



Document Identifier: DSP0222

Date: 2022-4-27

Version: 1.2WIP90

Network Controller Sideband Interface (NC-SI) Specification

Information for Work-in-Progress version:

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 This document is available for public review and comment until superseded.

Provide any comments through the DMTF Feedback Portal:

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

Supersedes: 1.1.1

Document Class: Work in Progress

Document Status: Published

Document Language: en-US

13

14 Copyright Notice

15 Copyright © 2009, 2013, 2022 DMTF. All rights reserved.

16 DMTF is a not-for-profit association of industry members dedicated to promoting enterprise and systems
17 management and interoperability. Members and non-members may reproduce DMTF specifications and
18 documents, provided that correct attribution is given. As DMTF specifications may be revised from time to
19 time, the particular version and release date should always be noted.

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30 withdrawn or modified after publication, and shall be indemnified and held harmless by any party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

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

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

37

CONTENTS

38	1	SCOPE	17
39	2	NORMATIVE REFERENCES	17
40	3	TERMS AND DEFINITIONS	18
41	3.1	WORDING INTERPRETATION.....	18
42	3.2	REQUIREMENT TERM DEFINITIONS	18
43	3.3	NC-SI TERM DEFINITIONS	20
44	3.4	NUMBERS AND NUMBER BASES.....	23
45	3.5	NETWORK ADDRESSES	23
46		NETWORK ADDRESSES IN THIS SPECIFICATION ARE WRITTEN AS FOLLOWS:	23
47	3.6	RESERVED FIELDS	23
48	4	ACRONYMS AND ABBREVIATIONS	23
49	5	NC-SI OVERVIEW	25
50	5.1	GENERAL.....	25
51	5.2	DEFINED TOPOLOGIES	27
52	5.3	SINGLE AND INTEGRATED NETWORK CONTROLLER IMPLEMENTATIONS.....	28
53	5.4	TRANSPORT STACK	30
54	5.5	TRANSPORT PROTOCOL	31
55	5.6	BYTE AND BIT ORDERING FOR TRANSMISSION	31
56	6	OPERATIONAL BEHAVIORS	31
57	6.1	TYPICAL OPERATIONAL MODEL	31
58		STATE DEFINITIONS -	32
59	6.1.1	<i>Defined States</i>	32
60	6.1.2	<i>NC-SI RBT pre-operational states</i>	33
61	6.1.3	<i>Package Ready state</i>	33
62	6.1.4	<i>Initial State</i>	34
63	6.1.5	<i>NC-SI Initial State recovery</i>	34
64	6.1.6	<i>State transition diagram</i>	35
65	6.1.7	<i>State diagram for NC-SI operation with hardware arbitration</i>	37
66	6.1.8	<i>Resets</i>	38
67	6.1.9	<i>Network Controller Channel ID</i>	38
68	6.1.10	<i>Configuration-related settings</i>	39
69	6.1.11	<i>Transmitting Pass-through packets from the Management Controller</i>	40
70	6.1.12	<i>Receiving Pass-through packets for the Management Controller</i>	41
71	6.1.13	<i>Pass-through operation in multiple medium implementations</i>	41
72	6.1.14	<i>Startup sequence examples</i>	41
73	6.2	NC-SI TRAFFIC TYPES	46
74	6.2.1	<i>Overview</i>	46
75	6.2.2	<i>Command protocol</i>	46
76	6.3	LINK CONFIGURATION AND CONTROL	49
77	6.3.1	<i>Link Configuration</i>	49
78		<i>The Management Controller should make link configuration changes only when the host network driver is</i>	
79		<i>absent or non-operational</i>	49
80	6.3.2	<i>Link Status</i>	49
81	6.4	FRAME FILTERING FOR PASS-THROUGH MODE.....	49

82	6.4.1	<i>Overview</i>	49
83	6.4.2	<i>Multicast filtering</i>	49
84	6.4.3	<i>Broadcast filtering</i>	49
85	6.4.4	<i>VLAN filtering</i>	49
86	6.5	OUTPUT BUFFERING BEHAVIOR	51
87	6.6	NC-SI FLOW CONTROL	51
88	6.7	ASYNCHRONOUS EVENT NOTIFICATION	51
89	6.8	AEN HANDLING IN MULTIPLE MEDIUM IMPLEMENTATIONS	52
90	6.9	ERROR HANDLING	52
91	6.9.1	<i>Overview</i>	52
92	6.9.2	<i>Transport errors</i>	52
93	6.9.3	<i>Missing responses</i>	53
94	6.9.4	<i>Detecting Pass-through traffic interruption</i>	53
95	6.10	SUPPORT FOR ADDITIONAL NETWORK FABRICS	54
96	6.10.1	<i>FC support</i>	54
97	6.10.2	<i>InfiniBand Support</i>	54
98	6.11	PLDM AND SPDM TRANSPORT	54
99	7	ARBITRATION IN CONFIGURATIONS WITH MULTIPLE NETWORK CONTROLLER PACKAGES	56
100	7.1	OVERVIEW	56
101	7.2	MULTI-CONTROLLER RBT	57
102	7.3	HARDWARE ARBITRATION	57
103	7.3.1	<i>General</i>	58
104	7.3.2	<i>Hardware arbitration opcodes</i>	59
105	7.3.3	<i>Opcode operations</i>	60
106	7.3.4	<i>Bypass mode</i>	62
107	7.3.5	<i>Hardware arbitration startup</i>	62
108	7.3.6	<i>ARB_MSTR assignment</i>	62
109	7.3.7	<i>Token timeout mechanism</i>	63
110	7.3.8	<i>Timing considerations</i>	63
111	7.3.9	<i>Example hardware arbitration state machine</i>	65
112	7.4	COMMAND-BASED ARBITRATION	67
113	8	PACKET DEFINITIONS	67
114	8.1	NC-SI PACKET ENCAPSULATION	67
115	8.1.1	<i>Ethernet frame header</i>	68
116	8.1.2	<i>Frame Check Sequence</i>	69
117	8.1.3	<i>Data length</i>	69
118	8.2	CONTROL PACKET DATA STRUCTURE	69
119	8.2.1	<i>Control Packet header</i>	69
120	8.2.2	<i>Control Packet payload</i>	70
121	8.2.3	<i>Command packet payload</i>	72
122	8.2.4	<i>Response packet payload</i>	72
123	8.2.5	<i>Response codes and reason codes</i>	73
124	8.2.6	<i>AEN packet format</i>	75
125	8.2.7	<i>Single OEM AEN packet format</i>	76
126	8.2.8	<i>Multiple OEMs AEN packet format</i>	76
127	8.3	CONTROL PACKET TYPE DEFINITIONS	77
128	8.4	COMMAND AND RESPONSE PACKET FORMATS	85
129	8.4.1	<i>NC-SI command frame format</i>	85
130	8.4.2	<i>NC-SI response packet format</i>	86
131	8.4.3	<i>Clear Initial State command (0x00)</i>	86
132	8.4.4	<i>Clear Initial State response (0x80)</i>	87

133	8.4.5	<i>Select Package command (0x01)</i>	87
134	8.4.6	<i>Select Package response (0x81)</i>	89
135	8.4.7	<i>Deselect Package command (0x02)</i>	89
136	8.4.8	<i>Deselect Package response (0x82)</i>	90
137	8.4.9	<i>Enable Channel command (0x03)</i>	90
138	8.4.10	<i>Enable Channel response (0x83)</i>	90
139	8.4.11	<i>Disable Channel command (0x04)</i>	91
140	8.4.12	<i>Disable Channel response (0x84)</i>	91
141	8.4.13	<i>Reset Channel command (0x05)</i>	92
142	8.4.14	<i>Reset Channel response (0x85)</i>	92
143	8.4.15	<i>Enable Channel Network TX command (0x06)</i>	92
144	8.4.16	<i>Enable Channel Network TX response (0x86)</i>	93
145	8.4.17	<i>Disable Channel Network TX command (0x07)</i>	93
146	8.4.18	<i>Disable Channel Network TX response (0x87)</i>	94
147	8.4.19	<i>AEN Enable command (0x08)</i>	94
148	8.4.20	<i>AEN Enable response (0x88)</i>	95
149	8.4.21	<i>Set Link command (0x09)</i>	96
150	8.4.22	<i>Set Link Response (0x89)</i>	99
151	8.4.23	<i>Get Link Status command (0x0A)</i>	100
152	8.4.24	<i>Get Link Status response (0x8A)</i>	100
153	8.4.25	<i>Set VLAN Filter command (0x0B)</i>	105
154	8.4.26	<i>Set VLAN Filter response (0x8B)</i>	107
155	8.4.27	<i>Enable VLAN command (0x0C)</i>	107
156	8.4.28	<i>Enable VLAN response (0x8C)</i>	108
157	8.4.29	<i>Disable VLAN command (0x0D)</i>	108
158	8.4.30	<i>Disable VLAN response (0x8D)</i>	109
159	8.4.31	<i>Set MAC Address command (0x0E)</i>	109
160	8.4.32	<i>Set MAC Address response (0x8E)</i>	111
161	8.4.33	<i>Enable Broadcast Filter command (0x10)</i>	111
162	8.4.34	<i>Enable Broadcast Filter response (0x90)</i>	113
163	8.4.35	<i>Disable Broadcast Filter command (0x11)</i>	114
164	8.4.36	<i>Disable Broadcast Filter response (0x91)</i>	114
165	8.4.37	<i>Enable Global Multicast Filter command (0x12)</i>	114
166	8.4.38	<i>Enable Global Multicast Filter response (0x92)</i>	119
167	8.4.39	<i>Disable Global Multicast Filter command (0x13)</i>	119
168	8.4.40	<i>Disable Global Multicast Filter response (0x93)</i>	119
169	8.4.41	<i>Set NC-SI Flow Control command (0x14)</i>	120
170	8.4.42	<i>Set NC-SI Flow Control response (0x94)</i>	121
171	8.4.43	<i>Get Version ID command (0x15)</i>	121
172	8.4.44	<i>Get Version ID Response (0x95)</i>	122
173	8.4.45	<i>Get Capabilities command (0x16)</i>	124
174	8.4.46	<i>Get Capabilities response (0x96)</i>	124
175	8.4.47	<i>Get Parameters command (0x17)</i>	127
176	8.4.48	<i>Get Parameters response (0x97)</i>	127
177	8.4.49	<i>Get Controller Packet Statistics command (0x18)</i>	130
178	8.4.50	<i>Get Controller Packet Statistics response (0x98)</i>	130
179	8.4.51	<i>Get NC-SI Statistics command (0x19)</i>	135
180	8.4.52	<i>Get NC-SI Statistics response (0x99)</i>	135
181	8.4.53	<i>Get NC-SI Pass-through Statistics command (0x1A)</i>	136
182	8.4.54	<i>Get NC-SI Pass-through Statistics response (0x9A)</i>	137
183	8.4.55	<i>Get Package Status command (0x1B)</i>	138

184	8.4.56	Get Package Status response (0x9B)	139
185	8.4.57	Get NC Capabilities and Settings command (0x25)	139
186	8.4.58	Get NC Capabilities and Settings response (0xA5)	140
187	8.4.59	Set NC Configuration command (0x26)	142
188	8.4.60	Set NC Configuration response (0xA6)	143
189	8.4.61	Get PF Assignment command (0x27)	143
190	8.4.62	Get PF Assignment Response (0xA7)	144
191	8.4.63	Set PF Assignment command (0x28)	147
192	8.4.64	Set PF Assignment Response (0xA8)	149
193	8.4.65	Get Port Configuration command (0x29)	150
194	8.4.66	Get Port Configuration response (0xA9)	150
195	8.4.67	Set Port Configuration command (0x2A)	151
196	8.4.68	Set Port Configuration response (0xAA)	153
197	8.4.69	Get Partition Configuration command (0x2B)	153
198	8.4.70	Get Partition Configuration response (0xAB)	154
199	8.4.71	Set Partition Configuration command (0x2C)	158
200	8.4.72	Set Partition Configuration response (0xAC)	161
201	8.4.73	Get Boot Config Command (0x2D)	161
202	8.4.74	Get Boot Config Response (0xAD)	162
203	8.4.75	Set Boot Config command (0x2E)	167
204	8.4.76	Set Boot Config Response (0xAE)	168
205	8.4.77	Get Partition Statistics command (0x2F)	169
206	8.4.78	Get Partition Statistics response for Ethernet (0xAF)	170
207	8.4.79	Get Partition Statistics response for FCoE (0xAF)	173
208	8.4.80	Get Partition Statistics response for iSCSI (0xAF)	174
209	8.4.81	Get Partition Statistics response for InfiniBand (0xAF)	176
210	8.4.82	Get Partition Statistics response for RDMA (0xAF)	178
211	8.4.83	Get Partition Statistics Response for Fibre Channel (0xAF)	180
212	8.4.84	Get FC Link Status command (0x31)	181
213	8.4.85	Get FC Link Status Response (0xB1)	182
214	8.4.86	Get Transceiver Management Data command (0x32)	185
215	8.4.87	Get Transceiver Management Data response (0xB2)	186
216	8.4.88	Get InfiniBand Link Status command (0x38)	187
217	8.4.89	Get InfiniBand Link Status Response (0xB8)	187
218	8.4.90	Get IB Statistics command (0x39)	190
219	8.4.91	Get IB Statistics Response (0xB9)	190
220	8.4.92	Settings Commit command (0x47)	192
221	8.4.93	Settings Commit response (0xC7)	193
222	8.4.94	Get ASIC Temperature (0x48)	193
223	8.4.95	Get ASIC Temperature Response (0xC8)	194
224	8.4.96	Get Ambient Temperature (0x49)	194
225	8.4.97	Get Ambient Temperature Response (0xC9)	195
226	8.4.98	Get Transceiver Temperature (0x4A)	195
227	8.4.99	Get Transceiver Temperature Response (0xCA)	195
228	8.4.100	Thermal Shutdown Control Command (0x4B)	196
229	8.4.101	Thermal Shutdown Control Response (0xCB)	197
230	8.4.102	Get Inventory Information command (0x4E)	198
231	8.4.103	Get Inventory Information response (0xCE)	198
232	8.5	SET PASS-THROUGH MODE CONTROL COMMAND (0x4F)	199
233	8.5.1	Pass-through Type Field	199
234	8.6	SET PASS-THROUGH MODE CONTROL RESPONSE (0xCF)	200
235	8.7	GET PASS-THROUGH MODE CONTROL COMMAND (0x50)	200

236	8.7.1	<i>Pass-through Type Field</i>	<i>Error! Bookmark not defined.</i>
237	8.8	GET PASS-THROUGH MODE CONTROL RESPONSE (0xD0)	201
238	8.8.1	<i>Feature Type Field</i>	<i>Error! Bookmark not defined.</i>
239	8.8.2	<i>Transmit Data to NC command (0x4C)</i>	202
240	8.8.3	<i>Transmit Data to NC response (0xCC)</i>	204
241	8.8.4	<i>Receive Data from NC command (0x4D)</i>	204
242	8.8.5	<i>Receive Data from NC response (0xCD)</i>	206
243	8.8.6	<i>Transfer SPDMM command (0x60)</i>	207
244	8.8.7	<i>Transfer SPDMM Response (0xE0)</i>	208
245	8.8.8	<i>Query Pending NC SPDMM Request (0x61)</i>	208
246	8.8.9	<i>Query Pending NC SPDMM Request Response (0xE1)</i>	208
247	8.8.10	<i>Send NC SPDMM Reply (0x62)</i>	209
248	8.8.11	<i>Send NC SPDMM Reply Response (0xE2)</i>	209
249	8.8.12	<i>Query and Set OEM AEN command (0x4E)</i>	210
250	8.8.13	<i>Query and Set OEM AEN Response (0xCE)</i>	211
251	8.8.14	<i>OEM command (0x50)</i>	212
252	8.8.15	<i>OEM response (0xD0)</i>	213
253	8.8.16	<i>PLDM Request (0x51)</i>	213
254	8.8.17	<i>PLDM Response (0xD1)</i>	214
255	8.8.18	<i>Query Pending NC PLDM Request (0x56)</i>	214
256	8.8.19	<i>Query Pending NC PLDM Request Response (0xD6)</i>	215
257	8.8.20	<i>Send NC PLDM Reply (0x57)</i>	215
258	8.8.21	<i>Send NC PLDM Reply Response (0xD7)</i>	216
259	8.8.22	<i>Transport-specific AEN Enable command (0x55)</i>	216
260	8.8.23	<i>Transport-specific AENs Enable Response (0xD5)</i>	217
261	8.8.24	<i>Get MC MAC Address command (0x58)</i>	218
262	8.8.25	<i>Get MC MAC Address response (0xD8)</i>	218
263	8.8.26	<i>Get Package UUID command (0x52)</i>	219
264	8.8.27	<i>Get Package UUID response (0xD2)</i>	220
265	8.9	AEN PACKET FORMATS	221
266	8.9.1	<i>Link Status Change AEN</i>	221
267	8.9.2	<i>Configuration Required AEN</i>	221
268	8.9.3	<i>Host Network Controller Driver Status Change AEN</i>	222
269	8.9.4	<i>Delayed Response Ready AEN</i>	222
270	8.9.5	<i>InfiniBand Link Status Change AEN</i>	223
271	8.9.6	<i>Fibre Channel Link Status Change AEN</i>	223
272	8.9.7	<i>Transceiver Event AEN</i>	224
273	8.9.8	<i>Request Data Transfer AEN</i>	226
274	8.9.9	<i>Partition Link Status Change AEN</i>	226
275	8.9.10	<i>Thermal Shutdown Event AEN</i>	227
276	8.9.11	<i>Pending PLDM Request AEN</i>	227
277	8.9.12	<i>Pending SPDMM Request AEN</i>	228
278	9	PACKET-BASED AND OPCODE TIMING.....	229
279	10	RBT ELECTRICAL SPECIFICATION	231
280	10.1	TOPOLOGIES	231
281	10.2	ELECTRICAL AND SIGNAL CHARACTERISTICS AND REQUIREMENTS	232
282	10.2.1	<i>Companion specifications</i>	232
283	10.2.2	<i>Full-duplex operation</i>	232
284	10.2.3	<i>Signals</i>	232
285	10.2.4	<i>High-impedance control</i>	233

286	10.2.5	<i>DC characteristics</i>	233
287	10.2.6	<i>AC characteristics</i>	235
288	10.2.7	<i>Interface power-up</i>	238
289	10.2.8	<i>REF_CLK startup</i>	239
290	10.3	RBT IMPLEMENTATION GUIDANCE	239
291	ANNEX A (NORMATIVE) EXTENDING THE MODEL		240
292	ANNEX B (INFORMATIVE) RELATIONSHIP TO RMII SPECIFICATION.....		241
293	ANNEX C		243
294	(INFORMATIVE)	CHANGE LOG	243
295			

296 **Figures**

297	Figure 1 – NC-SI functional block diagram	26
298	Figure 2 – NC-SI RBT traffic flow diagram.....	27
299	Figure 3 – Example topologies supported by the NC-SI.....	28
300	Figure 4 – Network Controller integration options.....	29
301	Figure 5 – NC-SI transport stack	31
302	Figure 6 – NC-SI package/channel operational state diagram	36
303	Figure 7 – NC-SI operational state diagram for hardware arbitration operation	37
304	Figure 8 – MC steps when the MC does not have prior knowledge of hardware arbitration	45
305	Figure 9 – NC-SI packet filtering flowchart	50
306	Figure 10 – Basic multi-drop block diagram.....	57
307	Figure 11 – Multiple Network Controllers in a ring format.....	58
308	Figure 12 – Opcode to RXD relationship	60
309	Figure 13 – Example TOKEN to transmit relationship	64
310	Figure 14 – Hardware arbitration state machine	65
311	Figure 15 – Ethernet frame encapsulation of NC-SI packet data without VLAN tag.....	68
312	Figure 16 – Example NC-SI RBT signal interconnect topology	231
313	Figure 17 – DC measurements	234
314	Figure 18 – AC measurements	235
315	Figure 19 – Overshoot measurement	237
316	Figure 20 – Undershoot measurement	238
317		

318 **Tables**

319	Table 1 – NC-SI operating state descriptions	32
320	Table 3 – Channel Ready state configuration settings	40
321	Table 4 – Hardware arbitration di-bit encoding	59
322	Table 5 – Hardware arbitration opcode format	59
323	Table 6 – Hardware arbitration states.....	66
324	Table 7 – Hardware arbitration events.....	67
325	Table 8 – Ethernet Header Format	68
326	Table 9 – Control Packet header format	69
327	Table 10 – Generic example of Control Packet payload.....	71
328	Table 11 – Generic example of Response packet payload format	72
329	Table 12 – Generic example of Delayed Response packet payload	73
330	Table 13 – Reason code ranges	73
331	Table 14 – Standard response code values	74
332	Table 15 – Standard Reason Code Values	74
333	Table 16 – AEN packet format.....	75
334	Table 17 – AEN Type Ranges	76
335	Table 18 – OEM AEN packet format.....	76
336	Table 19 – Multiple OEMs AEN packet format	76
337	Table 20 – Example of complete minimum-sized NC-SI command packet.....	85
338	Table 21 – Example of complete minimum-sized NC-SI response packet.....	86
339	Table 22 – Clear Initial State command packet format	87

340	Table 23 – Clear Initial State response packet format	87
341	Table 24 – Select Package command packet format	88
342	Table 25 – Features Control byte	88
343	Table 26 – Select package response packet format.....	89
344	Table 27 – Deselect Package command packet format	89
345	Table 28 – Deselect Package response packet format	90
346	Table 29 – Enable Channel command packet format.....	90
347	Table 30 – Enable Channel response packet format.....	90
348	Table 31 – Disable Channel command packet format	91
349	Table 32 – Disable Channel response packet format.....	91
350	Table 33 – Reset Channel command packet format.....	92
351	Table 34 – Reset Channel response packet format.....	92
352	Table 35 – Enable Channel Network TX command packet format	93
353	Table 36 – Enable Channel Network TX response packet format	93
354	Table 37 – Disable Channel Network TX command packet format	93
355	Table 38 – Disable Channel Network TX response packet format	94
356	Table 39 – AEN Enable command packet format.....	94
357	Table 40 – Format of AEN control	95
358	Table 41 – AEN Enable response packet format.....	96
359	Table 42 – Set Link command packet format	96
360	Table 43 – Set Link bit definitions	96
361	Table 44 – OEM Set Link bit definitions	99
362	Table 45 – Set Link response packet format	99
363	Table 46 – Set Link command-specific reason codes	99
364	Table 47 – Get Link Status command packet format.....	100
365	Table 48 – Get Link Status response packet format.....	100
366	Table 49 – Link Status field bit definitions.....	101
367	Table 50 – Other Indications field bit definitions	105
368	Table 51 – OEM Link Status field bit definitions (optional)	105
369	Table 52 – Get Link Status command-specific reason code.....	105
370	Table 53 – IEEE 802.1q VLAN Fields.....	106
371	Table 54 – Set VLAN Filter command packet format.....	106
372	Table 55 – Possible Settings for Filter Selector field (8-bit field)	106
373	Table 56 – Possible Settings for Enable (E) field (1-bit field).....	107
374	Table 57 – Set VLAN Filter response packet format.....	107
375	Table 58 – Set VLAN Filter command-specific reason code	107
376	Table 59 – Enable VLAN command packet format.....	107
377	Table 60 – VLAN Enable modes.....	108
378	Table 61 – Enable VLAN response packet format.....	108
379	Table 62 – Disable VLAN command packet format	109
380	Table 63 – Disable VLAN response packet format	109
381	Table 64 – Set MAC Address command packet format.....	110
382	Table 65 – Possible settings for MAC Address Number (8-bit field).....	110
383	Table 66 – Possible settings for Address Type (3-bit field).....	111
384	Table 67 – Possible settings for Enable Field (1-bit field).....	111
385	Table 68 – Set MAC Address response packet format.....	111
386	Table 69 – Set MAC Address command-specific reason code.....	111
387	Table 70 – Enable Broadcast Filter command packet format	112

388	Table 71 – Broadcast Packet Filter Settings field	112
389	Table 72 – Enable Broadcast Filter response packet format	113
390	Table 73 – Disable Broadcast Filter command packet format	114
391	Table 74 – Disable Broadcast Filter response packet format	114
392	Table 75 – Enable Global Multicast Filter command packet format.....	115
393	Table 76 – Bit Definitions for Multicast Packet Filter Settings field	115
394	Table 77 – Enable Global Multicast Filter response packet format.....	119
395	Table 78 – Disable Global Multicast Filter command packet format.....	119
396	Table 79 – Disable Global Multicast Filter response packet format.....	120
397	Table 80 – Set NC-SI Flow Control command packet format	120
398	Table 81 – Values for the Flow Control Enable field (8-bit field).....	120
399	Table 82 – Set NC-SI Flow Control response packet format	121
400	Table 83 – Set NC-SI Flow Control command-specific reason code	121
401	Table 84 – Get Version ID command packet format.....	121
402	Table 85 – Get Version ID response packet format.....	122
403	Table 86 – Get Capabilities command packet format.....	124
404	Table 87 – Get Capabilities response packet format	124
405	Table 88 – Capabilities Flags bit definitions.....	125
406	Table 89 – VLAN Mode Support bit definitions	126
407	Table 90 – Get Parameters command packet format	127
408	Table 91 – Get Parameters response packet format	128
409	Table 92 – Get Parameters data definition	128
410	Table 93 – MAC Address Flags bit definitions	129
411	Table 94 – VLAN Tag Flags bit definitions.....	129
412	Table 95 – Configuration Flags bit definitions	130
413	Table 96 – Get Controller Packet Statistics command packet format.....	130
414	Table 97 – Get Controller Packet Statistics response packet format.....	131
415	Table 98 – Get Controller Packet Statistics counters	132
416	Table 99 – Counters Cleared from Last Read Fields format	134
417	Table 100 – Get NC-SI Statistics command packet format	135
418	Table 101 – Get NC-SI Statistics response packet format	135
419	Table 102 – Get NC-SI Statistics counters	136
420	Table 103 – Get NC-SI Pass-through Statistics command packet format	136
421	Table 104 – Get NC-SI Pass-through Statistics response packet format	137
422	Table 105 – Get NC-SI Pass-through Statistics counters	137
423	Table 106 – Get Package Status packet format	138
424	Table 107 – Get Package Status response packet format	139
425	Table 108 – Package Status field bit definitions	139
426	Table 109 – Get NC Capabilities and Settings command packet format.....	139
427	Table 111 – Fabrics field bit definitions.....	141
428	Table 112 – Enabled Fabrics field bit definitions	141
429	Table 113 – Capabilities Flags bit definitions.....	142
430	Table 114 – Set NC Configuration command packet format	142
431	Table 115 – Set NC Configuration response packet format	143
432	Table 116 – Get PF Assignment Command Packet Format.....	144
433	Table 117 – Get PF Assignment Response packet format.....	144
434	Table 118 – Channel c Function Assignment bitmap field	145
435	Table 119 – Function Port Association bitmap field	145

436	Table 120 – Function Enablement bitmap field.....	146
437	Table 121 – PCI Bus b Assignment bitmap field.....	146
438	Table 122 – Set PF Assignment Command packet format.....	148
439	Table 123 – Channel Function Assignment bitmap field.....	148
440	Table 124 – Function Enablement bitmap field.....	149
441	Table 125 – PCI Bus Assignment bitmap field.....	149
442	Table 126 – Set PF Assignment Response packet format	149
443	Table 127 – Get Port Configuration command packet format.....	150
444	Table 128 – Get Port Configuration response packet format.....	150
445	Table 129 – Fabric Type bit definitions	150
446	Table 130 – Media Type bit definitions	151
447	Table 131 – bits field definitions.....	Error! Bookmark not defined.
448	Table 132 – Set Port Configuration command packet format	152
449	Table 133 – Fabric Type bit definitions	152
450	Table 136 – Set Port Configuration response packet format	153
451	Table 138 – Get Partition Configuration command packet format	153
452	Table 139 – Get Partition Configuration response packet format	154
453	Table 140 – Personality Cfg bit definitions.....	154
454	Table 141 – Personality Spt bit definitions.....	155
455	Table 142 – Configuration Flags bit definitions	155
456	Table 143 – Address Type-Length Field Bit Definitions	158
457	Table 144 – Set Partition Configuration command packet format	159
458	Table 145 – Personality Cfg bit definitions.....	159
459	Table 146 – Values for the Config flags field (8-bit field)	160
460	Table 148 – Address Type-Length field bit definitions	160
461	Table 149 – Set Partition Configuration response packet format	161
462	Table 151 – Get Boot Config command packet.....	161
463	Table 152 – Protocol Type field	162
464	Table 153 – Get Boot Config Response packet.....	162
465	Table 154 – Protocol Type field	163
466	Table 155 – PXE Boot Protocol Type-Length field	163
467	Table 156 – Get FC Boot Protocol Type-Length field	164
468	Table 157 – FCoE Boot Protocol Type-Length field	164
469	Table 158 – iSCSI Boot Protocol Type-Length field	165
470	Table 156 – Get NVMeoFC Boot Protocol Type-Length field	165
471	Table 159 – Set Boot Config command packet format	168
472	Table 160 – Set Boot Config Response packet format.....	168
473	Table 161 – TLV Error Reporting field	169
474	Table 164– Get Partition Statistics (Ethernet) response packet format	171
475	Table 166 – Counters Cleared from Last Read field format	172
476	Table 167 – Get Partition Statistics (FCoE) response packet format	173
477	Table 169 – Counters Cleared from Last Read field format	174
478	Table 170 – Get Partition Statistics (iSCSI) response packet format	175
479	Table 172 – Counters Cleared from Last Read field format	176
480	Table 173 – Get Partition Statistics (IB) response packet format	176
481	Table 175 – Counters Cleared from Last Read field format	177
482	Table 176 – Get Partition Statistics (RDMA) response packet format	178
483	Table 177 – Counter Sizes field format.....	179

484	Table 178 – Counters Cleared from Last Read field format	179
485	Table 179 – Get Partition Statistics (FC) Response packet.....	180
486	Table 180 – Counters Cleared from Last Read field format	180
487	Table 181 – FC Statistics.....	181
488	Table 182 – Get FC Link Status command packet format	182
489	Table 183 – Get FC Link Status Response packet format.....	182
490	Table 182 – FC Trunk Status field bit definitions	183
491	Table 183 – FC Link Status field bit definitions	183
492	Table 184 – Trunk Speeds field	184
493	Table 185 – FC Link Speed field.....	184
494	Table 186 – Get Transceiver Management Data command packet format	185
495	. Table 183 – Flag field bit definitions	186
496	Table 187 – Get Transceiver Management Data response packet format	186
497	Table 182 – Get InfiniBand Link Status command.....	187
498	Table 183 – Get InfiniBand Link Status Response packet.....	188
499	Table 184 – InfiniBand Link Status definitions	188
500	Table 185 – Get IB Statistics Command.....	190
501	Table 186 – Get IB Statistics Response packet.....	191
502	Table 187 – IB Statistics Counter definitions	191
503	Table 225 – Settings Commit command packet format	192
504	Table 226 – Settings Commit response packet format	193
505	Table 188 – Get ASIC Temperature Command packet	193
506	Table 189 – Get ASIC Temperature Response packet	194
507	Table 190 – Get Ambient Temperature command packet	194
508	Table 191 – Get Ambient Temperature Response packet.....	195
509	Table 192 – Get Transceiver Temperature Command Packet	195
510	Table 193 – Get Transceiver Temperature Response packet.....	196
511	Table 188 – Thermal Shutdown Control Command packet	196
512	Table 178 – Command field bit definitions	197
513	Table 189 – Thermal Shutdown Control Response packet	197
514	Table 178 – Status field bit definitions	197
515	Table 222 – Get Inventory Information command packet format.....	198
516	Table 223 – Get Inventory Information response packet format.....	198
517	Table 155 – Inventory Information Type-Length field	199
518	Table 58 – Pass-through Type definitions	200
519	Table 62 – Pass-through Type definitions	Error! Bookmark not defined.
520	Table 65 – Feature Type definitions	Error! Bookmark not defined.
521	Table 202 – Transmit Data to NC command packet format.....	203
522	Table 195 – Opcode field format.....	203
523	Table 196 – Transmit Data to NC response packet format.....	204
524	Table 197 – Transmit Data to NC command-specific reason codes.....	204
525	Table 198 – Receive Data from NC command packet format	205
526	Table 199 – Opcode field format.....	205
527	Table 200 illustrates the packet format of the Receive Data from NC command response.	206
528	Table 200 – Receive Data from NC response packet format	206
529	Table 195 – Opcode field format.....	206
530	Table 201 – Receive Data from NC command-specific reason codes	207
531	Table 202 – Transfer SPDM command packet.....	207

532	Table 203 – Transfer SPDm Response packet	208
533	Table 210 – Query Pending NC SPDm Request packet format	208
534	Table 211 – Query Pending NC SPDm Request Response Packet Format	208
535	Table 212 – Query Pending NC SPDm Request Response parameters.....	209
536	Table 213 – Send NC SPDm Reply packet format	209
537	Table 214 –Send NC SPDm Reply Response packet format.....	210
538	Table 215 – Reply NC SPDm Response parameters.....	210
539	Table 204 – Query and Set OEM AEN command packet.....	211
540	Table 205 – Query and Set OEM AEN Response packet	211
541	Table 206 – OEM command packet format	212
542	Table 207 – OEM response packet format	213
543	Table 208 – PLDM Request packet format.....	213
544	Table 209 – PLDM Response packet format.....	214
545	Table 210 – Query Pending NC PLDM Request packet format.....	214
546	Table 211 – Query Pending NC PLDM Request Response Packet Format.....	215
547	Table 212 – Query Pending NC PLDM Request Response parameters	215
548	Table 213 – Send NC PLDM Reply packet format	215
549	Table 214 –Send NC PLDM Reply Response packet format	216
550	Table 215 – Reply NC PLDM Response parameters	216
551	Table 216 – Transport-specific AEN Enable command packet format	217
552	Table 217 – Transport-specific AEN enable field format	217
553	Table 218 – Transport-specific AEN Enable Response packet format.....	217
554	Table 220 – Get MC MAC Address command packet format.....	218
555	Table 221 – Get MC MAC Address response packet format.....	219
556	Table 222 – Get Package UUID command packet format.....	219
557	Table 223 – Get Package UUID response packet format.....	220
558	Table 224 – UUID Format.....	220
559	Table 229 – Link Status Change AEN packet format	221
560	Table 230 – Configuration Required AEN packet format.....	221
561	Table 231 – Host Network Controller Driver Status Change AEN packet format	222
562	Table 232 – Host Network Controller Driver Status format.....	222
563	Table 234 – InfiniBand Link Status Change AEN packet format	223
564	Table 235 – Fibre Channel Link Status Change AEN packet format.....	223
565	Table 236 – Transceiver Event AEN packet format.....	224
566	Table 237 – Transceiver Event List format	224
567	Table 237 – Transceiver Presence format.....	225
568	Table 238 – Request Data Transfer AEN packet format	226
569	Table 250 – Partition Link Status Change AEN packet format	226
570	Table 251 – Partition Map Field	226
571	Table 252 – Partition Link Status	227
572	Table 230 – Thermal Shutdown Event AEN packet format	227
573	Table 253 – Pending PLDM Request AEN format.....	228
574	Table 254 – Pending SPDm Request AEN format	228
575	Table 255 – NC-SI packet-based and opcode timing parameters	229
576	Table 240 – Physical RBT signals	232
577	Table 241 – DC specifications	234
578	Table 242 – AC specifications	235
579		

580

Foreword

581 The *Network Controller Sideband Interface (NC-SI) Specification* (DSP0222) was prepared by the PMCI
582 Working Group.

583 .

584 This version supersedes version 1.1.1. For a list of changes, see the Change Log in ANNEX C.

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

587 Acknowledgments

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

589 Editors:

- 590 • Hemal Shah – Broadcom Inc.
- 591 • Bob Stevens – Dell Technologies

592 Contributors:

- 593 • Patrick Caporale - Lenovo
- 594 • Phil Chidester – Dell Inc.
- 595 • Yuval Itkin – NVIDIA Corporation
- 596 • Ira Kalman – Intel Corporation
- 597 • Patrick Kutch – Intel Corporation
- 598 • Eliel Louzoun – Intel Corporation
- 599 • Rob Mapes – Marvell Corporation
- 600 • Edward Newman – Hewlett Packard Enterprise
- 601 • Patrick Schoeller – Intel Corporation
- 602 • Tom Slaight – Intel Corporation

603

604

Introduction

605 In out-of-band management environments, the interface between the out-of-band Management Controller
606 and the Network Controller is critical. This interface is responsible for supporting communication between
607 the Management Controller and external management applications. Currently there are multiple such
608 proprietary interfaces in the industry, leading to inconsistencies in implementation of out-of-band
609 management.

610 The goal of this specification is to define an interoperable sideband communication interface standard to
611 enable the exchange of management data between the Management Controller and Network Controller.
612 The Sideband Interface is intended to provide network access for the Management Controller, and the
613 Management Controller is expected to perform all the required network functions.

614 This specification defines the protocol and commands necessary for the operation of the sideband
615 communication interface. This specification also defines physical and electrical characteristics of a
616 sideband binding interface that is a variant of RMII targeted specifically for sideband communication
617 traffic.

618 The specification is primarily intended for architects and engineers involved in the development of
619 network interface components and Management Controllers that will be used in providing out-of-band
620 management.

Network Controller Sideband Interface (NC-SI) Specification

1 Scope

This specification defines the functionality and behavior of the Sideband Interface responsible for connecting the Network Controller (including Ethernet, Fibre Channel, and InfiniBand controllers) to the Management Controller. It also outlines the behavioral model of the (Ethernet) network traffic destined for the Management Controller from the Network Controller.

This specification defines the following two aspects of the Network Controller Sideband Interface (NC-SI):

- behavior of the interface, which include its operational states as well as the states of the associated components
- the payloads and commands of the communication protocol supported over the interface

The scope of this specification is limited to addressing only a single Management Controller communicating with one or more Network Controllers.

This specification also defines the following aspects of a 3.3V RMI-Based Transport (RBT) based physical medium:

- transport binding for NC-SI over RBT
- electrical and timing requirements for the RBT
- an optional hardware arbitration mechanism for RBT

Only the topics that may affect the behavior of the Network Controller or Management Controller, as it pertains to the Sideband Interface operations, are discussed in this specification.

2 Normative references

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

DMTF DSP0261, *NC-SI over MCTP Binding Specification 1.2*
<https://www.dmtf.org/dsp/DSP0261>

DMTF DSP0240, *Platform Level Data Model (PLDM) Base Specification 1.0*
<https://www.dmtf.org/dsp/DSP0240>

DMTF DSP0274, *Security Protocol and Data Model (SPDM) Specification*
<https://www.dmtf.org/dsp/DSP0274>

IEEE 802.3, *IEEE Standard for Ethernet*, June 2018,
<http://www.ieee.org/portal/site>

IETF, RFC4122, *A Universally Unique Identifier (UUID) URN Namespace*, July 2005
<http://datatracker.ietf.org/doc/rfc4122/>

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

Reduced Media Independent Interface (RMII) Consortium, *RMII Specification*, revision 1.2, March 20, 1998, http://ebook.pldworld.com/eBook/-Telecommunications,Networks-/TCPIP/RMII/rmii_rev12.pdf

InfiniBand™ Architecture Specification
<https://www.infinibandta.org/ibta-specification/>

Fibre Channel Technical Committee (ANSI/INCITS TC T11)
<http://www.t11.org> and <http://www.incits.org>

SFF, SFF-8024, SFF Cross Reference to Industry Products
<https://www.snia.org/technology-communities/sff/specifications>

SFF, SFF-8472, Diagnostic Monitoring Interface for Optical Transceivers
<https://www.snia.org/technology-communities/sff/specifications>

SFF, SFF-8436, QSFP+ 10Gbs 4X Pluggable Transceiver
<https://www.snia.org/technology-communities/sff/specifications>

SFF, SFF-8636, Management Interface for Cabled Environments
<https://www.snia.org/technology-communities/sff/specifications>

CMIS, Common Management Interface Specification 4.0
<http://www.qsfp-dd.com/wp-content/uploads/2019/05/QSFP-DD-CMIS-rev4p0.pdf>

3 Terms and definitions

3.1 Wording Interpretation

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

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

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

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

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

3.2 Requirement term definitions

This clause defines key phrases and words that denote requirement levels in this specification.

- 693 **3.1.1**
694 **can**
695 indicates an ability or capability expressed by the specification or of the possibility of some outcome in the
696 context of the specification
- 697 **3.1.2**
698 **cannot**
699 indicates the inability or denial of the possibility of a certain outcome in the context of the specification
- 700 **3.1.3**
701 **conditional**
702 indicates that an item is required under specified conditions
- 703 **3.1.4**
704 **deprecated**
705 indicates that an element or profile behavior has been outdated by newer constructs
- 706 **3.1.5**
707 **mandatory**
708 indicates that an item is required under all conditions
- 709 **3.1.6**
710 **may**
711 a permission expressed by this specification
- 712 **3.1.7**
713 **may not**
714 an expression of permission in the negative; a lack of requirement
- 715 **3.1.8**
716 **not recommended**
717 indicates that valid reasons may exist in particular circumstances when the particular behavior is
718 acceptable or even useful, but the full implications should be understood and carefully weighed before
719 implementing any behavior described with this label
- 720 **3.1.9**
721 **obsolete**
722 indicates that an item was defined in prior specifications but has been removed from this specification
- 723 **3.1.10**
724 **optional**
725 indicates that an item is not mandatory, conditional, or prohibited
- 726 **3.1.11**
727 **recommended**
728 indicates that valid reasons may exist in particular circumstances to ignore a particular item, but the full
729 implications should be understood and carefully weighed before choosing a different course
- 730 **3.1.12**
731 **required**
732 indicates that the item is an absolute requirement of the specification

3.1.13**shall**

indicates that the item is an absolute requirement of the specification

3.1.14**shall not**

indicates that the item is an absolute prohibition of the specification

3.1.15**should**

indicates a recommendation of the specification, but the full implications should be understood and carefully weighed before choosing a different course

3.1.16**should not**

indicates a recommendation against, but the full implications should be understood and carefully weighed before implementing any behavior described with this label

3.3 NC-SI term definitions

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

3.2.1**frame**

a data packet of fixed or variable length that has been encoded for digital transmission over a node-to-node link

Frame is used in references to [IEEE 802.3 Frames](#). *Packet* is used in all other references.

3.2.2**packet**

a formatted block of information carried by a computer network

Frame is used in references to [IEEE 802.3 Frames](#). *Packet* is used in all other references.

3.2.3**external network interface**

the interface of the Network Controller that provides connectivity to the external network infrastructure; also known as *port*

3.2.4**internal host interface**

the interface of the Network Controller that provides connectivity to the host operating system running on the platform

3.2.5**Management Controller**

an intelligent entity composed of hardware/firmware/software that resides within a platform and is responsible for some or all of the management functions associated with the platform; also known as BMC and Service Processor

771	3.2.6
772	Network Controller
773	the component within a system that is responsible for providing connectivity to an external Ethernet, Fibre
774	Channel, or InfiniBand network
775	3.2.7
776	remote media
777	a manageability feature that enables remote media devices to appear as if they are attached locally to the
778	host
779	3.2.8
780	Network Controller Sideband Interface
781	NC-SI
782	The RBT interface of the Network Controller that provides network connectivity to a Management
783	Controller; also shown as <i>Sideband Interface</i> , <i>RBT</i> or <i>NC-SI</i> as appropriate in the context
784	3.2.9
785	integrated controller
786	a Network Controller device that supports two or more channels for the NC-SI that share a common
787	NC-SI physical interface (for example, a Network Controller that has two or more physical network ports
788	and a single NC-SI bus connection)
789	3.2.10
790	multi-drop
791	refers to the situation in which multiple physical communication devices share an electrically common bus
792	and a single device acts as the master of the bus and communicates with multiple “slave” or “target”
793	devices
794	Related to NC-SI, a Management Controller serves the role of the master, and the Network Controllers
795	are the target devices
796	3.2.11
797	point-to-point
798	refers to the situation in which only a single Management Controller and single Network Controller
799	package are used on the bus in a master/slave relationship, where the Management Controller is the
800	master
801	3.2.12
802	Channel
803	refers to the logical representation of a network port in a Network Controller that supports Control traffic
804	and may support Pass-through traffic
805	A Network Controller may have a 1:1 relationship of NC-SI channels to physical network ports, or Network
806	Controllers that support partitioning can have multiple channels on a given network port
807	3.2.13
808	Partition
809	one or more NC-SI channels in a Network Controller that share a common network port

3.2.14**Package**

one or more NC-SI channels in a Network Controller that share a common set of electrical buffers and common electrical buffer controls for the NC-SI bus

Typically a single, logical NC-SI package exists for a single physical Network Controller package (chip or module). However, this specification allows a single physical chip or module to hold multiple NC-SI logical packages

3.2.15**control traffic****Control Packets****control packets**

command, response, and asynchronous event notification packets transmitted between the Management Controller and Network Controllers for the purpose of managing the NC and NC-SI

3.2.16**Command**

Control Packet sent by the Management Controller to the Network Controller to request the Network Controller to perform an action, and/or return data

3.2.17**Response**

Control Packet sent by the Network Controller to the Management Controller as a positive acknowledgement of a command received from the Management Controller, and to provide the execution outcome of the command, as well as to return any required data

3.2.18**Asynchronous Event Notification**

Control Packet sent by the Network Controller to the Management Controller as an explicit notification of the occurrence of an event of interest to the Management Controller

3.2.19**pass-through traffic****pass-through packets**

network packets passed between the external network and the Management Controller through the Network Controller

3.2.20**RBT****RMII-Based Transport**

Electrical and timing specification for a 3.3V-signaling physical medium that is derived from [RMII](#)

3.2.21**PCI Endpoint**

Also PCI Port, physically the collection of Transmitters and Receivers located on the same chip that define a Link, logically the interface between a component and a PCI Express Link. For the purposes of this specification, it is a PCIe upstream port on the NC that is assigned a PCI Bus number when connecting to a PCI Switch or Root Complex

PCI Link

The collection of two Ports and their interconnecting Lanes. A Link is a dual-simplex communications path between two components.

3.4 Numbers and number bases

Numbers in this specification are written as follows:

- Hexadecimal numbers are written with a “0x” prefix (for example, 0xFF and 0x80).
- Binary numbers are written with a lowercase “b” suffix (for example, 1001b and 10b).
- Hexadecimal and binary numbers are formatted in the `Courier New` font.
- Uint8 describes an unsigned 8-bit integer value.

3.5 Network Addresses

Network addresses in this specification are written as follows:

- IPv4 addresses are written as decimal numbers with period (.) separators
- IPv6 addresses are written as hexadecimal numbers with colon (:) separators
- MAC addresses are written as 6 hexadecimal number pairs with colon (:) separators
- InfiniBand GUIDs are written as hexadecimal numbers with no separators
- Fibre Channel WWNs are written as hexadecimal numbers with no separators

3.6 Reserved fields

Unless otherwise specified, reserved fields (bytes, bits, etc.) are reserved for future use and should be written as zeros and ignored when read.

4 Acronyms and abbreviations

The following symbols and abbreviations are used in this document.

4.1**AC**

alternating current

4.2**AEN**

Asynchronous Event Notification

4.3**BMC**

Baseboard Management Controller (often used interchangeably with MC)

4.4**CRC**

cyclic redundancy check

4.5**CRS_DV**

a physical NC-SI signal used to indicate Carrier Sense/Received Data Valid

887	4.6
888	DC
889	direct current
890	4.7
891	DHCP
892	Dynamic Host Configuration Protocol
893	4.8
894	EEE
895	Energy Efficient Ethernet
896	4.9
897	FC
898	Fibre Channel
899	4.10
900	FCS
901	Frame Check Sequence
902	4.11
903	IB
904	InfiniBand
905	4.12
906	MC
907	Management Controller
908	4.13
909	NC
910	Network Controller
911	4.14
912	NC-SI
913	Network Controller Sideband Interface
914	4.15
915	NC-SI RX
916	the direction of traffic on RBT from the Network Controller to the Management Controller
917	4.16
918	NC-SI TX
919	the direction of traffic RBT to the Network Controller from the Management Controller
920	4.17
921	RMII
922	Reduced Media Independent Interface
923	4.18
924	RX
925	Receive

4.19**RXD**

physical NC-SI signals used to transmit data from the Network Controller to the Management Controller

4.20**RX_ER**

a physical NC-SI signal used to indicate a Receive Error

4.21**SerDes**

serializer/deserializer; an integrated circuit (IC or chip) transceiver that converts parallel data to serial data and vice-versa. This is used to support interfaces such as 1000Base-X and others.

4.22**TX**

Transmit

4.23**TXD**

physical NC-SI signals used to transmit data from the Management Controller to the Network Controller

4.24**VLAN**

Virtual LAN

5 NC-SI overview**5.1 General**

With the increasing emphasis on out-of-band manageability and functionality, such as Remote Media (R-Media) and Remote Keyboard-Video-Mouse (R-KVM), the need for defining an industry standard Network Controller Sideband Interface (NC-SI) has become clear. This specification enables a common interface definition between different Management Controller and Network Controller vendors. This specification addresses not only the electrical and protocol specifications, but also the system-level behaviors for the Network Controller and the Management Controller related to the NC-SI.

The NC-SI is defined as the interface (protocol, messages, and medium) between a Management Controller and one or multiple Network Controllers. This interface, referred to as a Sideband Interface in Figure 1, is responsible for providing external network connectivity for the Management Controller while also allowing the external network interface to be shared with traffic to and from the host.

The specification of how the NC-SI protocol and messages are implemented over a particular physical medium is referred to as a transport binding. This document, DSP0222, includes the definition of the transport binding, electrical, framing, and timing specifications for a physical interface called RBT (RMII-based Transport). Electrically, RBT, as described in clause 0, is similar to the Reduced Media Independent Interface™ (RMII) – see ANNEX B. Transport bindings for NC-SI over other media and transport protocols are defined through external transport binding specifications, such as [DSP0261](#), the *NC-SI over MCTP Transport Binding Specification*. That specification defines the Get Supported Media command (0x54) which is used to discover if the NC supports operation over multiple media. This command may be issued on any NC-SI transport including RBT.

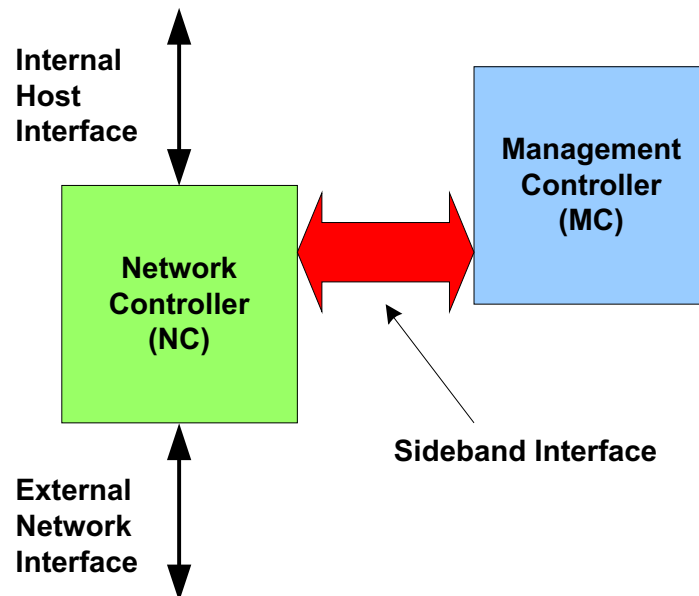


Figure 1 – NC-SI functional block diagram

NC-SI traffic flow is illustrated in Figure 2. Two classes of packet data can be delivered over the Sideband Interface:

- “Pass-through” packets that are transferred between the Management Controller and the external network
- “Control” packets that are transferred between the Management Controller and Network Controllers for control or configuration functionality. This specification defines NC-SI commands and responses as well as a mechanism to customize and extend functionality via OEM commands – see ANNEX A.

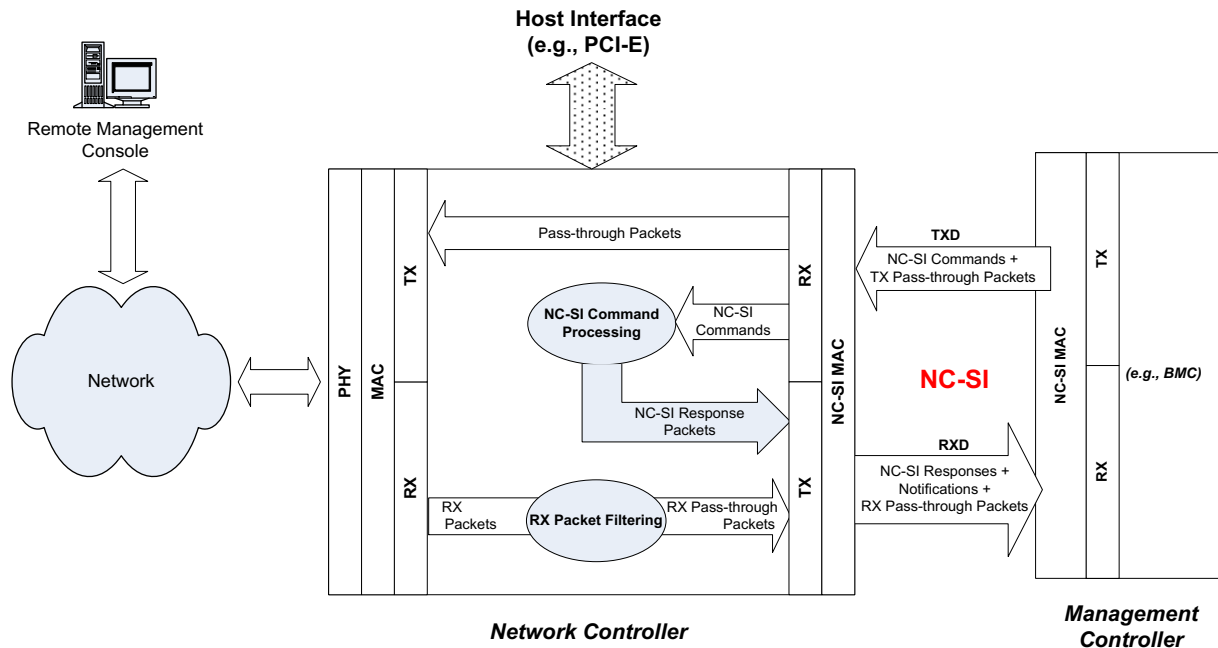


Figure 2 – NC-SI RBT traffic flow diagram



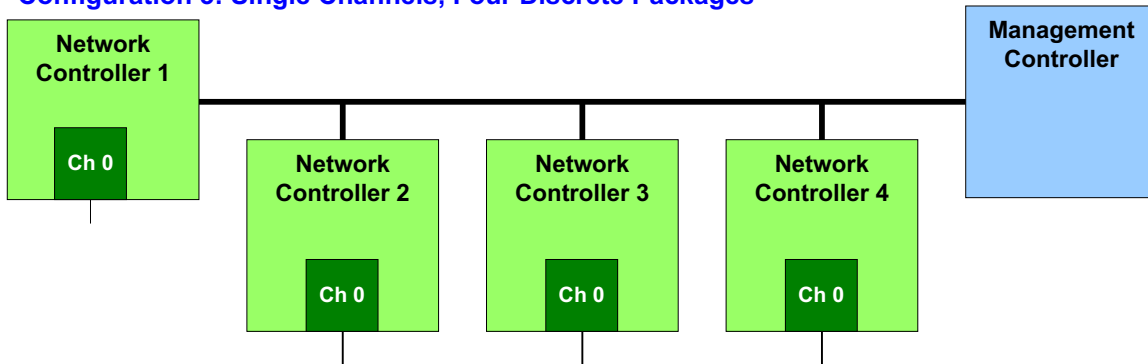
NC-SI is intended to operate independently from the in-band activities of the Network Controller. As such, the Sideband Interface is not specified to be accessible through the host interface of the Network Controller. From the external world, this interface should behave and operate like a standard Ethernet Interface.

5.2 Defined topologies

The topologies supported under this specification apply to the case in which a single Management Controller is actively communicating with one or more Network Controllers on the Sideband Interface over RBT. The RBT electrical specification is targeted to directly support up to four physical Network Controller packages. The protocol specification allows up to eight Network Controller packages, with up to 31 channels per package.

Figure 3 illustrates some examples of Network Controller configurations supported by the NC-SI in the current release:

- Configuration 1 shows a Management Controller connecting to a single Network Controller with a single external network connection.
- Configuration 2 shows a Management Controller connecting to a Network Controller package that supports two NC-SI channel connections.
- Configuration 3 shows a Management Controller connecting to four discrete Network Controllers.

Configuration 1: Single Channel, Single Package**Configuration 2: Integrated Dual Channel, Single Package****Configuration 3: Single Channels, Four Discrete Packages****Figure 3 – Example topologies supported by the NC-SI****5.3 Single and integrated Network Controller implementations**

This clause illustrates the general relationship between channels, packages, receive buffers, and bus buffers for different controller implementations.

An integrated controller is a Network Controller that connects to the NC-SI RBT (or other physical interfaces that support NC-SI) interface and provides NC-SI support for two or more network connections. A single controller is a controller that supports only a single NC-SI channel.

For the *NC-SI Specification*, an integrated controller can be logically implemented in one of three basic ways, as illustrated in Figure 4. Although only two channels are shown in the illustration, an integrated controller implementation can provide more than two channels. The example channel and package numbers (for example, channel 0, package 0) refer to the Internal Channel and Package ID subfields of the Channel ID. For more information, see 6.1.9.

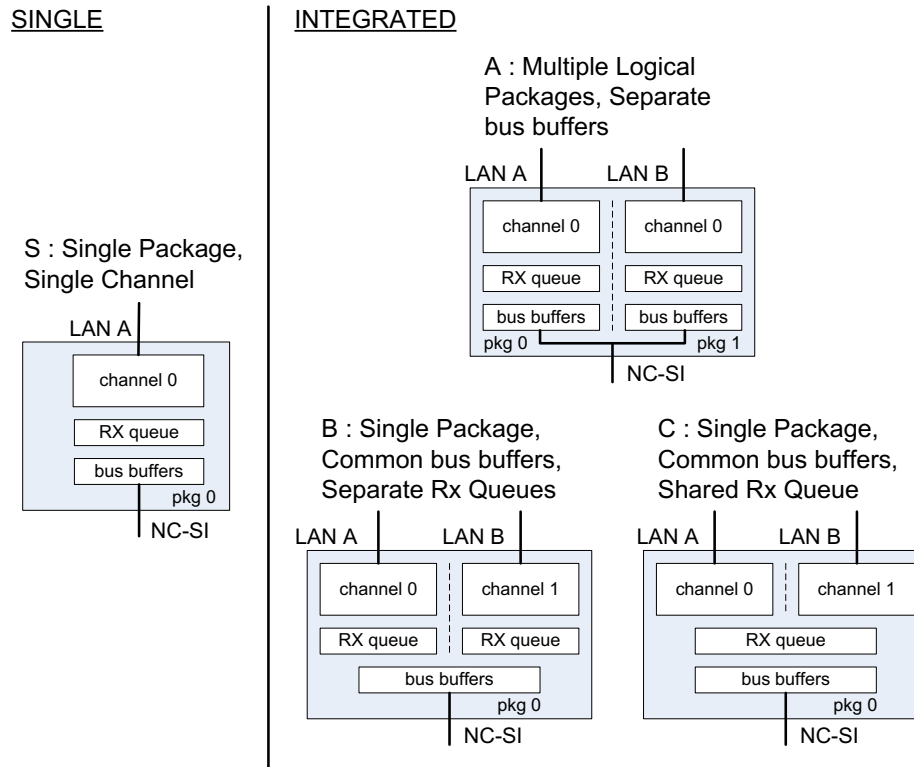


Figure 4 – Network Controller integration options

Packages that include multiple channels are required to handle internal arbitration between those channels and the Sideband Interface. The mechanism by which this occurs is vendor-specific and not specified in this document. This internal arbitration is always active by default. No NC-SI commands are defined for enabling or disabling internal arbitration between channels.

The following classifications refer to a logical definition. The different implementations are distinguished by their behavior with respect to the NC-SI bus and command operation. The actual physical and internal implementation can vary from the simple diagrams. For example, an implementation can act as if it has separate RX queues without having physically separated memory blocks for implementing those queues.

- **S: Single Package, Single Channel**

This implementation has a single NC-SI interface providing NC-SI support for a single LAN port, all contained within a package or module that has a single connection to the NC-SI physical bus. Note that FC Bonding is supported in this specification and thus multiple physical ports may be aggregated into one logical port.

- **A: Multiple Logical Packages, Separate Bus Buffers**

This implementation acts like two physically separate Network Controllers that happen to share a common overall physical container. Electrically, they behave as if they have separate electrical buffers connecting to the NC-SI bus. This behavior might be accomplished by means of a passive internal bus or by separate physical pins coming from the overall package. From the point of view of the Management Controller and the NC-SI command operation, this implementation behaves as if the logical controllers were implemented as physically separate controllers.

This type of implementation could include internal hardware arbitration between the two logical Network Controller packages. If hardware arbitration is provided external to the package, it shall meet the requirements for hardware arbitration described later in this specification. (For more information, see 7.3.)

- **B: Single Package, Common Bus Buffers, Separate RX Queues**

In this implementation, the two internal NC-SI channels share a common set of electrical bus buffers. A single Deselect Package command will deselect the entire package. The Channel Enable and Channel Disable commands to each channel control whether the channel can transmit Pass-through and AEN packets through the NC-SI interface. The Channel Enable command also determines whether the packets to be transmitted through the NC-SI interface will be queued up in an RX Queue for the channel while the channel is disabled or while the package is deselected. Because each channel has its own RX Queue, this queuing can be configured for each channel independently.

- **C: Single Package, Common Bus Buffers, Shared RX Queue**

This implementation is the same as described in the preceding implementation, except that the channels share a common RX Queue for holding Pass-through packets to be transmitted through the NC-SI interface. This queue could also queue up AEN or Response packets.

In addition to the general purpose architectures listed above, some Network Controllers support more advanced architectures that provide for multiple host interfaces that share a single channel/physical port, a single host interface that sends and receives traffic over multiple physical ports, but modeled as a single channel, and lastly an internally terminated channel that can be used to control some other functionality in the NC that requires a communication and control path to the MC.

5.4 Transport stack

The overall transport stack of the NC-SI is illustrated in Figure 5. The lowest level is the physical-level interface (for example, RBT), and the media-level interface is based on Ethernet. Above these interfaces are the two data-level protocols that are supported by the *NC-SI Specification*: NC-SI Command Protocol and the Network Data Protocol (for example, ARP, IP, DHCP, and NetBIOS) associated with Pass-through traffic for NCs supporting Ethernet. Both protocols are independent from binding to the underlying physical interface. This specification only defines the binding for NC-SI over RBT.

This document defines the necessary NC-SI command set and interface specification that allows the appropriate configuration of the Network Controller parameters and operation to enable network traffic to flow to and from external networks to the Management Controller for those devices that support it. As shown in Figure 5, the scope of the NC-SI Command Protocol is limited to the interface between the Network Controller and the Management Controller.

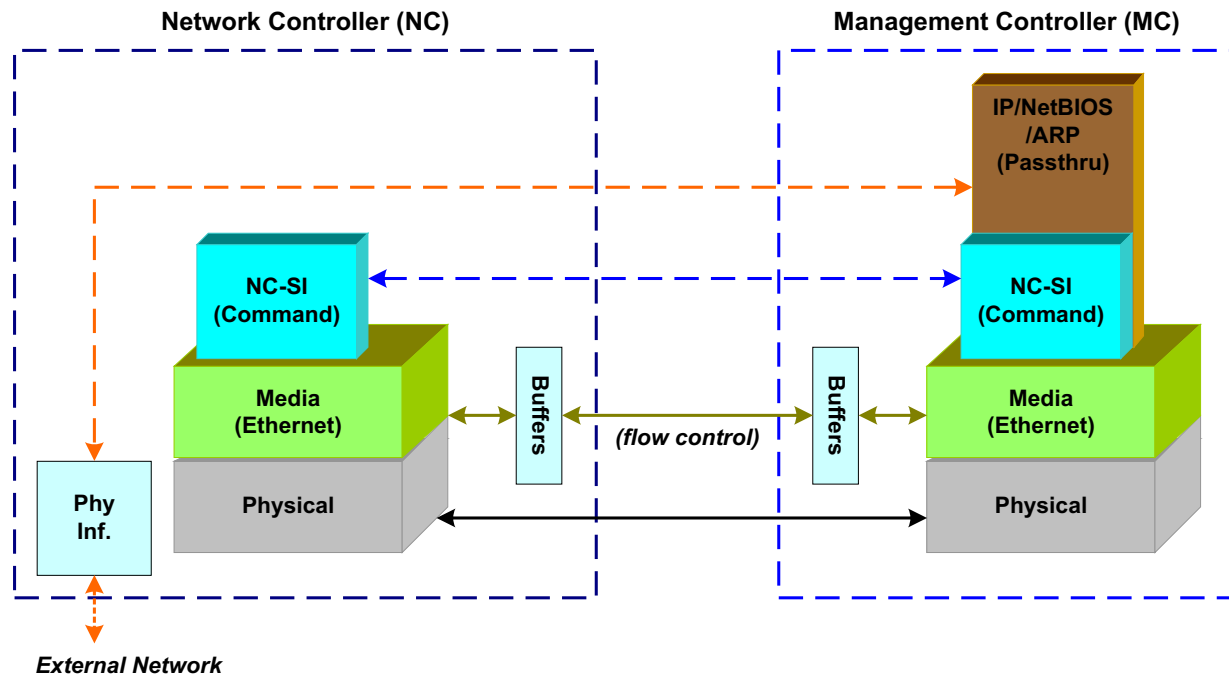


Figure 5 – NC-SI transport stack

5.5 Transport protocol

A simple transport protocol is used to track the reliable reception of command packets. The transport protocol is based upon a command/response paradigm and involves the use of unique Instance IDs (IIDs) in the packet headers to allow responses received to be matched to previously transmitted commands. The Management Controller is the generator of command packets sent to the Sideband Interface of one or more Network Controllers in the system, and it receives response packets from them. A response packet is expected to be received for every command packet successfully sent.

The transport protocol described here shall apply only to command and response packets sent between the Management Controller and the Network Controller.

5.6 Byte and bit ordering for transmission

Unless otherwise specified, the bytes for a multi-byte numeric field are transmitted most significant byte first and bits within a byte are transmitted most significant bit first.

6 Operational behaviors

6.1 Typical operational model

This clause describes the typical system-level operation of the NC-SI components.

The following tasks are associated with Management Controller use of the NC-SI:

- **Initial configuration**

When the NC-SI interface is first powered up, the Management Controller needs to discover and configure NC-SI devices as well as to enable pass-through operation. This task includes setting parameters such as MAC addresses, configuring Layer 2 filtering, setting Channel enables, and so on.

- **General Controller configuration and monitoring**

- The Management Controller may also configure and monitor aspects of Controller operation.

- **Pass-through**

The Management Controller handles transmitting and receiving Pass-through packets using the NC-SI. Pass-through packets can be delivered to and received from the network through the NC-SI based on the Network Controller's NC-SI configuration.

- **Asynchronous event handling**

In certain situations, a status change in the Network Controller, such as a Link State change, can generate an asynchronous event on the Sideband Interface. These event notifications are sent to the Management Controller where they are processed as appropriate.

- **Error handling**

The Management Controller handles errors that could occur during operation or configuration. For example, a Network Controller might have an internal state change that causes it to enter a state in which it requires a level of reconfiguration (this condition is called the "Initial State," described in more detail in 6.1.4); or a data glitch on the NC-SI could have caused an NC-SI command to be dropped by the Network Controller, requiring the Management Controller to retry the command.

6.1.1 State definitions - Defined States

Table 1 describes states related to whether and when the Network Controller is ready to handle NC-SI command packets, when it is allowed to transmit packets through the NC-SI interface, and when it has entered a state where it is expecting configuration by the Management Controller.

Table 1 – NC-SI operating state descriptions

State	Applies to	Description
Interface Power Down	Package	The NC-SI is in the power down state.
Interface Power Up	Package	The NC-SI is in the power up state, as defined in clause 0.
Package Selected (also referred to as the Selected state)	Package	A Selected package is allowed to turn on its electrical buffers and transmit through the NC-SI interface.
Package Deselected (also referred to as the Deselected state)	Package	A Deselected package is not allowed to turn on its electrical buffers and transmit through the NC-SI interface.
Hardware Arbitration Enabled	Package	When hardware arbitration is enabled, the package is allowed to transmit through the NC-SI interface only when it is Selected and has the TOKEN opcode.
Hardware Arbitration Disabled	Package	When hardware arbitration is disabled, the package is allowed to transmit through the NC-SI interface anytime that it is Selected, regardless of whether it has the TOKEN opcode.

State	Applies to	Description
Package Ready	Package	In the Package Ready state, the package is able to accept and respond to NC-SI commands for the package and be Selected.
Package Not Ready	Package	The Package Not Ready state is a transient state in which the package does not accept package-specific commands.
Channel Ready	Channel	In the Channel Ready state, a channel within the package is able to accept channel-specific NC-SI commands that are addressed to its Channel ID (Package ID + Internal Channel ID).
Channel Not Ready	Channel	The Channel Not Ready state is a transient state in which the channel does not accept channel-specific commands.
Initial State	Channel	In the Initial State, the channel is able to accept and respond to NC-SI commands, and one or more configuration settings for the channel need to be set or restored by the Management Controller (that is, the channel has not yet been initialized, or has encountered a condition where one or more settings have been lost and shall be restored). Refer to 6.1.4 for more information.
Channel Enabled	Channel	This is a sub-state of the Channel Ready state. When a channel is enabled, the channel is allowed to transmit unrequested packets (that is, packets that are not command responses—for example, AEN and Pass-through packets) through the NC-SI interface whenever the package is Selected.
Channel Disabled	Channel	This is a sub-state of the Channel Ready state. When a channel is disabled, the channel is not allowed to transmit unrequested packets (that is, packets that are not command responses—for example, AEN and Pass-through packets) through the NC-SI interface.

1111 6.1.2 NC-SI RBT pre-operational states

1112 There are two states defined on RBT before it becomes operational:

1113 • NC-SI Interface Power Down state

1114 In this state, the NC-SI Physical interface and the associated receive and transmit buffers in all
 1115 devices on the NC-SI RBT (that is, the NC-SI interfaces on the Network Controllers and
 1116 Management Controller) are not powered up.

1117 • NC-SI Power Up state

1118 In this state, the NC-SI Physical interface and the associated receive and transmit buffers in all
 1119 devices on the NC-SI RBT (that is, the Network Controller and Management Controller) are
 1120 powered up.

1121 NOTE: NC transmit I/O buffers should not be enabled in this state. The Network Controller is
 1122 expected to transition to the Initial State within T4 seconds after the Power Up state is entered.

1123 6.1.3 Package Ready state

1124 A Network Controller in the Package Ready state shall be able to respond to any NC-SI commands that
 1125 are directed to the ID for the overall package (versus being directed to a particular channel within the
 1126 package). Package-specific commands are identified by a particular set of Channel ID values delivered in
 1127 the command header (see 6.1.9).

6.1.4 Initial State

The Initial State for a channel corresponds to a condition in which the Sideband Interface is powered up and is able to accept NC-SI commands, and the channel has one or more configuration settings that need to be set or restored by the Management Controller. Unless default configuration settings are explicitly defined in this specification, the default values are implementation specific. The MC should not make any assumptions on any configuration settings that are not defined in this specification. Because this state may be entered at any time, the Initial State shall be acknowledged with a Clear Initial State command for the Initial State to be exited. This requirement helps to ensure that the Management Controller does not continue operating the interface unaware that the NC-SI configuration had autonomously changed in the Network Controller.

An NC-SI channel in the Initial State shall:

- be able to respond to NC-SI commands that are directed to the Channel ID for the particular channel (see 6.1.9)
 - respond to all non-OEM NC-SI command packets that are directed to the channel or partitions on the channel with a Response Packet that contains a Response Code of “Command Failed” and a Reason Code of “Initialization Required”
- NOTE This requirement does not apply to commands that are directed to the overall package, such as the Select Package and Deselect Package commands.
- place the channel into the Disabled state
 - set hardware arbitration (if supported) to “enabled” on Interface Power Up only; otherwise, the setting that was in effect before entry into the Initial State shall be preserved (that is, the hardware arbitration enable/disable configuration is preserved across entries into the Initial State)
 - set the enabled/disabled settings for the individual MAC and VLAN filters (typically set using the Set MAC Address, Set VLAN Filter, and Enable VLAN commands) to “disabled”
- NOTE It is recommended that global multicast and broadcast filters are also set to “disabled”.
- reset all counters defined in the various channel and partition level statistics commands, and the Get NC-SI Pass-Through Statistics command to 0x0
 - disable the Channel Network TX setting and transmission of Pass-through packets onto the network
 - clear any record of prior command instances received upon entry into the Initial State (that is, assume that the first command received after entering the Initial State is a new command and not a retried command, regardless of any Instance ID that it may have received before entering the Initial State)
 - disable transmission of AENs and reset any enabled AENs

Otherwise, there is no requirement that other NC-SI configuration settings be set, retained, or restored to particular values in the Initial State unless otherwise specified. Controller configuration settings that are identified as persistent and saved to NVRAM are one example of retained settings..

The Initial State is a NC-SI configuration state and therefore places no requirements on the NC's network link state.

6.1.5 NC-SI Initial State recovery

As described in 6.1.4, a channel in the Initial State shall receive the Clear Initial State command before other commands can be executed. This requirement ensures that if the Initial State is entered asynchronously, the Management Controller is made aware that one or more NC-SI settings may have

1172 changed without its involvement and blocks the Management Controller from issuing additional
1173 commands under that condition. Until the channel receives the Clear Initial State command, the Network
1174 Controller shall respond to any other received command (except the Select Package and Deselect
1175 Package commands) with a Command Failed response code and Interface Initialization Required reason
1176 code to indicate that the Clear Initial State command shall be sent. See response and reason code
1177 definitions in 8.2.5.2.

1178 NOTE Package commands (for example, Select Package and Deselect Package) are always accepted and
1179 responded to normally regardless of whether the Channel is in the Initial State.

1180 If the Management Controller, at any time, receives the response indicating that the Clear Initial State
1181 command is expected, it should interpret this response to mean that default settings have been restored
1182 for the channel (per the Initial State specification), and that one or more package/channel settings need to
1183 be restored by the Management Controller.

1184 **6.1.6** State transition diagram

1185 Figure 6 illustrates the general relationship between the package- and channel-related states described in
1186 Table 1 and the actions that cause transitions between the states. Each bubble in Figure 6 represents a
1187 particular combination of states as defined in Table 1.

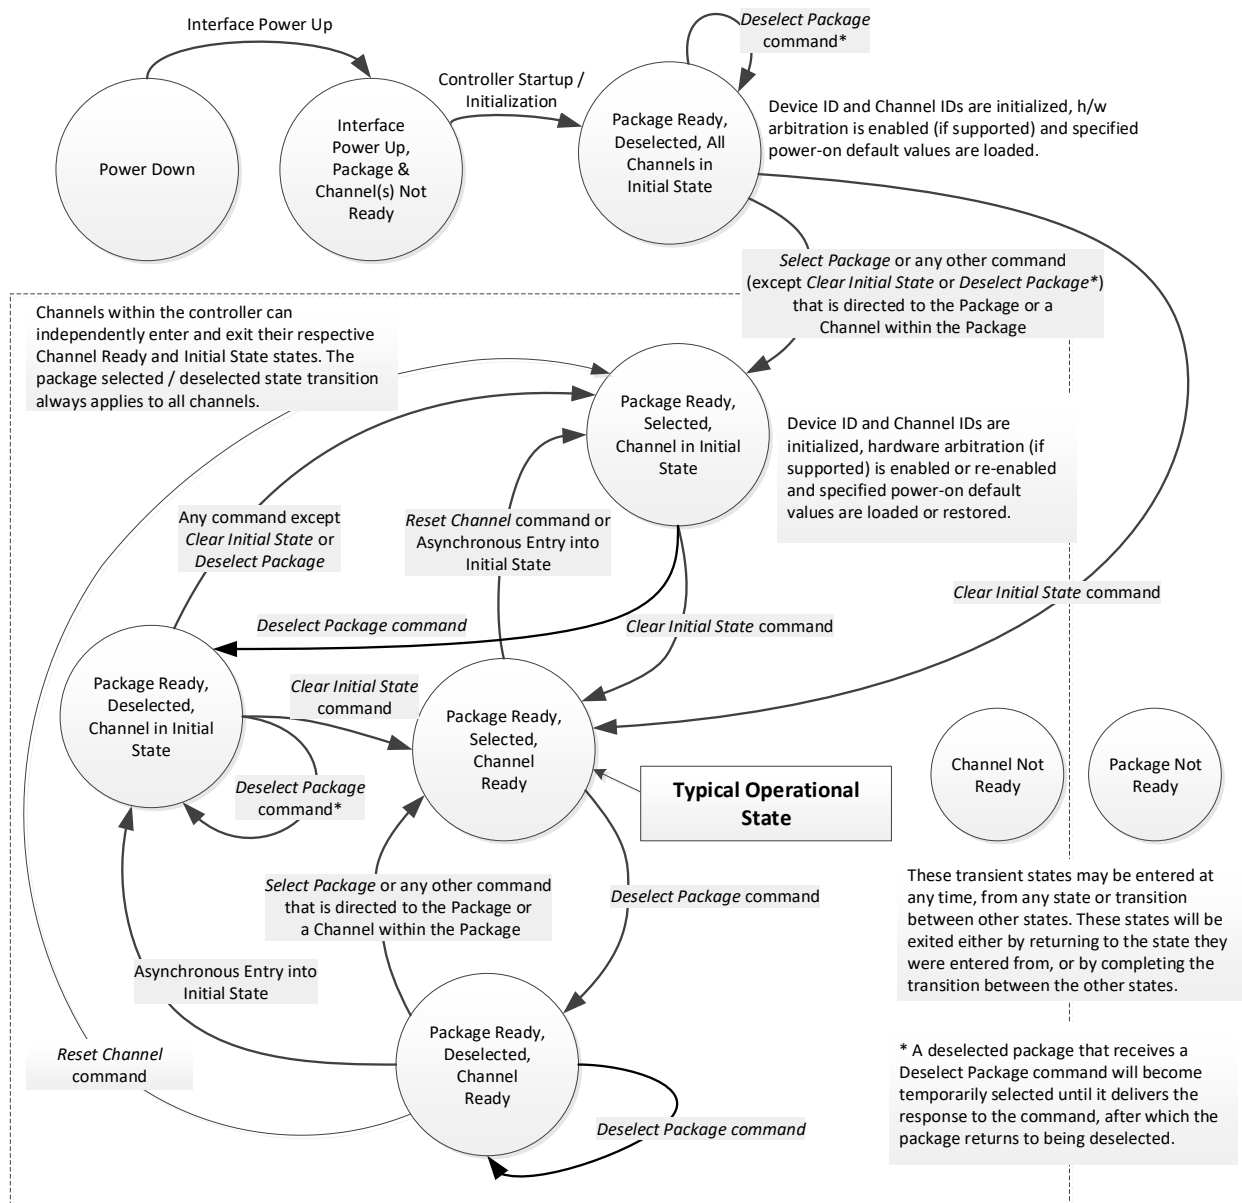


Figure 6 – NC-SI package/channel operational state diagram

6.1.7 State diagram for NC-SI operation with hardware arbitration

Figure 7 shows NC-SI operation in the hardware arbitration mode of operation. This is a sub-set of the general NC-SI operational state diagram (Figure 6) and has been included to illustrate the simplified sequence of package selection when this optional capability is used.

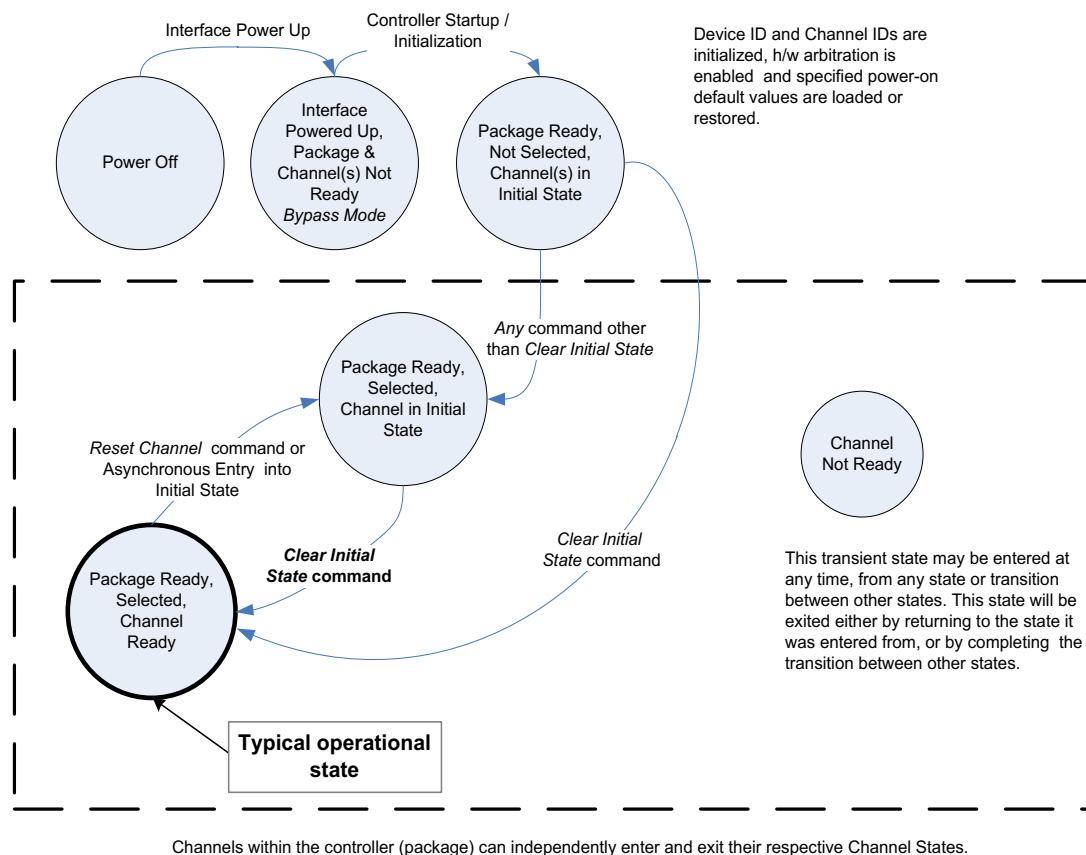


Figure 7 – NC-SI operational state diagram for hardware arbitration operation

While Select and Deselect package commands are not shown in Figure 7, these commands can be used with HW arbitration and will behave as specified in this specification.

Select and Deselect package commands can work together with HW arbitration. If HW arbitration is enabled, a package needs both the HW arbitration token and to be selected in order to transmit on the NC-SI RBT. If either the package is deselected, or the package does not have HW arbitration token, then the package is not allowed to transmit on the NC-SI RBT.

1202 6.1.8 Resets**1203 6.1.8.1 Asynchronous entry into Initial State**

1204 An Asynchronous Reset event is defined as an event that results in a Channel asynchronously entering
1205 the Initial State. This event could occur as a consequence of powering up, a System Reset, a Driver
1206 Reset, an internal firmware error, loss of configuration errors, internal hardware errors, and so on.
1207 Additionally, it is recommended that any event in the NC that causes a total or partial loss of configuration
1208 should be interpreted as an Asynchronous Reset event

1209 Unless otherwise specified, NC-SI configuration settings beyond those required by the Initial State may or
1210 may not be preserved following asynchronous entry into the Initial State, depending on the Network
1211 Controller implementation.

1212 There is no explicit definition of a Reset for an entire package. However, it is possible that an
1213 Asynchronous Reset condition may cause an asynchronous entry into the Initial State for all Channels in
1214 a package simultaneously.

1215 6.1.8.2 Synchronous Reset

1216 A Synchronous Reset event on the NC-SI is defined as a Reset Channel command issued by a
1217 Management Controller to a Channel. Upon the receipt of this command, the Network Controller shall
1218 place the Channel into the Initial State.

1219 Unless otherwise specified, NC-SI configuration settings beyond those required by the Initial State may or
1220 may not be preserved following a Synchronous Reset, depending on the Network Controller
1221 implementation.

1222 6.1.8.3 Other Resets

1223 Resets that do not affect NC-SI operation are outside the scope of this specification.

1224 6.1.9 Network Controller Channel ID

1225 Each channel in the Network Controller shall be physically assigned a Network Controller Channel ID that
1226 will be used by the Management Controller to specify which Network Controller channel, of possibly
1227 many, it is trying to communicate. The Network Controller Channel ID shall be physically assignable
1228 (configured) at system-integration time based on the following specification.

1229 It is the system integrator's or system designer's responsibility to correctly assign and provide these
1230 identifier values in single- and multi-port Network Controller configurations, and to ensure that Channel
1231 IDs do not conflict between devices sharing a common NC-SI RBT interconnect.

The Channel ID field is comprised of two subfields, Package ID and Internal Channel ID, as described in Table 2.

Table 2 – Channel ID format

Bits	Field Name	Description
[7..5]	Package ID	<p>The Package ID is required to be common across all channels within a single Network Controller that share a common NC-SI physical interconnect.</p> <p>The system integrator will typically configure the Package IDs starting from 0 and increasing sequentially for each physical Network Controller.</p> <p>The Network Controller shall allow the least significant two bits of this field to be configurable by the system integrator, with the most significant bit of this field = 0b. An implementation is allowed to have all 3 bits configurable.</p>
[4..0]	Internal Channel ID	<p>The Network Controller shall support Internal Channel IDs that are numbered starting from 0 and increasing sequentially for each channel supported by the Network Controller that is accessible by the Management Controller through the NC-SI using NC-SI commands.</p> <p>An implementation is allowed to support additional configuration options for the Internal Channel ID as long as the required numbering can be configured.</p> <p>An Internal Channel ID value of 0x1F applies to the entire Package.</p>

Channel IDs shall be completely decoded. Aliasing between values is not allowed (that is, the Network Controller is not allowed to have multiple IDs select the same channel on a given Sideband Interface).

Once configured, the settings of the Package ID and Internal Channel ID values shall be retained in a non-volatile manner. That is, they shall be retained across power-downs of the Sideband Interface and shall not be required to be restored by the Management Controller for NC-SI operation. This specification does not define the mechanism for configuring or retaining the Package ID or the Internal Channel ID (if configurable). Some implementations may use pins on the Network Controller for configuring the IDs, other implementations may use non-volatile storage logic such as electrically erasable memory or FLASH, while others may use a combination of pins and non-volatile storage logic.

6.1.10 Configuration-related settings

6.1.10.1 Package-specific operation

There are some NC-SI configuration settings that are package-specific:

- the enable/disable settings for hardware arbitration
- NC-SI flow control
- Package-related AENs

There may also be NC configuration settings that are controlled by NC-SI Commands addressed to the package. These commands specify this requirement in their command description.

Hardware arbitration is enabled or disabled through a parameter that is delivered using the Select Package command. If hardware arbitration is enabled on all Network Controller packages on the NC-SI RBT, more than one package can be in the Selected state simultaneously. Otherwise, only one package is allowed to be in the Selected state at a time in order to prevent electrical buffer conflicts (buffer fights) that can occur from more than one package being allowed to drive the bus.

NC-SI flow control is enabled or disabled using the Set NC-SI Flow Control command. The flow control setting applies to all channels in the package.

1259 Package-specific commands should only be allowed and executed when the Channel ID field is set to
1260 0x1F.

1261 There are some package-level AENs to allow the NC to alert the MC of controller-level events.

1262 6.1.10.2 Channel-specific operation

1263 Channel-specific commands should only be allowed to be executed when the Channel ID field is set to a
1264 value other than 0x1F. Channel-specific commands with Invalid Channel IDs are not allowed (see
1265 6.9.2.1).

1266 Table 3 shows the major categories of configuration settings that control channel operation when a
1267 channel is in the Channel Ready state. Channels that are not operating in Ethernet mode may not
1268 support Pass-through-related settings.

1269

1270 **Table 3 – Channel Ready state configuration settings**

Setting/Configuration Category	Description
"Channel Enable" settings	The Enable Channel and Disable Channel commands are used to control whether the channel is allowed to asynchronously transmit unrequested packets (AEN and Pass-through packets) through the NC-SI interface whenever the package is Selected. Note that channels are always allowed to transmit responses to commands sent to the channel.
"Channel Configuration" settings	Version 1.2 adds a number of commands for configuration setting of channels and their partitions (if supported) See Table 19
Pass-through Transmit Enable settings	The Enable Channel Network TX command is used to enable the channel to transmit any Pass-through packets that it receives through the NC-SI onto the network, provided that the source MAC address in those packets matches the Network Controller settings. Correspondingly, the Disable Channel Network TX command is used to direct the controller not to transmit Pass-through packets that it receives onto the network.
AEN Enable settings	The AEN Enable command is used to enable and disable the generation of the different AENs supported by the Network Controller.
MAC Address Filter settings and control	The Set MAC Address, Enable Broadcast Filter, and Enable Global Multicast Filter commands are used to configure the filters for unicast, broadcast, and multicast addresses that the controller uses in conjunction with the VLAN Filter settings for filtering incoming Pass-through packets.
VLAN Filter settings and control	The Set VLAN Filter command is used to configure VLAN Filters that the controller uses in conjunction with the MAC Address Filters for filtering incoming Pass-through packets. The Enable VLAN and Disable VLAN commands are used to configure VLAN filtering modes and enable or disable whether VLAN filtering is used.

1271 6.1.11 Transmitting Pass-through packets from the Management Controller

1272 Packets not recognized as command packets (that is, packets without the NC-SI Ethertype) that are
1273 received on the Network Controller's NC-SI interface shall be assumed to be Pass-through packets
1274 provided that the source MAC Address matches one of the unicast MAC addresses settings (as
1275 configured by the Set MAC Address command) for the channel in the Network Controller, and will be

1276 forwarded for transmission to the corresponding external network interface if Channel Network TX is
1277 enabled.

1278 **6.1.12** Receiving Pass-through packets for the Management Controller

1279 The Management Controller has control over and responsibility for configuring packet-filtering options,
1280 such as whether broadcast, multicast, or VLAN-tagged packets are accepted. Depending on the filter
1281 configurations, after the channel has been enabled, any packet that the Network Controller receives for
1282 the Management Controller shall be forwarded to the Management Controller through the NC-SI
1283 interface.

1284 **6.1.13** Pass-through operation in multiple medium implementations

1285 Pass-through operation is not restricted to certain physical interfaces, but a NC-SI channel shall support
1286 Pass-through on at most one physical interface at a time.

1287 **6.1.14** Startup sequence examples

1288 **6.1.14.1** Overview

1289 The following clauses show possible startup sequences that may be used by the Management Controller
1290 to start NC-SI operation. Depending upon the specific configuration of each system, there are many
1291 possible variations of startup sequences that may be used, and these examples are intended for
1292 reference only.

1293 **6.1.14.2** Typical non-hardware arbitration specific startup sequence

1294 The following sequence is provided as an example of one way a Management Controller can start up
1295 NC-SI operation. This sequence assumes that the Management Controller has no prior knowledge of how
1296 many Network Controllers are present on RBT, or what capabilities those controllers support. Note that
1297 this is not the only possible startup sequence. Alternative sequences can also be used to start up NC-SI
1298 operation. Some steps may be skipped if the Management Controller has prior knowledge of the Network
1299 Controller capabilities, such as whether Network Controllers are already connected and enabled for
1300 hardware arbitration.

1301 **1) Power up**

1302 The NC-SI is powered up (refer to 10.2.7 for the specification of this condition). The Network
1303 Controller packages are provided a Network Controller Power Up Ready Interval during which
1304 they can perform internal firmware startup and initialization to prepare their NC-SI to accept
1305 commands. The Management Controller first waits for the maximum Network Controller Power
1306 Up Ready Interval to expire (refer to Table 262). At this point, all the Network Controller
1307 packages and channels should be ready to accept commands through the NC-SI. (The
1308 Management Controller may also start sending commands before the Network Controller Power
1309 Up Ready Interval expires but will have to handle the case that Network Controller devices may
1310 be in a state in which they are unable to accept or respond to commands.)

1311 **2) Discover package**

1312 The Management Controller issues a Select Package command starting with the lowest
1313 Package ID (see 8.4.5 for more information). Because the Management Controller is assumed
1314 to have no prior knowledge of whether the Network Controller is enabled for hardware
1315 arbitration, the Select Package command is issued with the Hardware Arbitration parameter set
1316 to 'disable'.

If the Management Controller receives a response within the specified response time, it can record that it detected a package at that ID. If the Management Controller does not receive a response, it is recommended that the Management Controller retry sending the command. Three total tries are typical. (This same retry process should be used when sending all commands to the Network Controller and will be left out of the descriptions in the following steps.) If the retries fail, the Management Controller can assume that no Network Controller is at that Package ID and can immediately repeat this step 2) for the next Package ID in the sequence.

3) Discover and get capabilities for each channel in the package

The Management Controller can now discover how many channels are supported in the Network Controller package and their capabilities. To do this, the Management Controller issues the Clear Initial State command starting from the lowest Internal Channel ID (which selects a given channel within a package). If it receives a response, the Management Controller can then use the Get Version ID command to determine NC-SI specification compatibility, and the Get Capabilities command to collect information about the capabilities of the channel. The Management Controller can then repeat this step until the full number of internal channels has been discovered. (The Get Capabilities command includes a value that indicates the number of channels supported within the given package.)

NOTE The *NC-SI Specification* requires Network Controllers to be configurable to have their Internal Channel IDs be sequential starting from 0. If it is known that the Network Controller is configured this way, the Management Controller needs only to iterate sequentially starting from Internal Channel ID = 0 up to the number of channels reported in the first Get Capabilities response.

The Management Controller should temporarily retain the information from the Get Capabilities command, including the information that reports whether the overall package supports hardware arbitration. This information is used in later steps.

4) Repeat steps 2 and 3 for remaining packages

The Management Controller repeats steps 2) and 3) until it has gone through all the Package IDs.

IMPORTANT: Because hardware arbitration has not been enabled yet, the Management Controller shall issue a Deselect Package command to the present Package ID before issuing the Select Package command to the next Package ID. If hardware arbitration is not being used, only one package can be in the Selected state at a time. Otherwise, hardware electrical buffer conflicts (buffer fights) will occur between packages.

5) Initialize each channel in the package

Based on the number of packages and channels that were discovered, their capabilities, and the desired use of Pass-through communication, the Management Controller can initialize the settings for each channel. This process includes the following general steps for each package:

- a) Issue the Select Package command.
- b) For each channel in the package, depending on controller capabilities, perform the following actions. Refer to individual command descriptions for more information.
 - Use the Set MAC Address command to configure which unicast and multicast addresses are used for routing Pass-through packets to and from the Management Controller.
 - Use the Enable Broadcast Filter command to configure whether incoming broadcast Pass-through packets are accepted or rejected.

- 1362 • Use the Enable Global Multicast Filter command to configure how incoming multicast
1363 Pass-through packets are handled based on settings from the Set MAC Address
1364 command.
- 1365 • Use the Set VLAN Filter and Enable VLAN Filters commands to configure how
1366 incoming Pass-through packets with VLAN Tags are handled.
- 1367 • Use the Set NC-SI Flow Control command (if supported) to configure how Ethernet
1368 Pause Frames are used for flow control on RBT. Set NC-SI Flow Control is a
1369 package command and only needs to be issued once.
- 1370 • Use the AEN Enable command to configure what types of AEN packets the channel
1371 should send out on the NC-SI.
- 1372 • Use the Enable Channel Network TX command to configure whether the channel is
1373 enabled to deliver Pass-through packets from the NC-SI to the network (based on the
1374 MAC address settings) or is disabled from delivering any Pass-through packets to the
1375 network.

1376 c) Issue the Deselect Package command.

1377 6) **Start Pass-through packet and AEN operation on the channels**

1378 The channels should now have been initialized with the appropriate parameters for Pass-
1379 through packet reception and AEN operation. Pass-through operation can be started by issuing
1380 the Enable Channel command to each channel that is to be enabled for delivering Pass-through
1381 packets or generating AENs through the NC-SI interface.

1382 If hardware arbitration is not operational and it is necessary to switch operation over to another package,
1383 a Deselect Package command shall be issued to the presently selected package before a different
1384 package can be selected. Deselecting a package blocks all output from the package. Therefore, it is not
1385 necessary to issue Disable Channel commands before selecting another package. There is no restriction
1386 on enabling multiple channels within a package.

1387

1388 **6.1.14.3 Hardware arbitration-specific startup sequence**

1389 This clause applies when multiple NCs are used by the MC. This clause only applies to the NC-SI over
1390 RBT binding.

1391 The following is an example of the steps that a Management Controller may perform to start up NC-SI
1392 operation when Hardware Arbitration is specifically known to be used, present, and enabled on all
1393 Network Controllers. This example startup sequence assumes a high level of integration where the
1394 Management Controller knows the Network Controllers support and default to the use of Hardware
1395 Arbitration on startup but does not have prior knowledge of how many Network Controllers are present on
1396 RBT, or the full set of capabilities those controllers support, so discovery is still required.

1397 Although other startup examples may show a specific ordering of steps for the process of discovering,
1398 configuring and enabling channels, the Management Controller has almost total flexibility in choosing how
1399 these steps are performed once a channel in a package is discovered. In the end, it would be just as valid
1400 for a Management Controller to follow a breadth-first approach to discovery steps as it would be to follow
1401 a depth-first approach where each channel that is discovered is fully initialized and enabled before
1402 moving to the next.

1403 1) **Power up**

1404 No change from other startup scenarios.

2) **Discovery**

The process of discovery consists of identifying the number of packages that are available, the number of channels that are available in each package, and for each channel, the capabilities that are provided for Management Controller use. Because, in this startup scenario, the Management Controller knows Hardware Arbitration is used, it is not required to use the **Select Package** and **Deselect Package** commands for discovery but may elect to just use the **Clear Initial State** command for this purpose instead.

In this startup scenario, Packages and Channels are discovered by sending the **Clear Initial State** command starting with the lowest Package ID and Channel ID, then waiting for, and recording, the response event as previously described. Internal channel IDs are required to be numbered sequentially starting with 0, so when the Management Controller does not receive a response to repeated attempts at discovery, it knows this means no additional channels exist in the current package. If this happens when the internal channel ID is 0, the Management Controller knows a package is not available at the current package ID, and it continues with the next package ID in sequence. If the Management Controller receives a response to the **Clear Initial State** command, it records that the channel and package are available, and continues discovery.

During discovery, the Management Controller should interrogate the capabilities of each channel found to be available in each package by sending the **Get Capabilities** command appropriate package and channel ID values. However, it does not matter whether this is done as the very next step in the discovery process or performed for each channel after all packages and channels have been discovered, just as long as the Management Controller does interrogate each channel.

3) **Configure each channel and enable pass-through**

Once the existence of all packages and channels, and the capabilities of each channel, have been discovered and recorded, the Management Controller shall initialize and enable each channel as needed for use. The details of these steps remain essentially the same as have been previously stated, except to note that there are no restrictions on how they are performed. What this means is that the MC may perform these steps in any order across the channels in each package as it sees fit. The MC may fully initialize and enable each channel in each package one at a time or perform the same step on each channel in sequence before moving on to the next, or in a different order. The specific order of steps is not dictated by this specification.

6.1.14.4 Summary of scheme for the MC without prior knowledge of hardware arbitration

The following scheme describes the case when the MC does not have a priori knowledge of the hardware arbitration support across multiple NCs.

1. For each available NC,

- a. The MC checks whether a device supports the HW arbitration, using “**Get Capabilities**” command (this implicitly selects the package).
- b. The MC issues “**Deselect Package**” for the NC (needed as at this stage we do not know whether all the devices support HW arbitration).

2. If (all NCs support HW arbitration and HW arbitration is used by all NCs), then

the MC assumes that HW arbitration is active because according to clause 6.2.4 “set hardware arbitration (if supported) to *enabled* on Interface Power Up only”, and the MC can “Select” any number of packages at the same time.

1450 Otherwise (at least one NC reports that HW arbitration is not supported, or at least one NC
1451 reports that HW arbitration is not used, or at least one NC cannot report its support level) then
1452 HW arbitration is **not** active, and the MC can “Select” only single package at the any time.
1453 The MC configures every NC to disable HW arbitration, using the “**Select Package**”
1454 command.

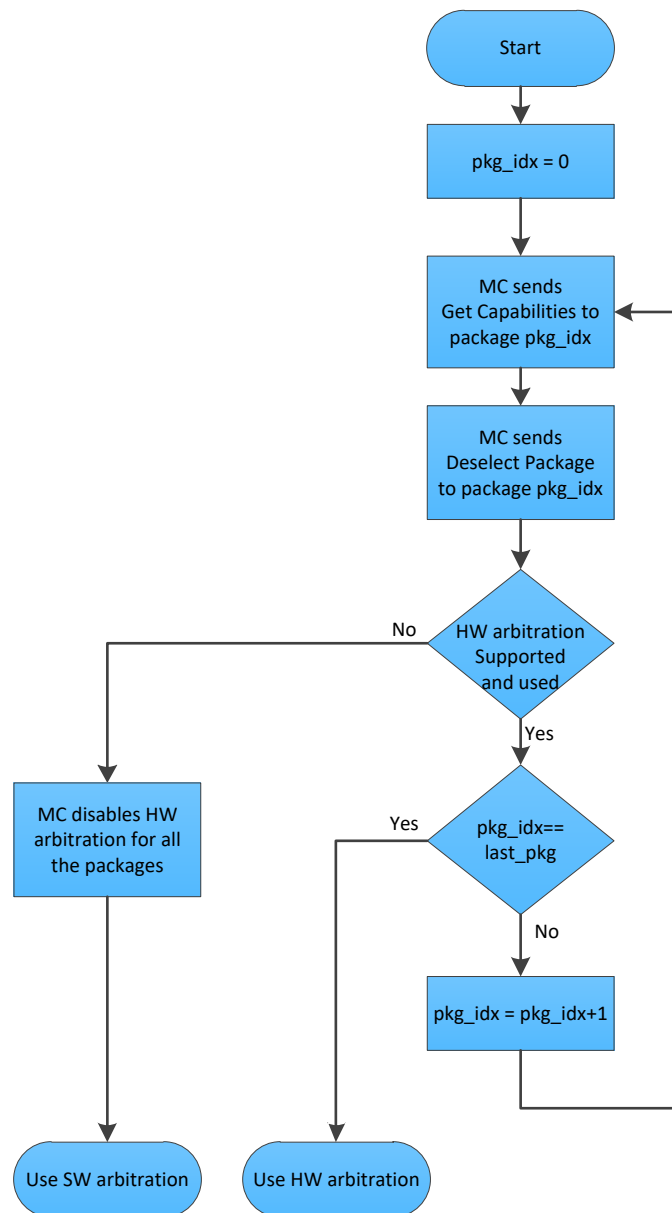


Figure 8 – MC steps when the MC does not have prior knowledge of hardware arbitration

6.2 NC-SI traffic types

6.2.1 Overview

Two types of traffic are defined by NC-SI, based on the network fabric type: Pass-through traffic and Control traffic.

- Pass-through traffic consists of packets that are transferred between the external network interface and the Management Controller using the Sideband Interface.
- Control traffic consists of commands (requests) and responses that support the inventory, configuration and control of the Network Controller, the Sideband Interface and Pass-through operation of the Network Controller, and AENs that support reporting various events to the Management Controller.

6.2.2 Command protocol

6.2.2.1 Overview

Commands are provided to allow a Management Controller to initialize, control, and regulate Management Controller packet flow across the sideband interface, configure channel filtering, and to interrogate the operational status of the Network Controller. As interface master, the Management Controller is the initiator of all commands, and the Network Controller responds to commands, but may also generated AENs if enabled.

6.2.2.2 Instance IDs

The command protocol uses a packet field called the Instance ID (IID). IID numbers are 8-bit values that shall range from 0x01 to 0xFF. IIDs are used to uniquely identify instances of a command, to improve the robustness of matching responses to commands, and to differentiate between new and retried commands. The Network Controller that receives a command handles the IID in the following ways:

- It returns the IID value from the command in the corresponding response.
- If the IID is the same as the IID for the previous command, it recognizes the command as a 'retried' command rather than as a new instance of the command. It is expected that the 'retried' command contains the same command type value in the Control Packet Type field. The NC behavior when a 'retried' command type does not match the original command type is outside the scope of this specification.
- If a retried command is received, the Network Controller shall return the previous response. Depending on the command, the Network Controller can accomplish this either by holding the previous response data so that it can be returned, or, if re-executing the command has no side effects (that is, the command is idempotent), by re-executing the command operation and returning that response.
- If the command IID is the same as the IID for the previous command, and the Poll Indication is set, the NC recognizes the command as a 'polling' command rather than as a new instance of the command.
 - When polling, the MC is expected to use the command type value of the original command in the Control Packet Type field. If there was no command in progress, the NC shall fail the 'polling' command and respond with an error. When the NC fails the 'polling' command, the outcome of the original command is indeterminate and is outside the scope of this specification.
 - If a command with Poll Indication set is received and the original command has been completed, then the Network Controller shall return the response of the completed command.

- If it is still processing the command, it shall return a “Delayed Response” reason code and optionally recommend a next polling time interval.
- When an IID value is received that is different from the one for the previous command, the Network Controller executes the command as a new command.
- When the NC-SI Channel first enters the Initial State, it shall clear any record of any prior requests. That is, it assumes that the first command after entering the Initial State is a new command and not a retried command, regardless of any IID that it may have received before entering the Initial State.

Thus, for single-threaded operation with idempotent commands, a responding Network Controller can simply execute the command and return the IID in the response that it received in the command. If it is necessary to not execute a retried command, the responding controller can use the IID to identify the retried command and return the response that was delivered for the original command.

The Management Controller that generates a command handles the IID in the following ways:

- The IID changes for each new instance of a command.
- If a command needs to be retried, the Management Controller uses the same value for the IID that it used for the initial command.
- The Management Controller can optionally elect to use the IID to provide additional confirmation that the response is being returned for a particular command.

Because an AEN is not a response, an AEN always uses a value of 0x00 for its IID.

NOTE: The Instance ID mechanism can be readily extended in the future to support multiple controllers and multiple outstanding commands. This extension would require having the responder track the IID on a per command and per requesting controller basis. For example, a retried command would be identified if the IID and command matched the IID and command for a prior command for the given originating controller's ID. That is, a match is made with the command, originating controller, and IID fields rather than on the IID field alone. A requester that generates multiple outstanding commands would correspondingly need to track responses based on both command and IID to match a given response with a given command. IIDs need to be unique for the number of different commands that can be concurrently outstanding.

6.2.2.3 Single-threaded operation

The Network Controller is required to support NC-SI commands only in a single-threaded manner. That is, the Network Controller is required to support processing only one command at a time and is not required to accept additional commands until after it has sent the response to the previous one.

Therefore, the Management Controller should issue NC-SI commands in a single-threaded manner. That is, the Management Controller should have only one command outstanding to a given Network Controller package at a time. Upon sending an NC-SI command packet, and before sending a subsequent command, the Management Controller should wait for the corresponding response packet to be received or a command timeout event to occur before attempting to send another command. For the full descriptions of command timeout, see 6.9.3.2.

Note: While NC implementations are only required to support single-threaded operations, they may choose to support more than one outstanding command. The use of unique IIDs is essential to properly match multiple outstanding commands and responses in such implementations.

6.2.2.4 Responses

The Network Controller shall process and acknowledge each validly formatted command received at the NC-SI interface by formatting and sending a valid response packet to the Management Controller through the NC-SI interface.

To allow the Management Controller to match responses to commands, the Network Controller shall copy the IID number of the Command into the Instance ID field of the corresponding response packet.

To allow for retransmission and error recovery, the Network Controller may re-execute the last command or maintain a copy of the response packet most recently transmitted to the Management Controller through its sideband interface. This “previous” response packet shall be updated every time a new response packet is transmitted to the Management Controller by replacing it with the one just sent.

The Network Controller shall return a “Command Unsupported” response code with an “Unknown Command Type” reason code for any command (standard or OEM) that the Network Controller does not support or recognize. If a command cannot be executed due to the processing of others, the response code Command Unavailable shall be returned.

6.2.2.5 Response and post-response processing

Typically, a Network Controller completes a requested operation before sending the response. In some situations, however, it may be useful for the controller to be allowed to queue up the requested operation and send the response assuming that the operation will complete correctly (for example, when the controller is requested to change link configuration). The following provisions support this process:

- A Network Controller is allowed to send a response before performing the requested action if the command is expected to complete normally and all parameters that are required to be returned with the response are provided.
- Temporal ordering of requested operations shall be preserved. For example, if one command updates a configuration parameter value and a following command reads back that parameter, the operation requested first shall complete so that the following operation returns the updated parameter.
- Under typical operation of the Network Controller, responses should be delivered within the Normal Execution Interval (T5) (see Table 262).
- Unless otherwise specified, all requested operations shall complete within the Asynchronous Reset/Asynchronous Not Ready interval (T6) following the response.
- If the Network Controller channel determines that the requested operation or configuration change has not been completed correctly after sending the response, the channel shall enter the Initial State.
- If the command response is dependent on the execution of the command and the command response cannot be provided within Normal Execution Interval (T5), then a “Delayed Response” response code may be returned. In this case, the MC can poll the command later with the “Poll Indication” set to retrieve the response. The decision on when the MC polls again can be based on one of the following criteria:
 - A fixed delay. In this case a delay greater than T5 is recommended.
 - If provided, based on the “recommended next polling time” in the original response
 - If the AEN is enabled, based on reception of a “Delayed Response Ready AEN”

When using delayed responses, the NC shall complete the command processing within T14 sec.

6.2.2.6 NC-SI traffic ordering

This specification does not require any ordering between AENs, NC-SI responses, and NC-SI Pass-through packets. Specific transport binding specifications may require ordering between AENs, NC-SI responses, and NC-SI Pass-through packets.

6.3 Link configuration and control

6.3.1 Link Configuration

The Network Controller provides commands to allow the Management Controller to specify the auto-negotiation, link speed, duplex settings, FEC algorithm, link training, SerDes lane configuration, and so on to be used on the network interface. For more information, see 8.4.21.

The Management Controller should make link configuration changes only when the host network driver is absent or non-operational.

6.3.2 Link Status

The Network Controller provides a Get Link Status command to allow the Management Controller to interrogate the configuration and operational status of the primary Ethernet links. The Management Controller may issue the Get Link Status command regardless of OS operational status.

6.4 Frame filtering for Pass-through mode

6.4.1 Overview

The Network Controller provides the option of configuring various types of filtering mechanisms for the purpose of controlling the delivery of received Ethernet frames to the Management Controller. These options include VLAN Tag filter, L2 address filters, MAC address support, and limited frame filtering using L3, L4 protocol header fields. All frames that pass frame filtering are forwarded to the Management Controller over the Sideband Interface. Refer to RFC2373, RFC2461 and RFC3315 for IPv6-related definitions.

6.4.2 Multicast filtering

The Network Controller may provide commands to allow the Management Controller to enable and disable global filtering of all multicast packets. The Network Controller may optionally provide one or more individual multicast filters, as well as DHCP v6, IPv6 Neighbor Advertisement, IPv6 Router Advertisement, IPv6 Neighbor Solicitation, IPv6 MLD, mDNSv4, mDNSv6 and LLDP filters.

6.4.3 Broadcast filtering

The Network Controller provides commands to allow the Management Controller to enable and disable forwarding of Broadcast and ARP packets. The Network Controller may optionally support selective forwarding of broadcast packets for specific protocols, such as DHCP (see RFC2131) and NetBIOS.

6.4.4 VLAN filtering

The Network Controller provides commands to allow the Management Controller to enable and disable VLAN filtering, configure one or more VLAN Filters, and to configure VLAN filtering modes.

Figure 9 illustrates the flow of frame filtering. Italicized text in the figure is used to identify NC-SI command names.

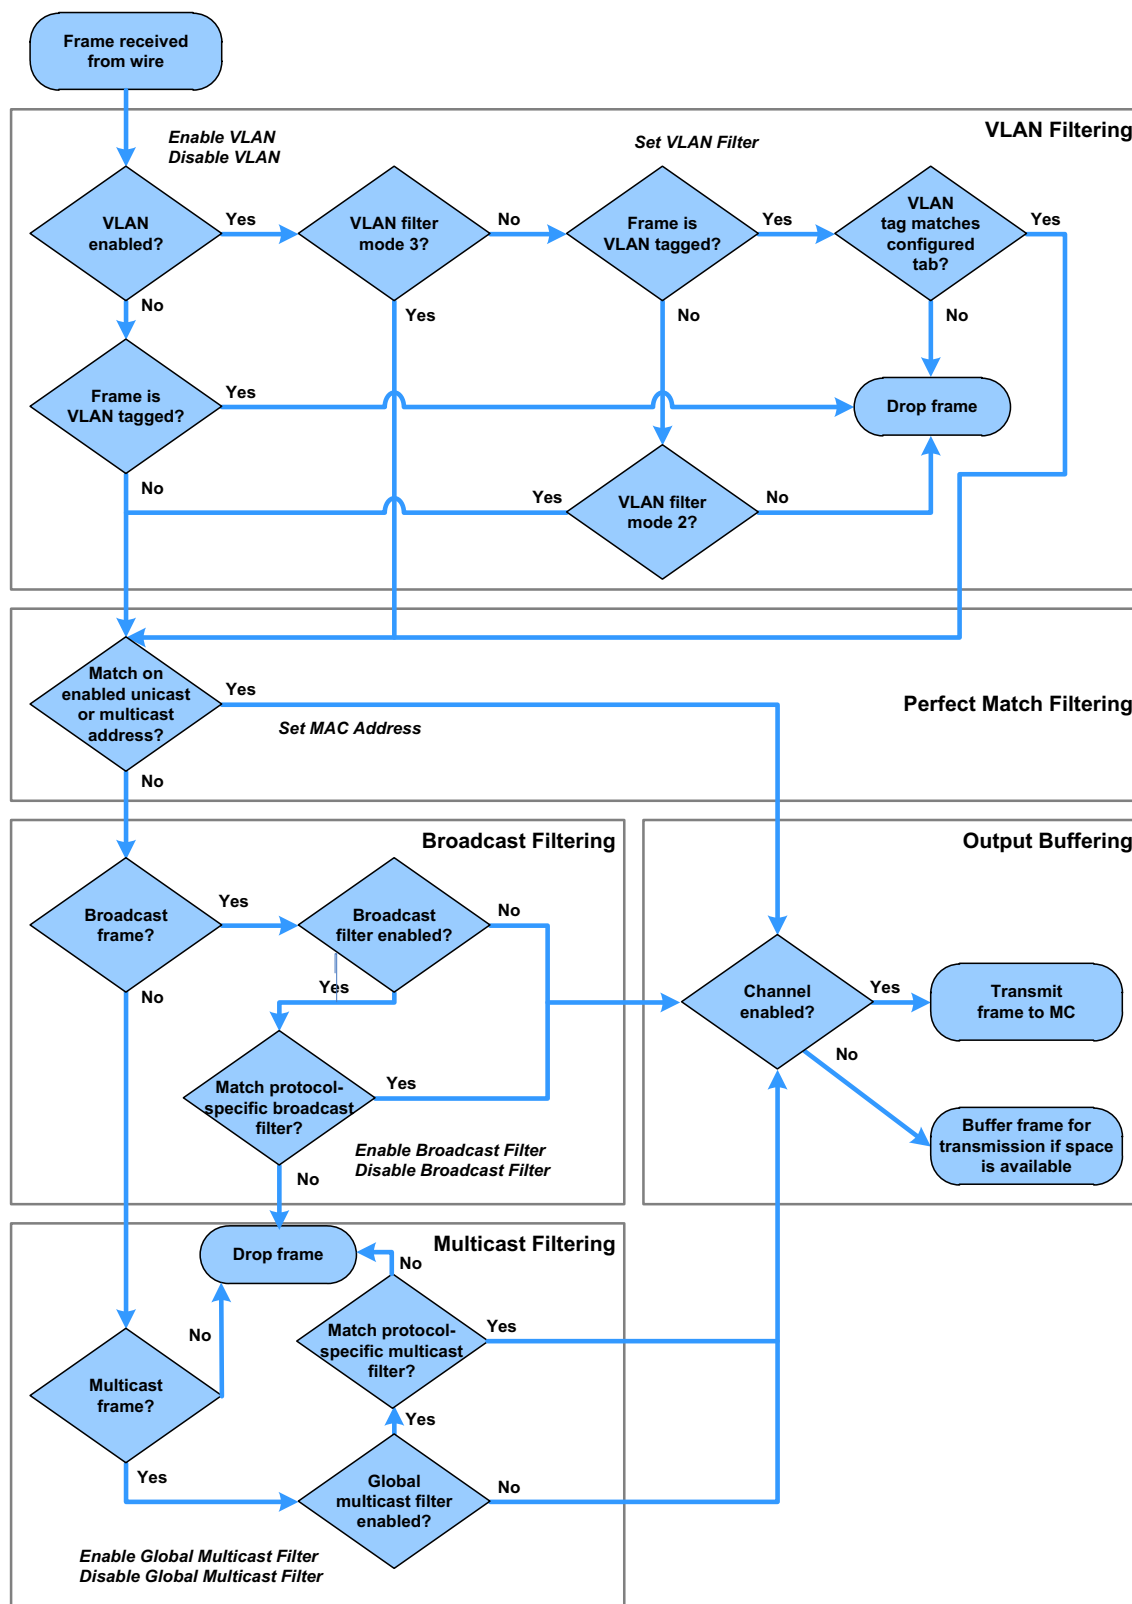


Figure 9 – NC-SI packet filtering flowchart

6.5 Output buffering behavior

There are times when the NC is not allowed to transmit Pass-through, AEN, or Control Packets onto the Sideband Interface.

The NC should buffer Pass-through frames to be transmitted to the MC under any of the following conditions:

- The package is deselected.
- For a channel within a package while that channel is disabled.
- When the hardware arbitration is enabled, and the NC does not have the token to transmit frames to the MC.

The NC may buffer AENs to the MC under any of the above conditions.

Control Packets (responses) are buffered when hardware arbitration is enabled, and the NC does not have the token to transmit frames to the MC.

Additionally, while an NC-SI channel is in the initial state, previously received Pass-through frames and AENs may or may not be buffered. This behavior is outside the scope of this specification.

6.6 NC-SI flow control

The Network Controller may provide commands to enable flow control on the RBT interface between the Network Controller and the Management Controller. The NC-SI flow control behavior follows the PAUSE frame behavior as defined in the [IEEE 802.3 specification](#). Flow control is configured using the Set NC-SI Flow command (see 8.4.41).

When enabled for flow control, a channel may direct the package to generate and renew 802.3x (XOFF) PAUSE Frames for a maximum interval of T12 for a single congestion condition. If the congestion condition remains in place after a second T12 interval expires, the congested channel shall enter the Initial State and remove its XOFF request to the package. Note that some implementations may have shared buffering arrangements where all channels within the package become congested simultaneously. Also note that if channels become congested independently, the package may not immediately go into the XON state after T12 if other channels within the package are still requesting XOFF.

6.7 Asynchronous Event Notification

Asynchronous Event Notification (AEN) packets enable the Network Controller to deliver unsolicited notifications to the Management Controller when certain status changes that could impact interface operation occur in the Network Controller. Because the NC-SI is a small part of the larger Network Controller, its operation can be affected by a variety of events that occur in the Network Controller. These events include link status changes, OS driver loads and unloads, and chip resets. This feature defines a set of notification packets that operate outside of the established command-response mechanism.

Control over the generation of the AEN packets is achieved by control bits in the AEN Enable command. Each type of notification is optional and can be independently enabled by the Management Controller.

AENs are not acknowledged, and there is no protection against the possible loss of an AEN packet. Each defined event has its own AEN packet. Because the AEN packets are generated asynchronously by the Network Controller, they cannot implement some of the features of the other Control Packets. AEN packets leverage the general packet format of Control Packets.

- The originating Network Controller channel shall fill in its Channel ID (Ch. ID) field in the command header to identify the source of notification.

- The IID field in an AEN shall be set to 0x00 to differentiate it from a response or command packet.
- The Network Controller shall copy the AEN MC ID field from the AEN Enable command into the MC ID field in every AEN sent to the Management Controller.

6.8 AEN handling in multiple medium implementations

Implementations that use NC-SI over physical interfaces other than RBT and enable Asynchronous Event Notifications (AEN) on those other media shall comply with the requirements in DSP-0261.

AENs that are enabled via RBT are specific to RBT-active operation and any AEN that is subsequently generated is only delivered over RBT and then only when RBT is active (maintained or restored operation).

AEN generation is suppressed and not cached when the media on which it was enabled is not active.

6.9 Error handling

6.9.1 Overview

This clause describes the error-handling methods that are supported over the NC-SI. Two types of error-handling methods are defined:

- Synchronous Error Handling
- Errors that trigger Asynchronous Entry into the Initial State

Synchronous Error Handling occurs when an Error (non-zero) Response/Reason Code is received in response to a command issued by the Management Controller. For information about response and reason codes, see 8.2.4.1.

Asynchronous Entry into the Initial State Error Handling occurs when the Network Controller asynchronously enters the Initial State because of an error condition that affects NC-SI configuration or a failure of a command that was already responded to. For more information, see 6.1.8.1.

6.9.2 Transport errors

6.9.2.1 Dropped Control Packets

A Network Controller with an active interface shall drop Control Packets received on the NC-SI interface under the following conditions:

- The packet has an invalid Frame Check Sequence (FCS) value.
- Frame length does not meet [IEEE 802.3](#) requirements (except for OEM commands, where accepting larger packets may be allowed as a vendor-specific option).
- The packet checksum (if provided) is invalid.
- The NC-SI Channel ID value in the packet does not match the expected value.
- The Network Controller does not have resources available to accept the packet.
- The Network Controller receives a command packet with an incorrect header revision.
- Control Packets may also be dropped if an event that triggers Asynchronous Entry into the Initial State causes packets to be dropped during the transition.

1701 6.9.2.2 Pass-through packet errors

1702 Handling of Pass-through packet errors, other than logging statistics, is out of scope of this specification.

1703 6.9.3 Missing responses**1704 6.9.3.1 Overview**

1705 There are typical scenarios in which the Management Controller does not receive the response to a
1706 command:

- 1707 • The Network Controller dropped the command and thus never sent the response.
- 1708 • The response was dropped by the Management Controller (for example, because of a CRC
1709 error in the response packet).
- 1710 • The Network Controller is in the process of being reset or is disabled.

1711 The Management Controller can detect a missing response packet as the occurrence of an NC-SI
1712 command timeout event.

1713 6.9.3.2 Command timeout

1714 The Management Controller may detect missing responses by implementing a command timeout interval.
1715 The timeout value chosen by the Management Controller shall not be less than Normal Execution
1716 Interval, T5. Upon detecting a timeout condition, the Management Controller should not make
1717 assumptions on the state of the unacknowledged command (for example, the command was dropped, or
1718 the response was dropped), but should retransmit (retry) the previous command using the same IID it
1719 used in the initial command.

1720 The Management Controller should try a command at least three times before assuming an error
1721 condition in the Network Controller.

1722 It is possible that a Network Controller could send a response to the original command at the same time a
1723 retried command is being delivered. Under this condition, the Management Controller could get more than
1724 one response to the same command. Thus, the Management Controller should be capable of determining
1725 that it has received a second instance of a previous response packet. Dropped commands may be
1726 detected by the Management Controller as a timeout event waiting for the response.

1727 6.9.3.3 Handling dropped commands or missing responses

1728 To recover from dropped commands or missing responses, the Management Controller can retransmit
1729 the unacknowledged command packet using the same IID that it used for the initial command.

1730 The Network Controller shall be capable of reprocessing retransmitted (retried) commands without error
1731 or undesirable side effects. The Network Controller can determine that the command has been
1732 retransmitted by verifying that the IID is unchanged from the previous command.

1733 6.9.4 Detecting Pass-through traffic interruption

1734 The Network Controller might asynchronously enter the Initial State because of a reset or other event. In
1735 this case, the Network Controller stops transmitting Pass-through traffic on the RXD lines. Similarly, Pass-
1736 through traffic sent to the Network Controller may be dropped. If the Management Controller is not in the
1737 state of sending or receiving Pass-through traffic, it may not notice this condition. Thus, the Management
1738 Controller should periodically issue a command to the Network Controller to test whether the Network
1739 Controller has entered the Initial State. How often this testing should be done is a choice of the
1740 Management Controller.

6.10 Support for additional network fabrics

6.10.1 FC support

NCs that support Fibre Channel connectivity can be inventoried, configured, and monitored. Fibre Channel-specific link speed, link status, boot configuration and statistics commands are provided. Fibre Channel over Ethernet (FCoE) support is also defined for Ethernet NCs that support it.

6.10.2 InfiniBand Support

NCs that support InfiniBand connectivity can be inventoried, configured, and monitored. InfiniBand-specific link speed, link status and statistics commands are provided.

6.11 PLDM and SPDM transport

NC-SI over RBT can be used to transport SPDM or PLDM messages over RBT. This transport supports the following modes:

- MC sends PLDM and/or SPDM commands to the NC.
- MC polls the NC for PLDM and/or SPDM commands originating at the NC.
- The NC indicates through an AEN that a PLDM/SPDM command is available for retrieval.

The following commands are used to implement these flows:

Command	PLDM	SPDM
Send command from MC	PLDM Request (0x51)	Transfer SPDM (0x60)
Poll for NC command	Query Pending NC PLDM Request (0x56)	Query Pending NC SPDM Request (0x61)
Respond to NC command	Send NC PLDM Reply (0x57)	Send NC SPDM Reply (0x62)
AEN	Pending PLDM AEN (0x71)	Pending SPDM AEN (0x72)

The PLDM and SPDM command flows are described in the UML diagrams below.





1763

1764

1765

1766

1767

1768

1769

1770

1771

1772

1773

1774

7

Arbitration in configurations with multiple Network Controller packages

7.1

Overview

This clause applies to NC-SI over RBT only.

More than one Network Controller package on a RBT interface can be enabled for transmitting packets to the Management Controller. This specification defines two mechanisms to accomplish Network Controller package arbitration operations. One mechanism uses software commands provided by the Network Controller for the Management Controller to control whose turn it is to transmit traffic. The other mechanism uses hardware arbitration to share the single RBT bus. Implementations are required to support command-based Device Selection operation; the hardware arbitration method is typically desired but is optional.

7.2 Multi-controller RBT

Figure 10 is a simplified block diagram of the Sideband Interface being used in a multi-drop configuration. The RMII (upon which NC-SI RBT is based) was originally designed for use as a point-to-point interconnect. Accordingly, only one party can transmit data onto the bus at any given time. There is no arbitration protocol intrinsic in the RMII specification to support managing multiple transmitters.

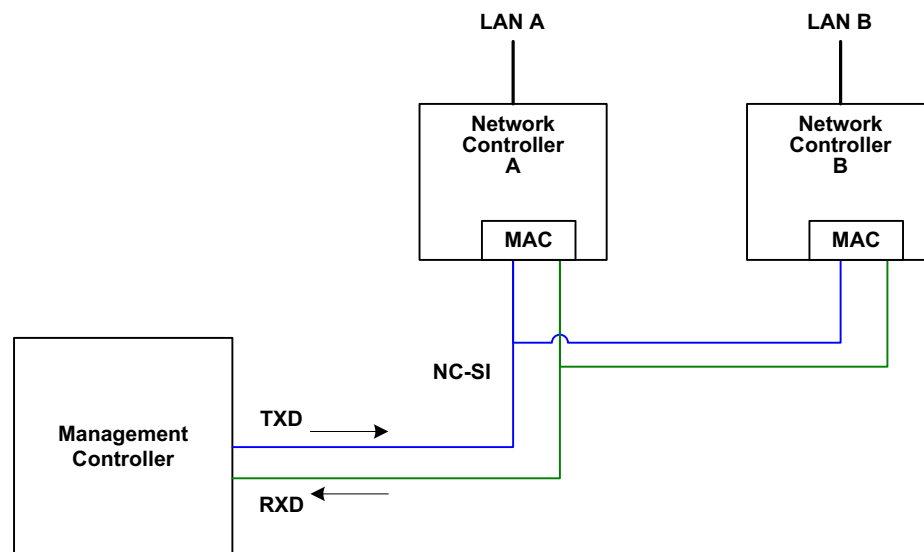


Figure 10 – Basic multi-drop block diagram

However, it is possible for multiple Network Controllers on the interface to be able to simultaneously *receive* traffic from the Management Controller that is being transmitted on the RBT TXD lines. The Network Controllers can receive commands from the Management Controller without having to arbitrate for the bus. This facilitates the Management Controller in delivering commands for setup and configuration of arbitration.

Arbitration allows multiple Network Controller packages that are attached to the interface to be enabled to share the RXD lines to deliver packets to the Management Controller.

This operation is summarized as follows:

- Only one Network Controller at a time can transmit packets on the RXD lines of the interface.
- Network Controllers can accept commands for configuring and controlling arbitration for the RXD lines.

7.3 Hardware arbitration

To prevent two or more NC-SI packages from transmitting at the same time, a hardware-based arbitration scheme was devised to allow only one Network Controller package to drive the RX lines of the shared interface at any given time. This scheme uses a mechanism of passing messages (opcodes) between Network Controller packages to coordinate when a controller is allowed to transmit through the RBT interface.

7.3.1 General

Three conceptual modes of hardware arbitration exist: arbitration master assignment, normal operation, and bypass. After a package is initialized and has its Channel IDs assigned, it enters the arbitration master assignment mode. This mode assigns one package the role of an Arbitration Master (ARB_Master) that is responsible for initially generating a TOKEN opcode that is required for the normal operating mode. In the normal operating mode, the TOKEN opcode is passed from one package to the next in the ring. The package is allowed to use the shared RXD signals and transmit if the package has received the TOKEN opcode and has a packet to send.

Bypass mode allows hardware arbitration opcodes to pass through a Network Controller package before it is initialized. Bypass mode shall be in effect while hardware arbitration is disabled. Bypass mode shall be exited, and arbitration master assignment mode shall be entered when the hardware arbitration becomes enabled or re-enabled.

Hardware-based arbitration requires two additional pins (ARB_IN and ARB_OUT) on the Network Controller. The ARB_OUT pin of one package is connected to the ARB_IN pin of the next package to form a ring configuration, as illustrated in Figure 11. The timing requirements for hardware arbitration are designed to accommodate a maximum of four Network Controller packages. If the implementation consists of a single Network Controller package, the ARB_OUT pin may be connected to the ARB_IN pin on the same package, or may be left disconnected, in which case hardware arbitration should be disabled by using the Select Package command. This specification optionally supports reporting of Hardware arbitration implementation status and hardware arbitration status using the **Get Capabilities** command.

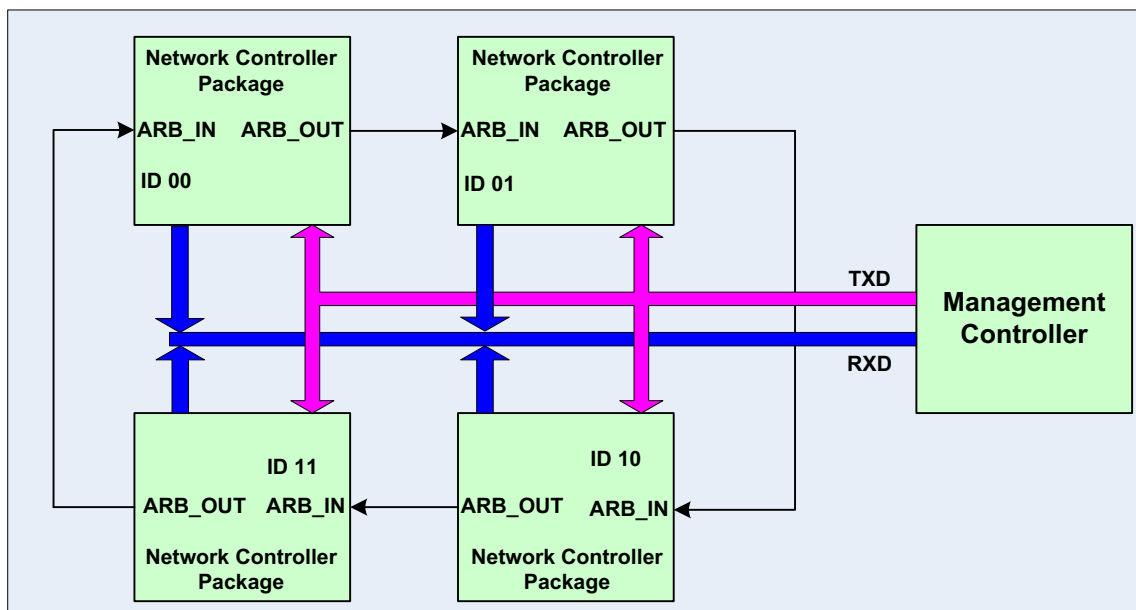


Figure 11 – Multiple Network Controllers in a ring format

Each Network Controller package sends out pulses on the ARB_OUT pin to create a series of symbols that form opcodes (commands) between Network Controllers. Each pulse is one clock wide and synchronized to REF_CLK. The hardware arbitration data bits follow the same timing specifications used

for the TXD and RXD data bits (see 10.2.6). The pulses are di-bit encoded to ensure that symbols are correctly decoded. The symbols have the values shown in Table 4.

While clause 7.3.2.1 allows for opcode to be truncated, it is recommended that the transmission of current opcode on ARB_OUT be completed if the HW arbitration mode is changed in the middle of an opcode transfer (or in the middle of a symbol).

Table 4 – Hardware arbitration di-bit encoding

Symbol Name	Encoded Value
Esync	11b
Ezero	00b
Eone	01b
Illegal symbol	10b

7.3.2 Hardware arbitration opcodes

The hardware-based arbitration feature has five defined opcodes: IDLE, TOKEN, FLUSH, XON, and XOFF. Each opcode starts with an Esync symbol and is followed by either E_{one} or E_{zero} symbols. The legal opcodes are listed in Table 5.

Table 5 – Hardware arbitration opcode format

Opcode	Format
IDLE	E _{sync} E _{zero} E _{zero} (110000b)
TOKEN	E _{sync} E _{one} E _{zero} (110100b)
FLUSH	E _{sync} E _{one} E _{one} E _{zero} E(Package_ID[2:0]) E _{zero} (11010100xxxxxx00b)
XOFF	E _{sync} E _{zero} E _{one} E _{zero} E _{zero} E _{zero} (110001000000b)
XON	E _{sync} E _{zero} E _{one} E _{one} E _{zero} E(Package_ID[2:0]) E _{zero} (1100010100uuuuuu00b)

7.3.2.1 Detecting truncated opcodes

A truncated opcode is detected when the number of clocks between E_{sync}s is less than the number of bits required for the opcode. Note that any additional bits clocked in after a legitimate opcode is detected do not indicate an error condition and are ignored until the next E_{sync}.

7.3.2.2 Handling truncated or illegal opcodes

When a Network Controller receives a truncated or illegal opcode, it should discard it.

7.3.2.3 Relationship of opcodes processing and driving the RX data lines

A Network Controller package shall take no more than T₉ REF_CLK times after receiving the last bit of the opcode to decode the incoming opcode and start generating the outgoing opcode. This time limit allows for decoding and processing of the incoming opcode under the condition that an outgoing opcode transmission is already in progress.

A package that has received a TOKEN and has packet data to transmit shall turn on its buffer and begin transmitting the packet data within T11 REF_CLK times of receiving the TOKEN, as illustrated in Figure 12. The package shall disable the RXD buffers before the last clock of the transmitted TOKEN.

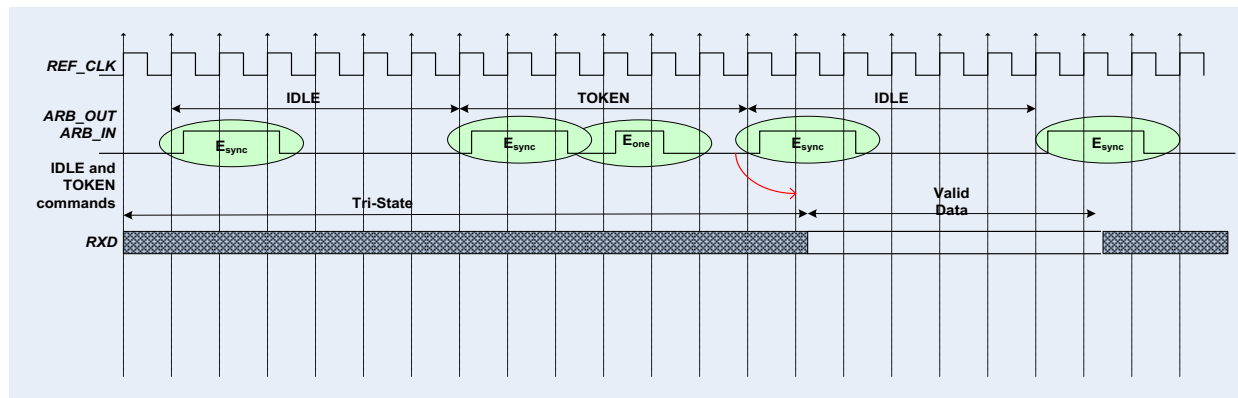


Figure 12 – Opcode to RXD relationship

7.3.3 Opcode operations

.

7.3.3.1 TOKEN opcode

When a TOKEN opcode is received, the Network Controller package may drive the RXD signals to send only one of the following items: a Pass-through packet, a command response, or an AEN. One [IEEE 802.3](#) PAUSE frame (XON or XOFF) may also be sent either before or after one of the previous packets, or on its own. While the Network Controller package is transmitting the data on the RXD signals of the interface, it shall generate IDLE opcodes on its ARB_OUT pin. Once a package completes its transmission, if any, it shall generate and send the TOKEN on its ARB_OUT pin.

7.3.3.2 IDLE opcode

A package that has no other opcode to send shall continuously generate IDLE opcodes. Typically, a received IDLE opcode indicates that the TOKEN is currently at another package in the ring. This opcode is also used in the ARB_Master assignment process (for details, see 7.3.5). An Idle opcode typically will also be generated when the package is transmitting on RBT

7.3.3.3 FLUSH opcode

A FLUSH opcode is used to establish an Arbitration Master for the ring when the package enters the Package Ready state or when the TOKEN is not received within the specified timeout, T8. This opcode is further explained in 7.3.5.

If the package receives a FLUSH opcode while it is in the middle of transmitting a packet onto NC-SI, it shall generate IDLE opcodes until the transmission is complete and then process the FLUSH opcode as described.

1873 **7.3.3.4 Flow Control opcodes**

1874 The XON and XOFF opcodes are used to manage the generation of [IEEE 802.3 PAUSE](#) frames on the
 1875 RBT interface. If the Network Controller supports flow control and flow control is enabled, the XOFF and
 1876 XON opcodes behave as described in this clause. If the Network Controller does not support flow control
 1877 or if flow control is not enabled, the Network Controller shall pass the opcodes to the next package.

1878 There may be a configuration where some NCs support flow control and others do not. In this
 1879 configuration, an NC sending an XOFF opcode may see the XOFF packet emission delayed by two or
 1880 more full size Pass-through packets, one for each package not supporting XOFF when it gets the token,
 1881 and one for the next package supporting XOFF before sending the XOFF packet. The NC is not required
 1882 to provide buffering to prevent packet loss in this configuration. No drop behavior should be expected by
 1883 an MC only if all NCs have flow control enabled.

1884 NOTE: There is a maximum amount of time that the Network Controller is allowed to maintain a PAUSE. For more
 1885 information, see 8.4.41.

1886 **7.3.3.4.1 XOFF opcode**

1887 A Network Controller package that becomes congested while receiving packets from the NC-SI shall
 1888 perform the following actions:

- 1889 • If it does not have a TOKEN, it sends the XOFF opcode to the next package.
- 1890 NOTE If it has the TOKEN and has not previously sent an XOFF frame for this instance of congestion, it
 1891 shall send a single XOFF frame (PAUSE frame with a pause time of 0xFFFF) and will not generate an
 1892 XOFF opcode.
- 1893 • A package may also regenerate an XOFF frame or opcode if it is still congested and determines
 1894 that the present PAUSE frame is about to expire.

1895 When a package on the ring receives an XOFF opcode, it shall perform one of the following actions:

- 1896 • If it does not have a TOKEN opcode, it passes the XOFF opcode to the next package in the
 1897 ring.
- 1898 • If it has the TOKEN, it shall send an XOFF frame (PAUSE frame with a pause time of 0xFFFF)
 1899 and will not regenerate the XOFF opcode. If it receives another XOFF opcode while sending the
 1900 XOFF frame or a regular network packet, it discards the received XOFF opcode.

1901 **7.3.3.4.2 XON opcode**

1902 XON frames (PAUSE frame with a pause time of 0x0000) are used to signal to the Management
 1903 Controller that the Network Controller packages are no longer congested and that normal traffic flow can
 1904 resume. XON opcodes are used between the packages to coordinate XON frame generation. The
 1905 package ID is included in this opcode to provide a mechanism to verify that every package is not
 1906 congested before sending an XON frame to the Management Controller.

1907 The XON opcode behaves as follows:

- 1908 • When a package is no longer congested, it generates an XON opcode with its own Package ID.
 1909 This puts the package into the 'waiting for its own XON' state.
- 1910 • A package that receives the XON opcode takes one of the following actions:
 - 1911 – If it is congested, it replaces the received XON opcode with the IDLE opcode. This action
 1912 causes the XON opcode to be discarded. Eventually, the congested package generates its
 1913 own XON opcode when it exits the congested state.
 - 1914 – If the package is not congested and is not waiting for the XON opcode with own Package
 1915 ID, it forwards the received XON opcode to the next package in the ring.

- 1916 – If the received XON opcode contains the package's own Package ID, the opcode should
1917 be discarded.
- 1918 – If the package is not congested and is waiting for its own XON opcode, it performs one of
1919 the following actions:
- 1920 • If it receives an XON opcode with a Package ID that is higher than its own, it replaces
1921 the XON opcode with its own Package ID.
 - 1922 • If it receives an XON opcode with a Package ID lower than its own, it passes that
1923 XON opcode to the next package and it exits the 'waiting for its own XON' state.
 - 1924 • If it receives an XON opcode with the Package ID equal to its own, it sends an XON
1925 frame on the NC-SI when it receives the TOKEN opcode and exits the 'waiting for its
1926 own XON' state.
- 1927 NOTE More than one XON opcode with the same Package ID can be received
1928 while waiting for the TOKEN and while sending the XON frame. These additional XON
1929 opcodes should be discarded.
- 1930 • If a package originates an XON opcode but receives an XOFF opcode, it terminates its XON
1931 request so that it does not output an XON frame when it receives the TOKEN.
- 1932 NOTE This behavior is not likely to occur because the Management Controller will be in the
1933 Pause state at this point.
- 1934 • A package that generated an XON opcode may receive its own XON opcode back while it has
1935 the TOKEN opcode. In this case, it may send a regular packet (Pass-through, command
1936 response, or AEN) to the Management Controller (if it has one to send), an XON frame, or both.

1937 7.3.4 Bypass mode

1938 When the Network Controller package is in bypass mode, data received on the ARB_IN pin is redirected
1939 to the ARB_OUT pin within the specified clock delay. This way, arbitration can continue between other
1940 devices in the ring.

1941 A package in bypass mode shall take no more than T10 REF_CLK times to forward data from the
1942 ARB_IN pin to the ARB_OUT pin. The transition in and out of bypass mode may result in a truncated
1943 opcode.

1944 A Network Controller package enters bypass mode immediately upon power up and transitions out of this
1945 mode after the Network Controller completes its startup/initialization sequence.

1946 7.3.5 Hardware arbitration startup

1947 Hardware arbitration startup works as follows:

- 1948 1) All the packages shall be in bypass mode within T_{pwrz} seconds of NC-SI power up.
- 1949 2) As each package is initialized, it shall continuously generate FLUSH opcodes with its own
1950 Package ID.
- 1951 3) The package then participates in the ARB_MSTR assignment process described in the
1952 following clause.

1953 7.3.6 ARB_MSTR assignment

1954 ARB_MSTR assignment works as follows:

- 1955 1) When a package receives a FLUSH opcode with a Package ID numerically smaller than its
1956 own, it shall forward on the received FLUSH opcode. If the received FLUSH opcode's Package
1957 ID is numerically larger than the local Package ID, the package shall continue to send its

- 1958 FLUSH opcode with its own Package ID. When a package receives a FLUSH opcode with its
 1959 own Package ID, it becomes the master of the ring (ARB_MSTR).
- 1960 2) The ARB_MSTR shall then send out IDLE opcodes until it receives an IDLE opcode.
- 1961 3) Upon receiving the IDLE opcode, the ARB_MSTR shall be considered to be in possession of
 1962 the TOKEN opcode (see 7.3.3.1).
- 1963 4) If the package receives a FLUSH opcode while it is in the middle of transmitting a packet onto
 1964 NC-SI, it shall generate IDLE opcodes until the transmission is complete and then process the
 1965 FLUSH opcode as described.

1966 **7.3.7** Token timeout mechanism

1967 Each Network Controller package that supports hardware-based arbitration control shall implement a
 1968 timeout mechanism in case the TOKEN opcode is not received. When a package has a packet to send, it
 1969 starts its timer. If it does not receive a TOKEN prior to the TOKEN timeout, the package shall send a
 1970 FLUSH opcode. This restarts the arbitration process.

1971 The timer may be programmable depending on the number of packages in the ring. The timeout value is
 1972 designed to accommodate up to four packages, each sending the largest packet (1536 bytes) plus
 1973 possible XON or XOFF frame transmission and opcode processing time. The timeout shall be no fewer
 1974 than T8 cycles of the REF_CLK.

1975 **7.3.8** Timing considerations

1976 The ARB_OUT and ARB_IN pins shall follow the timing specifications outlined in Clause 0.

1977 To improve the efficiency of the multi-drop NC-SI, TOKEN opcode generation may overlap the Inter
 1978 Packet Gap (IPG) defined by the [802.3](#) specification, as shown in Figure 13. The TOKEN opcode shall be
 1979 sent no earlier than the last T13 REF_CLK cycles of the IPG.

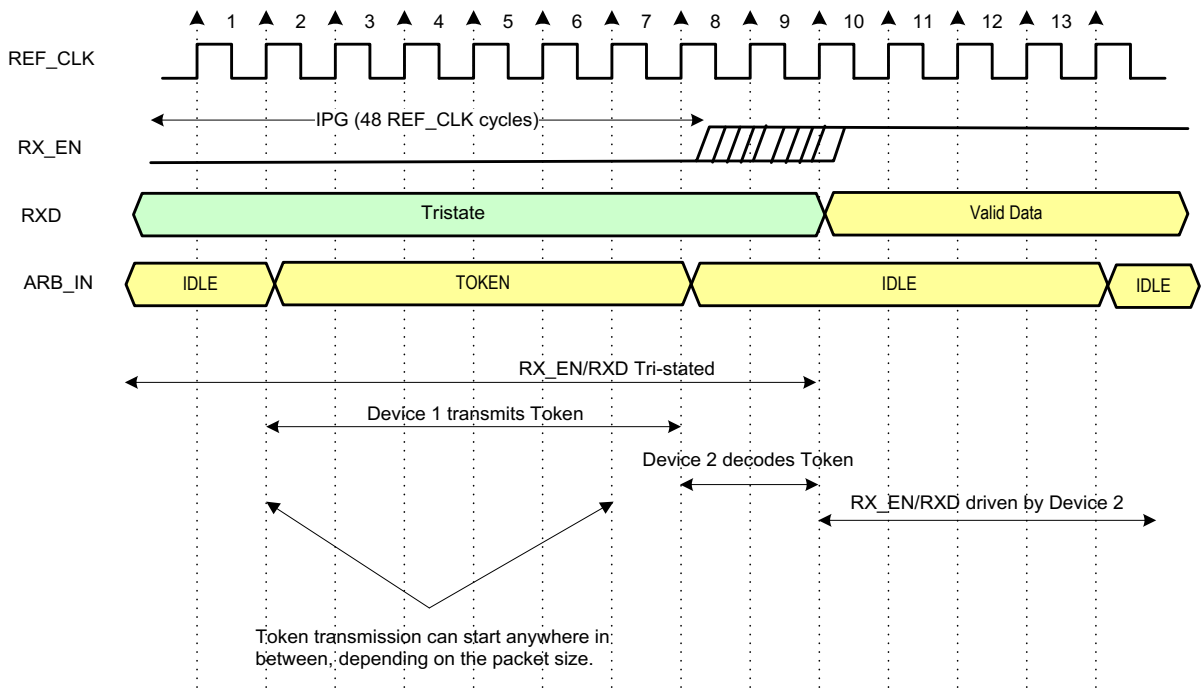


Figure 13 – Example TOKEN to transmit relationship

7.3.9 Example hardware arbitration state machine

The state machine diagram shown in Figure 14 is provided as a guideline to help illustrate the startup process and opcode operations described in the preceding clauses.

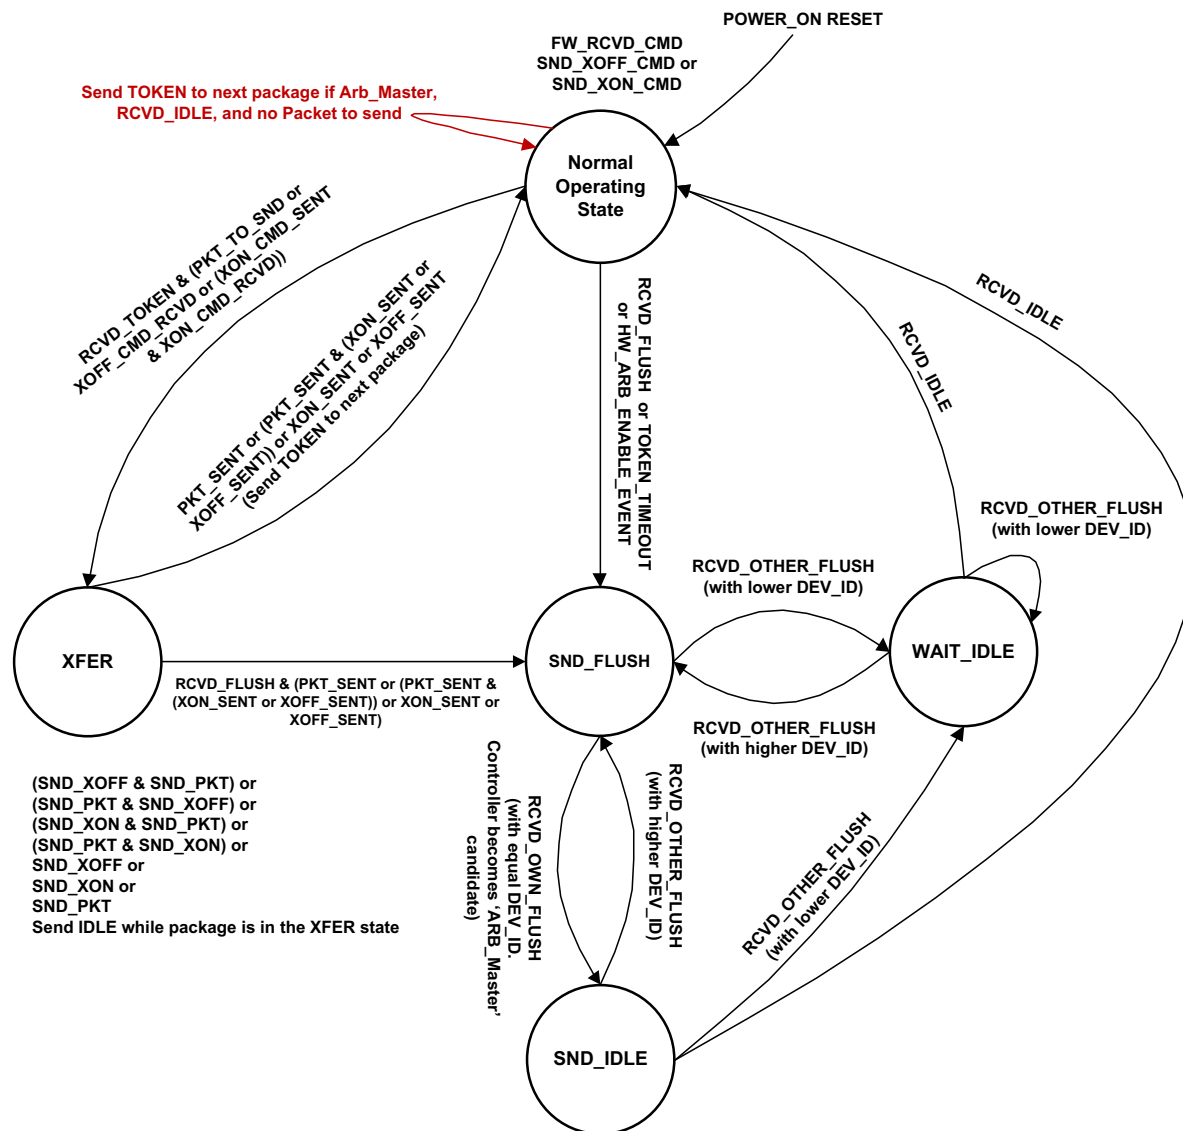


Figure 14 – Hardware arbitration state machine

1987 The states and events shown in Figure 14 are described in Table 6 and Table 7, respectively.

1988 **Table 6 – Hardware arbitration states**

State	Action
Normal Operating State	<p>This state is the normal operating state for hardware arbitration. The following actions happen in this state:</p> <ul style="list-style-type: none"> • FW_RCVD_CMD: Forward received command. As opcodes are received and acted upon, the resulting opcode is sent to the next package. For example, the TOKEN opcode is received, and no packet data is available to send, so the TOKEN opcode is sent to the next package in the ring. • SND_XOFF_CMD: Send the XOFF opcode to the next package. This action happens when the specific conditions are met as described in 7.3.3. • SND_XON_CMD: Send the XON opcode to the next package. This action happens when the specific conditions are met as described in 7.3.3. • If the Network Controller is ARB_Master, it generates the TOKEN opcode upon receiving an IDLE opcode at the end of the FLUSH process. • The RXD lines will be in a high-impedance condition in this state.
XFER	<p>In this state, data is sent on the RXD lines. This data will be a Pass-through packet, response packet, XON (Pause Off) packet, XOFF (Pause On) packet, or AEN. (An XON or XOFF packet can be sent in addition to a Pass-through packet, response packet, or AEN.) IDLE opcodes are sent to the next package while the device is in the XFER state.</p> <p>The following actions happen in this state:</p> <ul style="list-style-type: none"> • SND_XON: Transmit an XON frame (Pause Off) to the Management Controller. • SND_XOFF: Transmit an XOFF frame (Pause On) to the Management Controller. • SND_PKT: Transmit a Pass-through packet, response packet, or AEN to the Management Controller. • The TOKEN opcode is sent to the next package upon completion of the transfer.
SND_FLUSH	<p>This state is the entry point for determining the ARB_Master among the packages. In this state, the FLUSH opcode is continuously sent. This state is exited upon receiving a FLUSH opcode that has a DEV_ID that is equal to or lower than the package's own DEV_ID.</p>
SND_IDLE	<p>This is the final state for determining the ARB_Master, entered when a device's own FLUSH opcode is received. In this state, the IDLE opcode is continuously sent.</p>
WAIT_IDLE	<p>This state is entered when a FLUSH command is received from another package with a lower Device ID. When an IDLE opcode is received, the ARB_Master has been determined and the device transitions to the Normal Operating State.</p>

1989

Table 7 – Hardware arbitration events

Event	Description
RCVD_TOKEN	A TOKEN opcode was received, or the arbitration was just completed and won by this package.
RCVD_IDLE	An IDLE opcode was received.
XOFF_SENT	The Pause On frame was sent on the RXD interface.
XON_SENT	The Pause Off frame was sent on the RXD interface.
PKT_TO_SND	The Network Controller package has a Pass-through packet, command response packet, XON (Pause Off) frame, XOFF (Pause On) frame, or AEN to send.
XON_CMD_RCVD	A package received an XON opcode with its own Package ID.
XOFF_CMD_RCVD	An XOFF opcode was received.
XON_CMD_SENT	A package sent an XON opcode with its own Package ID.
RCVD_FLUSH	A FLUSH opcode was received.
TOKEN_TIMEOUT	The timeout limit expired while waiting for a TOKEN opcode.
HW_ARB_ENABLE_EVENT	This event begins ARB_MSTR assignment. This event occurs just after the Network Controller package initializes or when hardware arbitration is re-enabled through the Select Package command.
RCVD_OTHER_FLUSH	A package received a FLUSH opcode with a Package ID other than its own.
RCVD_OWN_FLUSH	A package received a FLUSH opcode with a Package ID equal to its own.

1990 7.4 Command-based arbitration

1991 If hardware arbitration is not being used, the **Select Package** and **Deselect Package** commands shall be
 1992 used to control which Network Controller package can transmit on the RXD lines. Because only one
 1993 Network Controller package is allowed to transmit on the RXD lines, the Management Controller shall
 1994 only have one package in the selected state at any given time. For more information, see 8.4.5 and 8.4.7.

1995 8 Packet definitions

1996 8.1 NC-SI packet encapsulation

1997 The RBT interface is an Ethernet interface adhering to the standard [IEEE 802.3](#) Ethernet frame format.
 1998 Whether or not the Network Controller accepts runt packets is unspecified.

1999 As shown in Figure 15, this L2, or data link layer, frame format encapsulates all NC-SI packets, including
 2000 Pass-through, command, and response packets, as the L2 frame payload data by adding a 14-byte
 2001 header to the front of the data and appending a 4-byte Frame Check Sequence (FCS) to the end.

2002 NC-SI Control Packets shall not include any VLAN tags. NC-SI Pass-through packets may include an
 2003 802.1Q VLAN tag.




Figure 15 – Ethernet frame encapsulation of NC-SI packet data without VLAN tag

8.1.1 Ethernet frame header

The Management Controller shall format the 14-byte Ethernet frame header so that when it is received, it shall be formatted in the big-endian byte order shown in Table 8.

Channels shall accept Pass-through packets that meet the [IEEE 802.3](#) frame requirements.

Table 8 – Ethernet Header Format

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..03	DA ₅ = 0xFF	DA ₄ = 0xFF	DA ₃ = 0xFF	DA ₂ = 0xFF
04..07	DA ₁ = 0xFF	DA ₀ = 0xFF	SA ₅	SA ₄
08..11	SA ₃	SA ₂	SA ₁	SA ₀
12..13	Ethertype = 0x88F8 (DMTF NC-SI)			

8.1.1.1 Destination Address (DA)

Bytes 0–5 of the header represent bytes 5–0 of the Ethernet Destination Address field of an L2 header.

The channel is not assigned a specific MAC address and the contents of this field are not interpreted as a MAC address by the Management Controller or the Network Controller. However, the DA field in all NC-SI Control Packets shall be set to the broadcast address (FF:FF:FF:FF:FF:FF) for consistency.

If the Network Controller receives a Control Packet with a Destination Address other than FF:FF:FF:FF:FF:FF, the Network Controller may elect to accept the packet, drop it, or return a response packet with an error response/reason code.

8.1.1.2 Source Address (SA)

Bytes 6–11 of the header represent bytes 5–0 of the Ethernet Source MAC Address field of the Ethernet header. The contents of this field may be set to any value. The Network **Controller should** use FF:FF:FF:FF:FF:FF as the source address for NC-SI Control Packets that it generates.

8.1.1.3 Ethertype

The final two bytes of the header, bytes 12..13, represent bytes 1..0 of the Ethertype field of the Ethernet header. For NC-SI Control Packets, this field shall be set to a fixed value of 0x88F8 as assigned to NC-SI by the IEEE. This value allows NC-SI Control Packets to be differentiated from other packets in the overall packet stream.

8.1.2 Frame Check Sequence

The Frame Check Sequence (FCS) shall be added at the end of the frame to provide detection of corruption of the frame. Any frame with an invalid FCS shall be discarded.

8.1.3 Data length

NC-SI Commands, Responses, and AENs do not carry any VLAN tag. NC-SI Commands, Responses and AENs shall have a payload data length between 46 and 1500 octets (bytes). This complies with the 802.3 specification. This means that the length of Ethernet frame shown in Figure 15 is between 64 octets (for a payload of 46 octets) and 1518 octets (for a payload with 1500 octets).

Pass-through packets also follow the 802.3 specification. The maximum payload size is 1500 octets; the minimum payload size shall be 42 octets when 802.1Q (VLAN) tag is present and 46 octets when the 802.1Q tag is not present. The Layer-2 Ethernet frame for an 802.1Q tagged frame shall be between 64 octets (for a payload of 42 octets) and 1522 octets (for a payload with 1500 octets). For Pass-through packets that are not 802.1Q tagged, the minimum Layer-2 Ethernet frame size is 64 octets (for a payload of 46 octets) and the maximum Layer-2 Ethernet frame size is 1518 octets (for a payload with 1500 octets).

8.2 Control Packet data structure

Each NC-SI Control Packet is made up of a 16-byte packet header and a payload section whose length is specific to the packet type.

8.2.1 Control Packet header

The 16-byte Control Packet header is used in command, response, and AEN packets, and contains data values intended to allow the packet to be identified, validated, and processed. The packet header is in big-endian byte order, as shown in Table 9.

Table 9 – Control Packet header format

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..03	MC ID	Header Revision	Reserved	IID
04..07	Control Packet Type	Ch. ID	Flags	Payload Length
08..11	Reserved			
12..15	Reserved			

8.2.1.1 Management Controller ID

In Control Packets, this 1-byte field identifies the Management Controller issuing the packet. For this version of the specification, Management Controllers should set this field to 0x00 (zero). This implies that only one management controller is supported for accessing the NC via NC-SI at any given time, Network Controllers responding to command packets should copy the Management Controller ID field from the command packet header into the response packet header. For AEN packets, this field should be copied from the parameter that was set using the AEN Enable command.

2060 8.2.1.2 Header revision

2061 This 1-byte field identifies the version of the Control Packet header in use by the sender. For this version
2062 of the specification, the header revision is 0x01.

2063 8.2.1.3 Instance ID (IID)

2064 This 1-byte field contains the IID of the command and associated response. The Network Controller can
2065 use it to differentiate retried commands from new instances of commands. The Management Controller
2066 can use this value to match a received response to the previously sent command. For more information,
2067 see 6.2.2.2.

2068 8.2.1.4 Control Packet type

2069 This 1-byte field contains the Identifier that is used to identify specific commands and responses, and to
2070 differentiate AENs from responses. Each NC-SI command is assigned a unique 7-bit command type
2071 value in the range 0x00 . . 0x60. The proper response type for each command type is formed by setting
2072 the most significant bit (bit 7) in the original 1-byte command value. This allows for a one-to-one
2073 correspondence between 96 unique response types and 96 unique command types.

2074 8.2.1.5 Channel ID

2075 This 1-byte field contains the Network Controller Channel Identifier. The Management Controller shall set
2076 this value to specify the package and internal channel ID for which the command is intended.

2077 In a multi-drop configuration, all commands are received by all NC-SI Network Controllers present in the
2078 configuration. The Channel ID is used by each receiving Network Controller to determine if it is the
2079 intended recipient of the command. In Responses and AENs, this field carries the ID of the channel from
2080 which the response of AEN was issued.

2081 8.2.1.6 Payload length

2082 This 12-bit field contains the length, in bytes, of any payload data present in the command or response
2083 frame following the NC-SI packet header. This value does not include the length of the NC-SI Control
2084 Packet Header, the checksum value, or any padding that might be present.

2085 8.2.1.7 Flags

2086 Bit 0: Poll Indication: If this bit is set, it indicates that this command instance is polling on a previously sent
2087 command that was responded with a "Delayed Response" response code. This bit is relevant only for
2088 commands and not for responses or AENs.

2089 Bits 3:1: Reserved

2090 8.2.1.8 Reserved

2091 These fields are reserved for future use and should be written as zeros and ignored when read.

2092 8.2.2 Control Packet payload

2093 The NC-SI packet payload may contain zero or more defined data values depending on whether the
2094 packet is a command or response packet, and on the specific type. The NC-SI packet payload is always
2095 formatted in big-endian byte order, as shown in Table 10.

2096

Table 10 – Generic example of Control Packet payload

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..03	Data0 ₃	Data0 ₂	Data0 ₁	Data0 ₀
04..07	Data1 ₇	Data1 ₆	Data1 ₅	Data1 ₄
08..11	Data1 ₃	Data1 ₂	Data1 ₁	Data1 ₀
..				
...	DataN-1 ₄	DataN-1 ₃	DataN-1 ₂	DataN-1 ₁
...	DataN-1 ₀	Payload Pad (as required)		
...	Checksum			
...	Ethernet Packet Pad (as required)			

2097 **8.2.2.1 Data**

2098 As shown in Table 10, the bytes following the NC-SI packet header may contain payload data fields of
 2099 varying sizes, and which may be aligned or require padding. In the case where data is defined in the
 2100 payload, all data-field byte layouts (Data0–Data-1) shall use big-endian byte ordering with the most
 2101 significant byte of the field in the lowest addressed byte position (that is, coming first).

2102 **8.2.2.2 Payload pad**

2103 If the payload is present and does not end on a 32-bit boundary, one to three padding bytes equal to
 2104 0x00 shall be present to align the checksum field to a 32-bit boundary.

2105 **8.2.2.3 Checksum**

2106 This 4-byte field contains the 32-bit checksum compensation value that may be included in each
 2107 command and response packet by the sender of the packet. When it is implemented, the checksum
 2108 compensation shall be computed as the 2's complement of the checksum, which shall be computed as
 2109 the 32-bit unsigned sum of the NC-SI packet header and NC-SI packet payload interpreted as a series of
 2110 16-bit unsigned integer values. A packet receiver supporting packet checksum verification shall use the
 2111 checksum compensation value to verify packet data integrity by computing the 32-bit checksum described
 2112 above, adding to it the checksum compensation value from the packet, and verifying that the result is 0.

2113 Verification of non-zero NC-SI packet checksum values is optional. An implementation may elect to
 2114 generate the checksums and may elect to verify checksums that it receives. The checksum field is
 2115 generated and handled according to the following rules:

- 2116 • A checksum field value of all zeros specifies that a header checksum is not being provided for
 2117 the NC-SI Control Packet, and that the checksum field value shall be ignored when processing
 2118 the packet.
- 2119 • If the originator of an NC-SI Control Packet is not generating a checksum, the originator shall
 2120 use a value of all zeros for the header checksum field.
- 2121 • If a non-zero checksum field is generated for an NC-SI Control Packet, that header checksum
 2122 field value shall be calculated using the specified algorithm.
- 2123 • All receivers of NC-SI Control Packets shall accept packets with all zeros as the checksum
 2124 value (provided that other fields and the CRC are correct).

- 2125 • The receiver of an NC-SI Control Packet may reject (silently discard) a packet that has an
2126 incorrect non-zero checksum.
- 2127 • The receiver of an NC-SI Control Packet may ignore any non-zero checksums that it receives
2128 and accept the packet, even if the checksum value is incorrect (that is, an implementation is not
2129 required to verify the checksum field).
- 2130 • A controller that generates checksums is not required to verify checksums that it receives.
- 2131 • A controller that verifies checksums is not required to generate checksums for NC-SI Control
2132 Packets that it originates.

2133 8.2.2.4 Ethernet packet pad

2134 Per [IEEE 802.3](#), all Ethernet frames shall be at least 64 bytes in length, from the DA through and
2135 including FCS. For NC-SI packets, this requirement applies to the Ethernet header and payload, which
2136 includes the NC-SI Control Packet header and payload. Most NC-SI Control Packets are less than the
2137 minimum Ethernet frame payload size of 46 bytes in length and require padding to comply with
2138 [IEEE 802.3](#).

2139 8.2.3 Command packet payload

2140 Command packets have no common fixed payload format.

2141 8.2.4 Response packet payload

2142 Unlike command packets that do not necessarily contain payload data, all response packets carry at least
2143 a 4-byte payload. This default payload carries the response codes and reason codes (described in
2144 8.2.4.1) that provide status on the outcome of processing the originating command packet and is present
2145 in all response packet payload definitions.

2146 The default payload occupies bytes 00..03 of the response packet payload, with any additional
2147 response-packet-specific payload defined to follow starting on the next word. All response packet payload
2148 fields are defined with big-endian byte ordering, as shown in Table 11.

2149 **Table 11 – Generic example of Response packet payload format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..03	Response Code		Reason Code	
..
...	DataN-1 ₄	DataN-1 ₃	DataN-1 ₂	DataN-1 ₁
...	DataN-1 ₀	Word Pad (as required)		
...	Checksum			
...	Ethernet Packet Pad (as required)			

2150 8.2.4.1 Response Packet in case of Delayed Response Code

2151 If a response includes a “Delayed Response” Code, then the response does not contain the payload of
2152 the original response. The Delayed Response shall contain a payload of a single word (uint16) including
2153 the recommended next polling time in milliseconds. If no polling time estimate is available, then the
2154 recommended next polling time shall be set to 0x0000.

Table 12 – Generic example of Delayed Response packet payload

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..03	Response Code = 0x0004		Reason Code = 0x0000	
04..07	Reserved		Next Polling time	
08...11	Checksum			
...	Ethernet Packet Pad (as required)			

8.2.5 Response codes and reason codes

8.2.5.1 General

Response codes and reason codes are status values that are returned in the responses to NC-SI commands. The response code values provide a general categorization of the status being returned. The reason code values provide additional detail related to a particular response code.

Response codes and reason codes are divided into numeric ranges that distinguish whether the values represent standard codes that are defined in this specification or are vendor/OEM-specific values that are defined by the vendor of the controller.

The response code is a 2-byte field where values from 0x00 through 0x7F are reserved for definition by this specification. Values from 0x80 through 0xFF are vendor/OEM-specific codes that are defined by the vendor of the controller.

The reason code is a 2-byte field. The ranges of values are defined in Table 13.

Table 13 – Reason code ranges

MS-byte	LS-byte	Description
00h	0x00–0x7F	Standard generic reason codes This range of values for the lower byte is used for reason codes that are not specific to a particular command but can be used as reason codes in responses for any command. The values in this range are reserved for definition by this specification.
	0x80–0xFF	Vendor/OEM generic reason codes This range of values for the lower byte is used for reason codes that are not specific to a particular command but can be used as reason codes in responses for any command. Values in this range are defined by the vendor of the controller.
Command Number Note: This means that Command	0x00–0x7F	Standard command-specific reason codes This range of values for the lower byte is used for reason codes that are specific to a particular command. The upper byte holds the value of the command for which the reason code is defined. The values in this range are reserved for definition by this specification.

MS-byte	LS-byte	Description
Number 00 cannot have any command-specific reason codes.	0x80–0xFF	Vendor/OEM command-specific reason codes This range of values for the lower byte is used for reason codes that are specific to a particular command. The upper byte holds the value of the command for which the reason code is defined. Values in this range are defined by the vendor of the controller.

2170 8.2.5.2 Response code and reason code values

2171 The standard response code values are defined in Table 14, and the standard reason code values are
 2172 defined in Table 15. Command-specific values, if any, are defined in the clauses that describe the
 2173 response data for the command. Unless otherwise specified, the standard reason codes may be used in
 2174 combination with any response code. There are scenarios where multiple combinations of response and
 2175 reason code values are valid. Unless otherwise specified, an implementation may return any valid
 2176 combination of response and reason code values for the condition.

2177 **Table 14 – Standard response code values**

Value	Description	Comment
0x0000	Command Completed	Returned for a successful command completion. When this response code is returned, the reason code shall be 0x0000 as described in Table 15
0x0001	Command Failed	Returned to report that a valid command could not be processed or failed to complete correctly
0x0002	Command Unavailable	Returned to report that a command is temporarily unavailable for execution because the controller is in a transient state, busy condition, or in need of external intervention.
0x0003	Command Unsupported	Returned to report that a command is not supported by the implementation. The reason code “Unknown / Unsupported Command Type” should be returned along with this response code for all unsupported commands.
0x0004	Delayed Response	Returned to report that the command was accepted, and the NC started to handle it, but it cannot respond within T5 seconds with a final answer. When this response code is provided, the reason code shall be 0x0000
0x8000–0xFFFF	Vendor/OEM-specific	Response codes defined by the vendor of the controller

2178 **Table 15 – Standard Reason Code Values**

Value	Description	Comment
0x0000	No Error/No Reason Code	When used with the Command Completed response code, indicates that the command completed normally. Otherwise this value indicates that no additional reason code information is being provided.
0x0001	Interface Initialization Required	Returned for all commands except Select/Deselect Package commands when the channel is in the Initial State, until the channel receives a Clear Initial State command

Value	Description	Comment
0x0002	Parameter Is Invalid, Unsupported, or Out-of-Range	Returned when a received parameter value is outside of the acceptable values for that parameter
0x0003	Channel Not Ready	Returned when the channel is in a transient state in which it is unable to process commands normally
0x0004	Package Not Ready	Returned when the package and channels within the package are in a transient state in which normal command processing cannot be done
0x0005	Invalid payload length	Returned when the payload length in the command is incorrect for the given command
0x0006	Information not available	Returned when the channel is unable to provide response data to a valid supported command.
0x0007	Intervention Required	May be returned for all commands, except for Select and Deselect Package, when the Package is not ready and requires intervention to restore its operational state. When this code is returned, the NC does not check if the command is otherwise valid and the defined response is not returned.
0x0008	Link Command Failed-Hardware Access Error	Returned when PHY R/W access fails on Link commands
0x0009	Command Timeout	Command execution has exceeded the allocated T5 time
0x7FFF	Unknown / Unsupported Command Type	Returned when the command type is unknown or unsupported. This reason code shall only be used when the response code is 0x0003 (Command Unsupported) as described in Table 14.
0x8000-0xFFFF	OEM Reason Code	Vendor-specific reason code defined by the vendor of the controller

2179 8.2.6 AEN packet format

2180 AEN packets shall follow the general packet format of Control Packets, with the IID field set to 0 because,
 2181 by definition, the Management Controller does not send a response packet to acknowledge an AEN
 2182 packet. The Control Packet Type field shall have the value 0xFF. The originating Network Controller shall
 2183 fill in the Channel ID (Ch. ID) field with its own ID to identify itself as the source of notification. The AEN
 2184 Type field contains the identifier of what condition caused the generation of the AEN packet. Table 16
 2185 represents the AEN packet format to be used for AENs defined in this specification.

2186
2187 **Table 16 – AEN packet format**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..03	MC ID = 0x0	0x01	Reserved	IID = 0x0
04..07	Control Packet Type = 0xFF	Originating Ch. ID	Reserved	Payload Length
08..11	Reserved			
12..15	Reserved			

16..19	Reserved	AEN Type
--------	----------	----------

Error! Reference source not found..

Table 17 – AEN Type Ranges

Value	AEN Type Allocation
0x0..0x6F	Specification-defined AENs see clause Error! Reference source not found.. , all others are Reserved
0x70..0x7F	Transport-specific AENs
0x80..0xFF	OEM-specific AENs

8.2.7 Single OEM AEN packet format

OEM AEN packets shall conform to the format shown in Table 18 below for NCs that only support AENs using a single OEM identifier including NCs that implement spec version 1.1 and lower.

Table 18 – OEM AEN packet format

Bits				
Bytes	31..24	23..16	15..08	07..00
00..03	MC ID = 0x0	0x01	Reserved	IID = 0x0
04..07	Control Packet Type = 0xFF	Originating Ch. ID	Reserved	Payload Length
08..11	Reserved			
12..15	Reserved			
16..19	Reserved			AEN Type
20..23	OPTIONAL AEN Data			
24..27	Checksum			

8.2.8 Multiple OEMs AEN packet format

OEM AEN packets shall conform to the format shown in Table 19 below for NCs that support multiple OEM AENs and implement the Query and Set OEM AEN command.

Table 19 – Multiple OEMs AEN packet format

Bits				
Bytes	31..24	23..16	15..08	07..00
00..03	MC ID = 0x0	0x01	Reserved	IID = 0x0
04..07	Control Packet Type = 0xFF	Originating Ch. ID	Reserved	Payload Length

08..11	Reserved		
12..15	Reserved		
16..19	Reserved	Multi field	AEN Type
	Manufacturer ID (IANA)		
20..23	OPTIONAL AEN Data		
24..27	Checksum		

2199

2200 **8.2.8.1** Multi field

2201 This field has a value of 0x01 to indicate the AEN contains a Manufacturer ID (IANA).

2202 **8.3 Control Packet type definitions**

2203 Command packet types are in the range of 0x00 to 0x7F. **Error! Reference source not found.**
 2204 describes each command, its corresponding response, and the type value for each. **Error! Reference**
 2205 **source not found.** includes commands addressed to either a package or a channel. The commands
 2206 addressed to a package are highlighted with gray background. PLDM and OEM-specific commands
 2207 carried over NC-SI may be package specific or channel specific or both.

2208 Mandatory (M), Optional (O), and Conditional (C) refer to command support requirements for the Network
 2209 Controller.

2210 Ethernet (E), Fibre Channel (FC) and InfiniBand (IB) columns under the Fabric Implementation heading
 2211 refer to the specific requirements of the NC implementing the network fabric type configured on the
 2212 channel.

2213

2214

Table 20 - Command and response types

Control Packet Type	Command Name	Description	Response Packet Type	Fabric Implementation		
				E	FC	IB
0x00	Clear Initial State	Used by the Management Controller to acknowledge that the Network Controller is in the Initial State	0x80	M	M	M
0x01	Select Package	Used to explicitly select a controller package to transmit packets through the NC-SI interface	0x81	M	M	M
0x02	Deselect Package	Used to explicitly instruct the controller package to stop transmitting packets through the NC-SI interface	0x82	M	M	M

Control Packet Type	Command Name	Description	Response Packet Type	Fabric Implementation		
				E	FC	IB
0x03	Enable Channel	Used to enable the NC-SI channel and to cause the forwarding of bidirectional Management Controller packets to start	0x83	M	M	M
0x04	Disable Channel	Used to disable the NC-SI channel and to cause the forwarding of bidirectional Management Controller packets to cease	0x84	M	M	M
0x05	Reset Channel	Used to synchronously put the Network Controller back to the Initial State	0x85	M	M	M
0x06	Enable Channel Network TX	Used to explicitly enable the channel to transmit Pass-through packets onto the network	0x86	M	N/A	N/A
0x07	Disable Channel Network TX	Used to explicitly disable the channel from transmitting Pass-through packets onto the network	0x87	M	N/A	N/A
0x08	AEN Enable	Used to control generating AENs	0x88	C	C	C
0x09	Set Link	Used during OS absence to force link settings, or to return to auto-negotiation mode	0x89	M	N/A	N/A
0x0A	Get Link Status	Used to get current link status information	0x8A	M	N/A	N/A
0x0B	Set VLAN Filter	Used to program VLAN IDs for VLAN filtering	0x8B	M	N/A	N/A
0x0C	Enable VLAN	Used to enable VLAN filtering of Management Controller RX packets	0x8C	M	N/A	N/A
0x0D	Disable VLAN	Used to disable VLAN filtering	0x8D	M	N/A	N/A
0x0E	Set MAC Address	Used to configure and enable unicast and multicast MAC address filters	0x8E	M	N/A	N/A
0x10	Enable Broadcast Filter	Used to enable selective broadcast packet filtering	0x90	M	N/A	N/A

Control Packet Type	Command Name	Description	Response Packet Type	Fabric Implementation		
				E	FC	IB
0x11	Disable Broadcast Filter	Used to disable all broadcast packet filtering, and to enable the forwarding of all broadcast packets	0x91	M	N/A	N/A
0x12	Enable Global Multicast Filter	Used to enable selective multicast packet filtering	0x92	C	N/A	N/A
0x13	Disable Global Multicast Filter	Used to disable all multicast packet filtering, and to enable forwarding of all multicast packets	0x93	C	N/A	N/A
0x14	Set NC-SI Flow Control	Used to configure IEEE 802.3 flow control on RBT	0x94	O	N/A	N/A
0x15	Get Version ID	Used to get controller-related version information	0x95	M	M	M
0x16	Get Capabilities	Used to get optional functions supported by the NC-SI	0x96	M	M	M
0x17	Get Parameters	Used to get configuration parameter values currently in effect on the controller	0x97	M	M	M
0x18	Get Controller Packet Statistics	Used to get current packet statistics for the Ethernet Controller	0x98	O	N/A	O
0x19	Get NC-SI Statistics	Used to request the packet statistics specific to the NC-SI	0x99	O	O	O
0x1A	Get NC-SI Pass-through Statistics	Used to request NC-SI Pass-through packet statistics	0x9A	O	N/A	O
0x1B	Get Package Status	Used to get current status of the package.	0x9B	O	O	O
0x25	Get NC Capabilities and Settings	Used to request device configuration information and capabilities	0xA5			
0x26	Set NC Configuration	Used to configure device interfaces	0xA6			
0x27	Get PF Assignment	Used to request Function assignment information	0xA7			

Control Packet Type	Command Name	Description	Response Packet Type	Fabric Implementation		
				E	FC	IB
0x28	Set PF Assignment	Used to configure and enable Functions	0xA8			
0x29	Get Port Configuration	Used to request port configuration information	0xA9			
0x2A	Set Port Configuration	Used to configure operational characteristics of the port	0xAA			
0x2B	Get Partition Configuration	Used to request partition configuration information	0xAB			
0x2C	Set Partition Configuration	Used to configure partition operational characteristics	0xAC			
0x2D	Get Boot Config	Used to request boot protocol configuration information	0xAD			
0x2E	Set Boot Config	Used to configure boot protocol attributes	0xAE			
0x2F	Get Partition Statistics	Used to request network link statistics for the partition	0xAF			
0x31	Get FC Link Status	Used to request link and trunk status and speed for Fibre Channel ports	0xB1		M	
0x38	Get InfiniBand Link Status	Used to request link status for InfiniBand ports	0xB8			M
0x39	Get InfiniBand Statistics	Used to request port level statistics for InfiniBand ports	0xB9			M
0x47	Settings Commit	Used to request the commit of certain settings to NVRAM	0xC7			
0x48	Get ASIC Temperature	Used to request current NC ASIC and other external device temperatures from the NC	0xC8			
0x49	Get Ambient Temperature	Used to request the current ambient temperature from the NC adapter	0xC9			
0x4A	Get Transceiver Temperature	Used to request the current optical module temperature and thresholds	0xCA			
0x4B	Thermal Shutdown Control	Used to control and query the state of the thermal-based self-shutdown feature	0xCB	C	C	C

Control Packet Type	Command Name	Description	Response Packet Type	Fabric Implementation		
				E	FC	IB
0X4C	Transmit Data to NC	Used by the MC to transfer a block of data to the NC	0xCC	O	O	O
0X4D	Retrieve Data from NC	Used by the MC to transfer a block of data from the NC	0xCD	O	O	O
0x50	OEM Command	Used to request vendor-specific data	0xD0			
0x51	PLDM Request	Used for PLDM request over NC-SI over RBT	0xD1			
0x52	Get Package UUID	Returns a universally unique identifier (UUID) for the package	0xD2	O	O	O
0x51–0x60	Reserved for Transport Protocol Oriented Commands	Used to define transport protocol-oriented commands (e.g., PLDM over NC-SI/RBT)	0xD1–0xE0	O	O	O
0x51	Reserved					
0x52	Get Package UUID	Returns a universally unique identifier (UUID) for the package	0xD2	O	O	O
0x53	PLDM	Used for PLDM request over NC-SI over RBT	0xD3	O	O	O
0x54	Get Supported Media	See MCTP DSP0261 for full definition This command may be used on any transport	0xD4			
0x55	Transport-specific AEN Enable	See MCTP DSP0261 for full definition	0xD5			
0x56	Query Pending NC PLDM Request	Used by the MC to see if the NC has any pending PLDM requests to be retrieved	0xD6	O	O	O
0x57	Send NC PLDM Reply	Used by the MC to provide a response to a previous SPDM request by the NC	0xD7	O	O	O
0x58	Get MC MAC Address	Used by the MC to retrieve MAC addresses provisioned for its use	0xD8	O	O	O

Control Packet Type	Command Name	Description	Response Packet Type	Fabric Implementation		
				E	FC	IB
0x60	Transfer SPDM	Used by the MC to transfer a SPDM payload to or from the NC	0xE0	O	O	O
0x61	Query Pending SPDM Request	Used by the MC to see if the NC has any pending SPDM requests to be retrieved	0xE1	O	O	O
0x62	Send NC SPDM Reply	Used by the MC to respond to a previously read SPDM command from the NC	0xE2	O	O	O

2215

[illegible]

[illegible]

2216

2217 **8.4 Command and response packet formats**

2218 This clause describes the format for each of the NC-SI commands and corresponding responses.

2219 The corresponding response packet format shall be mandatory when a given command is supported.

2220 **8.4.1 NC-SI command frame format**

2221 Table 21 illustrates the NC-SI frame format that shall be accepted by the Network Controller.

2222 **Table 21 – Example of complete minimum-sized NC-SI command packet**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..03	0xFF	0xFF	0xFF	0xFF
04..07	0xFF	0xFF	0xFF	0xFF
08..11	0xFF	0xFF	0xFF	0xFF
12..15	0x88F8		MC ID	Header Revision
16..19	Reserved	IID	Command Type	Ch. ID
20..23	Reserved	Payload Length	Reserved	
24..27	Reserved		Reserved	
28..31	Reserved		Checksum (3..2)	
32..35	Checksum (1..0)		Pad	
36..39	Pad			
40..43	Pad			
44..47	Pad			
48..51	Pad			
52..55	Pad			
56..59	Pad			
60..63	FCS			

8.4.2 NC-SI response packet format

Table 22 illustrates the NC-SI response packet format that shall be transmitted by the Network Controller.

Table 22 – Example of complete minimum-sized NC-SI response packet

	Bits				
Bytes	31..24		23..16	15..08	07..00
00..03	0xFF		0xFF	0xFF	0xFF
04..07	0xFF		0xFF	0xFF	0xFF
08..11	0xFF		0xFF	0xFF	0xFF
12..15	0x88F8			MC ID	Header Revision
16..19	Reserved		IID	Response Type	Ch. ID
20..23	Reserved	Payload Length		Reserved	
24..27	Reserved			Reserved	
28..31	Reserved			Response Code	
32..35	Reason Code			Checksum (3..2)	
36..39	Checksum (1..0)			Pad	
40..43	Pad				
44..47	Pad				
48..51	Pad				
52..55	Pad				
56..59	Pad				
60..63	FCS				

8.4.3 Clear Initial State command (0x00)

The Clear Initial State command provides the mechanism for the Management Controller to acknowledge that it considers a channel to be in the Initial State (typically because the Management Controller received an “Interface Initialization Required” reason code) and to direct the Network Controller to start accepting commands for initializing or recovering the NC-SI operation. When in the Initial State, the Network Controller shall return the “Interface Initialization Required” reason code for all channel commands until it receives the Clear Initial State command.

If the channel is in the Initial State when it receives the Clear Initial State command, the command shall cause the Network Controller to stop returning the “Interface Initialization Required” reason code. The channel shall also treat any subsequently received instance ID numbers as IDs for new command instances, not retries.

If the channel is not in the Initial State when it receives this command, it shall treat any subsequently received instance ID numbers as IDs for new command instances, not retries.

Table 23 illustrates the packet format of the Clear Initial State command.

2240 **Table 23 – Clear Initial State command packet format**

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

2241 **8.4.4 Clear Initial State response (0x80)**

2242 Currently no command-specific reason code is identified for this response (see Table 24).

2243 **Table 24 – Clear Initial State response packet format**

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

2244 **8.4.5 Select Package command (0x01)**

2245 A package is considered to be “selected” when its NC-SI output buffers are allowed to transmit packets
 2246 through the NC-SI interface. Conversely, a package is “deselected” when it is not allowed to transmit
 2247 packets through the NC-SI interface.

2248 The Select Package command provides a way for a Management Controller to explicitly take a package
 2249 out of the deselected state and to control whether hardware arbitration is enabled for the package.
 2250 (Similarly, the Deselect Package command allows a Management Controller to explicitly deselect a
 2251 package.)

2252 The NC-SI package in the Network Controller shall also become selected if the package receives any NC-
 2253 SI command (other than Deselect Package) that is directed to the package or to a channel within the
 2254 package.

2255 The Select Package command is addressed to the package, rather than to a channel (that is, the
 2256 command is sent with a Channel ID where the Package ID subfield matches the ID of the intended
 2257 package and the Internal Channel ID subfield is set to 0x1F).

2258 More than one package can be in the selected state simultaneously if hardware arbitration is used
 2259 between the selected packages and is active. The hardware arbitration logic ensures that buffer conflicts
 2260 will not occur between selected packages.

2261 If hardware arbitration is not active or is not used for a given package, only one package shall be selected
 2262 at a time. To switch between packages, the Deselect Package command is used by the Management
 2263 Controller to put the presently selected package into the deselected state before another package is
 2264 selected.

- 2265 A package shall stay in the selected state until it receives a Deselect Package command unless an
2266 internal condition causes all internal channels to enter the Initial State.
- 2267 A package that is not using hardware arbitration may leave its output buffers enabled for the time that it is
2268 selected, or it may place its output buffers into the high-impedance state between transmitting packets
2269 through the NC-SI interface. (Temporarily placing the output buffers into the high-impedance state is not
2270 the same as entering the deselected state.)
- 2271 For Type A integrated controllers: Because the RBT bus buffers are separately controlled, a separate
2272 Select Package command needs to be sent to each Package ID in the controller that is to be enabled to
2273 transmit through the NC-SI interface. If the internal packages do not support hardware arbitration, only
2274 one package shall be selected at a time; otherwise, a bus conflict will occur.
- 2275 For Type S single channel, and Types B and C integrated controllers: A single set of RBT bus buffers
2276 exists for the package. Sending a Select Package command selects the entire package and enables all
2277 channels within the package to transmit through the NC-SI interface. (Whether a particular channel in a
2278 selected package starts transmitting Pass-through and AEN packets depends on whether that channel
2279 was enabled or disabled using the Enable or Disable Channel commands and whether the package may
2280 have had packets queued up for transmission.)
- 2281 Implementation Note: the features control settings are only configurable via this command and are not
2282 altered by 'implicit' selection as described in 6.1.14.4.
- 2283
- 2284 Table 25 illustrates the packet format of the Select Package command. Table 26 illustrates the disable
2285 byte for hardware arbitration.

Table 25 – Select Package command packet format

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

Table 26 – Features Control byte

Bits	Description
0	<p>0b = Hardware arbitration between packages is enabled.</p> <p>1b = Disable hardware arbitration. Disabling hardware arbitration causes the package's arbitration logic to enter or remain in bypass mode.</p> <p>In the case that the Network Controller does not support hardware arbitration, this bit is ignored; the Network Controller shall not return an error if the Select Package command can otherwise be successfully processed.</p>
1	<p>Delayed Response Enable:</p> <p>0b = NC is not allowed to use the "Delayed Response" response code (default)</p> <p>1b = NC is allowed to use the "Delayed Response" response code</p>
7..2	Reserved

2288 8.4.6 Select Package response (0x81)

2289 Currently no command-specific reason code is identified for this response (see Table 27).

2290 **Table 27 – Select package response packet format**

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

2291 8.4.7 Deselect Package command (0x02)

2292 The Deselect Package command directs the controller package to stop transmitting packets through the
2293 NC-SI interface and to place the output buffers for the package into the high-impedance state.

2294 The Deselect Package command is addressed to the package, rather than to a particular channel (that is,
2295 the command is sent with a Channel ID where the Package ID subfield matches the ID of the intended
2296 package and the Internal Channel ID subfield is set to 0x1F).

2297 The controller package enters the deselected state after it has transmitted the response to the Deselect
2298 Package command and placed its buffers into the high-impedance state. The controller shall place its
2299 outputs into the high-impedance state within the Package Deselect to Hi-Z Interval (T1). (This interval
2300 gives the controller being deselected time to turn off its electrical output buffers after sending the
2301 response to the Deselect Package command.)

2302 If hardware arbitration is not supported or used, the Management Controller should wait for the Package
2303 Deselect to Hi-Z Interval (T1) to expire before selecting another controller.

2304 For Type A integrated controllers: Because the bus buffers are separately controlled, putting the overall
2305 controller package into the high-impedance state requires sending separate Deselect Package
2306 commands to each Package ID in the overall package.

2307 For Type S single channel, and Types B and C integrated controllers: A single set of bus buffers exists for
2308 the package. Sending a Deselect Package command deselects the entire NC-SI package and prevents
2309 all channels within the package from transmitting through the NC-SI interface.

2310 Table 28 illustrates the packet format of the Deselect Package command.

2311 **Table 28 – Deselect Package command packet format**

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

2312 8.4.8 Deselect Package response (0x82)

2313 The Network Controller shall always put the package into the deselected state after sending a Deselect
2314 Package Response.

2315 No command-specific reason code is identified for this response (see Table 29).

2316 **Table 29 – Deselect Package response packet format**

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

2317 8.4.9 Enable Channel command (0x03)

2318 The Enable Channel command shall enable the Network Controller to allow transmission of Pass-through
2319 and AEN packets to the Management Controller through the NC-SI.

2320 Table 30 illustrates the packet format of the Enable Channel command.

2321 **Table 30 – Enable Channel command packet format**

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

2322 8.4.10 Enable Channel response (0x83)

2323 No command-specific reason code is identified for this response (see Table 31).

2324 **Table 31 – Enable Channel response packet format**

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

2325 8.4.11 Disable Channel command (0x04)

2326 The Disable Channel command allows the Management Controller to disable the flow of packets,
2327 including Pass-through and AEN, to the Management Controller.

2328 A Network Controller implementation is not required to flush pending packets from its RX Queues when a
2329 channel becomes disabled. If queuing is subsequently disabled for a channel, it is possible that a number
2330 of packets from the disabled channel could still be pending in the RX Queues. These packets may
2331 continue to be transmitted through the NC-SI interface until the RX Queues are emptied of those packets.
2332 The Management Controller should be aware that it may receive a number of packets from the channel
2333 before receiving the response to the Disable Channel command.

2334 The 1-bit Allow Link Down (ALD) field can be used by the Management Controller to indicate that the link
2335 corresponding to the specified channel is not required after the channel is disabled. The Network
2336 Controller is allowed to take down the external network physical link if no other functionality (for example,
2337 host OS or WoL [Wake-on-LAN]) is active.

2338 Possible values for the 1-bit ALD field are as follows:

- 2339 • 0b = Keep link up (establish and/or keep a link established) while channel is disabled
- 2340 • 1b = Allow link to be taken down while channel is disabled

2341 Table 32 illustrates the packet format of the Disable Channel command.

2342 **Table 32 – Disable Channel command packet format**

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

2343 NOTE It is currently unspecified whether this command will cause the Network Controller to cease the passing
2344 through of traffic from the Management Controller to the network, or if this can only be done using the Disable
2345 Channel Network TX command.

2346 8.4.12 Disable Channel response (0x84)

2347 No command-specific reason code is identified for this response (see Table 33).

2348 **Table 33 – Disable Channel response packet format**

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

2349 8.4.13 Reset Channel command (0x05)

2350 The Reset Channel command allows the Management Controller to put the channel into the Initial State.
 2351 Packet transmission is not required to stop until the Reset Channel response has been sent. Thus, the
 2352 Management Controller should be aware that it may receive a number of packets from the channel before
 2353 receiving the response to the Reset Channel command.

2354 Table 34 illustrates the packet format of the Reset Channel command.

2355 **Table 34 – Reset Channel command packet format**

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

2356 8.4.14 Reset Channel response (0x85)

2357 Currently no command-specific reason code is identified for this response (see Table 35).

2358 **Table 35 – Reset Channel response packet format**

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

2359 8.4.15 Enable Channel Network TX command (0x06)

2360 The Enable Channel Network TX command shall enable the channel to transmit Pass-through packets
 2361 onto the network. After network transmission is enabled, this setting shall remain enabled until a Disable
 2362 Channel Network TX command is received, or the channel enters the Initial State.

2363 The intention of this command is to control which Network Controller ports are allowed to transmit to the
 2364 external network. The Network Controller compares the source MAC address in outgoing Pass-through
 2365 packets to the unicast MAC address(es) configured using the Set MAC Address command. If a match
 2366 exists, the packet is transmitted to the network.

2367 Table 36 illustrates the packet format of the Enable Channel Network TX command.

2368 **Table 36 – Enable Channel Network TX command packet format**

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

2369

2370 **8.4.16 Enable Channel Network TX response (0x86)**

2371 No command-specific reason code is identified for this response (see Table 37).

2372 **Table 37 – Enable Channel Network TX response packet format**

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

2373 **8.4.17 Disable Channel Network TX command (0x07)**

2374 The Disable Channel Network TX command disables the channel from transmitting Pass-through packets
 2375 onto the network. After network transmission is disabled, it shall remain disabled until an Enable Channel
 2376 Network TX command is received.

2377 Table 38 illustrates the packet format of the Disable Channel Network TX command.

2378 **Table 38 – Disable Channel Network TX command packet format**

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

2379 8.4.18 Disable Channel Network TX response (0x87)

2380 The NC-SI shall, in the absence of a checksum error or identifier mismatch, always accept the Disable
2381 Channel Network TX command and send a response.

2382 Currently no command-specific reason code is identified for this response (see Table 39).

2383 **Table 39 – Disable Channel Network TX response packet format**

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

2384 8.4.19 AEN Enable command (0x08)

2385 Network Controller implementations shall support this command on the condition that the Network
2386 Controller generates one or more standard AENs. The AEN Enable command enables and disables the
2387 different standard AENs supported by the Network Controller. The Network Controller shall copy the AEN
2388 MC ID field from the AEN Enable command into the MC ID field in every subsequent AEN sent to the
2389 Management Controller.

2390 For more information, see **Error! Reference source not found. ("Error! Reference source not found.")**
2391 and 8.2.1.1 ("Management Controller ID").

2392 Control of transport-specific AENs is outside the scope of this specification and should be defined by the
2393 transport binding specifications.

2394 Table 40 illustrates the packet format of the AEN Enable command.

2395 **Table 40 – AEN Enable command packet format**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Reserved			AEN MC ID
20..23	AEN Control			
24..27	Checksum			
28..45	Pad			

2396 The AEN Control field has the format shown in Table 41.

2397 **Table 41 – Format of AEN control**

Bit Position	Field Description	Value Description
0	Link Status Change AEN control	0b = Disable Link Status Change AEN 1b = Enable Link Status Change AEN
1	Configuration Required AEN control	0b = Disable Configuration Required AEN 1b = Enable Configuration Required AEN
2	Host NC Driver Status Change AEN control	0b = Disable Host NC Driver Status Change AEN 1b = Enable Host NC Driver Status Change AEN
3	Delayed Response Ready AEN control	0b = Disable Delayed Response Ready AEN 1b = Enable Delayed Response Ready AEN
4	InfiniBand Link Status Change AEN control	0b = Disable IB Link Status Change AEN 1b = Enable IB Link Status Change AEN
5	Fibre Channel Link Status Change AEN control	0b = Disable FC Link Status Change AEN 1b = Enable FC Link Status Change AEN
6	Transceiver Event AEN Control	0b = Disable Transceiver Event AEN 1b = Enable Transceiver Event AEN
7	Request Data Transfer AEN control	0b = Disable Request Data Transfer AEN 1b = Enable Request Data Transfer AEN
8	Partition Link Status Change AEN control	0b = Disable Partition Link Status Change AEN 1b = Enable Partition Link Status Change AEN
9	Thermal Shutdown Event AEN control	0b = Disable Thermal Shutdown Event AEN 1b = Enable Thermal Shutdown Event AEN
15..610	Reserved	Reserved
31..16	OEM-specific AEN control	OEM-specific control

2398 **8.4.20 AEN Enable response (0x88)**

2399 Currently no command-specific reason code is identified for this response (see Table 42). If the MC
2400 attempts to set an AEN type that is not supported, the NC shall reject the entire command even if it also

2401 includes valid AENs and respond with the “Command Failed” response and “Parameter Is Invalid...”
 2402 reason codes.

2403 **Table 42 – AEN Enable response packet format**

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

2404 **8.4.21 Set Link command (0x09)**

2405 The Set Link command may be used by the Management Controller to configure the external network
 2406 interface associated with the channel by using the provided settings. Upon receiving this command, while
 2407 the host NC driver is not operational, the channel shall attempt to set the link to the configuration
 2408 specified by the parameters. Upon successful completion of this command, link settings specified in the
 2409 command should be used by the network controller as long as the host NC driver does not overwrite the
 2410 link settings.

2411 In the absence of an operational host NC driver, the NC should attempt to make the requested link state
 2412 change even if it requires the NC to drop the current link. The channel shall send a response packet to
 2413 the Management Controller within the required response time. However, this specification does not
 2414 specify the amount of time the requested link state changes may take to complete.

2415 The actual link settings are controlled by the host NC driver when it is operational. When the host NC
 2416 driver is operational, link settings specified by the MC using the Set Link command may be overwritten by
 2417 the host NC driver. The link settings are not restored by the NC if the host NC driver becomes non-
 2418 operational.

2419 Table 43 illustrates the packet format of the Set Link command.

2420 **Table 43 – Set Link command packet format**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Link Settings			
20..23	OEM Link Settings			
24..27	Checksum			
28..45	Pad			

2421 Table 44 and Table 45 describe the Set Link bit definitions. Refer to [IEEE 802.3](#) for definitions of Auto
 2422 Negotiation, Duplex Setting, Pause Capability, and Asymmetric Pause Capability.

2423 **Table 44 – Set Link bit definitions**

Bit Position	Field Description	Value Description
00	Auto Negotiation If Auto Negotiation is not used, only one combination of single link speed, protocol and FEC settings is allowed to be configured, otherwise a Command Failed response code and Parameter Is Invalid, Unsupported, or Out-of-Range reason code shall be returned.	1b = enable 0b = disable
01..07	Link Speed Selection More than one speed can be selected when Auto Negotiation is set to 'enable'. If Auto Negotiation is not used, the channel attempts to force the link to the specified setting (in this case, if the setting is not supported or if multiple speeds are enabled, a Command Failed response code and Parameter Is Invalid, Unsupported, or Out-of-Range reason code shall be returned). If multiple settings are enabled, a Command Failed response code and Set Link Speed Conflict reason code shall be returned) NOTE Additional link speeds are defined below.	Bit 01: 1b = enable 10 Mbps
		Bit 02: 1b = enable 100 Mbps
		Bit 03: 1b = enable 1000 Mbps (1 Gbps)
		Bit 04: 1b = enable 10 Gbps
		Bit 05: 1b = enable 20 Gbps (optional for NC-SI 1.1, Reserved for NC-SI 1.0)
		Bit 06: 1b = enable 25 Gbps (optional for NC-SI 1.1, Reserved for NC-SI 1.0)
		Bit 07: 1b = enable 40 Gbps (optional for NC-SI 1.1, Reserved for NC-SI 1.0)
08..09	Duplex Setting (separate duplex setting bits) More than one duplex setting can be selected when Auto Negotiation is set to 'enable'. If Auto Negotiation is not used, the channel attempts to force the link to the specified setting (in this case, if the setting is not supported or if multiple settings are enabled, a Command Failed response code and Parameter Is Invalid, Unsupported, or Out-of-Range reason code shall be returned."	Bit 08: 1b = enable half-duplex
		Bit 09: 1b = enable full-duplex
10	Pause Capability If Auto Negotiation is not used, the channel should apply pause settings assuming the partner supports the same capability.	1b = disable 0b = enable
11	Asymmetric Pause Capability If Auto Negotiation is not used, the channel should apply asymmetric pause settings assuming the partner supports the same capability.	1b = enable 0b = disable
12	OEM Link Settings Field Valid (see Table 45)	1b = enable 0b = disable

13..19	Additional Link Speeds (see Link Speed Selection)	<p>Bit 13: 1b = enable 50 Gbps (optional for NC-SI 1.1, Reserved for NC-SI 1.0)</p> <p>Bit 14: 1b = enable 100 Gbps (optional for NC-SI 1.1, Reserved for NC-SI 1.0)</p> <p>Bit 15: 1b = enable 2.5 Gbps (optional for NC-SI 1.1, Reserved for NC-SI 1.0)</p> <p>Bit 16: 1b = enable 5 Gbps (optional for NC-SI 1.1, Reserved for NC-SI 1.0)</p> <p>Bit 17: 1b = enable 200 Gbps (optional for NC-SI 1.2, Reserved for NC-SI 1.1/1.0)</p> <p>Bit 18: 1b = enable 400 Gbps (optional for NC-SI 1.2, Reserved for NC-SI 1.1/1.0)</p> <p>Bit 19: 1b = enable 800 Gbps (optional for NC-SI 1.2, Reserved for NC-SI 1.1/1.0)</p>
20..21	Reserved	
22..23	Modulation Scheme	<p>Bit 22: 1b = NRZ (optional for NC-SI 1.2, Reserved for NC-SI 1.1/1.0)</p> <p>Bit 23: 1b = PAM-4 (optional for NC-SI 1.2, Reserved for NC-SI 1.1/1.0)</p> <p>Bit 23-22 Values:</p> <p>00 – Use default</p> <p>01 – Enable NRZ</p> <p>10 – Enable PAM-4</p> <p>11 – Enable NRZ and PAM-4</p>
24..27	Forward Error Correction (FEC) Algorithm	<p>Bit 24: 1b = BASE-R FEC (Firecode) (optional for NC-SI 1.2, Reserved for NC-SI 1.1/1.0)</p> <p>Bit 25: 1b = RS-FEC (Reed Solomon) (optional for NC-SI 1.2, Reserved for NC-SI 1.1/1.0)</p> <p>Bit 26..27 Reserved</p> <p>If all bits are set to 0, then no FEC algorithm shall be selected</p>
28	Energy Efficient Ethernet (EEE)	<p>1b = enable</p> <p>0b = disable</p>
29	Link Training (LT)	<p>1b = enable</p> <p>0b = disable</p>
30	Parallel Detect An auto-negotiation link partner's mechanism to establish links with non-negotiation, fixed-speed linked partners.	<p>1b = enable</p> <p>0b = disable</p>
31	Reserved	0

2424

Table 45 – OEM Set Link bit definitions

Bit Position	Field Description	Value Description
00..31	OEM Link Settings	Vendor specified

2425 **8.4.22 Set Link Response (0x89)**

2426 The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Set Link
 2427 command and send a response (see Table 46). In the presence of an operational Host NC driver, the NC
 2428 should not attempt to make link state changes and should send a response with reason code 0x1 (Set
 2429 Link Host OS/ Driver Conflict).

2430 If the Auto Negotiation field is set, the NC should ignore Link Speed Selection and Duplex Setting fields
 2431 that are not supported by the NC.

2432

Table 46 – Set Link response packet format

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

2433 Table 47 describes the reason codes that are specific to the Set Link command. Returning the following
 2434 command-specific codes is recommended, conditional upon Network Controller support for the related
 2435 capabilities.

2436

Table 47 – Set Link command-specific reason codes

Value	Description	Comment
0x0901	Set Link Host OS/ Driver Conflict	Returned when the Set Link command is received when the Host NC driver is operational
0x0902	Set Link Media Conflict	Returned when Set Link command parameters conflict with the media type (for example, Fiber Media)
0x0903	Set Link Parameter Conflict	Returned when Set Link parameters conflict with each other (for example, 1000 Mbps HD with copper media)
0x0904	Set Link Power Mode Conflict	Returned when Set Link parameters conflict with current low-power levels by exceeding capability
0x0905	Set Link Speed Conflict	Returned when Set Link parameters attempt to force more than one speed at the same time when Auto Negotiation is disabled
0x0906	Link Command Failed-Hardware Access Error	Returned when PHY R/W access fails to complete normally while executing the Set Link or Get Link Status command
0x0907	Set Link SerDes Conflict	Returned when Set Link parameters attempt to force an unsupported SerDes configuration

Value	Description	Comment
0x0908	Set Link FEC Conflict	Returned when Set Link parameters attempt to force an unsupported FEC algorithm
0x0909	Set Link EEE Conflict	Returned when Set Link parameters attempt to force an unsupported EEE configuration
0x090A	Set Link LT Conflict	Returned when Set Link parameters attempt to force an unsupported link training configuration
0x090B	Set Link Parallel Detection Conflict	Returned when Set Link parameters attempt to force an unsupported parallel detection configuration

2437 8.4.23 Get Link Status command (0x0A)

2438 The Get Link Status command allows the Management Controller to query the channel for potential link
 2439 status and error conditions (see Table 48).

2440 **Table 48 – Get Link Status command packet format**

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

2441 8.4.24 Get Link Status response (0x8A)

2442 The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Get Link
 2443 Status command and send a response (see Table 49).

2444 **Table 49 – Get Link Status response packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	Link Status			
24..27	Other Indications			
28..31	OEM Link Status			
32..35	Checksum			
36..45	Pad			

2445 Table 50 describes the Link Status bit definitions.

2446 **Table 50 – Link Status field bit definitions**

Bit Position	Field Description	Value Description
00	Link Flag	<p>0b = Link is down 1b = Link is up (including Low Power Idle state in EEE)</p> <p>This field is mandatory.</p>
04..01	Speed and duplex	<p>0x0 = Auto-negotiate not complete [per IEEE 802.3], or SerDes Flag = 1b, or no Highest Common Denominator (HCD) from the following options (0x1 through 0xF) was found.</p> <p>0x1 = 10BASE-T half-duplex 0x2 = 10BASE-T full-duplex 0x3 = 100BASE-TX half-duplex 0x4 = 100BASE-T4 0x5 = 100BASE-TX full-duplex 0x6 = 1000BASE-T half-duplex 0x7 = 1000BASE-T full-duplex 0x8 = 10G-BASE-T support or 10 Gbps 0x9 = 20 Gbps (optional for NC-SI 1.1, Reserved for NC-SI 1.0) 0xA = 25 Gbps (optional for NC-SI 1.1, Reserved for NC-SI 1.0) 0xB = 40 Gbps (optional for NC-SI 1.1, Reserved for NC-SI 1.0) 0xC = 50 Gbps (optional for NC-SI 1.1, Reserved for NC-SI 1.0) 0xD = 100 Gbps (optional for NC-SI 1.1, Reserved for NC-SI 1.0) 0xE = 2.5 Gbps (optional for NC-SI 1.1, Reserved for NC-SI 1.0) 0xF = Use values defined in Extended Speed and Duplex field starting at bit 24 (optional for NC-SI 1.1, Reserved for NC-SI 1.0)</p> <p>When SerDes Flag = 0b, the value may reflect forced link setting.</p> <p>NOTE For the physical medium and/or speed/duplex not listed above, the closest speed and duplex option can be reported by the NC. This field does not infer any media type information.</p>
05	Auto Negotiate Flag	<p>1b = Auto-negotiation is enabled.</p> <p>This field always returns 0b if auto-negotiation is not supported, or not enabled.</p> <p>This field is mandatory if supported by the controller.</p>
06	Auto Negotiate Complete	<p>1b = Auto-negotiation has completed.</p> <p>This includes if auto-negotiation was completed using Parallel Detection. Always returns 0b if auto-negotiation is not supported or is not enabled.</p> <p>This field is mandatory if the Auto Negotiate Flag is supported.</p>

Bit Position	Field Description	Value Description
07	Parallel Detection Flag	1b = Link partner did not support auto-negotiation and parallel detection was used to get link. This field contains 0b if Parallel Detection was not used to obtain link.
08	Reserved	None
09	Link Partner Advertised Speed and Duplex 1000TFD	1b = Link Partner is 1000BASE-T full-duplex capable. Valid when: SerDes Flag = 0b Auto-Negotiate Flag = 1b Auto-Negotiate Complete = 1b This field is mandatory.
10	Link Partner Advertised Speed and Duplex 1000THD	1b = Link Partner is 1000BASE-T half-duplex capable. Valid when: SerDes Flag = 0b Auto-Negotiate Flag = 1b Auto-Negotiate Complete = 1b This field is mandatory.
11	Link Partner Advertised Speed 100T4	1b = Link Partner is 100BASE-T4 capable. Valid when: SerDes Flag = 0b Auto-Negotiate Flag = 1b Auto-Negotiate Complete = 1b This field is mandatory.
12	Link Partner Advertised Speed and Duplex 100TXFD	1b = Link Partner is 100BASE-TX full-duplex capable. Valid when: SerDes Flag = 0b Auto-Negotiate Flag = 1b Auto-Negotiate Complete = 1b This field is mandatory.
13	Link Partner Advertised Speed and Duplex 100TXHD	1b = Link Partner is 100BASE-TX half-duplex capable. Valid when: SerDes Flag = 0b Auto-Negotiate Flag = 1b Auto-Negotiate Complete = 1b This field is mandatory.

Bit Position	Field Description	Value Description
14	Link Partner Advertised Speed and Duplex 10TFD	<p>1b = Link Partner is 10BASE-T full-duplex capable.</p> <p>Valid when:</p> <p>SerDes Flag = 0b</p> <p>Auto-Negotiate Flag = 1b</p> <p>Auto-Negotiate Complete = 1b</p> <p>This field is mandatory.</p>
15	Link Partner Advertised Speed and Duplex 10THD	<p>1b = Link Partner is 10BASE-T half-duplex capable.</p> <p>Valid when:</p> <p>SerDes Flag = 0b</p> <p>Auto-Negotiate Flag = 1b</p> <p>Auto-Negotiate Complete = 1b</p> <p>This field is mandatory.</p>
16	TX Flow Control Flag	<p>0b = Transmission of Pause frames by the NC onto the external network interface is disabled.</p> <p>1b = Transmission of Pause frames by the NC onto the external network interface is enabled.</p> <p>This field is mandatory.</p>
17	RX Flow Control Flag	<p>0b = Reception of Pause frames by the NC from the external network interface is disabled.</p> <p>1b = Reception of Pause frames by the NC from the external network interface is enabled.</p> <p>This field is mandatory.</p>
19..18	Link Partner Advertised Flow Control	<p>00b = Link partner is not pause capable.</p> <p>01b = Link partner supports symmetric pause.</p> <p>10b = Link partner supports asymmetric pause toward link partner.</p> <p>11b = Link partner supports both symmetric and asymmetric pause.</p> <p>Valid when:</p> <p>SerDes Flag = 0b</p> <p>Auto-Negotiate = 1b</p> <p>Auto-Negotiate Complete = 1b</p> <p>This field is mandatory.</p>
20	SerDes Link	<p>SerDes status (See 4.21.)</p> <p>0b = SerDes is not used or used to connect to an external PHY</p> <p>1b = SerDes is used as a direct attach interface</p> <p>This field is mandatory.</p>

Bit Position	Field Description	Value Description
21	OEM Link Speed Valid	0b = OEM link settings are invalid. 1b = OEM link settings are valid.
23..22	Modulation Scheme	00b = Reserved 01b = NRZ is used. 10b = PAM-4 is used. 11b = Reserved This field is optional for NC-SI 1.2, reserved for NC-SI 1.1/1.0.
31..24	Extended Speed and duplex	Optional for NC-SI 1.2/1.1, Reserved for NC-SI 1.0 0x0 = Auto-negotiation not complete [per IEEE 802.3], or SerDes Flag = 1b, or no highest common denominator speed from the following options (0x01 through 0x0F) was found. 0x01 = 10BASE-T half-duplex 0x02 = 10BASE-T full-duplex 0x03 = 100BASE-TX half-duplex 0x04 = 100BASE-T4 0x05 = 100BASE-TX full-duplex 0x06 = 1000BASE-T half-duplex 0x07 = 1000BASE-T full-duplex 0x08 = 10G-BASE-T support or 10 Gbps 0x09 = 20 Gbps 0x0A = 25 Gbps 0x0B = 40 Gbps 0x0C = 50 Gbps 0x0D = 100 Gbps 0x0E = 2.5 Gbps 0x10 = 1 Gbps (for non Base-T) 0x11 = 200 Gbps 0x12 = 400 Gbps 0x13 = 800 Gbps 0x14-0xFF = Reserved When SerDes Flag = 0b, the value may reflect forced link setting. NOTE For the physical medium and/or speed/duplex not listed above, the closest speed and duplex option can be reported by the NC. This field does not infer any media type information.

2447 Table 51 describes the Other Indications field bit definitions.

2448 **Table 51 – Other Indications field bit definitions**

Bits	Description	Values
00	Host NC Driver Status Indication	<p>0b = The Network Controller driver for the host external network interface associated with this channel is not operational (not running), unknown, or not supported.</p> <p>1b = The Network Controller driver for the host external network interface associated with this channel (or when partitioned, at least one partition driver) is being reported as operational (running).</p> <p>This bit always returns 0b if the Host NC Driver Status Indication is not supported.</p>
01	Energy Efficient Ethernet (EEE)	<p>1b = enabled</p> <p>0b = disabled</p>
02	Link Training (LT)	<p>1b = enabled</p> <p>0b = disabled</p>
03	Parallel Detect	<p>1b = enabled</p> <p>0b = disabled</p>
04	OEM Link Status Field	<p>1b = enabled</p> <p>0b = disabled</p>
05..31	Reserved	

2449 Table 52 describes the OEM Link Status field bit definitions.

2450 **Table 52 – OEM Link Status field bit definitions (optional)**

Bits	Description	Values
00..31	OEM Link Status	OEM specific

2451 Table 53 describes the reason code that is specific to the Get Link Status command.

2452 **Table 53 – Get Link Status command-specific reason code**

Value	Description	Comment
0x0A06	Link Command Failed-Hardware Access Error	Returned when PHY R/W access fails to complete normally while executing the Set Link or Get Link Status command

2453 **8.4.25 Set VLAN Filter command (0x0B)**

2454 The Set VLAN Filter command is used by the Management Controller to program one or more VLAN IDs
2455 that are used for VLAN filtering.

2456 Incoming packets that match both a VLAN ID filter and a MAC address filter are forwarded to the
2457 Management Controller. Other packets may be dropped based on the VLAN filtering mode per the Enable
2458 VLAN command.

2459 The quantity of each filter type that is supported by the channel can be discovered by means of the Get
 2460 Capabilities command. Up to 15 filters can be supported per channel. A Network Controller
 2461 implementation shall support at least one VLAN filter per channel.

2462 To configure a VLAN filter, the Management Controller issues a Set VLAN Filter command with the Filter
 2463 Selector field indicating which filter is to be configured, the VLAN ID field set to the VLAN TAG values to
 2464 be used by the filter, and the Enable field set to either enable or disable the selected filter.

2465 The VLAN-related fields are specified per [IEEE 802.1q](#). When VLAN Tagging is used, the packet includes
 2466 a Tag Protocol Identifier (TPID) field and VLAN Tag fields, as shown in Table 54.

2467 **Table 54 – IEEE 802.1q VLAN Fields**

Field	Size	Description
TPI	2 bytes	Tag Protocol Identifier = 8100h
VLAN TAG – user priority	3 bits	User Priority (typical value = 000b)
VLAN TAG – CFI	1 bit	Canonical Format Indicator = 0b
VLAN TAG – VLAN ID	12 bits	Zeros = no VLAN

2468 When checking VLAN field values, the Network Controller shall match against the enabled VLAN Tag
 2469 Filter values that were configured with the Set VLAN Filter command. The Network Controller shall also
 2470 match on the TPI value of 8100h, as specified by [IEEE 802.1q](#). Matching against the User Priority/CFI
 2471 bits is optional. An implementation may elect to ignore the setting of those fields.

2472 Table 55 illustrates the packet format of the Set VLAN Filter command.

2473 **Table 55 – Set VLAN Filter command packet format**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Reserved		User Priority/CFI	VLAN ID
20..23	Reserved		Filter Selector	Reserved E
24..27	Checksum			
28..45	Pad			

2474 Table 56 provides possible settings for the Filter Selector field. Table 57 provides possible settings for the
 2475 Enable (E) field.

2476 **Table 56 – Possible Settings for Filter Selector field (8-bit field)**

Value	Description
1	Settings for VLAN filter number 1
2	Settings for VLAN filter number 2
..	

Value	Description
N	Settings for VLAN filter number <i>N</i>

2477 **Table 57 – Possible Settings for Enable (E) field (1-bit field)**

Value	Description
0b	Disable this VLAN filter
1b	Enable this VLAN filter

2478 **8.4.26 Set VLAN Filter response (0x8B)**

2479 The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Set
2480 VLAN Filter command and send a response (see Table 58).

2481 **Table 58 – Set VLAN Filter response packet format**

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

2482 Table 59 describes the reason code that is specific to the Set VLAN Filter command.

2483 **Table 59 – Set VLAN Filter command-specific reason code**

Value	Description	Comment
0x0B07	VLAN Tag Is Invalid	Returned when the VLAN ID is invalid (VLAN ID = 0)

2484 **8.4.27 Enable VLAN command (0x0C)**

2485 The Enable VLAN command may be used by the Management Controller to enable the channel to accept
2486 VLAN-tagged packets from the network for NC-SI Pass-through operation (see Table 60).

2487 **Table 60 – Enable VLAN command packet format**

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

2488 Table 61 describes the modes for the Enable VLAN command.

2489 **Table 61 – VLAN Enable modes**

Mode	#	O/M	Description
Reserved	0x00	N/A	Reserved
VLAN only	0x01	M	Only VLAN-tagged packets that match the enabled VLAN Filter settings (and also match the MAC Address Filtering configuration) are accepted. Non-VLAN-tagged packets are not accepted.
VLAN + non-VLAN	0x02	O	VLAN-tagged packets that match the enabled VLAN Filter settings (and also match the MAC Address Filtering configuration) are accepted. Non-VLAN-tagged packets (that also match the MAC Address Filtering configuration) are also accepted.
Any VLAN + non-VLAN	0x03	O	Any VLAN-tagged packets that also match the MAC Address Filtering configuration are accepted, regardless of the VLAN Filter settings. Non-VLAN-tagged packets (that also match the MAC Address Filtering configuration) are also accepted.
Reserved	0x04 – 0xFF	N/A	Reserved

2490 **8.4.28 Enable VLAN response (0x8C)**

2491 The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Enable
2492 VLAN command and send a response.

2493 Currently no command-specific reason code is identified for this response (see Table 62).

2494 **Table 62 – Enable VLAN response packet format**

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

2495 **8.4.29 Disable VLAN command (0x0D)**

2496 The Disable VLAN command may be used by the Management Controller to disable VLAN filtering. In the
2497 disabled state, only non-VLAN-tagged packets (that also match the MAC Address Filtering configuration)
2498 are accepted. VLAN-tagged packets are not accepted.

2499 Table 63 illustrates the packet format of the Disable VLAN command.

2500

Table 63 – Disable VLAN command packet format

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

2501 **8.4.30 Disable VLAN response (0x8D)**

2502 The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Disable
 2503 VLAN command and send a response.

2504 Currently no command-specific reason code is identified for this response (see Table 64).

2505

Table 64 – Disable VLAN response packet format

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

2506 **8.4.31 Set MAC Address command (0x0E)**

2507 The Set MAC Address command is used by the Management Controller to program the channel's unicast
 2508 or multicast MAC address filters.

2509 The channel supports one or more “perfect match” MAC address filters that are used to selectively
 2510 forward inbound frames to the Management Controller. Assuming that a packet passes any VLAN filtering
 2511 that may be active, it will be forwarded to the Management Controller if its 48-bit destination MAC address
 2512 exactly matches an active MAC address filter.

2513 MAC address filters may be configured as unicast or multicast addresses, depending on the capability of
 2514 the channel. The channel may implement three distinct types of filter:

- 2515 • **Unicast filters** support exact matching on 48-bit unicast MAC addresses (AT = 0x0 only).
- 2516 • **Multicast filters** support exact matching on 48-bit multicast MAC addresses (AT = 0x1 only).
- 2517 • **Mixed filters** support matching on both unicast and multicast MAC addresses. (AT=0x0 or
 2518 AT=0x1)

2519 The number of each type of filter that is supported by the channel can be discovered by means of the Get
 2520 Capabilities command. The channel shall support at least one unicast address filter or one mixed filter, so
 2521 that at least one unicast MAC address filter may be configured on the channel. Support for any
 2522 combination of unicast, multicast, or mixed filters beyond this basic requirement is vendor specific. The
 2523 total number of all filters shall be less than or equal to 8.

2524 To configure an address filter, the Management Controller issues a Set MAC Address command with the
 2525 Address Type field indicating the type of address to be programmed (unicast or multicast) and the MAC
 2526 Address Num field indicating the specific filter to be programmed.

2527 Filters are addressed using a 1-based index ordered over the unicast, multicast, and mixed filters
 2528 reported by means of the Get Capabilities command. For example, if the interface reports four unicast
 2529 filters, two multicast filters, and two mixed filters, then MAC Address numbers 1 through 4 refer to the
 2530 interface's unicast filters, 5 and 6 refer to the multicast filters, and 7 and 8 refer to the mixed filters.
 2531 Similarly, if the interface reports two unicast filters, no multicast filters, and six mixed filters, then MAC
 2532 address numbers 1 and 2 refer to the unicast filters, and 3 through 8 refer to the mixed filters.

2533 The filter type of the filter to be programmed (unicast, multicast, or mixed) shall be compatible with the
 2534 Address Type being programmed. For example, programming a mixed filter to a unicast address is
 2535 allowed, but programming a multicast filter to a unicast address is an error.

2536 The Enable field determines whether the indicated filter is to be enabled or disabled. When a filter is
 2537 programmed to be enabled, the filter is loaded with the 48-bit MAC address in the MAC Address field of
 2538 the command, and the channel enables forwarding of frames that match the configured address. If the
 2539 specified filter was already enabled, it is updated with the new address provided.

2540 When a filter is programmed to be disabled, the contents of the MAC Address field are ignored. Any
 2541 previous MAC address programmed in the filter is discarded and the channel no longer uses this filter in
 2542 its packet-forwarding function.

2543 Only unicast MAC addresses, specified with AT set to 0x0, should be used in source MAC address
 2544 checking and for determining the NC-SI channel for Pass-through transmit traffic.

2545 Table 65 illustrates the packet format of the Set MAC Address command.

Table 65 – Set MAC Address command packet format

		Bits						
Bytes	31..24	23..16		15..08		07..00		
00..15	NC-SI Control Packet Header							
16..19	MAC Address byte 5	MAC Address byte 4		MAC Address byte 3		MAC Address byte 2		
20..23	MAC Address byte 1	MAC Address byte 0		MAC Address Num		AT	Rsvd	E
24..27	Checksum							
28..45	Pad							
NOTE AT = Address Type, E = Enable.								

2547 Table 66 provides possible settings for the MAC Address Number field. Table 67 provides possible
 2548 settings for the Address Type (AT) field. Table 68 provides possible settings for the Enable (E) field.

Table 66 – Possible settings for MAC Address Number (8-bit field)

Value	Description
0x01	Configure MAC address filter number 1
0x02	Configure MAC address filter number 2
..	

Value	Description
N	Configure MAC address filter number <i>N</i>

2550

Table 67 – Possible settings for Address Type (3-bit field)

Value	Description
0x0	Unicast MAC address
0x1	Multicast MAC address
0x2–0x7	Reserved

2551

Table 68 – Possible settings for Enable Field (1-bit field)

Value	Description
0b	Disable this MAC address filter
1b	Enable this MAC address filter

2552 **8.4.32 Set MAC Address response (0x8E)**

2553 The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Set MAC
 2554 Address command and send a response (see Table 69).

2555

Table 69 – Set MAC Address response packet format

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

2556 Table 70 describes the reason code that is specific to the Set MAC Address command.

2557

Table 70 – Set MAC Address command-specific reason code

Value	Description	Comment
0x0E08	MAC Address Is Zero	Returned when the Set MAC Address command is received with the MAC address set to 0

2558 **8.4.33 Enable Broadcast Filter command (0x10)**

2559 The Enable Broadcast Filter command allows the Management Controller to control the forwarding of
 2560 broadcast frames to the Management Controller. The channel, upon receiving and processing this
 2561 command, shall filter all received broadcast frames based on the broadcast packet filtering settings
 2562 specified in the payload. If no broadcast packet types are specified for forwarding, all broadcast packets
 2563 shall be filtered out.

2564 The Broadcast Packet Filter Settings field is used to specify those protocol-specific broadcast filters that
 2565 should be activated. The channel indicates which broadcast filters it supports in the Broadcast Filter
 2566 Capabilities field of the Get Capabilities Response frame defined in 8.4.46.

2567 Table 71 illustrates the packet format of the Enable Broadcast Filter command.

2568 **Table 71 – Enable Broadcast Filter command packet format**

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

2569 Table 72 describes the Broadcast Packet Filter Settings field bit definitions.

2570 **Table 72 – Broadcast Packet Filter Settings field**

Bit Position	Field Description	Value Description
0	ARP Packets	<p>1b = Forward this packet type to the Management Controller. 0b = Filter out this packet type.</p> <p>For the purposes of this specification, an ARP broadcast packet is defined to be any packet that meets all of the following requirements:</p> <ul style="list-style-type: none"> The destination MAC address field is set to the layer 2 broadcast address (FF:FF:FF:FF:FF:FF). The Ethertype field set to 0x0806. <p>This field is mandatory.</p>
1	DHCP Client Packets	<p>1b = Forward this packet type to the Management Controller. 0b = Filter out this packet type.</p> <p>For the purposes of this filter, a DHCP client broadcast packet is defined to be any packet that meets all of the following requirements:</p> <ul style="list-style-type: none"> The destination MAC address field is set to the layer 2 broadcast address (FF:FF:FF:FF:FF:FF). The Ethertype field is set to 0x0800 (IPv4). The IP header's Protocol field is set to 17 (UDP). The UDP destination port number is set to 68. <p>This field is optional. If unsupported, broadcast DHCP client packets will be blocked when broadcast filtering is enabled. The value shall be set to 0 if unsupported.</p>

Bit Position	Field Description	Value Description
2	DHCP Server Packets	<p>1b = Forward this packet type to the Management Controller. 0b = Filter out this packet type.</p> <p>For the purposes of this filter, a DHCP server broadcast packet is defined to be any packet that meets all of the following requirements:</p> <ul style="list-style-type: none"> The destination MAC address field is set to the layer 2 broadcast address (FF:FF:FF:FF:FF:FF). The Ethertype field is set to 0x0800 (IPv4). The IP header's Protocol field is set to 17 (UDP). The UDP destination port number is set to 67. <p>This field is optional. If unsupported, broadcast DHCP packets will be blocked when broadcast filtering is enabled. The value shall be set to 0b if unsupported.</p>
3	NetBIOS Packets	<p>1b = Forward this packet type to the Management Controller. 0b = Filter out this packet type.</p> <p>For the purposes of this filter, NetBIOS broadcast packets are defined to be any packet that meets all of the following requirements:</p> <ul style="list-style-type: none"> The destination MAC address field is set to the layer 2 broadcast address (FF:FF:FF:FF:FF:FF). The Ethertype field is set to 0x0800 (IPv4). The IP header's Protocol field is set to 17 (UDP). The UDP destination port number is set to 137 for NetBIOS Name Service or 138 for NetBIOS Datagram Service, per the assignment of IANA well-known ports. <p>This field is optional. If unsupported, broadcast NetBIOS packets will be blocked when broadcast filtering is enabled. The value shall be set to 0b if unsupported.</p>
4..31	Reserved	None

2571 8.4.34 Enable Broadcast Filter response (0x90)

2572 The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Enable
2573 Broadcast Filter command and send a response.

2574 Currently no command-specific reason code is identified for this response (see Table 73).

2575 **Table 73 – Enable Broadcast Filter response packet format**

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

2576 8.4.35 Disable Broadcast Filter command (0x11)

2577 The Disable Broadcast Filter command may be used by the Management Controller to disable the
2578 broadcast filter feature and enable the reception of all broadcast frames. Upon processing this command,
2579 the channel shall discontinue the filtering of received broadcast frames.

2580 Table 74 illustrates the packet format of the Disable Broadcast Filter command.

2581 **Table 74 – Disable Broadcast Filter command packet format**

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

2582 8.4.36 Disable Broadcast Filter response (0x91)

2583 The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Disable
2584 Broadcast Filter command and send a response.

2585 Currently no command-specific reason code is identified for this response (see Table 75).

2586 **Table 75 – Disable Broadcast Filter response packet format**

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

2587 8.4.37 Enable Global Multicast Filter command (0x12)

2588 The Enable Global Multicast Filter command is used to activate global filtering of multicast frames with
2589 optional filtering of specific multicast protocols. Upon receiving and processing this command, the
2590 channel shall only deliver multicast frames that match specific multicast MAC addresses enabled for
2591 Pass-through using this command or the Set MAC Address command.

2592 The Multicast Packet Filter Settings field is used to specify optional, protocol-specific multicast filters that
2593 should be activated. The channel indicates which optional multicast filters it supports in the Multicast Filter
2594 Capabilities field of the Get Capabilities Response frame defined in 8.4.46. The Management Controller
2595 should not set bits in the Multicast Packet Filter Settings field that are not indicated as supported in the
2596 Multicast Filter Capabilities field.

2597 Neighbor Solicitation messages are sent to a Solicited Node multicast address that is derived from the
2598 target node's IPv6 address. This command may be used to enable forwarding of solicited node
2599 multicasts.

The IPv6 neighbor solicitation filter, as defined in this command, may not be supported by the Network Controller. In this case, the Management Controller may configure a multicast or mixed MAC address filter for the specific Solicited Node multicast address using the Set MAC Address command to enable forwarding of Solicited Node multicasts.

This command shall be implemented if the channel implementation supports accepting all multicast addresses. An implementation that does not support accepting all multicast addresses shall not implement these commands. Pass-through packets with multicast addresses can still be accepted depending on multicast address filter support provided by the Set MAC Address command. Multicast filter entries that are set to be enabled in the Set MAC Address command are accepted; all others are rejected. Table 76 illustrates the packet format of the Enable Global Multicast Filter command. Unsupported fields should be treated as reserved fields unless otherwise specified.

Table 76 – Enable Global Multicast Filter command packet format

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

Table 77 describes the bit definitions for the Multicast Packet Filter Settings field.

Table 77 – Bit Definitions for Multicast Packet Filter Settings field

Bit Position	Field Description	Value Description
0	IPv6 Neighbor Advertisement	<p>1b = Forward this packet type to the Management Controller. 0b = Filter out this packet type.</p> <p>For the purposes of this specification, an IPv6 Neighbor Advertisement multicast packet is defined to be any packet that meets all of the following requirements:</p> <ul style="list-style-type: none"> The destination MAC address field is set to a layer 2 multicast address of the form 33:33:00:00:00:01. This address corresponds to the all-nodes multicast address (FF02::1). The Ethertype field is set to 0x86DD (IPv6). The IPv6 header's Next Header field is set to 58 (ICMPv6). The ICMPv6 header's Message Type field is set to the following value: 136 – Neighbor Advertisement. <p>This field is optional.</p>

Bit Position	Field Description	Value Description
1	IPv6 Router Advertisement	<p>1b = Forward this packet type to the Management Controller. 0b = Filter out this packet type.</p> <p>For the purposes of this specification, an IPv6 Router Advertisement multicast packet is defined to be any packet that meets all of the following requirements:</p> <ul style="list-style-type: none"> The destination MAC address field is set to a layer 2 multicast address of the form 33:33:00:00:00:01. This corresponds to the all-nodes multicast address (FF02::1). The Ethertype field is set to 0x86DD (IPv6). The IPv6 header's Next Header field is set to 58 (ICMPv6). The ICMPv6 header's Message Type field is set to 134. <p>This field is optional.</p>
2	DHCPv6 relay and server multicast	<p>1b = Forward this packet type to the Management Controller. 0b = Filter out this packet type.</p> <p>For the purposes of this filter, a DHCPv6 multicast packet is defined to be any packet that meets all of the following requirements:</p> <ul style="list-style-type: none"> The destination MAC address field is set to the layer 2 multicast address 33:33:00:01:00:02 or 33:33:00:01:00:03. These correspond to the IPv6 multicast addresses FF02::1:2 (All_DHCP_Relay_Agents_and_Servers) and FF05::1:3 (All_DHCP_Servers). The Ethertype field is set to 0x86DD (IPv6). The IPv6 header's Next Header field is set to 17 (UDP). The UDP destination port number is set to 547. <p>This field is optional.</p>
3	DHCPv6 multicasts from server to clients listening on well-known UDP ports	<p>1b = Forward this packet type to the Management Controller. 0b = Filter out this packet type.</p> <p>For the purposes of this filter, a DHCPv6 multicast packet is defined to be any packet that meets all of the following requirements:</p> <ul style="list-style-type: none"> The destination MAC address field is set to the layer 2 multicast address 33:33:00:01:00:02. These correspond to the IPv6 multicast addresses FF02::1:2 (All_DHCP_Relay_Agents_and_Servers). The Ethertype field is set to 0x86DD (IPv6). The IPv6 header's Next Header field is set to 17 (UDP). The UDP destination port number is set to 546. <p>This field is optional.</p>

Bit Position	Field Description	Value Description
4	IPv6 MLD	<p>1b = Forward this packet type to the Management Controller. 0b = Filter out this packet type.</p> <p>For the purposes of this specification, an IPv6 MLD packet is defined to be any packet that meets all of the following requirements:</p> <ul style="list-style-type: none"> The destination MAC address field is set to a layer 2 multicast address of the form 33:33:00:00:00:01. This address corresponds to the All_Nodes (FF02::1) multicast address. The Ethertype field is set to 0x86DD (IPv6). The IPv6 header's Next Header field is set to 58 (ICMPv6). The ICMPv6 header's Message Type field is set to one of the following values: 130 (Multicast Listener Query), 131 (Multicast Listener Report), 132 (Multicast Listener Done) <p>This field is optional.</p>
5	IPv6 Neighbor Solicitation	<p>1b = Forward this packet type to the Management Controller. 0b = Filter out this packet type.</p> <p>For the purposes of this specification, an IPv6 MLD packet is defined to be any packet that meets all of the following requirements:</p> <ul style="list-style-type: none"> The destination MAC address field is set to a layer 2 multicast address of the form 33:33:FF:XX:XX:XX. This address corresponds to the Solicited Node multicast address where the last three bytes of the destination MAC address are ignored for this filter. The Ethertype field is set to 0x86DD (IPv6). The IPv6 header's Next Header field is set to 58 (ICMPv6). The ICMPv6 header's Message Type field is set to one of the following values: 135 <p>This field is optional.</p> <p>IMPLEMENTATION NOTE Enabling of this filter results in receiving all IPv6 neighbor solicitation traffic on this channel. If IPv6 neighbor solicitation traffic for a specific multicast address is of interest, then it is recommended that the MC uses a multicast address filter (configured for the multicast address using the Set MAC Address command) instead of this filter.</p>

Bit Position	Field Description	Value Description
6	LLDP	<p>1b = Forward this packet type to the Management Controller. 0b = Filter out this packet type.</p> <p>For the purposes of this specification, a LLDP packet is defined to be any packet that meets all of the following requirements:</p> <ul style="list-style-type: none"> The destination MAC address field is set to a layer 2 multicast address of the form 01:80:C2:00:00:00, or 01:80:C2:00:00:03, or 01:80:C2:00:00:0E. The Ethertype field is set to 0x88CC. <p>This field is optional.</p> <p>Implementation Note: Enabling of this filter results in receiving a copy of all LLDP traffic on this channel. If LLDP traffic for a specific LLDP multicast address is of interest, then it is recommended that the MC uses a multicast address filter (configured for the multicast address using the Set MAC Address command) instead of this filter.</p> <p>The intent of this filter is to allow the MC to snoop the received LLDP frame by the port, not to achieve ownership of any contained protocols.</p>
7	mDNSv4	<p>1b = Forward this packet type to the Management Controller. 0b = Filter out this packet type.</p> <p>For the purposes of this specification, a mDNS/IPv4 packet is defined to be any packet that meets all the following requirements:</p> <ul style="list-style-type: none"> The destination MAC address field is set to a layer 2 multicast address of the form 01:00:5E:00:00:FB. The Ethertype field is set to 0x0800. The IPv4 address is 224.0.0.251. The IPv4 header's Protocol field is set to 17 (UDP). The UDP destination port number is set to 5353. <p>This field is optional.</p>
8	mDNSv6	<p>1b = Forward this packet type to the Management Controller. 0b = Filter out this packet type.</p> <p>For the purposes of this specification, a mDNS/IPv6 packet is defined to be any packet that meets all the following requirements:</p> <ul style="list-style-type: none"> The destination MAC address field is set to a layer 2 multicast address of the form 33:33:00:00:00:FB. This corresponds to the All Nodes IPv6 multicast address, FF02::FB. The Ethertype field is set to 0x086DD. The IPv6 header's Next Header field is set to 17 (UDP). The UDP destination port number is set to 5353. <p>This field is optional.</p>
31..9	Reserved	None

2615 8.4.38 Enable Global Multicast Filter response (0x92)

2616 The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Enable
2617 Global Multicast Filter command and send a response.

2618 Currently no command-specific reason code is identified for this response (see Table 78).

2619 **Table 78 – Enable Global Multicast Filter response packet format**

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

2620 8.4.39 Disable Global Multicast Filter command (0x13)

2621 The Disable Global Multicast Filter command is used to disable global filtering of multicast frames. Upon
2622 receiving and processing this command, and regardless of the current state of multicast filtering, the
2623 channel shall forward all multicast frames to the Management Controller.

2624 This command shall be implemented on the condition that the channel implementation supports accepting
2625 all multicast addresses. An implementation that does not support accepting all multicast addresses shall
2626 not implement these commands. Pass-through packets with multicast addresses can still be accepted
2627 depending on multicast address filter support provided by the Set MAC Address command. Packets with
2628 destination addresses matching multicast filter entries that are set to enabled in the Set MAC Address
2629 command are accepted; all others are rejected.

2630 Table 79 illustrates the packet format of the Disable Global Multicast Filter command.

2631 **Table 79 – Disable Global Multicast Filter command packet format**

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

2632 8.4.40 Disable Global Multicast Filter response (0x93)

2633 In the absence of any errors, the channel shall process and respond to the Disable Global Multicast Filter
2634 command by sending the response packet shown in Table 80.

2635 Currently no command-specific reason code is identified for this response.

2636

Table 80 – Disable Global Multicast Filter response packet format

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

2637 8.4.41 Set NC-SI Flow Control command (0x14)

2638 The Set NC-SI Flow Control command allows the Management Controller to configure [IEEE 802.3](#) pause
 2639 packet flow control on the NC-SI.

2640 The Set NC-SI Flow Control command is addressed to the package, rather than to a particular channel
 2641 (that is, the command is sent with a Channel ID where the Package ID subfield matches the ID of the
 2642 intended package and the Internal Channel ID subfield is set to 0x1F).

2643 The setting of [IEEE 802.3](#) Pause packet flow control on RBT is independent from any arbitration scheme,
 2644 if any is used.

2645 Table 81 illustrates the packet format of the Set NC-SI Flow Control command.

2646

Table 81 – Set NC-SI Flow Control command packet format

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

2647 Table 82 describes the values for the Flow Control Enable field.

2648

Table 82 – Values for the Flow Control Enable field (8-bit field)

Value	Description
0x0	Disables NC-SI flow control
0x1	Enables Network Controller to Management Controller flow control frames (Network Controller generates flow control frames) This field is optional.
0x2	Enables Management Controller to Network Controller flow control frames (Network Controller accepts flow control frames) This field is optional.

Value	Description
0x3	Enables bi-directional flow control frames This field is optional.
0x4..0xFF	Reserved

2649 8.4.42 Set NC-SI Flow Control response (0x94)

2650 The package shall, in the absence of a checksum error or identifier mismatch, always accept the Set
2651 NC-SI Flow Control command and send a response (see Table 83).

2652 **Table 83 – Set NC-SI Flow Control response packet format**

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

2653 Table 84 describes the reason code that is specific to the Set NC-SI Flow Control command.

2654 **Table 84 – Set NC-SI Flow Control command-specific reason code**

Value	Description	Comment
0x1409	Independent transmit and receive enable/disable control is not supported	Returned when the implementation requires that both transmit and receive flow control be enabled and disabled simultaneously

2655 8.4.43 Get Version ID command (0x15)

2656 The Get Version ID command may be used by the Management Controller to request the channel to
2657 provide the controller and firmware type and version strings listed in the response payload description.

2658 Table 85 illustrates the packet format of the Get Version ID command.

2659 **Table 85 – Get Version ID command packet format**

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

2660 **8.4.44 Get Version ID Response (0x95)**

2661 The channel shall, in the absence of an error, always accept the Get Version ID command and send the
 2662 response packet shown in Table 86. Currently no command-specific reason code is identified for this
 2663 response.

2664 Note: When multiple Physical Functions are enabled on the channel, the PCI ID that is returned shall be
 2665 that of the lowest numbered Function on the channel.

2666 **Table 86 – Get Version ID response packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	NC-SI Version			
	Major	Minor	Update	Alpha1
24..27	reserved	reserved	reserved	Alpha2
28..31	Firmware Name String (11-08)			
32..35	Firmware Name String (07-04)			
36..39	Firmware Name String (03-00)			
40..43	Firmware Version			
	MS-byte (3)	Byte (2)	Byte (1)	LS-byte (0)
44..47	PCI DID		PCI VID	
48..51	PCI SSID		PCI SVID	
52..55	Manufacturer ID (IANA)			
56..59	Checksum			

2667 **8.4.44.1 NC-SI Version encoding**

2668 The NC-SI Version field holds the version number of the NC-SI specification with which the controller is
 2669 compatible. The version field shall be encoded as follows:

- 2670 • The ‘major’, ‘minor’, and ‘update’ bytes are BCD-encoded, and each byte holds two BCD digits.
- 2671 • The ‘alpha’ byte holds an optional alphanumeric character extension that is encoded using the
 2672 ISO/IEC 8859-1 Character Set.
- 2673 • The semantics of these fields follow the semantics specified in [DSP4014](#).
- 2674 • The value 0x00 in the Alpha1 or Alpha2 fields means that the corresponding alpha field is not
 2675 used. The Alpha1 field shall be used first.
- 2676 • The value 0xF in the most-significant nibble of a BCD-encoded value indicates that the most-
 2677 significant nibble should be ignored and the overall field treated as a single digit value.
- 2678 • A value of 0xFF in the update field indicates that the entire field is not present. 0xFF is not
 2679 allowed as a value for the major or minor fields.

2680 EXAMPLE: Version 3.7.10a → 0xF3F7106100
 2681 Version 10.01.7 → 0x1001F70000
 2682 Version 3.1 → 0xF3F1FF0000
 2683 Version 1.0a → 0xF1F0FF4100
 2684 Version 1.0ab → 0xF1F0FF4142 (Alpha1 = 0x41, Alpha2 = 0x42)

2685 8.4.44.2 Firmware Name encoding

2686 The Firmware Name String shall be encoded using the ISO/IEC 8859-1 Character Set. Strings are left-
 2687 justified where the leftmost character of the string occupies the most-significant byte position of the
 2688 Firmware Name String field, and characters are populated starting from that byte position. The string is
 2689 null terminated if the string is smaller than the field size. That is, the delimiter value, 0x00, follows the last
 2690 character of the string if the string occupies fewer bytes than the size of the field allows. A delimiter is not
 2691 required if the string occupies the full size of the field. Bytes following the delimiter (if any) should be
 2692 ignored and can be any value.

2693 8.4.44.3 Firmware Version encoding

2694 To facilitate a common way of representing and displaying firmware version numbers across different
 2695 vendors, each byte is hexadecimal encoded where each byte in the field holds two hexadecimal digits.
 2696 The Firmware Version field shall be encoded as follows. The bytes are collected into a single 32-bit field
 2697 where each byte represents a different 'point number' of the overall version. The selection of values that
 2698 represent a particular version of firmware is specific to the Network Controller vendor.

2699 Software displaying these numbers should not suppress leading zeros, which should help avoid user
 2700 confusion in interpreting the numbers. For example, consider the two values 0x05 and 0x31.
 2701 Numerically, the byte 0x31 is greater than 0x05, but if leading zeros were incorrectly suppressed, the two
 2702 displayed values would be ".5" and ".31", respectively, and a user would generally interpret 0.5 as
 2703 representing a greater value than 0.31 instead of 0.05 being smaller than 0.31. Similarly, if leading zeros
 2704 were incorrectly suppressed, the value 0x01 and 0x10 would be displayed as 0.1 and 0.10, which could
 2705 potentially be misinterpreted as representing the same version instead of 0.01 and 0.10 versions.

2706 EXAMPLE: 0x00030217 → Version 00.03.02.17
 2707 0x010100A0 → Version 01.01.00.A0

2708 8.4.44.4 PCI ID fields

2709 These fields (PCI DID, PCI VID, PCI SSID, PCI SVID) hold the PCI ID information for the Network
 2710 Controller when the Network Controller incorporates a PCI or PCI Express™ interface that provides a
 2711 host network interface connection that is shared with the NC-SI connection to the network.

2712 If this field is not used, the values shall all be set to zeros (0000h). Otherwise, the fields shall hold the
 2713 PCI ID information for the host interface as defined by the version of the PCI/PCI Express™ specification
 2714 to which the device's interface was designed.

2715 If multiple partitions are enabled on the channel, the values should represent the PCI ID of the lowest
 2716 Function number assigned to the channel by the Set PF Assignment command (0x28).

2717 8.4.44.5 Manufacturer ID (IANA) field

2718 The Manufacturer ID holds the [IANA Enterprise Number](#) for the manufacturer of the Network Controller as
 2719 a 32-bit binary number. If the field is unused, the value shall be set to 0xFFFFFFFF.

2720 **8.4.45 Get Capabilities command (0x16)**

2721 The Get Capabilities command is used to discover additional optional functions supported by the channel,
 2722 such as the number of unicast/multicast addresses supported, the amount of buffering in bytes available
 2723 for packets bound for the Management Controller, and so on.

2724 Table 87 illustrates the packet format for the Get Capabilities command.

2725 **Table 87 – Get Capabilities command packet format**

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

2726 **8.4.46 Get Capabilities response (0x96)**

2727 In the absence of any errors, the channel shall process and respond to the Get Capabilities Command
 2728 and send the response packet shown in Table 88. Currently no command-specific reason code is
 2729 identified for this response.

2730 **Table 88 – Get Capabilities response packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	Capabilities Flags			
24..27	Broadcast Packet Filter Capabilities			
28..31	Multicast Packet Filter Capabilities			
32..35	Buffering Capability			
36..39	AEN Control Support			
40..43	VLAN Filter Count	Mixed Filter Count	Multicast Filter Count	Unicast Filter Count
44..47	Reserved		VLAN Mode Support	Channel Count
48..51	Checksum			

2731 **8.4.46.1 Capabilities Flags field**

2732 The Capabilities Flags field indicates which optional features of this specification the channel supports, as
 2733 described in Table 89.

2734

Table 89 – Capabilities Flags bit definitions

Bit Position	Field Description	Value Description
0	Hardware Arbitration Capability	0b = Hardware arbitration capability is not supported by the package. 1b = Hardware arbitration capability is supported by the package.
1	Host NC Driver Status	0b = Host NC Driver Indication status is not supported. 1b = Host NC Driver Indication status is supported. See Table 51 for the definition of Host NC Driver Indication Status.
2	Network Controller to Management Controller Flow Control Support	0b = Network Controller to Management Controller flow control is not supported. 1b = Network Controller to Management Controller flow control is supported.
3	Management Controller to Network Controller Flow Control Support	0b = Management Controller to Network Controller flow control is not supported. 1b = Management Controller to Network Controller flow control is supported.
4	All multicast addresses support	0b = The channel cannot accept all multicast addresses. The channel does not support enable/disable global multicast commands. 1b = The channel can accept all multicast addresses. The channel supports enable/disable global multicast commands.
6..5	Hardware Arbitration Implementation Status	00b = Unknown 01b = Hardware arbitration capability is not implemented for the package on the given system. 10b = Hardware arbitration capability is implemented for the package on the given system. 11b = Reserved.
7	Thermal shutdown Implementation Status	0b = The thermal self-shutdown capability is not supported by the channel (package). 1b = The thermal self-shutdown capability is supported by the channel (package).
8	Delayed Response Support	0b = Delayed response operation and signaling is not supported by the channel (package). 1b = Delayed response operation and signaling is supported by the channel (package).
9..31	Reserved	Reserved

2735 **8.4.46.2 Broadcast Packet Filter Capabilities field**

2736 The Broadcast Packet Filter Capabilities field defines the optional broadcast packet filtering capabilities
 2737 that the channel supports. The bit definitions for this field correspond directly with the bit definitions for the
 2738 Broadcast Packet Filter Settings field defined for the Enable Broadcast Filter command in Table 72. A bit

2739 set to 1 indicates that the channel supports the filter associated with that bit position; otherwise, the
 2740 channel does not support that filter.

2741 **8.4.46.3 Multicast Packet Filter Capabilities field**

2742 The Multicast Packet Filter Capabilities field defines the optional multicast packet filtering capabilities that
 2743 the channel supports. The bit definitions for this field correspond directly with the bit definitions for the
 2744 Multicast Packet Filter Settings field defined for the Enable Global Multicast Filter command in Table 77.
 2745 A bit set to 1 indicates that the channel supports the filter associated with that bit position; otherwise, the
 2746 channel does not support that filter.

2747 **8.4.46.4 Buffering Capability field**

2748 The Buffering Capability field defines the amount of buffering in bytes that the channel provides for
 2749 inbound packets destined for the Management Controller. The Management Controller may make use of
 2750 this value in software-based Device Selection implementations to determine the relative time for which a
 2751 specific channel may be disabled before it is likely to start dropping packets. A value of 0 indicates that
 2752 the amount of buffering is unspecified.

2753 **8.4.46.5 AEN Control Support field**

2754 The AEN Control Support field indicates various standard AENs supported by the implementation. The
 2755 format of the field is shown in Table 41.

2756 **8.4.46.6 VLAN Filter Count field**

2757 The VLAN Filter Count field indicates the number of VLAN filters, up to 15, that the channel supports, as
 2758 defined by the Set VLAN Filter command.

2759 **8.4.46.7 Mixed, Multicast, and Unicast Filter Count fields**

2760 The Mixed Filter Count field indicates the number of mixed address filters that the channel supports. A
 2761 mixed address filter can be used to filter on specific unicast or multicast MAC addresses.

2762 The Multicast Filter Count field indicates the number of multicast MAC address filters that the channel
 2763 supports.

2764 The Unicast Filter Count field indicates the number of unicast MAC address filters that the channel
 2765 supports.

2766 The channel is required to support at least one unicast or mixed filter, such that at least one unicast MAC
 2767 address can be configured on the interface. The total number of unicast, multicast, and mixed filters shall
 2768 not exceed 8.

2769 **8.4.46.8 VLAN Mode Support field**

2770 The VLAN Mode Support field indicates various modes supported by the implementation. The format of
 2771 field is defined in Table 90.

2772 **Table 90 – VLAN Mode Support bit definitions**

Bit Position	Field Description	Value Description
0	VLAN only	1 = VLAN shall be supported in the implementation.

1	VLAN + non-VLAN	0 = Filtering 'VLAN + non-VLAN' traffic is not supported in the implementation. 1 = Filtering 'VLAN + non-VLAN' traffic is supported in the implementation.
2	Any VLAN + non-VLAN	0 = Filtering 'Any VLAN + non-VLAN' traffic is not supported in the implementation. 1 = Filtering 'Any VLAN + non-VLAN' traffic is supported in the implementation.
3..7	Reserved	0

2773 8.4.46.9 Channel Count field

2774 The Channel Count field indicates the number of channels supported by the Network Controller.

2775 8.4.47 Get Parameters command (0x17)

2776 The Get Parameters command can be used by the Management Controller to request that the channel
2777 send the Management Controller a copy of all of the currently stored parameter settings that have been
2778 put into effect by the Management Controller, plus "other" Host/Channel parameter values that may be
2779 added to the Get Parameters Response Payload.

2780 Table 91 illustrates the packet format for the Get Parameters command.

2781 **Table 91 – Get Parameters command packet format**

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

2782 8.4.48 Get Parameters response (0x97)

2783 The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Get
2784 Parameters command and send a response. As shown in Table 92, each parameter shall return the value
2785 that was set by the Management Controller. If the parameter is not supported, 0 is returned. Currently no
2786 command-specific reason code is identified for this response.

2787 The payload length of this response packet will vary according to how many MAC address filters or VLAN
2788 filters the channel supports. All supported MAC addresses are returned at the end of the packet, without
2789 any intervening padding between MAC addresses.

2790 MAC addresses are returned in the following order: unicast filtered addresses first, followed by multicast
2791 filtered addresses, followed by mixed filtered addresses, with the number of each corresponding to those
2792 reported through the Get Capabilities command. For example, if the interface reports four unicast filters,
2793 two multicast filters, and two mixed filters, then MAC addresses 1 through 4 are those currently
2794 configured through the interface's unicast filters, MAC addresses 5 and 6 are those configured through
2795 the multicast filters, and 7 and 8 are those configured through the mixed filters. Similarly, if the interface
2796 reports two unicast filters, no multicast filters, and six mixed filters, then MAC addresses 1 and 2 are

2797 those currently configured through the unicast filters, and 3 through 8 are those configured through the
 2798 mixed filters.

2799 **Table 92 – Get Parameters response packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	MAC Address Count	Reserved		MAC Address Flags
24..27	VLAN Tag Count	Reserved	VLAN Tag Flags	
28..31	Link Settings			
32..35	Broadcast Packet Filter Settings			
36..39	Configuration Flags			
40..43	VLAN Mode	Flow Control Enable	Reserved	
44..47	AEN Control			
48..51	MAC Address 1 byte 5	MAC Address 1 byte 4	MAC Address 1 byte 3	MAC Address 1 byte 2
52..55 ^a	MAC Address 1 byte 1	MAC Address 1 byte 0	MAC Address 2 byte 5	MAC Address 2 byte 4
56..59	MAC Address 2 byte 3	MAC Address 2 byte 2	MAC Address 2 byte 1	MAC Address 2 byte 0
variable	...			
	VLAN Tag 1		VLAN Tag 2	
	...			
	...		Pad (if needed)	
	Checksum			

^a Variable fields can start at this byte offset.

2800 Table 93 lists the parameters for which values are returned in this response packet.

2801 **Table 93 – Get Parameters data definition**

Parameter Field Name	Description
MAC Address Count	The number of MAC addresses supported by the channel
MAC Address Flags	The enable/disable state for each supported MAC address See Table 94.
VLAN Tag Count	The number of VLAN Tags supported by the channel
VLAN Tag Flags	The enable/disable state for each supported VLAN Tag See Table 95.

Parameter Field Name	Description
Link Settings	The 32-bit Link Settings value as defined in the Set Link command. See Table 44.
Broadcast Packet Filter Settings	The current 32-bit Broadcast Packet Filter Settings value
Configuration Flags	See Table 96.
VLAN Mode	See Table 61.
Flow Control Enable	See Table 82.
AEN Control	See Table 41.
MAC Address 1..8	The current contents of up to eight 6-byte MAC address filter values.
VLAN Tag 1..15	The current contents of up to 15 16-bit VLAN Tag filter values
.	

2802 The format of the MAC Address Flags field is defined in Table 94.

2803 **Table 94 – MAC Address Flags bit definitions**

Bit Position	Field Description	Value Description
0	MAC address 1 status	0b = Default or unsupported or disabled 1b = Enabled
1	MAC address 2 status, or Reserved	0b = Default or unsupported or disabled 1b = Enabled
2	MAC address 3 status, or Reserved	0b = Default or unsupported or disabled 1b = Enabled
...
7	MAC address 8 status, or Reserved	0b = Default or unsupported or disabled 1b = Enabled

2804 The format of the VLAN Tag Flags field is defined in Table 95.

2805 **Table 95 – VLAN Tag Flags bit definitions**

Bit Position	Field Description	Value Description
0	VLAN Tag 1 status	0b = Default or unsupported or disabled 1b = Enabled
1	VLAN Tag 2 status, or Reserved	0b = Default or unsupported or disabled 1b = Enabled
2	VLAN Tag 3 status, or Reserved	0b = Default or unsupported or disabled 1b = Enabled
...
14	VLAN Tag 15 status, or Reserved	0b = Default or unsupported or disabled 1b = Enabled

2806 The format of the Configuration Flags field is defined in Table 96.

2807

Table 96 – Configuration Flags bit definitions

Bit Position	Field Description	Value Description
0	Broadcast Packet Filter status	0b = Disabled 1b = Enabled
1	Channel Enabled	0b = Disabled 1b = Enabled
2	Channel Network TX Enabled	0b = Disabled 1b = Enabled
3	Global Multicast Packet Filter Status	0b = Disabled 1b = Enabled
4..31	Reserved	Reserved

2808 **8.4.49 Get Controller Packet Statistics command (0x18)**

2809 The Get Controller Packet Statistics command may be used by the Management Controller to request a
 2810 copy of the aggregated Ethernet packet statistics that the channel maintains for its external interface to
 2811 the LAN network. The statistics are an aggregation of statistics for both the host side traffic and the NC-SI
 2812 Pass-through traffic.

2813

Table 97 – Get Controller Packet Statistics command packet format

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

2814 **8.4.50 Get Controller Packet Statistics response (0x98)**

2815 The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Get
 2816 Controller Packet Statistics command and send the response packet shown in Table 98.

2817 The Get Controller Packet Statistics Response frame contains a set of Ethernet statistics counters that
 2818 monitor the LAN traffic in the Network Controller. Implementation of the counters listed in Table 99 is
 2819 optional. The Network Controller shall return any unsupported counter with a value of 0xFFFFFFFF for
 2820 32-bit counters and 0xFFFFFFFFFFFFFFFF for 64-bit counters.

2821

Table 98 – Get Controller Packet Statistics response packet format

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	Counters Cleared from Last Read (MS Bits)			
24..27	Counters Cleared from Last Read (LS Bits)			
28..35	Total Bytes Received			
36..43	Total Bytes Transmitted			
44..51	Total Unicast Packets Received			
52..59	Total Multicast Packets Received			
60..67	Total Broadcast Packets Received			
68..75	Total Unicast Packets Transmitted			
76..83	Total Multicast Packets Transmitted			
84..91	Total Broadcast Packets Transmitted			
92..95	FCS Receive Errors			
96..99	Alignment Errors			
100..103	False Carrier Detections			
104..107	Runt Packets Received			
108..111	Jabber Packets Received			
112..115	Pause XON Frames Received			
116..119	Pause XOFF Frames Received			
120..123	Pause XON Frames Transmitted			
124..127	Pause XOFF Frames Transmitted			
128..131	Single Collision Transmit Frames			
132..135	Multiple Collision Transmit Frames			
136..139	Late Collision Frames			
140..143	Excessive Collision Frames			
144..147	Control Frames Received For version 1.2, this counter may include Priority flow control packets			
148..151	64-Byte Frames Received			
152..155	65–127 Byte Frames Received			
156..159	128–255 Byte Frames Received			
160..163	256–511 Byte Frames Received			
164..167	512–1023 Byte Frames Received			
168..171	1024–1522 Byte Frames Received			
172..175	1523–9022 Byte Frames Received			

Bytes	Bits			
	31..24	23..16	15..08	07..00
176..179	64-Byte Frames Transmitted			
180..183	65–127 Byte Frames Transmitted			
184..187	128–255 Byte Frames Transmitted			
188..191	256–511 Byte Frames Transmitted			
192..195	512–1023 Byte Frames Transmitted			
196..199	1024–1522 Byte Frames Transmitted			
200..203	1523–9022 Byte Frames Transmitted			
204..211	Valid Bytes Received			
212..215	Error Runt Packets Received			
216..219	Error Jabber Packets Received			
220..223	Checksum			

2822

Table 99 – Get Controller Packet Statistics counters

Counter Number	Name	Meaning
0	Total Bytes Received	Counts the number of bytes received
1	Total Bytes Transmitted	Counts the number of bytes transmitted
2	Total Unicast Packets Received	Counts the number of good (FCS valid) packets received that passed L2 filtering by a specific MAC address
3	Total Multicast Packets Received	Counts the number of good (FCS valid) multicast packets received
4	Total Broadcast Packets Received	Counts the number of good (FCS valid) broadcast packets received
5	Total Unicast Packets Transmitted	Counts the number of good (FCS valid) packets transmitted that passed L2 filtering by a specific MAC address
6	Total Multicast Packets Transmitted	Counts the number of good (FCS valid) multicast packets transmitted
7	Total Broadcast Packets Transmitted	Counts the number of good (FCS valid) broadcast packets transmitted
8	FCS Receive Errors	Counts the number of receive packets with FCS errors
9	Alignment Errors	Counts the number of receive packets with alignment errors
10	False Carrier Detections	Counts the false carrier errors reported by the PHY
11	Runt Packets Received	Counts the number of received frames that passed address filtering, were less than minimum size (64 bytes from <Destination Address> through <FCS>, inclusively), and had a valid FCS

Counter Number	Name	Meaning
12	Jabber Packets Received	Counts the number of received frames that passed address filtering, were greater than the maximum size, and had a valid FCS
13	Pause XON Frames Received	Counts the number of XON packets received from the network
14	Pause XOFF Frames Received	Counts the number of XOFF packets received from the network
15	Pause XOFF Frames Transmitted	Counts the number of XON packets transmitted to the network
16	Pause XOFF Frames Transmitted	Counts the number of XOFF packets transmitted to the network
17	Single Collision Transmit Frames	Counts the number of times that a successfully transmitted packet encountered a single collision
18	Multiple Collision Transmit Frames	Counts the number of times that a transmitted packet encountered more than one collision but fewer than 16
19	Late Collision Frames	Counts the number of collisions that occurred after one slot time (defined by IEEE 802.3)
20	Excessive Collision Frames	Counts the number of times that 16 or more collisions occurred on a single transmit packet
21	Control Frames Received	Counts the number of MAC control frames received that are <i>not</i> XON or XOFF flow control frames
22	64 Byte Frames Received	Counts the number of good packets received that are exactly 64 bytes (from <Destination Address> through <FCS>, inclusively) in length
23	65–127 Byte Frames Received	Counts the number of good packets received that are 65–127 bytes (from <Destination Address> through <FCS>, inclusively) in length
24	128–255 Byte Frames Received	Counts the number of good packets received that are 128–255 bytes (from <Destination Address> through <FCS>, inclusively) in length
25	256–511 Byte Frames Received	Counts the number of good packets received that are 256–511 bytes (from <Destination Address> through <FCS>, inclusively) in length
26	512–1023 Byte Frames Received	Counts the number of good packets received that are 512–1023 bytes (from <Destination Address> through <FCS>, inclusively) in length
27	1024–1522 Byte Frames Received	Counts the number of good packets received that are 1024–1522 bytes (from <Destination Address> through <FCS>, inclusively) in length
28	1523–9022 Byte Frames Received	Counts the number of received frames that passed address filtering and were greater than 1523 bytes in length
29	64 Byte Frames Transmitted	Counts the number of good packets transmitted that are exactly 64 bytes (from <Destination Address> through <FCS>, inclusively) in length

Counter Number	Name	Meaning
30	65–127 Byte Frames Transmitted	Counts the number of good packets transmitted that are 65–127 bytes (from <Destination Address> through <FCS>, inclusively) in length
31	128–255 Byte Frames Transmitted	Counts the number of good packets transmitted that are 128–255 bytes (from <Destination Address> through <FCS>, inclusively) in length
32	256–511 Byte Frames Transmitted	Counts the number of good packets transmitted that are 256–511 bytes (from <Destination Address> through <FCS>, inclusively) in length
33	512–1023 Byte Frames Transmitted	Counts the number of good packets transmitted that are 512–1023 bytes (from <Destination Address> through <FCS>, inclusively) in length
34	1024–1522 Byte Frames Transmitted	Counts the number of good packets transmitted that are 1024–1522 bytes (from <Destination Address> through <FCS>, inclusively) in length
35	1523–9022 Byte Frames Transmitted	Counts the number of transmitted frames that passed address filtering and were greater than 1523 in length
36	Valid Bytes Received	Counts the bytes received in all packets that did not manifest any type of error
37	Error Runt Packets Received	Counts the number of invalid frames that were less than the minimum size (64 bytes from <Destination Address> through <FCS>, inclusively)
38	Error Jabber Packets Received	Counts Jabber packets, which are defined as packets that exceed the programmed MTU size <i>and</i> have a bad FCS value

2823 The Network Controller shall also indicate in the Counters Cleared from Last Read fields whether the
 2824 corresponding field has been cleared by means other than NC-SI (possibly by the host) since it was last
 2825 read by means of the NC-SI. Counting shall resume from 0 after a counter has been cleared. The
 2826 Counters Cleared from Last Read field's format is shown in Table 100.

2827 Currently no command-specific reason code is identified for this response.

2828 **Table 100 – Counters Cleared from Last Read Fields format**

Field	Bits	Mapped to Counter Numbers
MS Bits	0..6	32..38
	7..31	Reserved
LS Bits	0..31	0..31

2829 IMPLEMENTATION NOTE The Get Controller Packet Statistics response contains the following counters related
 2830 to flow control: Pause XON Frames Received, Pause XOFF Frames Received, Pause
 2831 XON Frames Transmitted, and Pause XOFF Frames Transmitted. An implementation
 2832 can optionally include Priority-Based Flow Control (PFC) packets in these counters.

2833 8.4.51 Get NC-SI Statistics command (0x19)

2834 In addition to the packet statistics accumulated on the LAN network interface, the channel separately
 2835 accumulates a variety of NC-SI specific packet statistics for the channel. The Get NC-SI Statistics
 2836 command may be used by the Management Controller to request that the channel send a copy of all
 2837 current NC-SI packet statistic values for the channel. The implementation may or may not include
 2838 statistics for commands that are directed to the package.

2839 Table 101 illustrates the packet format of the Get NC-SI Statistics command.

2840 **Table 101 – Get NC-SI Statistics command packet format**

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

2841 8.4.52 Get NC-SI Statistics response (0x99)

2842 In the absence of any error, the channel shall process and respond to the Get NC-SI Statistics command
 2843 by sending the response packet and payload shown in Table 102.

2844 **Table 102 – Get NC-SI Statistics response packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	NC-SI Commands Received			
24..27	NC-SI Control Packets Dropped			
28..31	NC-SI Command Type Errors			
32..35	NC-SI Command Checksum Errors			
36..39	NC-SI Receive Packets			
40..43	NC-SI Transmit Packets			
44..47	AENs Sent			
48..51	Checksum			

2845 The Get NC-SI Statistics Response frame contains a set of statistics counters that monitor the NC-SI
 2846 traffic in the Network Controller. Counters that are supported shall be reset to 0x0 when entering the
 2847 Initial State and after being read. Implementation of the counters shown in Table 103 is optional. The
 2848 Network Controller shall return any unsupported counter with a value of 0xFFFFFFFF. Counters may
 2849 wraparound or stop if they reach 0xFFFFFFFFE. It is vendor-specific how NC-SI commands that are sent
 2850 to the package ID are included in the NC-SI statistics.

2851 Currently no command-specific reason code is identified for this response.

2852 **Table 103 – Get NC-SI Statistics counters**

Counter Number	Name	Meaning
1	NC-SI Commands Received	For packets that are not dropped, this field returns the number of NC-SI Control Packets received and identified as NC-SI commands.
2	NC-SI Control Packets Dropped	Counts the number of NC-SI Control Packets that were received and dropped (Packets with correct FCS and Ethertype, but are dropped for one of the other reasons listed in 6.9.2.1). NC-SI Control Packets that were dropped because the channel ID was not valid may not be included in this statistics counter.
3	NC-SI Unsupported Commands Received	Counts the number of NC-SI command packets that were received but are not supported. (Network controller responded to the command with a Command Unsupported response code).
4	NC-SI Command Checksum Errors	Counts the number of NC-SI Control Packets that were received but dropped because of an invalid checksum (if checksum is provided and checksum validation is supported by the channel)
5	NC-SI Receive Packets	Counts the total number of NC-SI Control Packets received. This count is the sum of NC-SI Commands Received and NC-SI Control Packets Dropped.
6	NC-SI Transmit Packets	Counts the total number of NC-SI Control Packets transmitted to the Management Controller. This count is the sum of NC-SI responses sent and AENs sent.
7	AENs Sent	Counts the total number of AEN packets transmitted to the Management Controller

2853 **8.4.53 Get NC-SI Pass-through Statistics command (0x1A)**

2854 The Get NC-SI Pass-through Statistics command may be used by the Management Controller to request
 2855 that the channel send a copy of all current NC-SI Pass-through packet statistic values.

2856 Table 104 illustrates the packet format of the Get NC-SI Pass-through Statistics command.

2857 **Table 104 – Get NC-SI Pass-through Statistics command packet format**

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

2858 **8.4.54** Get NC-SI Pass-through Statistics response (0x9A)

2859 In the absence of any error, the channel shall process and respond to the Get NC-SI Pass-through
2860 Statistics command by sending the response packet and payload shown in Table 105.

2861 **Table 105 – Get NC-SI Pass-through Statistics response packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..27	Pass-through TX Packets Received on NC-SI Interface (Management Controller to Network Controller)			
28..31	Pass-through TX Packets Dropped			
32..35	Pass-through TX Packet Channel State Errors			
36..39	Pass-through TX Packet Undersized Errors			
40..43	Pass-through TX Packet Oversized Errors			
44..47	Pass-through RX Packets Received on LAN Interface			
48..51	Total Pass-through RX Packets Dropped			
52..55	Pass-through RX Packet Channel State Errors			
56..59	Pass-through RX Packet Undersized Errors			
60..63	Pass-through RX Packet Oversized Errors			
64..67	Checksum			

2862 The Get NC-SI Statistics Response frame contains a set of statistics counters that monitor the NC-SI
2863 Pass-through traffic in the Network Controller. Supported counters shall be reset to 0x0 when entering
2864 the Initial State and after being read. Implementation of the counters shown in Table 106 is optional. The
2865 Network Controller shall return any unsupported counter with a value of 0xFFFFFFFF for 32-bit counters
2866 and 0xFFFFFFFFFFFFFFFF for 64-bit counters. Counters may wraparound or stop if they reach
2867 0xFFFFFFFF for 32-bit counters and 0xFFFFFFFFFFFFFFFF for 64-bit counters.

2868 **Table 106 – Get NC-SI Pass-through Statistics counters**

Counter Number	Name	Meaning
1	Total Pass-through TX Packets Received (Management Controller to Channel)	Counts the number of Pass-through packets forwarded by the channel to the LAN
2	Total Pass-through TX Packets Dropped (Management Controller to Channel)	Counts the number of Pass-through packets from the Management Controller that were dropped by the Network Controller
3	Pass-through TX Packet Channel State Errors (Management Controller to Channel)	Counts the number of egress management packets (Management Controller to Network Controller) that were dropped because the channel was in the disabled state when the packet was received

Counter Number	Name	Meaning
4	Pass-through TX Packet Undersized Errors (Management Controller to Channel)	Counts the number of Pass-through packets from the Management Controller that were undersized (under 64 bytes, including FCS)
5	Pass-through TX Packet Oversized Errors (Management Controller to Channel)	Counts the number of Pass-through packets from the Management Controller that were oversized (over 1522 bytes, including FCS)
6	Total Pass-through RX Packets Received on the LAN Interface (LAN to Channel)	Counts the number of Pass-through packets that were received on the LAN interface of the channel. This counter does not necessarily count the number of packets that were transmitted to the Management Controller, because some of the packets might have been dropped due to RX queue overflow.
7	Total Pass-through RX Packets Dropped (LAN to Channel)	Counts the number of Pass-through packets that were received on the LAN interface of the channel but were dropped and not transmitted to the Management Controller
8	Pass-through RX Packet Channel State Errors (LAN to Channel)	Counts the number of ingress management packets (channel to Management Controller) that were dropped because the channel was in the disabled state when the packet was received. The NC may also count packets that were dropped because the package was in the deselected state.
9	Pass-through RX Packet Undersized Errors (LAN to Channel)	Counts the number of Pass-through packets from the LAN that were undersized (under 64 bytes, including FCS)
10	Pass-through RX Packet Oversized Errors (LAN to Channel)	Counts the number of Pass-through packets from the LAN that were oversized (over 1522 bytes, including FCS)

2869 Currently no command-specific reason code is identified for this response.

2870 8.4.55 Get Package Status command (0x1B)

2871 The Get Package Status command provides a way for a Management Controller to explicitly query the
 2872 status of a package. The Get Package Status command is addressed to the package, rather than to a
 2873 particular channel (that is, the command is sent with a Channel ID where the Package ID subfield
 2874 matches the ID of the intended package and the Internal Channel ID subfield is set to 0x1F).

2875 Table 107 illustrates the packet format of the Get Package Status command.

2876 **Table 107 – Get Package Status packet format**

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

2877 **8.4.56** Get Package Status response (0x9B)

2878 In the absence of any errors, the package shall process and respond to the Get Package Status
 2879 Command and send the response packet shown in Table 108.

2880 Currently no command-specific reason code is identified for this response.

2881 **Table 108 – Get Package Status response packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	Package Status			
24..27	Checksum			
28..45	Pad			

2882 **Table 109 – Package Status field bit definitions**

Bit Position	Field Description	Value Description
0	Hardware Arbitration Status	0b = Hardware arbitration is non-operational (inactive) or unsupported. NOTE This means that hardware arbitration tokens are not flowing through this NC. 1b = Hardware arbitration is supported, active, and implemented for the package on the given system.
1	Delayed Response Status	0b = Delayed Response handling is disabled. 1b = Delayed Response handling is enabled.
31.. 2	Reserved	Reserved

2883

2884 **8.4.57** Get NC Capabilities and Settings command (0x25)

2885 The Get NC Capabilities and Settings command is sent only as a package command. It is used to
 2886 discover the supported architectural and currently configured (active) parameters of the NC.

2887 Table 110 illustrates the packet format for the Get NC Capabilities and Settings command.

2888 **Table 110 – Get NC Capabilities and Settings command packet format**

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

2889 **8.4.58** Get NC Capabilities and Settings response (0xA5)

2890 In the absence of any errors, the package shall process and respond to the Get NC Capabilities and
 2891 Settings Command and send the response packet shown in **Error! Reference source not found..**

2892 Currently no command-specific reason code is identified for this response.

2893 **Table 111 - Get NC Capabilities and Settings response packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	Max Ports	Enabled Ports	Max PCI Endpoints	Enabled PCI Endpoints
24..27	Max PFs	Enabled PFs	Max VFs	
28..31	Fabrics	Enabled Fabrics	Other Capabilities	
32..35	Checksum			
36..45	Pad			

2894

2895 **8.4.58.1** Max Ports field

2896 The Max Ports field indicates the maximum number of network ports that can be supported by the
 2897 implementation (uint8).

2898 **8.4.58.2** Enabled Ports field

2899 The Enabled Ports field indicates the current number of network ports that are currently configured
 2900 (uint8).

2901 **8.4.58.3** Max PCI Endpoints field

2902 The Max PCI Endpoints field indicates the maximum number of PCI Endpoints that can be supported by
 2903 the implementation (uint8).

2904 **8.4.58.4** Enabled PCI Endpoints field

2905 The Enabled PCI Endpoints field indicates the current number of PCI Endpoints that are currently
 2906 configured (uint8).

2907 **8.4.58.5** Max PFs field

2908 The Max PFs field indicates the maximum number of PCI Physical Functions that can be supported by
 2909 the implementation (uint8).

2910 **8.4.58.6** Enabled PFs field

2911 The Enabled PFs field indicates the current number of PCI Physical Functions that are currently
 2912 configured (uint8).

2913 Max VFs field

2914 The Max VFs field indicates the maximum number of PCI Virtual Functions that can be supported by the
 2915 implementation (uint8).

2916 **8.4.58.7** Fabrics field

2917 The Fabrics field indicates the network fabrics that can be supported by the implementation.

2918 **Table 112 – Fabrics field bit definitions**

Bit Position	Field Description	Value Description
0	Ethernet	0b0 = Ethernet Fabric is not supported 0b1 = Ethernet Fabric is supported
1	Fibre Channel	0b0 = Fibre Channel Fabric is not supported 0b1 = Fibre Channel Fabric is supported
2	InfiniBand	0b0 = InfiniBand Fabric is not supported 0b1 = InfiniBand Fabric is supported
3..7	Reserved	Reserved

2919

2920 **8.4.58.8** Enabled Fabrics field

2921 The Enabled Fabrics field indicates the currently configured fabrics.

2922 **Table 113 – Enabled Fabrics field bit definitions**

Bit Position	Field Description	Value Description
0	Ethernet	0b0 = Ethernet Fabric is not enabled 0b1 = Ethernet Fabric is enabled
1	Fibre Channel	0b0 = Fibre Channel Fabric is not enabled 0b1 = Fibre Channel Fabric is enabled
2	InfiniBand	0b0 = InfiniBand Fabric is not enabled 0b1 = InfiniBand Fabric is enabled
3..7	Reserved	Reserved

2923

2924 **8.4.58.9 Other Capabilities field**

2925 The Other Capabilities field indicates which features of this specification the NC supports, as described in
 2926 Table 114.

2927 **Table 114 – Capabilities Flags bit definitions**

Bit Position	Field Description	Value Description
0	VF allocation	0b = The Max VFs field is interpreted as per port 1b = The Max VFs field is interpreted as per device
1	Enabled Ports	0b = The number of Enabled Ports is fixed 1b = The number of Enabled Ports is programmable
2	Enabled Buses	0b = The number of Enabled Buses is fixed 1b = The number of Enabled Buses is programmable
3	Enabled PFs	0b = The number of Enabled PFs is fixed 1b = The number of Enabled PFs is programmable
4..15	Reserved	Reserved

2928 **8.4.59 Set NC Configuration command (0x26)**

2929 The Set NC Configuration command allows the Management Controller to configure the number of active
 2930 Physical functions and PCI (host) and network interfaces, where allowed (generally if the reported max
 2931 value of the respective entity is greater than one). The values (programmed or fixed) are used in the PF
 2932 Assignment command where the associations are made between the physical ports, partitions and host
 2933 buses. If the implementation or controller architecture does not allow any configuration of these
 2934 parameters, this command shall not be implemented.

2935 The values configured by this command are held by the NC and only take effect at the next PCI reset.

2936 The Set NC Configuration command is addressed to the package, rather than to a channel (that is, the
 2937 command is sent with a Channel ID where the Package ID subfield matches the ID of the intended
 2938 package and the Internal Channel ID subfield is set to 0x1F).

2939 Table 115 illustrates the packet format of the Set NC Configuration command.

2940 **Table 115 – Set NC Configuration command packet format**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Enable Ports	Enable PCI Endpoints	Enable PFs	Reserved
20..23	Checksum			
24..45	Pad			

2941 **8.4.59.1 Enable Ports field**

2942 The Enable Ports field indicates the number of network ports to be enabled at the next PCI reset(uint8).

2943 **8.4.59.2** Enable PCI Endpoints field

2944 The Enable PCI Endpoints field indicates the number of PCI Endpoints to be enabled at the next PCI
 2945 reset(uint8). In some implementation architectures, this is not settable by NC-SI; in those cases this field
 2946 becomes read-only and the value is ignored.

2947 **8.4.59.3** Enable PFs field

2948 The Enable PFs field indicates the number of PCI Physical Functions to be enabled at the next PCI
 2949 reset(uint8).

2950

2951 **8.4.60** Set NC Configuration response (0xA6)

2952 The package shall, in the absence of a checksum error or identifier mismatch, always accept the Set NC
 2953 Configuration command and send a response (see Table 116).

2954 **Table 116 – Set NC Configuration response packet format**

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

2955

2956 **8.4.61** Get PF Assignment command (0x27)

2957 The Get PF Assignment command is a Package command that allows the Management controller to
 2958 receive the list of PCI Physical Functions (partitions) currently assigned to channels in the package, their
 2959 enablement state and conditionally what PCI Endpoint they are assigned to if the NC supports multiple
 2960 host interfaces.

2961 See the Set PF Assignment command description for additional information.

2962 Table 117 – Get PF Assignment Command Packet Format illustrates the packet format of the Get PF
 2963 Assignment Command.

2964 **Table 117 – Get PF Assignment Command Packet Format**

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

2965

2966 **8.4.62** Get PF Assignment Response (0xA7)

2967 In the absence of any errors, the channel shall process and respond to the Get PF Assignment Command
 2968 and send the response packet shown in the table below.

2969 Note: Braces {} denote fields that depend on device capabilities.

2970 **Table 118 – Get PF Assignment Response packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Channel 0 Function Assignment bitmap			
24..27	{Channel 1 Function Assignment bitmap}			
	...			
	{Channel c-1 Function Assignment bitmap}			
	Function - Port Association			
	Function Enablement bitmap			
	{ PCI Endpoint 0 Function Assignment bitmap}			
	{ PCI Endpoint 1 Function Assignment bitmap}			
	...			
	{ PCI Endpoint b-1 Function Assignment bitmap}			
	Checksum			
	Pad			

2971

2972 **8.4.62.1** Channel c Function Assignment bitmap fields

2973 The number of Channel Function Assignment bitmaps returned in the response is equal to 'c', the number
 2974 returned in the Get NC Capabilities and Settings Command Enabled Ports field. The Channel c Function
 2975 Assignment bitmaps are 32-bit fields in which each bit position corresponds to a PCI physical function in

the NC on the specified channel. If the physical function is assigned to the c^{th} channel, even if it not currently enabled, the bit value shall be set to 1b; otherwise, the bit is set to 0b.

Table 119 – Channel c Function Assignment bitmap field

Bit Position	Field Description	Value Description
0	F0 status	0b = F0 is not assigned on the channel. 1b = F0 is assigned on the channel.
1	F1 status	0b = F1 is not assigned on the channel. 1b = F1 is assigned on the channel.
...
15	F15 status	0b = F15 is not assigned on the channel. 1b = F15 is assigned on the channel

2979

2980 **8.4.62.2** Function Port Association bitmap field

2981 The Function Assignment bitmap is a 32-bit field in which each bit position corresponds to a physical
2982 function in the device. Unused bits are Reserved.

Table 120 – Function Port Association bitmap field

Bit Position	Field Description	Value Description
0	F0 association	0b = F0 is fixed to the specified channel. 1b = F0 may be assigned to any channel.
1	F1 association	0b = F1 is fixed to the specified channel. 1b = F1 may be assigned to any channel.
...
15	F15 association	0b = F15 is fixed to the specified channel. 1b = F15 may be assigned to any channel.

2984

2985 **8.4.62.3** Function Enablement bitmap field

2986 The Function Assignment bitmap is a 32-bit field in which each bit position corresponds to a physical
2987 function in the NC. The number of functions shown as enabled in this field shall be equal to the number
2988 shown in the Get/Set NC Configuration command. A function may be assigned to a PCI Endpoint and be
2989 enabled and not be assigned to a channel in some implementations (i.e., a non-networking function).

2990

Table 121 – Function Enablement bitmap field

Bit Position	Field Description	Value Description
0	F0 status	0b = F0 is not enabled 1b = F0 is enabled
1	F1 status	0b = F1 is not enabled. 1b = F1 is enabled.
...
31	F31 status	0b = F31 is not enabled. 1b = F31 is enabled

2991

2992 8.4.62.4 PCI Endpoint b Assignment bitmap field

2993 The number of PCI Endpoint Assignment bitmaps returned in the response is equal to 'b', the number
 2994 returned in the Get NC Capabilities and Settings Command Enabled PCI Endpoints field. The PCI
 2995 Endpoint b Assignment bitmaps are 32-bit fields in which each bit position corresponds to a physical
 2996 function in the NC on the specified host bus. If the physical function is assigned to the bth Endpoint , even
 2997 if it not currently enabled, the bit value shall be set to 1b, otherwise the bit is set to 0b.

2998

Table 122 – PCI Bus b Assignment bitmap field

Bit Position	Field Description	Value Description
0	F0 status	0b = F0 is not assigned on the specified PCI Endpoint. 1b = F0 is assigned on the specified PCI Endpoint.
1	F1 status	0b = F1 is not assigned on the specified PCI Endpoint. 1b = F1 is assigned on the specified PCI Endpoint.
...
31	F15 status	0b = F31 is not assigned on the specified PCI Endpoint. 1b = F31 is assigned on the specified PCI Endpoint

2999

3000 8.4.62.5 Calculation of Partition ID

3001 When multiple functions are assigned to a channel, they are addressed by a value called the Partition ID.
 3002 The Partition ID is created by taking the set of Functions that are assigned to a channel and assigning
 3003 each an index value starting with the lowest numbered Function. A Function assigned to a channel has a
 3004 Partition ID even if it is not enabled. Partition numbering starts at 0. For example, if F2 and F6 are
 3005 assigned to channel 3, but only F2 is enabled, then F2 has Partition ID = 0 and F6 has Partition ID = 1 on
 3006 that channel.

3007 **8.4.63 Set PF Assignment command (0x28)**

3008 The Set PF Assignment command is a Package command that allows the Management controller to
3009 enable, disable, and assign PCI Physical Functions (partitions) in the controller to the channels, and, if
3010 applicable, to different PCI Endpoints in multi-home or multi-host configurations.

3011 The format of the command payload is dependent on the numbers of Physical Functions, Channels and
3012 PCI Endpoints supported by the controller:

- 3013 1) The number of Function Assignments bitmap fields shall be determined by the value (c) of the
3014 Channel Count field in the Get Capabilities response.
- 3015 2) The number of Physical Functions allowed to be configured in the Function Assignment and
3016 Enablement bitmap fields shall be determined by the value of the Physical Function Count field
3017 in the Get NC Capabilities and Settings command response. Assignment in all bitmaps starts
3018 at bit 0 and continues sequentially for the number of Functions supported. To support various
3019 implementation architectures, the definition of assignment/enablement rules is beyond the
3020 scope of this specification.
- 3021 3) If the value (b) of the <PCI Bus Count> field in the <Get Device Capabilities and Settings
3022 command> response is greater than 1, the Controller shall also include that number of PCI
3023 Endpoint Function Assignment bitmap fields in the command. Controllers that do not support
3024 multiple PCI interfaces shall not implement PCI Endpoint Host Function Assignment bitmap
3025 fields. PCI Endpoint 0 shall be used if the Controller is configured for single bus operation.

3026 The values configured by this command are held by the controller and only take effect at the next PCI
3027 reset. The configuration is persistent unless changed by another Set PF Assignment command or other
3028 mechanism.

3029 Table 123 illustrates the packet format of the Set PF Assignment Command.

3030 **Table 123 – Set PF Assignment Command packet format**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Channel 0 Function Assignment bitmap			
	{Channel 1 Function Assignment bitmap}			
	...			
	{Channel c-1 Function Assignment bitmap}			
	Function Enablement bitmap			
	{ PCI Bus 0 Function Assignment bitmap}			
	{ PCI Bus 1 Function Assignment bitmap}			
	...			
	{ PCI Bus b-1 Function Assignment bitmap}			
	Checksum			
	Pad			

3031 8.4.63.1 Channel Function Assignment bitmap field

3032 The Channel Function Assignment bitmap is a 32-bit field in which each bit position corresponds to a
 3033 physical function in the device. If the physical function is assigned to the channel, even if it not currently
 3034 enabled, the bit value shall be set to 0b1. This allows for a partition ID to be assigned and partition
 3035 commands to be sent to the function even if it is not enabled.

3036 **Table 124 – Channel Function Assignment bitmap field**

Bit Position	Field Description	Value Description
0	F0 status	0b = F0 is not assigned on the channel. 1b = F0 is assigned on the channel.
1	F1 status	0b = F1 is not assigned on the channel. 1b = F1 is assigned on the channel.
...
15	F15 status	0b = F15 is not assigned on the channel. 1b = F15 is assigned on the channel

3037 8.4.63.2 Function Enablement bitmap field

3038 The Function Assignment bitmap is a 32-bit field in which each bit position corresponds to a physical
 3039 function in the device.

3040 **Table 125 – Function Enablement bitmap field**

Bit Position	Field Description	Value Description
0	F0 status	0b = F0 is not enabled on the specified channel. 1b = F0 is enabled on the specified channel.
1	F1 status	0b = F1 is not enabled on the specified channel. 1b = F1 is enabled on the specified channel.
...
15	F15 status	0b = F15 is not enabled on the specified channel. 1b = F15 is enabled on the specified channel

3041

3042 **8.4.63.3 PCI Endpoint Assignment bitmap field**

3043 The PCI Endpoint Assignment bitmap is a 32-bit field in which each bit position corresponds to a physical
 3044 function in the device.

3045 **Table 126 – PCI Bus Assignment bitmap field**

Bit Position	Field Description	Value Description
0	F0 status	0b = F0 is not assigned on the specified PCI Endpoint. 1b = F0 is assigned on the specified PCI Endpoint.
1	F1 status	0b = F1 is not assigned on the specified PCI Endpoint. 1b = F1 is assigned on the specified PCI Endpoint.
...
15	F15 status	0b = F15 is not assigned on the specified PCI Endpoint. 1b = F15 is assigned on the specified PCI Endpoint

3046 **8.4.64 Set PF Assignment Response (0xA8)**

3047 In the absence of any errors, the channel shall process and respond to the Get PF Assignment Command
 3048 and send the response packet shown in Table 127.

3049 **Table 127 – Set PF Assignment Response packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
24..27	Checksum			
36..39	Pad			

3050

3051 8.4.65 Get Channel Configuration command (0x29)

3052 The Get Port Configuration command is used to discover the currently configured settings of the channel,
3053 including the fabric type, the implemented media type, the number of enabled partitions, if any, and their
3054 bandwidth allocation settings where applicable..

3055 Table 128 illustrates the packet format for the Get Port Configuration command.

3056 **Table 128 – Get Port Configuration command packet format**

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

3057 8.4.66 Get Channel Configuration response (0xA9)

3058 In the absence of any errors, the channel shall process and respond to the Get Channel Configuration
3059 Command and send the response packet shown in Table 129.

3060 Currently no command-specific reason code is identified for this response.

3061 **Table 129 – Get Channel Configuration response packet format**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	Fabric Type	Media Type	Max MTU	
24..27	Reserved		Reserved	# Enabled Partitions
28..31	P1 Max TX BW	P1 Min TX BW	P2 Max TX BW	P2 Min TX BW
32..35	Checksum			

3062 8.4.66.1 Fabric Type field

3063 The Fabric Type field indicates which personality types are currently enabled on the channel, as
3064 described in Table 133.

3065 **Table 130 – Fabric Type bit definitions**

Value	Fabric Type	Value Description
1	Ethernet Mode	Ethernet operation is enabled
2	Fibre Channel Mode	Fibre Channel operation is enabled
3	InfiniBand Mode	InfiniBand operation is enabled
All others	Reserved	Reserved

3066 **8.4.66.2** Max MTU field

3067 The Max MTU field is used to report the maximum allowed MTU size (Bytes) when the port is configured
3068 for Ethernet.

3069 **8.4.66.3** Media Type field

3070 The Media Type field indicates the physical interface type used on the port implementation and if that port
3071 supports one or more than one NC-SI channels (for example, some designs may support up to 4
3072 independent ports in a QSFP interface), as described in **Error! Reference source not found.** Table 131.

3073 NOTE An implementation that implements a SFF cage interface into which a RJ-45 transceiver is plugged shall
3074 return 'SFF cage' as the media type.

3075 **Table 131 – Media Type bit definitions**

Bit Position	Field Description	Value Description
0	Backplane	0b = The port does not have a backplane interface 1b = The port has a backplane interface
1	Base-T (RJ-45 style)	0b = The port does not have a Base-T interface 1b = The port has a Base-T (RJ-45 style) interface
2	SFF cage	0b = The port does not have an SFF-style interface 1b = The port has an SFF-style interface
3..6	Reserved	Reserved
7	Shared Interface	0b = The port is dedicated to one NC-SI channel 1b = The port is shared between multiple channels

3076

3077 **8.4.66.4** P(n) Max TX BW Fields

3078 These fields contain the Maximum TX bandwidth allocation of the nth enabled partition expressed in % of
3079 the physical port link speed.

3080 **8.4.66.5** P(n) Min TX BW Fields

3081 These fields contain the Minimum TX bandwidth allocation of the nth enabled partition expressed in % of
3082 the physical port link speed.

3083 **8.4.67** Set Channel Configuration command (0x2A)

3084 The Set Channel Configuration command allows the Management Controller to configure characteristics
3085 of the channel. The TX Bandwidth fields must be set for each enabled partition, but their values may be
3086 overridden during operation by data from protocols such as DCB.

3087 Table 132 illustrates the packet format of the Set Channel Configuration command.

3088

Table 132 – Set Channel Configuration command packet format

Bits				
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Fabric Type	# Partitions	Max MTU	
20..23	P1 Max TX BW	P1 Min TX BW	P2 Max TX BW	P2 Min TX BW
	...			
	Checksum			
	Pad			

3089 **8.4.67.1 Fabric Type field**

3090 The Fabric Type field indicates the personality type to be enabled on the channel, as described in Table
 3091 133. The contents of this field may be ignored if the channel only supports one fabric type. The Fabric
 3092 type is a channel property shared by all partitions assigned to the channel.

3093

Table 133 – Fabric Type bit definitions

Value	Fabric Type	Value Description
1	Ethernet Mode	Enable Ethernet operation
2	Fibre Channel Mode	Enable Fibre Channel operation
3	InfiniBand Mode	Enable InfiniBand operation
all others	Reserved	Reserved

3094 **8.4.67.2 Max MTU field**

3095 The Max MTU field is used to configure the maximum allowed MTU size (Bytes) when the port is
 3096 configured for Ethernet.

3097 **8.4.67.3 # Partitions**

3098 The Number of Partitions field indicates the number of Functions that have been assigned to the
 3099 channel/port in the Set PF Assignment command. This field is used only to provide the number of
 3100 partitions present in the bandwidth fields and does not have the ability to change the number of assigned
 3101 partitions on the channel. Each assigned partition must be allocated min and max TX bandwidth values
 3102 when enabled.

3103 The initial value is generally expected to be one partition enabled per port and if modified, the new value
 3104 should persist across system boot and power cycles.

3105 **8.4.67.4 P(n) Max TX BW fields**

3106 These fields contain the Maximum TX bandwidth allocation of the nth enabled partition expressed in % of
 3107 the physical port link speed. Oversubscription of partition maximum bandwidth is allowed. The field value
 3108 is an integer ranging from 0 to 100₁₀.

3109 The initial value is generally expected to be 100% per partition, allowing each enabled partition full use of
 3110 the channel bandwidth if no other partition has traffic. If modified, the new value should persist across
 3111 system boot and power cycles.

3112 8.4.67.5 P(n) Min TX BW field

3113 These fields contain the Minimum TX bandwidth allocation of the n^{th} enabled partition expressed in % of
 3114 the physical port link speed. This is interpreted as committed bandwidth to the partition and as such the
 3115 Min TX BW fields of all enabled partitions on the port must sum to 100%. The field value is an integer
 3116 ranging from 0 to 100₁₀.

3117 The initial value is generally expected to be equal weighting among all enabled partitions, allowing each
 3118 enabled partition equal use of the channel bandwidth. If modified, the new value should persist across
 3119 system boot and power cycles

3120 8.4.68 Set Channel Configuration response (0xAA)

3121 The package shall, in the absence of a checksum error or identifier mismatch, always accept the Set
 3122 Channel Configuration command and send a response (see Table 134).

3123 **Table 134 – Set Channel Configuration response packet format**

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

3124 8.4.69 Get Partition Configuration command (0x2B)

3125 The Get Partition Configuration command is used to discover additional optional functions supported by
 3126 the channel, such as the number of unicast/multicast addresses supported, the amount of buffering in
 3127 bytes available for packets bound for the Management Controller, and so on.

3128 Table 135 illustrates the packet format for the Get Partition Configuration command.

3129 **Table 135 – Get Partition Configuration command packet format**

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

3130 8.4.69.1 Partition ID field

3131 The Partition ID field is the identifier for the function on the channel as defined in clause 8.4.63

3132 **8.4.70 Get Partition Configuration response (0xAB)**

3133 In the absence of any errors, the channel shall process and respond to the Get Partition Configuration
 3134 Command and send the response packet shown in Table 136.

3135 Currently no command-specific reason code is identified for this response.

3136 **Table 136 – Get Partition Configuration response packet format**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	Personality Cfg	Personality Spt	Configuration Flags	
24..27	Max TX BW	Min TX BW	Advertised VF Count	
28..31	PCI DID		PCI VID	
32..35	PCI SSID		PCI SVID	
36..39	PCI Endpoint #	PCI Bus #	PCI Device #	PCI Function #
40..43	FCoE Cfg	Address Count	Address TLVs	
44..47	Address (MSB)	Address

	Checksum			

3137 **8.4.70.1 Personality Cfg field**

3138 The Personality Configured field indicates which personality type(s) are currently enabled on the partition,
 3139 as described in Table 137.

3140 Note: Some implementations may support multiple personalities being simultaneously enabled.

3141 **Table 137 – Personality Cfg bit definitions**

Bit Position	Field Description	Value Description
0	Ethernet Status	0b = Ethernet operation is not enabled 1b = Ethernet operation is enabled
1	Fibre Channel Status	0b = Fibre Channel operation is not enabled 1b = Fibre Channel operation is enabled
2	Fibre Channel over Ethernet Status	0b = Fibre Channel over Ethernet operation is not enabled 1b = Fibre Channel over Ethernet operation is enabled
3	InfiniBand Status	0b = InfiniBand operation is not enabled 1b = InfiniBand operation is enabled
4	iSCSI Offload Status	0b = iSCSI Offload operation is not enabled 1b = iSCSI Offload operation is enabled

Bit Position	Field Description	Value Description
5	RDMA Status	0b = RDMA operation is not enabled 1b = RDMA operation is enabled
6	NVMe	0b = NVMe operation is not enabled 1b = NVMe operation is enabled
7	Reserved	Reserved

3142 8.4.70.2 Personality Spt field

3143 The Personality Supported field indicates which personality types the partition supports, as described in
3144 Table 138.

3145 **Table 138 – Personality Spt bit definitions**

Bit Position	Field Description	Value Description
0	Ethernet Support	0b = Ethernet operation is not supported 1b = Ethernet operation is supported
1	Fibre Channel Support	0b = Fibre Channel operation is not supported 1b = Fibre Channel operation is supported
2	Fibre Channel over Ethernet Support	0b = Fibre Channel over Ethernet operation isn't supported 1b = Fibre Channel over Ethernet operation is supported
3	InfiniBand Support	0b = InfiniBand operation is not supported 1b = InfiniBand operation is supported
4	iSCSI Offload Support	0b = iSCSI Offload operation is not supported 1b = iSCSI Offload operation is supported
5	RDMA Support	0b = RDMA operation is not supported 1b = RDMA operation is supported
6	NVMe	0b = NVMe Offload operation is not supported 1b = NVMe Offload operation is supported
7	Reserved	Reserved

3146 8.4.70.3 Configuration Flags field

3147 The Configuration Flags field indicates which optional features of this specification the channel supports,
3148 as described in Table 139.

3149 **Table 139 – Configuration Flags bit definitions**

Bit Position	Field Description	Value Description
0	Host Driver Status	0b = When reporting is supported, Host driver is not present 1b = When reporting is supported, Host driver is present
1	Host Driver Status Reporting	0b = Host Driver status reporting is not supported. 1b = Host Driver status reporting (bit 0) is supported.

Bit Position	Field Description	Value Description
2	Partition Link Status	0b = When reporting is supported, Partition Link is down 1b = When reporting is supported, Partition Link is up
3	Partition Link Status Reporting	0b = Partition Link Status reporting is not supported. 1b = Partition Link Status reporting (bit 2) is supported.
4	Boot Status	0b = The partition is not configured for boot. 1b = The partition is configured for boot.
5	Bootable	0b = The partition supports boot and reporting 1b = The partition does not support boot
7..31	Reserved	Reserved

3150

3151 **8.4.70.4 Partition Link fields**

3152 This fields describe the ability of a partition to support traffic when the partition is assigned to a PCI bus
 3153 and NC-SI channel and either its associated physical port link is up or the implementation supports
 3154 internal communication between partitions when the physical port link is down.

3155 **8.4.70.5 Max TX BW field**

3156 This field contains the Maximum TX bandwidth allocation of the partition expressed in % of the physical
 3157 port link speed. The % value ranges from 0 to 100 represented as an integer.

3158 **8.4.70.6 Min TX BW field**

3159 This field contains the Minimum TX bandwidth allocation of the partition expressed in % of the physical
 3160 port link speed. This is interpreted as committed bandwidth to the partition and as such the Min TX BW
 3161 fields of all enabled partitions on the port must sum to 100%. The % value ranges from 0 to 100
 3162 represented as an integer.

3163 **8.4.70.7 Advertised VF Count field**

3164 The Advertised VF Count field indicates the number of Virtual Functions that shall be advertised by the
 3165 partition's PF.

3166 **8.4.70.8 PCI DID**

3167 The current PCI Device ID of the Partition

3168 **8.4.70.9 PCI VID**

3169 The current PCI Vendor ID of the Partition

3170 **8.4.70.10 PCI SSID**

3171 The current PCI Subsystem ID of the Partition

3172 **8.4.70.11 PCI SVID**

3173 The current PCI Subvendor ID of the Partition

3174 8.4.70.12 PCI Endpoint #

3175 The identifier indicating which PCI Endpoint on the NC the partition is associated with

3176 8.4.70.13 PCI Bus #

3177 The assigned PCI Bus number assigned to the partition in the host system's bus enumeration process

3178 8.4.70.14 PCI Device #

3179 The assigned PCI Device number assigned to the partition in the host system's bus enumeration process
3180 except in the cases of ARI mode operation when it shall contain the arbitrary value of 0xFF

3181 8.4.70.15 PCI Function #

3182 The assigned PCI Function number assigned to the partition in the host system's bus enumeration
3183 process

3184

3185 8.4.70.16 FC/FCoE Cfg

3186 This field contains nothing right now.

3187 8.4.70.17 Address Count field

3188 This field indicates the number of permanent and virtual addresses reported by the partition.

3189

3190 8.4.70.18 Address TLVs

3191 These TLVs show the permanently programmed and current addresses being used by the partition.

3192

3193

3194

Table 140 – Address Type-Length Field Bit Definitions

Bit Position	Field Description	Value Description
7..0	Address Type	<p>The following type encodings shall be used to indicate the address values that are permanently assigned to the partition. The response shall include all types whether or not that mode of operation is active, or the partition is enabled:</p> <p>0x0 = Reserved</p> <p>0x1 = Ethernet MAC</p> <p>0x2 = iSCSI Offload (Ethernet MAC)</p> <p>0x3 = Fibre Channel World Wide Node Name</p> <p>0x4 = Fibre Channel World Wide Port Name</p> <p>0x5 = FCoE-FIP MAC</p> <p>0x6 = InfiniBand Node GUID</p> <p>0x7 = InfiniBand Port GUID</p> <p>0x8 = InfiniBand VPort/LID</p> <p>The following type encodings shall be used to indicate all address values that are currently in use by the partition based on configured mode of operation. These may be the permanent address or a programmatically assigned address.</p> <p>:</p> <p>0xF1 = Ethernet MAC</p> <p>0xF2 = iSCSI Offload (Ethernet MAC)</p> <p>0xF3 = Fibre Channel World Wide Node Name</p> <p>0xF4 = Fibre Channel World Wide Port Name</p> <p>0xF5 = FCoE-FIP MAC</p> <p>0xF6 = InfiniBand Node GUID</p> <p>0xF7 = InfiniBand Port GUID</p> <p>0xF8 = InfiniBand VPort/LID</p> <p>all others = Reserved</p>
15..8	Address Length	The length indicates the number of bytes used in the address

3195 **8.4.71 Set Partition Configuration command (0x2C)**

3196 The Set Partition Configuration command allows the Management Controller to configure various settings
 3197 of the partition including virtual addresses, VF allocation and other parameters.

3198 The Set Partition Configuration command is addressed to the channel with the Partition ID field set to the
 3199 index/ordinal of the target PF on the channel.

3200 The partition's personality configuration and VF count settings may be made persistent if written to the
 3201 NVRAM via the Commit command. These settings take effect at the next PCI Reset.

3202 Table 141 illustrates the packet format of the Set Partition Configuration command.

3203 **Table 141 – Set Partition Configuration command packet format**

Bits				
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Partition ID	Personality Cfg	VF Count	
20..23	Partition Link Control	Reserved	Address Count	Address TLV
24..27	Checksum			
28..45	Pad			

3204 8.4.71.1 Personality Cfg field

3205 The Personality Configuration field indicates which personality type(s) shall be enabled on the partition,
 3206 as described in Table 142. Any attempt to enable a personality not shown as supported in clause 8.4.70.2
 3207 shall be cause the command to fail. In some implementations it may be appropriate to select more than
 3208 one personality at a time, for instance Ethernet and RDMA.

3209 **Table 142 – Personality Cfg bit definitions**

Bit Position	Field Description	Value Description
0	Ethernet Status	0b = Disable Ethernet operation 1b = Enable Ethernet operation
1	Fibre Channel Status	0b = Disable Fibre Channel operation 1b = Enable Fibre Channel operation
2	Fibre Channel over Ethernet Status	0b = Disable Fibre Channel over Ethernet operation 1b = Enable Fibre Channel over Ethernet operation
3	InfiniBand Status	0b = Disable InfiniBand operation 1b = Enable InfiniBand operation
4	iSCSI Offload Status	0b = Disable iSCSI Offload operation 1b = Enable iSCSI Offload operation
5	RDMA Status	0b = Disable RDMA operation 1b = Enable RDMA operation
6	NVMe	0b = Disable NVMe operation 1b = Enable NVMe operation
7	Reserved	Reserved

3210 8.4.71.2 VF Count

3211 The VF Count field contains the number of VFs to be advertised in PCI Configuration Space by the
 3212 partition.

3213 **8.4.71.3 Partition Link Control**

3214 Table 143 describes the values for the Partition Link Control field.

3215 **Table 143 – Values for the Config flags field (8-bit field)**

Value	Description
0x0	Partition Link is down
0x1	Partition Link is up
0x4..0xFF	Reserved

3216 **8.4.71.4 Address Count field**

3217 The Address Count field contains the number of partition virtual addresses to be configured as specified
 3218 in the Address TLV field.

3219 **8.4.71.5 Address TLV**3220 **Table 144 – Address Type-Length field bit definitions**

Bit Position	Field Description	Value Description
7..0	Address Type	<p>Addresses specified herein override the permanent or factory-programmed network address to be used by the partition based on configured mode of operation. To return to using the permanent address, supply either an address of 0 or the permanent address in this field or remove power from the NC.</p> <p>:</p> <p>0xF1 = Ethernet MAC</p> <p>0xF2 = iSCSI Offload (Ethernet MAC)</p> <p>0xF3 = Fibre Channel World Wide Node Name</p> <p>0xF4 = Fibre Channel World Wide Port Name</p> <p>0xF5 = FCoE-FIP MAC</p> <p>0xF6 = InfiniBand Node GUID</p> <p>0xF7 = InfiniBand Port GUID</p> <p>0xF8 = InfiniBand VPort/LID</p> <p>All others = Reserved</p>
15..8	Address Length	The length indicates the number of bytes used in the address

3221 8.4.72 Set Partition Configuration response (0xAC)

3222 The package shall, in the absence of a checksum error or identifier mismatch, always accept the Set
3223 Partition Configuration command and send a response (see Table 145).

3224 **Table 145 – Set Partition Configuration response packet format**

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

3225

3226 8.4.73 Get Boot Config Command (0x2D)

3227 The Get Boot Config Command allows the Management Controller to query for the Boot Initiator settings
3228 of a given Boot Protocol type configured on the channel/PF/partition and stored in the NVRAM of the
3229 controller.

3230 If the command is sent to a destination that exists but that does not support the specified Boot Protocol
3231 type, the command execution shall fail with a reason code indicating a Parameter Is Invalid, Unsupported,
3232 or Out-of-Range.

3233 Table 146 illustrates the packet format of the Get Boot Config command.

3234 **Table 146 – Get Boot Config command packet**

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

3235

3236 8.4.73.1 Protocol Type field

3237 The Protocol Type field specifies the boot protocol for which configuration data is requested.

3238

Table 147 – Protocol Type field

Bit Position	Field Description	Value Description
7..0	Boot Protocol Type	0x0 = PXE (legacy) 0x1 = iSCSI Offload 0x2 = FCoE Offload 0x3 = FC 0x4 = NVMe (independent of fabric type) 0x5-0xFF = Reserved

3239

3240 Note: Selection of protocol type NVMe covers NVMeoF, NVMe over RDMA, NVMeoFC, and NVMeoIB
 3241 depending on the configured fabric type of the channel.

3242

3243 8.4.74 Get Boot Config Response (0xAD)

3244 The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Get Boot
 3245 Config command and send a response.

3246 The Get Boot Config Response frame contains the currently stored settings for the specified Boot
 3247 Protocol type contained in the controller's NVRAM that the channel/PF/partition will use in a boot
 3248 operation done locally by the adapter. Settings that the Controller supports but does not have a value for
 3249 (e.g., have no initial or current value) should be included in the Response and have a length of 0.

3250 All attribute values returned by this command shall be in unterminated ASCII string format.

3251 Table 148 illustrates the packet format of the Get Boot Config Response.

3252

Table 148 – Get Boot Config Response packet

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23			Protocol Type	Number of TLVs
28..	Type-Length Field #1		Value Field #1	
...	Type-Length Field #2		Value Field #2	
...	...			
....	Checksum			

3253

3254 8.4.74.1 Protocol Type field

3255 The Protocol Type field specifies the boot protocol for which boot attributes are being returned.

3256

Table 149 – Protocol Type field

Bit Position	Field Description	Value Description
7..0	Boot Protocol Type	0x0 = PXE 0x1 = iSCSI 0x2 = FCoE 0x3 = FC 0x4 = NVMe (independent of fabric type) 0x5-0xFF = Reserved

3257 Note: Selection of protocol type NVMe covers NVMeoF, NVMe over RDMA, NVMeoFC, and NVMeoIB
 3258 depending on the configured fabric type of the channel.

3259

3260 8.4.74.2 Boot Protocol Type-Length-Value fields

3261 The set of boot attributes (one of the following 4 tables) that correspond to the specified Protocol Type in
 3262 the Command are returned as TLVs in the Response.

3263

Table 150 – PXE Boot Protocol Type-Length field

Bit Position	Field Description	Value Description
7..0	Attribute Name/Type	0x0 = VLAN ID 0x1 = VLAN enable 0x2-0xFF = Reserved
15..8	Length	
	Attribute Value	Value data

3264

3265

Table 151 – Get FC Boot Protocol Type-Length field

Bit Position	Field Description	Value Description
7..0	Attribute Name/Type	0x0 = FCInitiatorBootSelection 0x1 = FirstFCTargetWWPN 0x2 = FirstFCTargetLUN 0x3 = SecondFCTargetWWPN 0x4 = SecondFCTargetLUN 0x5 = ThirdFCTargetWWPN 0x6 = ThirdFCTargetLUN 0x7 = FourthFCTargetWWPN 0x8 = FourthFCTargetLUN 0x9 = FifthFCTargetWWPN 0xA = FifthFCTargetLUN 0xB = SixthFCTargetWWPN 0xC = SixthFCTargetLUN 0xD = SeventhFCTargetWWPN 0xE = SeventhFCTargetLUN 0xF = EighthFCTargetWWPN 0x10 = EighthFCTargetLUN 0x11-0xFF = Reserved
15..8	Length	
	Attribute Value	Value data

3266

Table 152 – FCoE Boot Protocol Type-Length field

Bit Position	Field Description	Value Description
7..0	Attribute Name/Type	0x0 = FCoEInitiatorBootSelection 0x1 = FirstFCoEWWPNTarget 0x2 = FirstFCoEBootTargetLUN 0x3 = FirstFCoEFCFVLANID 0x4 = FCoETgTBoot 0x5-0xF = Reserved
15..8	Length	
	Attribute Value	Value data

3267

3268

Table 153 – iSCSI Boot Protocol Type-Length field

Bit Position	Field Description	Value Description
7..0	Attribute Name/Type	0x0 = IscsiInitiatorIPAddrType 0x1 = IscsiInitiatorAddr 0x2 = IscsiInitiatorName 0x3 = IscsiInitiatorSubnet 0x4 = IscsiInitiatorSubnetPrefix 0x5 = IscsiInitiatorGateway 0x6 = IscsiInitiatorFirstDNS 0x7 = IscsiInitiatorSecondDNS 0x10 = ConnectFirstTgt 0x11 = FirstTgtIpAddress 0x12 = FirstTgtTcpPort 0x13 = FirstTgtBootLun 0x14 = FirstTgtIscsiName 0x15 = FirstTgtChapId 0x16 = FirstTgtChapPwd 0x17 = FirstTgtVLANEnable *bool 0x18 = FirstTgtVLAN 0x20 = ConnectSecondTgt 0x21 = SecondTgtIpAddress 0x22 = SecondTgtTcpPort 0x23 = SecondTgtBootLun 0x24 = SecondTgtIscsiName 0x25 = SecondTgtChapId 0x26 = SecondTgtChapPwd 0x27 = SecondTgtVLANEnable *bool 0x28 = SecondTgtVLAN All others = Reserved
15..8	Length	
	Attribute Value	Value data

3269

Table 154 – Get NVMeoFC Boot Protocol Type-Length field

Bit PositionField DescriptionValue Description7..0Attribute Name/Type 0x0 = FirstNVMeTargetNQN

0x1 = FirstNVMeTargetWWN

0x2 = FirstNVMeTargetWWPN

0x3 = FirstNVMeTgtConn

0x4 = FirstNVMeTgtCntlrlID

0x5 = FirstNVMeTgtNSID

0x6-0x7 = Reserved

0x8 = SecondNVMeTargetNQN

0x9 = SecondNVMeTargetWWN

0xA = SecondNVMeTargetWWPN

0xB = SecondNVMeTgtConn

0xC = SecondNVMeTgtCntlrlID

0xD = SecondNVMeTgtNSID

0xE-0xF = Reserved

0x10 = ThirdNVMeTargetNQN

0x11 = ThirdNVMeTargetWWN

0x12 = ThirdNVMeTargetWWPN

0x13 = ThirdNVMeTgtConn

0x14 = ThirdNVMeTgtCntlrlID

0x15 = ThirdNVMeTgtNSID

0x16-0x17 = Reserved

0x18 = FourthNVMeTargetNQN

0x19 = FourthNVMeTargetWWN

0x1A = FourthNVMeTargetWWPN

0x1B = FourthNVMeTgtConn

0x1C = FourthNVMeTgtCntlrlID

0x1D = FourthNVMeTgtNSID

0x1E-0x1F = Reserved

0x20 = FifthNVMeTargetNQN

0x21 = FifthNVMeTargetWWN

0x22 = FifthNVMeTargetWWPN

0x23 = FifthNVMeTgtConn

0x24 = FifthNVMeTgtCntlrlID

0x25 = FifthNVMeTgtNSID

0x26-0x27 = Reserved

0x28 = SixthNVMeTargetNQN
 0x29 = SixthNVMeTargetWWN
 0x2A = SixthNVMeTargetWWPN
 0x2B = SixthNVMeTgtConn
 0x2C = SixthNVMeTgtCntlrlID
 0x2D = SixthNVMeTgtNSID
 0x2E-0x2F = Reserved

0x30 = SeventhNVMeTargetNQN
 0x31 = SeventhNVMeTargetWWN
 0x32 = SeventhNVMeTargetWWPN
 0x33 = SeventhNVMeTgtConn
 0x34 = SeventhNVMeTgtCntlrlID
 0x35 = SeventhNVMeTgtNSID
 0x36-0x37 = Reserved

0x38 = EighthNVMeTargetNQN
 0x39 = EighthNVMeTargetWWN
 0x3A = EighthNVMeTargetWWPN
 0x3B = EighthNVMeTgtConn
 0x3C = EighthNVMeTgtCntlrlID
 0x3D = EighthNVMeTgtNSID
 0x3E-0xFF = Reserved

3270 **8.4.75** 15..8Length Attribute ValueValue dataSet Boot Config command (0x2E)

3271 The Set Boot Config command allows the Management Controller to send to the channel/PF/partition the
 3272 Boot settings to be used by the channel/PF/partition in conducting boot operations of the specified type.

3273 The Network Controller shall apply the attribute values in the order received in this command (e.g., TLV1
 3274 before TLV2, etc.) so that any dependency relationships are maintained.

3275 See the Get Boot Config Command for the definition of the **command** fields.

3276 All string values specified in this command shall be in unterminated ASCII string format.

3277 A NC that does not support or is not in partitioning mode shall have the Partition ID field programmed as
 3278 0x00.

3279 A TLV length value of 0 indicates the clearing of the current value of the attribute to null or no value.

3280 A maximum of 32 TLVs may be sent in any one instance of the Set Boot Config command.

3281 If the command is sent to a destination that exists but that does not support the specified Boot Protocol
 3282 type, the command execution shall fail with a reason code of Parameter Is Invalid, Unsupported, or Out-
 3283 of-Range.

3284

Table 155 – Set Boot Config command packet format

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Partition ID		Protocol Type	Number of TLVs
24..	Type-Length Field #1.		Value Field #1.	
....	Type-Length Field #2		Value Field #2	
....			
....	Checksum (3..2)		Checksum (1..0)	
....	Pad			

3285 8.4.76 Set Boot Config Response (0xAE)

3286 The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Set Boot
3287 Config command and send a response.

3288 Only if all the TLVs are accepted without error then the Command Completed/No Error response/reason
3289 code shall be returned with the TLV Error Reporting field set to all 0's.

3290 If the command is sent to a destination that exists but that does not support the specified Boot Protocol
3291 type, the command response shall return the Parameter Is Invalid, Unsupported, or Out-of-Range reason
3292 code.

3293 If there are errors in any of the TLVs included in the Set command, the entire command is deemed to fail,
3294 and no configuration changes are to be made by the controller. The TLV Error Reporting field shall be
3295 used to provide individual status reporting on the TLVs received.

3296

Table 156 – Set Boot Config Response packet format

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	TLV Error Reporting			
28..31	Checksum			
32..45	Pad			

3297 8.4.76.1 TLV Error Reporting field

3298 The TLV Error Reporting field is a bitmap indicating which TLVs were processed successfully and which
3299 were not in the incoming Set command. The bit order corresponds to the order of TLVs in the incoming
3300 Set command. There is a 1:1 correspondence between incoming TLVs and the active bits in this field. If
3301 fewer than 32 TLVs are transmitted, the bits corresponding to the unsent TLVs shall be set to 0.

3302

Table 157 – TLV Error Reporting field

Bit Position	Field Description	Value Description
0	TLV #1 status	0b = 0 No error detected in TLV1 0b = 1 Error detected in TLV1
n	TLV n+1 status	1b = 0 No error detected in TLV n+1 or TLV n+1 not present 1b = 1 Error detected in TLV n+1 all others = Reserved

3303

3304 **8.4.77 Get Partition Statistics command (0x2F)**

3305 The Get Partition Statistics command is used to retrieve network statistics relevant to the partition from
 3306 the NC. For example, the MC should only request Ethernet statistics from a partition configured for
 3307 Ethernet operation. The defined responses are customized for each personality type.

3308 Implementation of this command is conditional and is required only for NCs that support partitioning.
 3309 Implementation of each response type is conditional based on the NC supporting the specified type of
 3310 operation on the partition.

3311 The NC shall return in the response a value of 0xFFFFFFFF for unsupported 32-bit counters and
 3312 0xFFFFFFFFFFFFFFFF for unsupported 64-bit counters. For implementations that declare a particular
 3313 counter only occupies 32 bits in a defined 64-bit (upper/lower) field, the lower field shall be used to
 3314 provide the count and the upper field shall be set to 0xFFFFFFFF.

3315 As the intent of the command is to retrieve live statistics from enabled partitions, if the command is sent to
 3316 a Partition ID that doesn't exist in the current configuration or if the Stats type does not match the
 3317 configured personality of the partition, the command shall fail with the Parameter is Invalid reason code.

3318

3319

3320

Table 158 – Get Partition Statistics command packet format

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Partition ID	Reserved		Stats Type
20..23	Checksum			
24..45	Pad			

3321 **8.4.77.1 Stats Type field**

3322 The Stats Type field is the identifier for the type of statistics to be queried.

3323

3324

Table 159 – Stats Type Field

Bit Position	Field Description	Value Description
7..0	Stats Type	0x01 = Ethernet 0x02 = iSCSI 0x04 = FCoE 0x08 = RDMA 0x10 = IB All others = Reserved

3325

3326 **8.4.78** Get Partition Statistics response for Ethernet (0xAF)

3327 In the absence of any errors, the channel shall process and respond to the Get Partition Statistics
3328 Command and send the response packet shown below when the Stats Type indicates Ethernet.

3329 Currently no command-specific reason code is identified for this response.

3330

Table 160– Get Partition Statistics (Ethernet) response packet format

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	Stats Type	Counter Sizes	Counters Cleared	
24..27	Total Bytes Received (upper)			
28..31	Total Bytes Received (lower)			
32..35	Total Bytes Transmitted (upper)			
36..39	Total Bytes Transmitted (lower)			
40..43	Total Unicast Packets Received			
44..47	Total Multicast Packets Received			
48..51	Total Broadcast Packets Received			
52..55	Total Unicast Packets Transmitted			
56..59	Total Multicast Packets Transmitted			
60..63	Total Broadcast Packets Transmitted			
64..67	Total Unicast Bytes Received (upper)			
68..71	Total Unicast Bytes Received (lower)			
72..75	Total Multicast Bytes Received (upper)			
76..79	Total Multicast Bytes Received (lower)			
80..83	Total Broadcast Bytes Received (upper)			
84..87	Total Broadcast Bytes Received (lower)			
88..91	Total Unicast Bytes Transmitted (upper)			
92..95	Total Unicast Bytes Transmitted (lower)			
96..99	Total Multicast Bytes Transmitted (upper)			
100..103	Total Multicast Bytes Transmitted (lower)			
104..107	Total Broadcast Bytes Transmitted (upper)			
108..111	Total Broadcast Bytes Transmitted (lower)			
112..115	Checksum			

3331

3332 **8.4.78.1** Counter Sizes field

3333 The NC shall indicate in the Counter Sizes field whether the implementation uses 32-bit counters or 64-bit
 3334 counters in those counter fields above that are defined as 64-bit.

3335

Table 161 – Counter Sizes field format

Bit Position	Field Description	Value Description
0	Total Bytes Received	0b = 32-bit 1b = 64-bit
1	Total Bytes Transmitted	0b = 32-bit 1b = 64-bit
2	Total Unicast Bytes Received	0b = 32-bit 1b = 64-bit
3	Total Multicast Bytes Received	0b = 32-bit 1b = 64-bit
4	Total Broadcast Bytes Received	0b = 32-bit 1b = 64-bit
5	Total Unicast Bytes Transmitted	0b = 32-bit 1b = 64-bit
6	Total Multicast Bytes Transmitted	0b = 32-bit 1b = 64-bit
7	Total Broadcast Bytes Transmitted	0b = 32-bit 1b = 64-bit

3336

3337 **8.4.78.2** Counters Cleared from Last Read field

3338 The NC shall indicate in the Counters Cleared from Last Read field whether the corresponding fields have
 3339 been cleared since they were last read over NC-SI.

3340

Table 162 – Counters Cleared from Last Read field format

Bit Position	Field Description	Value Description
0	Total Bytes Received	0b = Not Cleared 1b = Cleared
1	Total Bytes Transmitted	0b = Not Cleared 1b = Cleared
2	Total Unicast Packets Received	0b = Not Cleared 1b = Cleared
3	Total Multicast Packets Received	0b = Not Cleared 1b = Cleared
4	Total Broadcast Packets Received	0b = Not Cleared 1b = Cleared
5	Total Unicast Packets Transmitted	0b = Not Cleared 1b = Cleared
6	Total Multicast Packets Transmitted	0b = Not Cleared 1b = Cleared
7	Total Broadcast Packets Transmitted	0b = Not Cleared 1b = Cleared

Bit Position	Field Description	Value Description
8	Total Unicast Bytes Received	0b = Not Cleared 1b = Cleared
9	Total Multicast Bytes Received	0b = Not Cleared 1b = Cleared
10	Total Broadcast Bytes Received	0b = Not Cleared 1b = Cleared
11	Total Unicast Bytes Transmitted	0b = Not Cleared 1b = Cleared
12	Total Multicast Bytes Transmitted	0b = Not Cleared 1b = Cleared
13	Total Broadcast Bytes Transmitted	0b = Not Cleared 1b = Cleared
15..14	Reserved	

3341

3342 **8.4.79 Get Partition Statistics response for FCoE (0xAF)**

3343 In the absence of any errors, the channel shall process and respond to the Get Partition Statistics
 3344 Command and send the response packet shown below when the Stats Type indicates FCoE.

3345 Currently no command-specific reason code is identified for this response.

3346 **Table 163 – Get Partition Statistics (FCoE) response packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	Stats Type	Counter Sizes	Counters Cleared	
24..27	Total FCoE Bytes Received (upper)			
	Total FCoE Bytes Received (lower)			
	Total FCoE Bytes Transmitted (upper)			
	Total FCoE Bytes Transmitted (lower)			
	Total FCoE Packets Received (upper)			
	Total FCoE Packets Received (lower)			
	Total FCoE Packets Transmitted (upper)			
	Total FCoE Packets Transmitted (lower)			
	Checksum			

3347

3348 **8.4.79.1** Counter Sizes field

3349 The NC shall indicate in the Counter Sizes field whether the implementation uses 32-bit counters or 64-bit
 3350 counters in those counter fields above that are defined as 64-bit.

3351 **Table 164 – Counter Sizes field format**

Bit Position	Field Description	Value Description
0	Total FCoE Bytes Received	0b = 32-bit 1b = 64-bit
1	Total FCoE Bytes Transmitted	0b = 32-bit 1b = 64-bit
2	Total FCoE Packets Received	0b = 32-bit 1b = 64-bit
3	Total FCoE Packets Transmitted	0b = 32-bit 1b = 64-bit
4..7	Reserved	Reserved

3352 **8.4.79.2** Counters Cleared from Last Read

3353 The NC shall indicate in the Counters Cleared from Last Read field whether the corresponding fields have
 3354 been cleared since they were last read over NC-SI.

3355 **Table 165 – Counters Cleared from Last Read field format**

Bit Position	Field Description	Value Description
0	Total FCoE Bytes Received	0b = Not Cleared 1b = Cleared
1	Total FCoE Packets Transmitted	0b = Not Cleared 1b = Cleared
2	Total FCoE Packets Received	0b = Not Cleared 1b = Cleared
3	Total FCoE Packets Transmitted	0b = Not Cleared 1b = Cleared
15..4	Reserved	Reserved

3356

3357 **8.4.80** Get Partition Statistics response for iSCSI (0xAF)

3358 In the absence of any errors, the channel shall process and respond to the Get Partition Statistics
 3359 Command and send the response packet shown below when the Stats Type indicates iSCSI.

3360 Currently no command-specific reason code is identified for this response.

3361

Table 166 – Get Partition Statistics (iSCSI) response packet format

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	Stats Type	Counter Sizes	Counters Cleared	
24..27	Total iSCSI Offload Bytes Received (upper)			
	Total iSCSI Offload Bytes Received (lower)			
	Total iSCSI Offload Bytes Transmitted (upper)			
	Total iSCSI Offload Bytes Transmitted (lower)			
	Total iSCSI Offload PDUs Received (upper)			
	Total iSCSI Offload PDUs Received (lower)			
	Total iSCSI Offload PDUs Transmitted (upper)			
	Total iSCSI Offload PDUs Transmitted (lower)			
	Checksum			

3362

3363 **8.4.80.1 Counter Sizes field**

3364 The NC shall indicate in the Counter Sizes field whether the implementation uses 32-bit counters or 64-bit
 3365 counters in those counter fields above that are defined as 64-bit.

3366

Table 167 – Counter Sizes field format

Bit Position	Field Description	Value Description
0	Total iSCSI Offload Bytes Received	0b = 32-bit 1b = 64-bit
1	Total iSCSI Offload Bytes Transmitted	0b = 32-bit 1b = 64-bit
2	Total iSCSI Offload PDUs Received	0b = 32-bit 1b = 64-bit
3	Total iSCSI Offload PDUs Transmitted	0b = 32-bit 1b = 64-bit
4..7	Reserved	Reserved

3367 **8.4.80.2 Counters Cleared from Last Read**

3368 The NC shall indicate in the Counters Cleared from Last Read field whether the corresponding fields have
 3369 been cleared since they were last read over NC-SI.

3370 **Table 168 – Counters Cleared from Last Read field format**

Bit Position	Field Description	Value Description
0	Total iSCSI Offload Bytes Received	0b = Not Cleared 1b = Cleared
1	Total iSCSI Offload Bytes Transmitted	0b = Not Cleared 1b = Cleared
2	Total iSCSI Offload PDUs Received	0b = Not Cleared 1b = Cleared
3	Total iSCSI Offload PDUs Transmitted	0b = Not Cleared 1b = Cleared
15..4	Reserved	Reserved

3371

3372 **8.4.81 Get Partition Statistics response for InfiniBand (0xAF)**

3373 In the absence of any errors, the channel shall process and respond to the Get Partition Statistics
 3374 Command and send the response packet shown below when the Stats Type indicates InfiniBand.

3375 Currently no command-specific reason code is identified for this response.

3376 **Table 169 – Get Partition Statistics (IB) response packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	Stats Type	Counter Sizes	Counters Cleared	
24..27	Total Unicast Packets Received			
	Total Multicast Packets Received			
	Total Unicast Packets Transmitted			
	Total Multicast Packets Transmitted			
	Total Unicast Bytes Received			
	Total Multicast Bytes Received			
	Total Unicast Bytes Transmitted			
	Total Multicast Bytes Transmitted			
	Checksum			

3377

3378 **8.4.81.1 Counter Sizes field**

3379 The NC shall indicate in the Counter Sizes field whether the implementation uses 32-bit counters or 64-bit
 3380 counters in those counter fields above that are defined as 64-bit.

3381

3382

Table 170 – Counter Sizes field format

Bit Position	Field Description	Value Description
0	Total Unicast Packets Received	0b = 32-bit 1b = 64-bit
1	Total Unicast Packets Transmitted	0b = 32-bit 1b = 64-bit
2	Total Multicast Packets Received	0b = 32-bit 1b = 64-bit
3	Total Multicast Packets Transmitted	0b = 32-bit 1b = 64-bit
4	Total Unicast Bytes Received	0b = 32-bit 1b = 64-bit
5	Total Unicast Bytes Transmitted	0b = 32-bit 1b = 64-bit
6	Total Multicast Bytes Received	0b = 32-bit 1b = 64-bit
7	Total Broadcast Bytes Transmitted	0b = 32-bit 1b = 64-bit

3383 **8.4.81.2** Counters Cleared from Last Read

3384 The NC shall indicate in the Counters Cleared from Last Read field whether the corresponding fields have
 3385 been cleared since they were last read over NC-SI.

3386

Table 171 – Counters Cleared from Last Read field format

Bit Position	Field Description	Value Description
0	Total Unicast Packets Received	0b = Not Cleared 1b = Cleared
1	Total Multicast Packets Received	0b = Not Cleared 1b = Cleared
2	Total Unicast Packets Transmitted	0b = Not Cleared 1b = Cleared
3	Total Multicast Packets Transmitted	0b = Not Cleared 1b = Cleared
4	Total Unicast Bytes Received	0b = Not Cleared 1b = Cleared
5	Total Multicast Bytes Received	0b = Not Cleared 1b = Cleared
6	Total Unicast Bytes Transmitted	0b = Not Cleared 1b = Cleared

Bit Position	Field Description	Value Description
7	Total Multicast Bytes Transmitted	0b = Not Cleared 1b = Cleared
15..8	Reserved	

3387

3388 **8.4.82** Get Partition Statistics response for RDMA (0xAF)

3389 In the absence of any errors, the channel shall process and respond to the Get Partition Statistics
3390 Command and send the response packet shown below when the Stats Type indicates RDMA.

3391 Currently no command-specific reason code is identified for this response.

3392 **Table 172 – Get Partition Statistics (RDMA) response packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	Stats Type	Counter Sizes	Counters Cleared	
24..27	Total RDMA Bytes Received (upper)			
28..31	Total RDMA Bytes Received (lower)			
32..35	Total RDMA Bytes Transmitted (upper)			
36..39	Total RDMA Bytes Transmitted (lower)			
40..43	Total RDMA Packets Received (upper)			
44..47	Total RDMA Packets Received (lower)			
48..51	Total RDMA Packets Transmitted (upper)			
52..55	Total RDMA Packets Transmitted (lower)			
56..59	Total Read Request Packets Transmitted (upper)			
60..63	Total Read Request Packets Transmitted (lower)			
64..67	Total Send Packets Transmitted (upper)			
68..71	Total Send Packets Transmitted (lower)			
72..75	Total Write Packets Transmitted (upper)			
76..79	Total Write Packets Transmitted (lower)			
80..83	Checksum			

3393

3394 **8.4.82.1** Counter Sizes

3395 The NC shall indicate in the Counter Sizes field whether the implementation uses 32-bit counters or 64-bit
3396 counters in those counter fields above that are defined as 64-bit.

3397

Table 173 – Counter Sizes field format

Bit Position	Field Description	Value Description
0	Total RDMA Bytes Received	0b = 32-bit 1b = 64-bit
1	Total RDMA Bytes Transmitted	0b = 32-bit 1b = 64-bit
2	Total RDMA Packets Received	0b = 32-bit 1b = 64-bit
3	Total RDMA Packets Transmitted	0b = 32-bit 1b = 64-bit
4	Total Read Request Packets Transmitted	0b = 32-bit 1b = 64-bit
5	Total Send Packets Transmitted	0b = 32-bit 1b = 64-bit
6	Total Write Packets Transmitted	0b = 32-bit 1b = 64-bit
7	Reserved	

3398

3399

3400 **8.4.82.2** Counters Cleared from Last Read

3401 The NC shall indicate in the Counters Cleared from Last Read field whether the corresponding fields have
 3402 been cleared since they were last read over NC-SI.

3403

Table 174 – Counters Cleared from Last Read field format

Bit Position	Field Description	Value Description
0	Total RDMA Bytes Received	0b = Not Cleared 1b = Cleared
1	Total RDMA Bytes Transmitted	0b = Not Cleared 1b = Cleared
2	Total RDMA Packets Received	0b = Not Cleared 1b = Cleared
3	Total RDMA Packets Transmitted	0b = Not Cleared 1b = Cleared
4	Total Read Request Packets Transmitted	0b = Not Cleared 1b = Cleared
5	Total Send Packets Transmitted	0b = Not Cleared 1b = Cleared
6	Total Write Packets Transmitted	0b = Not Cleared 1b = Cleared

Bit Position	Field Description	Value Description
15..7	Reserved	

3404

3405

3406 **8.4.83 Get Partition Statistics Response for Fibre Channel (0xAF)**

3407 The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Get
 3408 Partition Statistics command and send a response when the Stats Type indicates FC.

3409 Table 175 illustrates the packet format of the Get FC Statistics Response.

3410 **Table 175 – Get Partition Statistics (FC) Response packet**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Stats Type	Reserved	Counters Cleared from Last Read	
24..27	Total FC Frames Received			
28..31	Total FC Frames Transmitted			
32..35	Receive KB Count			
36..39	Transmit KB Count			
40..43	FC Sequences Received			
44..47	FC Sequences Transmitted			
48..51	Link Failures			
52..55	Loss of Signal			
56..59	Invalid CRCs			
60..63	Checksum (3..2)		Checksum (1..0)	

3411 **8.4.83.1 Counters Cleared from Last Read field**

3412 The FC Controller shall also indicate in the Counters Cleared from Last Read field whether the
 3413 corresponding fields has been cleared since it was last read via NC-SI. The Counters Cleared from Last
 3414 Read fields should have the format shown in Table 176.

3415 **Table 176 – Counters Cleared from Last Read field format**

Bit Position	Field Description	Value Description
0	Total FC Frames Received	0b = Not Cleared 1b = Cleared
1	Total FC Frames Transmitted	0b = Not Cleared 1b = Cleared

Bit Position	Field Description	Value Description
2	Receive KB Count	0b = Not Cleared 1b = Cleared
3	Transmit KB Count	0b = Not Cleared 1b = Cleared
4	FC Sequences Received	0b = Not Cleared 1b = Cleared
5	FC Sequences Transmitted	0b = Not Cleared 1b = Cleared
6	Link Failures	0b = Not Cleared 1b = Cleared
7	Loss of Signal	0b = Not Cleared 1b = Cleared
8	Invalid CRCs	0b = Not Cleared 1b = Cleared
15..9	Reserved	

3416 8.4.83.2 FC Statistics Counter definitions

3417 **Table 177 – FC Statistics**

Name	Meaning
Total FC Frames Received	Counts the number of FC frames received by the port
Total FC Frames Transmitted	Counts the number of FC frames transmitted by the port
Receive KB Count	Counts the number of kilobytes transmitted by the port
Transmit KB Count	Counts the number of kilobytes transmitted by the port
FC Sequences Received	Counts the number of FC sequences received by the port
FC Sequences Transmitted	Counts the number of FC sequences transmitted by the port
Link Failures	Counts the number of times the link has failed.
Loss of Signal	Counts the number of times the signal was lost.
Invalid CRCs	Counts the number of CRC errors detected.

3418

3419 8.4.84 Get FC Link Status command (0x31)

3420 The Get FC Link Status command allows the Management Controller to query the channel for potential
 3421 link status and error conditions (see Table 178).

3422 Implementation of this command is conditional and is required only for controllers supporting native Fibre
3423 Channel.

3424 Implementation note:

3425 Some controllers may include a port trunking (bonding) capability in which one (or more) channels will
3426 map to multiple physical ports. FC trunking (bonding) is based on the following rules:

- 3427 • FC controllers provide a maximum of 4 physical ports
- 3428 • All ports are configured to the same speed
- 3429 • If trunking is enabled, all ports become involved in a bond, no standalone ports remain
- 3430 • Ports may bond in pairs or all together
 - 3431 ○ Dual port controllers bond Ports 1&2 and present one channel to the MC
 - 3432 ○ Quad port controllers bond Ports (1&2) [trunk 1] and {3&4} [trunk2] or {1&2&3&4} and
 - 3433 present two or one channel(s) respectively

3434 **Table 178 – Get FC Link Status command packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Reserved	Reserved	Reserved	Reserved
20..23	Checksum (3..2)		Checksum (1..0)	
24..27	Pad			

3435

3436 **8.4.85 Get FC Link Status Response (0xB1)**

3437 The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Get FC
3438 Link Status command and send a response (see Table 179).

3439 **Table 179 – Get FC Link Status Response packet format**

Bits				
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	# of FC Ports	FC Trunk Status	FC Link Status	Trunk Speeds
24..27	Channel 1 Link Speed	Channel 2 Link Speed	Channel 3 Link Speed	Channel 4 Link Speed
28..31	Checksum			
33..36	Pad			

3440 **8.4.85.1 # of FC Ports field**

3441 This is an integer value that specifies the total number of physical ports on the Package

3442 **8.4.85.2 FC Trunk Status field**

3443 This field indicates if the physical port is a member of a FC trunk.

3444 **Table 180 – FC Trunk Status field bit definitions**

Bit Position	Field Description	Value Description
0	Port 1 Trunk Flag	0b = Physical Port 1 Is not a member of a trunk 1b = Physical Port 1 Is a member of a trunk
1	Port 2 Trunk Flag	0b = Physical Port 2 Is not a member of a trunk 1b = Physical Port 2 Is a member of a trunk
2	Port 3 Trunk Flag	0b = Physical Port 3 Is not a member of a trunk 1b = Physical Port 3 Is a member of a trunk
3	Port 4 Trunk Flag	0b = Physical Port 4 Is not a member of a trunk 1b = Physical Port 4 Is a member of a trunk
7..4	Reserved	None

3445

3446 **8.4.85.3 FC Link Status field**

3447 Table 181 describes the FC Link Status field bit definitions.

3448 **Table 181 – FC Link Status field bit definitions**

Bit Position	Field Description	Value Description
0	Port 1 Link Flag	0b = Physical Port 1 Link is down 1b = Physical Port 1 Link is up
1	Port 2 Link Flag	0b = Physical Port 2 Link is down 1b = Physical Port 2 Link is up
2	Port 3 Link Flag	0b = Physical Port 3 Link is down 1b = Physical Port 3 Link is up
3	Port 4 Link Flag	0b = Physical Port 4 Link is down 1b = Physical Port 4 Link is up
7..5	Reserved	None

3449

3450 **8.4.85.4 Trunk Speeds field**

3451 The percentage of the configured trunk speed that is currently available represented as an integer.

3452 **Error! Reference source not found.** describes the Trunk Speeds field.

3453 **Table 182 – Trunk Speeds field**

Bit Position	Field Description	Value Description
3..0	Trunk 1 Percentage Speed	Percentage of the Trunk 1 configured link speed that is available expressed as hex value. Not applicable if no Trunks are configured. 0x0 = 0% 0x1 = 25% 0x2 = 50% 0x3 = 75% 0x4 = 100%
7..4	Trunk 2 Percentage Speed	Percentage of the Trunk 2 configured link speed that is available (expressed as hex value. Not applicable if two Trunks are not configured. 0x0 = 0% 0x2 = 50% 0x4 = 100%

3454

3455 **8.4.85.5 FC Link Speed field**

3456 The Link Speed field provides a link speed based on NC-SI Channel configuration. If the number of FC
 3457 ports is equal to the number of reported NC-SI channels, then trunking is not active, and the reported
 3458 speed is the speed of the channel on the port. In two- or four-port trunking modes, the number of FC
 3459 ports will be twice or four times the number of reported NC-SI channels and the reported configured link
 3460 speed is the sum of the individual link speeds in the trunk. If one or more of the member links goes down
 3461 the reported link speed will not change, but the FC Link Status and Trunk Speed fields will provide the
 3462 indication that the trunk is not operating at its stated speed.

3463 **Error! Reference source not found.** describes the FC Link Speed field bit definitions.

3464 **Table 183 – FC Link Speed field**

Value	Field Description	Value Description
0	Link Speed	0x0 = No link speed established 0x1 = FC2 0x2 = FC4 0x3 = FC8 0x4 = FC16 0x5 = FC32 0x6 = FC64 0x7 = FC128 0x8 = FC256

Value	Field Description	Value Description
Others	Reserved	None

3465

3466 **8.4.86** Get Transceiver Management Data command (0x32)

3467 The Get Transceiver Management Data command is used to retrieve 128-byte blocks of management
 3468 and inventory data stored in the passive copper cable or optical transceiver module associated with the
 3469 channel. Different standards and specifications exist in the industry for this management data, but they
 3470 share common data access methods allowing this command to successfully operate with the known
 3471 variety of module interface specifications.

3472 A two-byte Type identifier is used to specify the bank and page index of the target data to be returned.
 3473 Some devices only support 1 bank and therefore will only respond with data with the bank index set to
 3474 0x00.

3475 The lower 128 bytes of page 00h typically contains more important time-critical data. The upper 128
 3476 bytes of page 00h contains static inventory information. The implementation may read and cache the
 3477 upper 128 bytes once upon power on or module insertion to expedite processing of requests for page 00h
 3478 data.

3479 This command should fail as unsupported on backplane and RJ-45 implementations.

3480 Table 135 illustrates the packet format for the Get Transceiver Management Data command.

3481 **Table 184 – Get Transceiver Management Data command packet format**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Requested Bank	Requested Page	Reserved	Flags
20..23	Checksum			
24..45	Pad			

3482 **8.4.86.1** Requested Bank field

3483 The Requested Bank field is the value of the Bank data being requested.

3484 **8.4.86.2** Requested Bank field

3485 The Requested Bank field is the value of the Bank data being requested.

8.4.86.3 Flags field**Table 185 – Flag field bit definitions**

Bit Position	Field Description	Value Description
0	Page Upper Flag	0b = Requesting lower page data 1b = Requesting upper page data
7..1	Reserved	None

8.4.87 Get Transceiver Management Data response (0xB2)

In the absence of any errors, the channel shall process and respond to the Get Transceiver Management Data Command and send the response packet shown in Table 136.

Currently no command-specific reason code is identified for this response.

If there is no module installed, then use response/reason codes Command Unavailable/Information not available

Use the Command Failed reason code with the following conditions:

If the Requested Bank or Page number does not exist, then use reason code Parameter Out-of-Range

If the module is resetting or powering up, then use reason code Information Not Available

If the module cannot respond with data in the allocated time, either use Command Timeout or Delayed Response as supported by the implementation.

Table 186 – Get Transceiver Management Data response packet format

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	Max Bank	Max Page	Bank Number	Page Number
24..27	Data ₀	Data ₁	...	
48..51	Checksum			

8.4.87.1 Max Bank field

The Max Bank field contains the value of the highest Bank number supported by the module. If the module type does not support Banks, the field shall be set to 0x00.

8.4.87.2 Max Page field

The Max Page field contains the value of the highest Page number in the current Bank supported by the module.

8.4.87.3 Bank Number field

The Bank Number field contains the value of the Bank number requested by the command.

8.4.87.4 Page Number field

The Page Number field contains the value of the Page number requested by the command.

8.4.88 Get InfiniBand Link Status command (0x38)

The Get InfiniBand Link Status command allows the Management Controller to query the channel for the IB Statistics.

Implementation of this command is conditional and is required only for controllers supporting InfiniBand.

Table 187 illustrates the packet format of the InfiniBand Link Status command.

Table 187 – Get InfiniBand Link Status command

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Checksum (3..2)		Checksum (1..0)	
20..45	Pad			

8.4.89 Get InfiniBand Link Status Response (0xB8)

The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Get InfiniBand Link Status command and send a response.

The Get InfiniBand Link Status Response frame reports link width, logical and physical link states, and the supported and the configured link speed of the port.

Table 188 illustrates the packet format of the Get InfiniBand Link Status Response.

3527

Table 188 – Get InfiniBand Link Status Response packet

	Bits				
Bytes	31..24	23..16	15..08	07..00	
00..15	NC-SI Header				
16..19	Response Code		Reason Code		
28..31	IB Link Active Width	IB Link Supported Width	Link Type	Phys State	Log State
32..35	Reserved	IB Link Active Speed	Reserved	IB Link Supported Speed	
36..47	Checksum (3..2)		Checksum (1..0)		

3528

3529

Table 189 – InfiniBand Link Status definitions

Name	Direction	Description
IB Link Active Width	TX	<p>When Link Type is InfiniBand and physical link is up, this field reflects the active link width. Otherwise this field returns 0b.</p> <p>Bit 0 – 1b = 1X link width</p> <p>Bit 1 - 1b = 2X link width</p> <p>Bit 2 - 1b = 4X link width</p> <p>Bit 3 - 1b = 8X link width</p> <p>Bits 7:4 Reserved</p>
IB Link Supported Width	RX	<p>When Link Type is InfiniBand, this field reflects the supported link widths. When Link Type is Ethernet, this field returns 0.</p> <p>Bit 0 - 1b = 1X link width is supported</p> <p>Bit 1 - 1b = 2X link width is supported</p> <p>Bit 2 - 1b = 4X link width is supported</p> <p>Bit 3 - 1b = 8X link width is supported</p> <p>Bits 7:4 Reserved</p>
Link Type	TX	<p>Reflects the configured link type.</p> <p>Bit 0 - 0b = Ethernet</p> <p>1b = InfiniBand</p>

Name	Direction	Description
Phys State	RX	<p>The physical link state as specified in IB spec (PortInfoPortPhysicalState)</p> <p>0x0 = Used when Link Type is Ethernet</p> <p>0x1 = Sleep</p> <p>0x2 = Polling</p> <p>0x3 = Disabled</p> <p>0x4 = PortConfigurationTraining</p> <p>0x5 = LinkUp</p> <p>0x6 = LinkErrorRecovery</p> <p>0x7 = PhyTest</p>
Logical Port State	TX	<p>The logical port state of the physical port as specified in IB spec (PortInfo.PortState)</p> <p>0x0: Used when Link Type is Ethernet</p> <p>0x1: Down</p> <p>0x2: Init</p> <p>0x3: Arm</p> <p>0x4: Active</p>
IB Link Active Speed	TX	<p>When Link Type is InfiniBand and the physical link is up, this field reflects the active link speed. Otherwise this field returns 0x00.</p> <p>Bit 0 – 1b = SDR</p> <p>Bit 1 - 1b = DDR</p> <p>Bit 2 - 1b = QDR</p> <p>Bit 3 - 1b = FDR10</p> <p>Bit 4 - 1b = FDR</p> <p>Bit 5 - 1b = EDR</p> <p>Bit 6 - 1b = HDR</p> <p>Bit 7 - 1b = NDR</p>

Name	Direction	Description
IB Link Supported Speed	RX	<p>When Link Type is InfiniBand, this field reflects the supported link speeds. When Link Type is Ethernet this field returns 0x00.</p> <p>Bit 0 - 1b = SDR</p> <p>Bit 1 - 1b = DDR</p> <p>Bit 2 - 1b = QDR</p> <p>Bit 3 - 1b = FDR10</p> <p>Bit 4 - 1b = FDR</p> <p>Bit 5 - 1b = EDR</p> <p>Bit 6 - 1b = HDR</p> <p>Bit 7 - 1b = NDR</p>

3530

3531 **8.4.90 Get IB Statistics command (0x39)**

3532 The Get IB Statistics command allows the Management Controller to query the channel for the IB
 3533 Statistics.

3534 Implementation of this command is conditional and is required only for controllers supporting InfiniBand.

3535 Table 190 illustrates the packet format of the Get IB Statistics Command.

3536

3537 **Table 190 – Get IB Statistics Command**

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

3538

3539 **8.4.91 Get IB Statistics Response (0xB9)**

3540 The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Get IB
 3541 Statistics command and send a response.

3542 The Get IB Statistics Response frame reports a set of IB statistics from the channel. A value of
 3543 0xFFFFFFFF shall be used for any unsupported counter.

3544 All counters are reset on Controller resets or power-cycles only.

3545 Table 191 illustrates the packet format of the Get IB Statistics Response.

3546

Table 191 – Get IB Statistics Response packet

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	PortXmitData			
24..27	PortRcvData			
28..31	PortXmitPkts			
32..35	PortRcvPkts			
36..39	PortXmitWait			
40..43	PortXmitDiscard			
44..47	SymbolErrorCounter			
48..51	LinkErrorRecoveryCounter			
52..55	LinkDownedCounter			
56..59	PortRcvErrors			
60..63	PortRcvRemotePhysicalErrors			
64..67	PortRcvSwitchRelayErrors			
68..71	LocalLinkIntegrityErrors			
72..75	ExcessiveBufferOverrun			
76..79	VL15Dropped			
80..83	Checksum (3..2)		Checksum (1..0)	

3547

3548

Table 192 – IB Statistics Counter definitions

Name	Direction	Description
PortXmitData	TX	Total number of data octets, divided by 4 (lanes), transmitted on all VLs.
PortRcvData	RX	Total number of data octets, divided by 4 (lanes), received on all VLs.
PortXmitPkts	TX	Total number of packets transmitted on all VLs from this port. This may include packets with errors.
PortRcvPkts	RX	Total number of packets (this may include packets containing Errors).
PortXmitWait	TX	Number of ticks during which the port had data to transmit but no data was sent during the entire tick (either because of insufficient credits or because of lack of arbitration).
PortXmitDiscard	TX	Total number of outbound packets discarded by the port because the port is down or congested.

Name	Direction	Description
SymbolErrorCounter	RX	Total number of minor link errors detected on one or more physical lanes.
LinkErrorRecoveryCounter	RX	Total number of times the Port Training state machine has successfully completed the link error recovery process.
LinkDownedCounter	RX	Total number of times the Port Training state machine has failed the link error recovery process and downed the link.
PortRcvErrors	RX	Total number of packets containing an error that were received on the port.
PortRcvRemotePhysicalErrors	RX	Total number of packets marked with the EBP delimiter received on the port.
PortRcvSwitchRelayErrors	RX	Total number of packets received on the port that were discarded because they could not be forwarded by the switch relay.
LocalLinkIntegrityErrors	RX	Number of times that the count of local physical errors exceeded the threshold specified by LocalPhyErrors.
ExcessiveBufferOverrun	RX	Number of times that OverrunErrors consecutive flow control update periods occurred, each having at least one overrun error.
VL15Dropped	RX	Number of incoming VL15 packets dropped due to resource limitations (e.g., lack of buffers) of the port.

3549

3550 **8.4.92 Settings Commit command (0x47)**

3551 The Settings Commit command is a package command used by the Management Controller to indicate
 3552 that those previously programmed settings defined as persistent must now be written to non-volatile
 3553 storage. It also indicates that any previously programmed individual settings that have dependencies on
 3554 other settings (e.g., partition bandwidth) have been fully programmed and can be finalized and/or
 3555 validated. Only those settings in commands that returned successful response/reason codes will be
 3556 written to non-volatile storage.

3557 The MC can only be assured that settings have been persisted when this commit command has a
 3558 successful completion. It is highly likely that execution of this command will result in a Delayed
 3559 Response. The MC should assume that all settings that were sent but not committed are lost on losses
 3560 of power, various types of resets as defined by the NC, return to initial states of any affected channel, etc.
 3561 and must be resent after the interruption.

3562 **Error! Reference source not found.** illustrates the packet format of the Settings Commit command.

3563 **Table 193 – Settings Commit command packet format**

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

8.4.93 Settings Commit response (0xC7)

The package shall, in the absence of an error, always accept the Settings Commit command and send the response packet shown in Table 246.

Currently no command-specific reason code is identified for this response.

Table 194 – Settings Commit response packet format

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

8.4.94 Get ASIC Temperature (0x48)

The Get ASIC Temperature command allows the Management controller to query for temperature values from the Controller's on-chip thermal sensor(s) or alternately from attached (external) devices.

The Get ASIC Temperature command is defined as both a package level command and a channel command. This means the command can be either addressed to the package (that is, the command is sent with a Channel ID set to 0x1F) or addressed to a specific channel in the package.

When sent as a package command, the internal temperature of the controller is returned. If the controller has multiple internal temperature sensors, the highest measured temperature with respect to its threshold shall be returned.

In cases where there are other devices connected to the controller that can also report silicon temperature via the controller (such as one or more external PHYs), then the channel version of the command is used and the response contains the temperature data and threshold from the external device on that channel. Multiple sensor implementations in the external device shall be handled as described above.

Table 195 illustrates the packet format of the Get ASIC Temperature Command.

Table 195 – Get ASIC Temperature Command packet

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

3587 8.4.95 Get ASIC Temperature Response (0xC8)

3588 The package shall, in the absence of a checksum error or identifier mismatch, always accept the Get
3589 ASIC Temperature Command and send a response.

3590 Table 196 illustrates the packet format of the Get ASIC Temperature Response.

3591 **Table 196 – Get ASIC Temperature Response packet**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Maximum temperature		Current temperature	
24..27	Checksum			
28..45	Pad			

3592 8.4.95.1 Maximum Temperature Value

3593 This value is the maximum T-Diode temperature limit in degrees Celsius at which the controller can
3594 operate at full load for its rated service lifetime. The value should be derated to take measurement
3595 tolerance into account. The value shall be reported as a signed 16-bit integer.

3596 8.4.95.2 Current Temperature Value

3597 This value is the highest current real-time temperature of the ASIC sensors in degrees Celsius. The value
3598 shall be reported as a signed 16-bit integer.

3599 8.4.96 Get Ambient Temperature (0x49)

3600 The Get Ambient Temperature command allows the Management controller to query for temperature
3601 values from ambient temperature sensor(s) attached to the Controller.

3602 The Get Ambient Temperature command is defined as a package command.

3603 Controllers that do not support ambient temperature sensors should not implement this command.

3604 Table 197 illustrates the packet format of the Get Ambient Temperature command.

3605 **Table 197 – Get Ambient Temperature command packet**

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

3606

3607 8.4.97 Get Ambient Temperature Response (0xC9)

3608 The Package shall, in the absence of a checksum error or identifier mismatch, always accept the Get
3609 Ambient Temperature Command and send a response.

3610 Table 198 illustrates the packet format of the Get Ambient Temperature Response.

3611 **Table 198 – Get Ambient Temperature Response packet**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Temperature3 value	Temperature2 Value	Temperature1 Value	Number of sensors
24..27	Checksum			
28..45	Pad			

3612 8.4.97.1 Temperature Value

3613 This value (zero or more as specified by the Number of sensors field) is the real time ambient
3614 temperature reported in degrees Celsius. The value shall be reported as a signed 8-bit integer.

3615 8.4.98 Get Transceiver Temperature (0x4A)

3616 The Get Transceiver Temperature command allows the Management controller to query for the real time
3617 temperature value and thresholds of the (optical) transceiver attached to the channel. Implementations
3618 that do not support any type of temperature reporting module, such as a Base-T or backplane Ethernet
3619 adapter, should not implement this command.

3620 Table 199 illustrates the packet format of the Get Transceiver Temperature Command.

3621 **Table 199 – Get Transceiver Temperature Command Packet**

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

3622 8.4.99 Get Transceiver Temperature Response (0xCA)

3623 The channel shall, in the absence of a checksum error or identifier mismatch, always accept the Get
3624 Transceiver Temperature command and send a response.

3625 The Get Transceiver Temperature Response frame contains the current temperature of the attached
3626 module and the high side temperature thresholds.

3627 Definitions and interpretation of the data fields in the response are defined in the relevant SFF or MSA
3628 specification (e.g., SFF-8472, SFF-8436, SFF-8636, CMIS etc.) for the transceiver. 16-bit values are

3629 encoded as one contiguous entity with the most significant bit in bit 15 (or 31) and least significant bit in
 3630 bit 0 (or 16) in the response packet. The Controller is not expected to modify the data read from the
 3631 transceiver.

3632 In cases where the transceiver supports more than one channel, each channel shall provide a response
 3633 when queried.

3634 The reason code - *Information not available* - shall be used if the transceiver is not present, does not
 3635 provide temperature data or if the command is issued before the transceiver has not yet achieved power
 3636 up state.

3637 Table 200 illustrates the packet format of the Get Transceiver Temperature Response.

3638 **Table 200 – Get Transceiver Temperature Response packet**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Temp High Alarm Threshold		Temp High Warning Threshold	
24..27	Temperature Value		Reserved	
28..31	Checksum			

3639

3640 **8.4.100 Thermal Shutdown Control Command (0x4B)**

3641 The Thermal Shutdown Control command allows the Management controller to query for the state of or
 3642 alternatively set or reset the enablement state of the NC's thermal self-shutdown feature. NCs shall
 3643 indicate the implementation state of this feature in the Get Capabilities command response bit 7 and
 3644 implement this command/response appropriately. .

3645 The Thermal Shutdown Control command is defined as a package level command and is sent with the
 3646 Channel ID set to 0x1F.

3647 Table 195 illustrates the packet format of the Thermal Shutdown Control Command.

3648 **Table 201 – Thermal Shutdown Control Command packet**

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

3649

3650 **8.4.100.1 Command Field**

3651 The value specified in this field defines the action required for the NC's shutdown feature.

3652

Table 202 – Command field bit definitions

Value	Description	Value Description
0	Disable	Thermal self-shutdown shall be disabled on the device
1	Enable	Thermal self-shutdown shall be enabled on the device
2	Query	The currently configured shutdown setting shall be returned
others	Reserved	None

3653

3654 8.4.101 Thermal Shutdown Control Response (0xCB)

3655 The package shall, in the absence of a checksum error or identifier mismatch, always accept the Thermal
 3656 Shutdown Control Command and send a response.

3657 The Operating State status provided in the response shall be confirming the state after the execution of
 3658 the command. If the Config Control state is set to Read-only, any command to enable or disable the
 3659 feature shall be failed with the Parameter Is Invalid reason code. The other fields shall be included in the
 3660 response with their current setting.

3661 Table 196 illustrates the packet format of the Thermal Shutdown Control Response.

3662

Table 203 – Thermal Shutdown Control Response packet

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Reserved	Reserved	Status	Shutdown temperature
24..27	Checksum			
28..45	Pad			

3663 8.4.101.1 Shutdown Temperature Value

3664 This value is the integer temperature value in degrees Celsius at which the NC will shut itself down when
 3665 reached.

3666 8.4.101.2 Status Field

3667 The value returned in this field is the enablement status of the shutdown feature.

3668

Table 204 – Status field bit definitions

Bit	Description	Value Description
0	Operating State	0b = Thermal self-shutdown is disabled on the device 1b = Thermal self-shutdown is enabled on the device

Bit	Description	Value Description
1	Config Control	0b = Thermal self-shutdown setting is read-only 1b = Thermal self-shutdown setting is configurable
others	Reserved	None

3669

3670 **8.4.102** Get Inventory Information command (0x4E)

3671 The Get Inventory Information command may be used by the Management Controller to query the
3672 Network Controller for defined inventory information about the NC.

3673 This command is defined as a package command.

3674 Table 242 illustrates the packet format of the Inventory Information command.

3675 **Table 205 – Get Inventory Information command packet format**

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

3676 **8.4.103** Get Inventory Information response (0xCE)

3677 The package shall, in the absence of an error, always accept the Get Inventory Information command and
3678 send the response packet shown in Table 243. The value fields are defined as non-terminated ASCII
3679 strings.

3680 Currently no command-specific reason code is identified for this response.

3681 **Table 206 – Get Inventory Information response packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..35	Number of TLVs	Type-Length Field #1		Value Field #1
	...			
	Checksum			
	Pad			

3682

3683 **8.4.103.1** Inventory Information Type-Length-Value fields

3684 The Type definitions for the inventory elements are defined below.

Table 207 – Inventory Information Type-Length field

Bit Position	Field Description	Value Description
7..0	Attribute Name/Type	0x0 = Product Number 0x1 = Serial Number 0x2-0x7F = Reserved 0x80-0xFF = Reserved for OEM use
15..8	Length	Length in bytes of the field

8.5 Set Pass-through Mode Control Command (0x4F)

The Set Pass-through Mode Control command allows the Management controller to enable and disable specified data paths for Pass-through data on the channel when supported by the NC.

Implementation of this command is conditional depending on the support of Host-BMC Pass-through and embedded CPU-BMC Pass-through functionality.

The Host-BMC Pass-through, Network-BMC Pass-through and embedded CPU-BMC Pass-through controls specified in this command act as masks in conjunction with the existing Enable Channel and Enable Channel TX commands. The existing Pass-through MAC address and filtering control methods are simply extended to all defined data paths when configured. No additional filters or MACs are provided.

Table 57 illustrates the packet format for the Set Pass-through Mode Control Command.

Table 57 – Set Pass-through Mode Control Command

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

8.5.1 Pass-through Type Field

The Pass-through Type field indicates which Pass-through data path is to be enabled or disabled as described in Table 58.

3704

Table 58 – Pass-through Type definitions

Bit	Field Description	Value Description
0	Network-BMC Pass-through traffic	0b = Disallowed 1b = Allowed (default)
1	Host-BMC Pass-through traffic	0b = Disallowed (default) 1b = Allowed
2	embedded CPU -BMC Pass-through traffic	0b = Disallowed (default) 1b = Allowed
7..3	Reserved	0b

3705

3706 8.6 Set Pass-through Mode Control Response (0xCF)

3707 In the absence of any errors, the channel shall process and respond to the Set Pass-through Mode
 3708 Control command and send the response packet shown in Table 8859.

3709

Table 59 – Set Pass-through Mode Control Response Packet

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
28..31	Checksum			
32..45	Pad			

3710

3711 8.7 Get Pass-through Mode Command (0x50)

3712 The Get Pass-through Mode command allows the Management controller to query the Ethernet Controller
 3713 for the current state of the Pass-through data paths are supported by the channel.

Table 61 illustrates the packet format for the Get Pass-through Mode Control command.

Table 61 – Get Pass-through Mode Command Packet

Bits				
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Reserved	Reserved	Pass-through Mode Status	Pass-through Mode Capability
20..23	Checksum			
24..45	Pad			

8.8 Get Pass-through Mode Control Response (0xD0)

In the absence of any errors, the channel shall process and respond to the Get Pass-through Mode Control command and send the response packet shown in Table 8863.

Table 63 – Get Pass-through Mode Response Packet

Bits				
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Reserved	Reserved	Pass-through Mode Status	Pass-through Mode Capability
24..27	Checksum			
28..45	Pad			

8.8.1 Pass-through Mode Status Field

The Pass-through Mode Status field indicates which Pass-through data path(s) are currently allowed.

Table 58 – Pass-through Type definitions

Bit	Field Description	Value Description
0	Network-BMC Pass-through traffic	0b = Currently Disallowed 1b = Currently Allowed (default)
1	Host-BMC Pass-through traffic	0b = Currently Disallowed (default) 1b = Currently Allowed

Bit	Field Description	Value Description
2	embedded CPU -BMC Pass-through traffic	0b = Currently Disallowed (default) 1b = Currently Allowed
7..3	Reserved	0b

3727

3728 **8.8.2 Pass-through Mode Capability Field**

3729 The Pass-through Mode Capability field indicates which Pass-through Mode data path(s) are supported
3730 by the implementation.

3731

Table 58 – Pass-through Type definitions

Bit	Field Description	Value Description
0	Network-BMC Pass- through traffic	0b = Not Supported 1b = Supported
1	Host-BMC Pass- through traffic	0b = Not Supported 1b = Supported
2	embedded CPU -BMC Pass-through traffic	0b = Not Supported 1b = Supported
7..3	Reserved	0b

3732

3733 **8.8.3 Transmit Data to NC command (0x4C)**

3734 The Transmit Data to NC command is a package command that allows the MC to transfer an opaque
3735 block of data of up to 16 MB to the NC. The transfer can be initiated by the MC itself or in response to the
3736 reception of the Transfer Data AEN. In the latter case, the Total Length of Transfer and Data Handle
3737 fields (if provided) should be populated from the AEN fields. If the requested Data Handle is not
3738 supported, then the Abort opcode shall be used. Blocks of data that exceed the data space available in
3739 one NC-SI frame will be broken down into multiple transfers that comply with NC-SI RBT frame size.
3740 When multiple transfers are used:

- 3741 • Transmission ordering shall be maintained
- 3742 • All chunks shall be an integer multiple of 32 bits, (i.e., double-word aligned), except for the last
- 3743 which may include padding to make it double-word aligned
- 3744 • If the NC detects a transfer error it may request a retransmission of the active chunk, but no other
- 3745 • Any processing of the block of data will only after the successful reception of all transmitted
- 3746 chunks

3747 The MC and the NC both have the ability to abort the transfer at any time during the transfer by use of the
3748 proper opcode or reason code respectively. If the NC loses transfer context due to being reset or other
3749 event, or if it detects an out of order chunk number being specified in the command, it shall abort the
3750 transfer. Any data transfer that is aborted is deemed to have failed and cannot be resumed. The MC
3751 may attempt to repeat the transfer as a new transfer sequence.

3752 Only one active transfer sequence (transmit or receive) is supported at a given time.

3753 Table 208 illustrates the packet format of the Transmit Data to NC command.

3754 **Table 208 – Transmit Data to NC command packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Total Length of Transfer (Bytes)			Opcode
20..23	Offset		Chunk Length	
24..27	Data Handle/Chunk Number			
	Chunk or Part of Data			
	Checksum			
	Pad			

3755

3756 8.8.3.1 Total Length of Transfer field

3757 Length in bytes of the entire data block to be transferred.

3758 8.8.3.2 Opcode field

3759 **Table 209 – Opcode field format**

Value	Description	Value Description
0	Initial Chunk	First block of data in the transfer
1	Final Chunk	Last block of data in the transfer
2	Middle Chunk	Intermediate block of data in the transfer
3	Abort Transfer	Terminate the transfer
others	Reserved	

3760 8.8.3.3 Offset

3761 Offset of the current transfer within the larger data block.

3762 8.8.3.4 Chunk Length

3763 The length in bytes of the chunk being transferred with this command.

3764 8.8.3.5 Data Handle/Chunk number

3765 For the first chunk being transferred (Initial Chunk Opcode), this is an identifier of the block of data being transferred.

3766 For subsequent chunk transfers it is a sequentially incrementing count for the chunk being transferred (equal to 2 for the second chunk transfer, 3 for the third, etc.).

3768

3769 8.8.4 Transmit Data to NC response (0xCC)

3770 The package shall, in the absence of a checksum error or identifier mismatch, always accept the Transmit
3771 Data to NC command and send a response.

3772 Table 210 illustrates the packet format of the Transmit Data to NC command response.

3773 There are command-specific reason codes identified for this response (see Table 211 – Transmit Data to
3774 NC command-specific reason codes).

3775 **Table 210 – Transmit Data to NC response packet format**

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

3776

3777 **Table 211 – Transmit Data to NC command-specific reason codes**

Value	Description	Comment
0x4C01	Abort Transfer	Returned when the NC is terminating the transfer for unspecified reason
0x4C02	Unknown Data Handle	Specified Data Handle is not supported
0x4C03	Sequence count error	Chunk Number received is not consecutive with the previous number received. Also results in an aborted transfer.
0x4C04	Length error	Incorrect chunk length
0x4C05	Insufficient Storage	NC cannot process or store a data block of Total Length
0x4C06	Invalid Handle Value	Data Handle is invalid or not supported

3778

3779 8.8.5 Receive Data from NC command (0x4D)

3780 The Receive Data from NC command is a package command that allows the MC to receive an opaque
3781 block of data of up to 16 MB from the NC. Blocks of data that exceed the data space available in one NC-
3782 SI frame will be broken down into multiple transfers that comply with NC-SI RBT frame size. When
3783 multiple transfers are used:

- 3784 • Reception ordering shall be maintained
- 3785 • All chunks shall be an integer multiple of 32 bits, (i.e., double-word aligned), except for the last
- 3786 which may include padding to make it double-word aligned
- 3787 • If the MC detects a transfer error it may request a retransmission of the active chunk, but no other
- 3788 • Any processing of the block of data will only after the successful reception of all transmitted
- 3789 chunks

The MC and the NC both have the ability to abort the transfer at any time during the transfer by use of the proper opcode or reason code respectively. If the NC loses transfer context due to being reset or other event, or if it detects an out of order chunk number being specified in the command, it shall abort the transfer. Any data transfer that is aborted is deemed to have failed and cannot be resumed. The MC may attempt to repeat the transfer as a new transfer sequence.

3795

3796 Only one active transfer sequence (transmit or receive) is supported at a given time.

3797 Table 212 illustrates the packet format of the Receive Data from NC command.

3798 **Table 212 – Receive Data from NC command packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
	Reserved			Opcode
	Offset		Reserved	
	Data Handle/Chunk Number			
16..19	Checksum			
20..45	Pad			

3799

3800 **8.8.5.1** Total Length of Transfer field

3801 Length in bytes of the entire data block to be transferred.

3802 **8.8.5.2** Opcode field

3803 **Table 213 – Opcode field format**

Value	Description	Value Description
0	Initial Chunk	Request for the first chunk of the transfer to be returned
1	Reserved	
2	Next Chunk	Request for the next chunk of the transfer to be returned
3	Abort Transfer	Termination of transfer by MC
others	Reserved	

3804 **8.8.5.3** Offset field

3805 Offset of the current transfer within the larger data block

3806 **8.8.5.4** Chunk Length field

3807 The length in bytes of the chunk being requested by this command.

8.8.5.5 Data Handle/Chunk number field

For the first chunk being requested (Initial Chunk Opcode), this is an identifier of the block of data being requested .
 For subsequent chunk transfers it is a sequentially incrementing count for the chunk being transferred (equal to 2 for the second chunk transfer, 3 for the third, etc.)

8.8.6 Receive Data from NC response (0xCD)

The package shall, in the absence of a checksum error or identifier mismatch, always accept the Receive Data from NC command and send a response.

Table 214 illustrates the packet format of the Receive Data from NC command response.

Table 214 – Receive Data from NC response packet format

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	Total Length of Transfer (Bytes)			Opcode
24..27	Offset		Chunk Length	
	Data Handle/Chunk Number			
	Data			
	Checksum			
	Pad			

8.8.6.1 Total Length of Transfer field

Length in bytes of the entire data block to be transferred

8.8.6.2 Opcode field**Table 215 – Opcode field format**

Value	Description	Value Description
0	Initial Chunk	First block of data in the transfer
1	Final Chunk	Last block of data in the transfer
2	Middle Chunk	Intermediate block of data in the transfer
3	Abort Transfer	Terminate the transfer
others	Reserved	

3824 **8.8.6.3** Offset field

3825 Offset of the current transfer within the larger data block

3826

3827 **8.8.6.4** Chunk Length field

3828 The length in bytes of the chunk being requested by this command.

3829 **Table 216 – Receive Data from NC command-specific reason codes**

Value	Description	Comment
0x4D01	Abort Transfer	NC cannot proceed with transfer
0x4D02	Sequence count error	Chunk Number requested is not consecutive with the previous number transmitted
0x4D03	Final Chunk of Transfer	Sent with Response Code 0000 to indicate the last chunk of the transfer
0x4C06	Invalid Handle Value	Data Handle is invalid or not supported

3830

3831 **8.8.7** Transfer SPDM command (0x60)

3832 The Transfer SPDM command is used by the Management controller in RBT implementations to
 3833 encapsulate and send a SPDM payload as defined in DSP0274 to the NC or alternately receive an
 3834 encapsulated SPDM payload from the NC.

3835 The SPDM payload must be smaller than the maximum NC-SI payload allowed over RBT. Payloads that
 3836 exceed the RBT limits shall use SPDM's native multi-part transfer mechanism. Polling mode shall be
 3837 used to transfer each part of a multi-part transfer from the NC.

3838 The command response may be a long running command due to the nature of some SPDM tasks..

3839 The Transfer SPDM command is defined as a package command.

3840 This command and response is not supported on NC-SI over MCTP.

3841 Table 217 illustrates the packet format of Transfer SPDM command.

3842 **Table 217 – Transfer SPDM command packet**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	SPDM Version	Request Code	Param 1	Param 2
20..	SPDM Message Payload			
	Checksum			
	Pad			

3843

8.8.8 Transfer SPDM Response (0xE0)

The Package shall, in the absence of a checksum error or identifier mismatch, always accept the Transfer SPDM Command and send a response.

Table 218 illustrates the packet format of the Transfer SPDM Response.

Table 218 – Transfer SPDM Response packet

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	SPDM Version	Completion Code	Param 1	Param 2
24..	SPDM Response Payload			
	Checksum			
	Pad			

3849

8.8.9 Query Pending NC SPDM Request (0x61)

The Query Pending NC SPDM Request may be used by the Management Controller in RBT implementations to read the status of pending SPDM commands which the NC needs to send to the MC. Only one SPDM request can be handled by a Pending SPDM Request instance. When multiple requests are pending in the NC, each will be handled independently and the order at which requests are provided to the MC is decided by the NC.

The Query Pending NC SPDM command is defined as a package command.

This command and response is not supported on NC-SI over MCTP.

Table 231 illustrates the packet format of the Query Pending NC SPDM Request command.

Table 219 – Query Pending NC SPDM Request packet format

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

3860

8.8.10 Query Pending NC SPDM Request Response (0xE1)

In the event there are no pending requests, the command shall execute successfully and return with no SPDM payload. Currently no command-specific reason code is identified for this response (see Table 232).

Table 232 illustrates the packet format of the Query Pending NC SPDM Request Response.

Table 220 – Query Pending NC SPDM Request Response Packet Format

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..	SPDM Version	Request Code	Param 1	Param 2
	SPDM Message Payload + Payload Pad (zero or more bytes)			
	Checksum			
	Pad			

3867

3868

Table 221 – Query Pending NC SPDM Request Response parameters

Name	Meaning
SPDM Version	Optional, included only when there is a pending request
Request Code	Optional, included only when there is a pending request
Param1	Optional, included only when there is a pending request
Param2	Optional, included only when there is a pending request
SPDM Message Payload	Optional, included only when there is a pending request

3869

8.8.11 Send NC SPDM Reply (0x62)

3871 The Reply Pending SPDM command may be used by the Management Controller to provide the SPDM
 3872 command response to previously read SPDM command from the NC. The response to this command
 3873 further provides indication to the MC regarding additional pending SPDM NC commands.

3874 Table 234 illustrates the packet format of the Send NC SPDM Reply command.

3875

Table 222 – Send NC SPDM Reply packet format

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	SPDM Version	Completion Code	Param 1	Param 2
20..	SPDM Message Payload (zero or more bytes) + Payload Pad			
	Checksum			
	Pad			

3876

8.8.12 Send NC SPDM Reply Response (0xE2)

3878 Currently no command-specific reason code is identified for this response.

3879 Table 235 illustrates the packet format of the Send NC SPDM Reply command.

3880 **Table 223 –Send NC SPDM Reply Response packet format**

Bits				
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Reserved			Flags
24..27	Checksum			
28..45	Pad			

3881

3882 **Table 224 – Reply NC SPDM Response parameters**

Name	Meaning
Flags bit 0 – Pending request	0b – No additional pending SPDM command from NC to MC 1b – The NC has additional pending SPDM command to the MC
Flags bits 7:1 - Reserved	Reserved, always return 0.

3883

3884

3885 **8.8.13** Query and Set OEM AEN command (0x4E)

3886 The channel command Query and Set OEM AEN is used by the Management controller when sets of
 3887 different OEM AENs, identified by the OEM's IANA value, are simultaneously supported by a NC. It
 3888 allows the MC to query the channel for the active OEM AEN set as well as the other OEM AEN sets that
 3889 are supported. The MC can then configure a particular IANA as the active one for subsequent issues of
 3890 the Enable AEN command.

- 3891 • Implementation of this command is optional for those NCs that support only one set of OEM
 3892 AENs
- 3893 • Implementation of this command is required when the NC has implemented multiple sets of OEM
 3894 AENs and allows the MC to select a set that is different than the default
- 3895 • The NC may allow AENs from multiple sets to be simultaneously enabled through the successive
 3896 uses of this command and AEN Enable
- 3897 • The NC shall interpret a null IANA in the received command as a request for the list of OEM AEN
 3898 sets and shall not change the active set.

3899 The Query and Set OEM AEN command is defined as a channel command.

3900 Table 225 illustrates the packet format of Query and Set OEM AEN command.

3901 **Table 225 – Query and Set OEM AEN command packet**

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

3902

3903 **8.8.14 Query and Set OEM AEN Response (0xCE)**

3904 The Channel shall, in the absence of a checksum error or identifier mismatch, always accept the Query
3905 and Set OEM AEN Command and send a response.

3906 For each supported OEM IANA, #1 through #n, three fields are required: the identifying IANA field, and
3907 the 16-bit Enabled AENs and Supported AENs fields that correspond 1:1 to bits 31..16 in the AEN Control
3908 Field of the AEN Enable command

3909

3910 Table 226 illustrates the packet format of the Query and Set OEM AEN Response.

3911 **Table 226 – Query and Set OEM AEN Response packet**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Reserved	Reserved	Reserved	# of IANAs
24..27	Configured IANA			
28..31	IANA # 1			
32..35	IANA # 1 Enabled AENs		IANA # 1 Supported AENs	
	IANA # 2			
	...			
	Checksum			
	Pad			

3912

3913 **8.8.14.1 # of IANAs field**

3914 An integer value representing the number of OEM AEN sets supported by the NC.

3915 **8.8.14.2** Configured IANA field

3916 The IANA representing the currently enabled OEM AEN set for configuration by subsequent Enable OEM
 3917 AEN commands. If a valid IANA was sent in the command, the response shall confirm the change to that
 3918 IANA set. If the sent IANA was not valid, the previously configured IANA set shall remain active.

3919 **8.8.14.3** IANA #n field

3920 The identifier for the nth OEM AEN set supported by the NC.

3921 **8.8.14.4** IANA #n Enabled AENs field

3922 A bitmap showing the currently enabled AENs from the IANA #n's set of supported AENs.

3923 **8.8.14.5** IANA #n Supported AENs field

3924 A bitmap showing the supported OEM AENs in the IANA #n's AEN set.

3925

3926 **8.8.15** OEM command (0x50)

3927 The OEM command may be used by the Management Controller to request that the channel provide
 3928 vendor-specific information. The [Vendor Enterprise Number](#) is the unique MIB/SNMP Private Enterprise
 3929 number assigned by IANA per organization. Vendors are free to define their own internal data structures
 3930 in the vendor data fields.

3931 Table 227 illustrates the packet format of the OEM command.

3932

Table 227 – OEM command packet format

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Manufacturer ID (IANA)			
20...	Vendor-Data			
	NOTE: The optional checksum is unspecified for the OEM command. OEMs supporting checksum validation for NC-SI commands may include the checksum in the OEM specific payload for the command and response.			

3933 **8.8.16 OEM response (0xD0)**

3934 The channel shall return the “Unknown Command Type” reason code for any unrecognized enterprise
 3935 number, using the packet format shown in Table 228. If the command is valid, the response, if any, is
 3936 allowed to be vendor specific. The 0x8000 range is recommended for vendor-specific code.

3937 Table 228 illustrates the packet format of the OEM command response.

3938

3939 **Table 228 – OEM response packet format**

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	Manufacturer ID (IANA)			
24...	Return Data (Optional)			

3940 **8.8.17 PLDM Request (0x51)**

3941 The PLDM Request Packet may be used by the Management Controller to send PLDM commands over
 3942 NC-SI/RBT. This command may be targeted at the entire package or a specific channel. It is expected
 3943 that the MC will use PLDM Request command 0x51 to query the supported PLDM commands, before
 3944 using Query Pending NC PLDM Request command.

3945 Table 229 illustrates the packet format of the PLDM Request Packet over NC-SI/RBT.

3946 **Table 229 – PLDM Request packet format**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	PLDM Message Common Fields			
20..	PLDM Message Payload (zero or more bytes) + Payload Pad)			
..	Checksum			
..	Pad			

3947 Refer to the PLDM Base specification (DSP0240) for details on the PLDM messaging control and
 3948 discovery commands.

3949 8.8.18 PLDM Response (0xD1)

3950 The PLDM Response Packet may be used by the Network Controller to send PLDM responses over NC-
3951 SI/RBT. The package shall, in the absence of a checksum error or identifier mismatch, always accept the
3952 PLDM Request Command and send a response.

3953 Table 230 illustrates the packet format of the PLDM command response.

3954 **Table 230 – PLDM Response packet format**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..23	PLDM Message Common Fields			PLDM Completion Code
24..	PLDM Message Payload (zero or more bytes) + Payload Pad			
..	Checksum			
..	Ethernet Packet Pad			

3955 Refer to the PLDM Base specification (DSP0240) for details on the PLDM Response Messages.

3956 Note that the NC-SI PLDM Response (0xD1) response/reason codes are only used to report the support,
3957 success, or failure of the PLDM Request command (0x51) at the NC-SI over RBT messaging layer. The
3958 PLDM Completion Code is used for determining the success or failure of the encapsulated PLDM
3959 Commands at the PLDM messaging layer.

3960

3961 8.8.19 Query Pending NC PLDM Request (0x56)

3962 The Query Pending NC PLDM Request may be used by the Management Controller to read the status of
3963 pending PLDM commands which the NC needs to send to the MC. Only one PLDM request can be
3964 handled by a Pending PLDM Request instance. When multiple requests are pending in the NC, each will
3965 be handled independently and the order at which requests are provided to the MC is decided by the NC.

3966 Implementations using PLDM over RBT, where the NC has to send PLDM commands to the MC, shall
3967 support this command.

3968 Table 231 illustrates the packet format of the Query Pending NC PLDM Request command.

3969 **Table 231 – Query Pending NC PLDM Request packet format**

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

3970

8.8.20 Query Pending NC PLDM Request Response (0xD6)

In the event there are no pending requests, the command shall execute successfully and return with no PLDM payload. Currently no command-specific reason code is identified for this response (see Table 232).

Table 232 illustrates the packet format of the Query Pending NC PLDM Request Response.

Table 232 – Query Pending NC PLDM Request 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..	PLDM Message Common Fields			PLDM Message Payload
	PLDM Message Payload + Payload Pad (zero or more bytes)			
	Checksum			
	Pad			

Table 233 – Query Pending NC PLDM Request Response parameters

Name	Meaning
PLDM Message Common fields	Optional, included only when there is a pending request
PLDM Message Payload	Optional, included only when there is a pending request

8.8.21 Send NC PLDM Reply (0x57)

The Reply Pending PLDM command may be used by the Management Controller to provide the PLDM command response to previously read PLDM command from the NC that requires a response (Rq = 1, D = 0 in PLDM Message Common Fields). The response to this command further provides indication to the MC regarding additional pending PLDM NC commands.

Table 234 illustrates the packet format of the Send NC PLDM Reply command.

Table 234 – Send NC PLDM Reply packet format

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	PLDM Message Common Fields			PLDM Completion Code
20..	PLDM Message Payload (zero or more bytes) + Payload Pad			

	Checksum
	Pad

3987

3988 **8.8.22 Send NC PLDM Reply Response (0xD7)**

3989 Currently no command-specific reason code is identified for this response.

3990 Table 235 illustrates the packet format of the Send NC PLDM Reply command.

3991

Table 235 –Send NC PLDM Reply 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			Flags
24..27	Checksum			
28..45	Pad			

3992

3993

Table 236 – Reply NC PLDM Response parameters

Name	Meaning
Flags bit 0 – Pending request	0b – No additional pending PLDM command from NC to MC 1b – The NC has additional pending PLDM command to the MC
Flags bits 7:1 - Reserved	Reserved, always return 0.

3994

3995 **8.8.23 Transport-specific AEN Enable command (0x55)**

3996 Network Controller implementations shall support this command on the condition that the Network
 3997 Controller generates one or more RBT-specific AENs defined in this specification or other NC-SI bindings
 3998 such as DSP0261. The AEN Enable command enables and disables the different transport specific AENs
 3999 supported by the Network Controller. The Network Controller shall copy the AEN MC ID field from the
 4000 AEN Enable command into the MC ID field in every subsequent AEN sent to the Management Controller
 4001 as defined in AEN Enable command

4002 Table 237 illustrates the packet format of the Enable Transport-specific AENs command.

Table 237 – Transport-specific AEN Enable command packet format

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

Table 238 – Transport-specific AEN 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 Relevant only for NC-SI/MCTP
1	Pending PLDM Request AEN (0x71)	0b = Disable Pending PLDM Request AEN 1b = Enable Pending PLDM Request AEN Relevant only for PLDM over NC-SI control over RBT
2	Pending SPDM Request AEN (0x72)	0b = Disable Pending SPDM Request AEN 1b = Enable Pending SPDM Request AEN Relevant only for SPDM over NC-SI control over RBT
3..15	Reserved	Reserved

8.8.24 Transport-specific AENs Enable Response (0xD5)

In the absence of any error, the package shall process and respond to the Transport-specific AEN Enable command by sending the response packet and payload shown in Table 239.

Table 239 – Transport-specific AEN Enable Response packet format

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

8.8.25 Get MC MAC Address command (0x58)

A network controller may provision MAC addresses for Out-Of-Band (OOB) management traffic. These MAC addresses are not visible to the host(s). Get MC MAC Address is used to discover MAC addresses provisioned on the network controller for the MC. Get MC MAC Address is a channel-specific command. For multiport devices, it is expected that the MC queries provisioned MC MAC Addresses on each channel individually.

Table 240 illustrates the packet format of the Get MC Address Command.

Table 240 – Get MC MAC Address command packet format

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

8.8.26 Get MC MAC Address response (0xD8)

In the response of Get MC MAC Address command, the network controller provides the information about the provisioned MAC address(es) for the MC on that channel. The NC shall, in the absence of an error, always accept the Get MC MAC Address command and send the response packet shown in Table 241. Currently no command-specific reason code is identified for this response.

4025

Table 241 – Get MC MAC Address response packet format

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Response Code		Reason Code	
20..23	Address Count	Reserved		
Variable	Addr 1 Byte 5	Addr 1 Byte 4	Addr 1 Byte 3	Addr 1 Byte 2
	Addr 1 Byte 1	Addr 1 Byte 0	Addr 2 Byte 5	Addr 2 Byte 4
	...			
	...		Pad (if needed)	

4026 **8.8.26.1** Address Count

4027 This field shall be set to the number of MC MAC addresses provisioned on the channel.

4028 **8.8.26.2** Reserved

4029 This field shall be set to 0 by the network controller and shall be ignored by the management controller.

4030 **8.8.26.3** Addr i Byte j4031 This field shall be set to the value of j^{th} byte ($1 \leq j \leq 6$) of i^{th} provisioned MC MAC address.4032 **8.8.26.4** Pad

4033 If the number of MC MAC addresses is an odd number, then 2 bytes of the Pad field shall be present at
 4034 the end of the payload to align the payload on a 32-bit boundary. If present, each byte of the Pad field
 4035 shall be set to 0x00.

4036 If the number of MC MAC addresses is an even number, then 0 bytes of Pad shall be present.

4037 **8.8.27** Get Package UUID command (0x52)

4038 The Get Package UUID command may be used by the Management Controller to query Universally
 4039 Unique Identifier (UUID), also referred to as a globally unique ID (GUID), of the Network Controller over
 4040 NC-SI/RBT. This command is targeted at the package. This command can be used by the MC to
 4041 correlate endpoints used on different NC-SI transports (e.g. RBT, MCTP).

4042 Table 242 illustrates the packet format of the Get Package UUID Command over NC-SI/RBT.

4043 **Table 242 – Get Package UUID command packet format**

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

8.8.28 Get Package UUID response (0xD2)

The package shall, in the absence of an error, always accept the Get Package UUID command and send the response packet shown in Table 243. Currently no command-specific reason code is identified for this response.

Table 243 – Get Package UUID response packet format

	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Control Packet Header			
16..19	Response Code		Reason Code	
20..35	UUID bytes 1:16, respectively			
36..39	Checksum			
40..45	Pad			

The individual fields within the UUID are stored most-significant byte (MSB) first per the convention described in RFC4122. RFC4122 specifies four different versions of UUID formats and generation algorithms suitable for use for a UUID. These are version 1 (0001b) "time based", and three "name-based" versions: version 3 (0011b) "MD5 hash", version 4 (0100b) "Pseudo-random", and version 5 "SHA1 hash". The version 1 format is recommended, however versions 3, 4, or 5 formats are also allowed to be used. See Table 244 for the UUID format version 1.

Table 244 – UUID Format

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

8.9 AEN packet formats

This clause defines the formats for the different types of AEN packets. For a list of the AEN types, see Table 17.

8.9.1 Link Status Change AEN

The Link Status Change AEN indicates to the Management Controller any changes in the channel's external Ethernet interface link status.

This AEN should be sent if any change occurred in the link status (that is, the actual link mode was changed). The Link Status and OEM Link Status fields reproduce the bit definitions defined in the Get Link Status Response Packet (see Table 50).

Table 245 illustrates the packet format of the Link Status Change AEN.

Table 245 – Link Status Change AEN packet format

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	AEN Header			
16..19	Reserved			AEN Type = 0x00
20..23	Link Status			
24..27	OEM Link Status			
28..31	Checksum			

8.9.2 Configuration Required AEN

The Configuration Required AEN indicates to the Management Controller that the channel is transitioning into the Initial State. (This AEN is not sent if the channel enters the Initial State because of a Reset Channel command.)

NOTE This AEN may not be generated in some situations in which the channel goes into the Initial State. For example, some types of hardware resets may not accommodate generating the AEN.

Table 246 illustrates the packet format of the Configuration Required AEN.

Table 246 – Configuration Required AEN packet format

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	AEN Header			
16..19	Reserved			AEN Type = 0x01
20..23	Checksum			

8.9.3 Host Network Controller Driver Status Change AEN

This AEN indicates a change of the Host Network Controller Driver Status. Table 247 illustrates the packet format of the AEN.

Table 247 – Host Network Controller Driver Status Change AEN packet format

Bits				
Bytes	31..24	23..16	15..08	07..00
00..15	AEN Header			
16..19	Reserved			AEN Type = 0x02
20..23	Host Network Controller Driver Status			
24..27	Checksum			

The Host Network Controller Driver Status field has the format shown in Table 248.

Table 248 – Host Network Controller Driver Status format

Bit Position	Name	Description
0	Host Network Controller Driver Status	<p>0b = The Network Controller driver for the host external network interface associated with this channel is not operational (not running).</p> <p>1b = The Network Controller driver for the host external network interface associated with this channel is being reported as operational (running).</p>
1..31	Reserved	Reserved

8.9.4 Delayed Response Ready AEN

This AEN indicates the response to a delayed command is ready. **Error! Reference source not found.** illustrates the packet format of the AEN.

Note: This AEN does not deliver the delayed command response, it must be retrieved separately.

Table 249 - Delayed Response Ready AEN packet format

Table 2-16 Delayed Response Ready AEN packet format				
	Bits			
Bytes	31..24	23..16	15..08	07..00
00..15	AEN Header			
16..19	Reserved			AEN Type = 0x03
20..23	Original Command Type	Original Command IID	Padding	
24..27	Checksum			

The Original Command Type includes the Control Packet Type field of the completed command and the Original Command IID includes the IID field of the original command.

8.9.5 InfiniBand Link Status Change AEN

The InfiniBand Link Status Change AEN indicates to the Management Controller any changes in the channel's external InfiniBand interface link status.

This AEN should be sent if any change occurred in the link status (that is, the actual link mode was changed). The Link Status and OEM Link Status fields reproduce the bit definitions defined in the Get IB Link Status Response Packet (see Table 50).

Table 255 illustrates the packet format of the InfiniBand Link Status Change AEN.

Table 250 – InfiniBand Link Status Change AEN packet format

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	AEN Header			
16..19	Reserved			AEN Type = 0x04
20..23	Link Status			
24..27	OEM Link Status			
28..31	Checksum			

8.9.6 Fibre Channel Link Status Change AEN

The Fibre Channel Link Status Change AEN indicates to the Management Controller any changes in the channel's external Fibre Channel interface link status including when trunked.

This AEN should be sent if any change occurred in the link status (that is, the actual link mode was changed). The Link Status and OEM Link Status fields reproduce the bit definitions defined in the Get FC Link Status Response Packet (see Table 50).

Table 262 illustrates the packet format of the FC Link Status Change AEN.

Table 251 – Fibre Channel Link Status Change AEN packet format

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	AEN Header			
16..19	Reserved			AEN Type = 0x05
20..23	Link Status			
24..27	OEM Link Status			
28..31	Checksum			

8.9.7 Transceiver Event AEN

This indicates to the Management Controller that a change in presence status or a thermal threshold in the SFF-compliant Transceiver attached to the channel has occurred.

Since some SFF cages have multiple TX and RX lanes, it is possible that multiple NC-SI channels are handled by a single transceiver module or copper cable assembly. Only one instance of the Transceiver Event AEN sent to one of the channels involved is required to enable reporting for all channels. The NC shall send the Transceiver Event AEN on all affected channels if one or more alerts are triggered.

In the case of FC port trunking (bonding), the 1:1 relationship of NC-SI channel to transceiver is lost and multiple transceivers will handle the aggregated traffic. When operating in trunking mode, one enablement of the AEN will cover all transceivers that are members of the trunk. AENs will be generated individually for members in the trunk and use the SFF Cage number field to identify the transceiver generating the AEN.

Table 263 illustrates the packet format of the AEN.

Table 252 – Transceiver Event AEN packet format

Bits				
Bytes	31..24	23..16	15..08	07..00
00..15	AEN Header			
16..19	Reserved	Transceiver Presence	SFF Cage Number	AEN Type = 0x06
20..23	Transceiver Event List			
24..27	Reserved			
28..31	Checksum			

8.9.7.1 SFF Cage Number field

SFF cage numbers are assigned to SFF cages in the implementation based on the NC-SI channel they are associated with (when not trunked) offset by one. Thus, the SFF cage associated with NC-SI channel 0 is #1, channel 1 has cage 2, etc.

8.9.7.2 Transceiver Event List field

The Transceiver Event List field has the format shown in Table 264

Table 253 – Transceiver Event List format

Bit Position	Name	Description
0	Low Temp Warning	0b = no alert 1b = The Transceiver's low temperature warning threshold has been exceeded
1	High Temp Warning	0b = no alert 1b = The Transceiver's high temperature warning threshold has been exceeded

Bit Position	Name	Description
2	Low Temp Alarm	0b = no alert 1b = The Transceiver's low temperature alarm threshold has been exceeded
3	High Temp Alarm	0b = no alert 1b = The Transceiver's high temperature alarm threshold has been exceeded
4	Low Voltage Warning	0b = no alert 1b = The Transceiver's low voltage warning threshold has been exceeded
5	High Voltage Warning	0b = no alert 1b = The Transceiver's high voltage warning threshold has been exceeded
6	Low Voltage Alarm	0b = no alert 1b = The Transceiver's low voltage alarm threshold has been exceeded
7	High Voltage Alarm	0b = no alert 1b = The Transceiver's high voltage alarm threshold has been exceeded
15..8	8 x RX Power Levels	0b = no alert 1b = The Transceiver's RX Power alarm threshold has been exceeded. lsb is lane 1 thru msb is lane8
23..16	8 x TX Power Levels	0b = no alert 1b = The Transceiver's TX Power alarm threshold has been exceeded. lsb is lane 1 thru msb is lane8
31..24	8 x TX Bias Levels	0b = no alert 1b = The Transceiver's TX Bias Current alarm threshold has been exceeded. lsb is lane 1 thru msb is lane8

4129 8.9.7.3 Transceiver Presence field

4130 **Table 254 – Transceiver Presence format**

Bit Position	Name	Description
0	Transceiver Presence Change	0b = No change in presence detected 1b = The Transceiver was either removed or inserted. The insertion event reporting shall occur only after the Transceiver has completed its initialization stage
7..1	Reserved	

4131

8.9.8 Request Data Transfer AEN

This AEN indicates to the Management Controller that the NC is requesting the MC initiate a transfer of an opaque data package from the NC to the MC.

Table 255 illustrates the packet format of the AEN.

Table 255 – Request Data Transfer AEN packet format

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	AEN Header			
16..19	Reserved			AEN Type = 0x07
20..23	Total Length of Transfer (Bytes)			
	Data Handle			
24..27	Checksum			

4137

8.9.9 Partition Link Status Change AEN

The Partition Link Status Change AEN indicates to the Management Controller any change in the internal link status of any partition on the channel. This AEN is only valid when the NC supports partitioning and it is enabled.

This AEN should be sent if any change occurred in the internal link status of any enabled partition on the channel.

Table 245 illustrates the packet format of the Partition Link Status Change AEN.

Table 256 – Partition Link Status Change AEN packet format

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	AEN Header			
16..19	Reserved			AEN Type = 0x08
20..23	Reserved		Partition Map	Link Status
24..27	Checksum			

4146

Table 257 – Partition Map Field

Bit	Description
0	0b = Partition 1 on channel link state has not changed 1b = Partition 1 on channel link state has changed

1	0b = Partition 2 on channel link state has not changed 1b = Partition 2 on channel link state has changed
...	...
7	0b = Partition 8 on channel link state has not changed 1b = Partition 8 on channel link state has changed

4148 **Table 258 – Partition Link Status**

Bit	Description
0	0b = Partition 1 on channel link is down 1b = Partition 1 on channel link is up
1	0b = Partition 2 on channel link is down 1b = Partition 2 on channel link is up
...	...
7	0b = Partition 8 on channel link is down 1b = Partition 8 on channel link is up

4149

4150 **8.9.10 Thermal Shutdown Event AEN**

4151 The Thermal Shutdown Event AEN indicates to the Management Controller that NC device shutdown is
4152 imminent due to the defined thermal threshold being reached.

4153 Table 246 illustrates the packet format of the Thermal Shutdown Event AEN.

4154 **Table 259 – Thermal Shutdown Event AEN packet format**

Bytes	Bits			
	31..24	23..16	15..08	07..00
00..15	AEN Header			
16..19	Reserved			AEN Type = 0x09
20..23	Checksum			

4155

4156 **8.9.11 Pending PLDM Request AEN**

4157 The Pending PLDM Request AEN is an RBT-specific AEN used to alert the MC that there is a pending
4158 PLDM request for the MC in the NC. This AEN allows for the MC to poll for pending PLDM request on the
4159 NC at a lower rate.

4160 As a transport-specific AEN, this AEN is enabled using the transport-specific AEN enable command and
4161 is controlled by bit 1 in Transport Specific AEN's enable field.

4162 This AEN should be sent if there is a new pending PLDM command that is available in the NC designated
 4163 to the MC, which was not reported to the MC through **Send NC PLDM Reply Response (0xD7)**. A
 4164 Pending PLDM Request AEN should not be sent from the time the NC recognizes an incoming **Query**
 4165 **Pending NC PLDM Request (0x56)** until the NC sends **Send NC PLDM Reply Response (0xD7)** for the
 4166 PLDM request.

Table 260 – Pending PLDM Request AEN format

Bits				
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Reserved			AEN Type = 0x71
20..23	Checksum			
24..45	Pad			

4168

4169 **8.9.12 Pending SPDM Request AEN**

4170 The Pending SPDM Request AEN is an RBT-specific AEN used to alert the MC that there is a pending
 4171 SPDM command request for the MC in the NC.

4172 As a transport-specific AEN, this AEN is enabled using the transport-specific AEN enable command and
 4173 is controlled by bit 2 in Transport Specific AEN's enable field.

4174 This AEN should be sent if there is a new pending SPDM command that is generated in the NC
 4175 designated for the MC, which was not reported to the MC through **Send NC PLDM Reply Response**
 4176 **(0xD7)**. A Pending SPDM Request AEN should not be sent from the time the NC recognizes an incoming
 4177 **Query Pending NC PLDM Request (0x56)** until the NC sends **Send NC PLDM Reply Response (0xD7)**
 4178 for the SPDM request.

Table 261 – Pending SPDM Request AEN format

Bits				
Bytes	31..24	23..16	15..08	07..00
00..15	NC-SI Header			
16..19	Reserved			AEN Type = 0x72
20..23	Checksum			
24..45	Pad			

4180

9 Packet-based and opcode timing

Table 262 presents the timing specifications for a variety of packet-to-electrical-buffer interactions, inter-packet timings, and opcode processing requirements. The following timing parameters shall apply to NC-SI over RBT binding defined in this specification.

Table 262 – NC-SI packet-based and opcode timing parameters

Name	Symbol	Value	Description
Package Deselect to Hi-Z Interval	T1	200 μ s, max	Maximum time interval from when a Network Controller completes transmitting the response to a Deselect Package command to when the Network Controller outputs are in the high-impedance state Measured from the rising edge of the first clock that follows the last bit of the packet to when the output is in the high-impedance state as defined in clause 0
Package Output to Data	T2	2 clocks, min	Minimum time interval after powering up the output drivers before a Network Controller starts transmitting a packet through the NC-SI interface Measured from the rising edge of the first clock of the packet
Network Controller Power Up Ready Interval	T4	2 s, max	Time interval from when the NC-SI on a Network Controller is powered up to when the Network Controller is able to respond to commands over the NC-SI Measured from when V_{ref} becomes available
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
Asynchronous Reset Interval	T6	2 s, max	Interval during which a controller may not recognize or respond to commands or handle Pass-through traffic due to an Asynchronous Reset event. See clause 6.1.8 For a Management Controller, this means that a Network Controller could become unresponsive for up to T6 seconds if an Asynchronous Reset event occurs. This is not an error condition. The Management Controller retry behavior should be designed to accommodate this possibility.
Synchronous Reset Interval	T7	2 s, max	Interval during which a controller may not recognize or respond to commands or handle Pass-through traffic due to a Synchronous Reset event. See clause 6.1.8 Measured from the rising edge of the first clock following the last bit of the Reset Channel response packet
Token Timeout	T8	32,000 REF_CLK min	Number of REF_CLKs before timing out while waiting for a TOKEN to be received

Name	Symbol	Value	Description
Opcode Processing	T9	32 REF_CLK max	Number of REF_CLKs after receiving an opcode on ARB_IN to decode the opcode and generate the next opcode on ARB_OUT Measured from the falling edge of the last bit of the opcode received on ARB_IN to the rising edge of the next opcode on ARB_OUT
Opcode Bypass Delay	T10	32 REF_CLK max	Number of REF_CLK delays between a bit received on ARB_IN and the corresponding bit passed on to ARB_OUT while in Bypass Mode Measured from the falling edge of the last bit of the opcode received on ARB_IN to the rising edge of the next opcode on ARB_OUT
TOKEN to RXD	T11	T2 min, 32 REF_CLK max	Number of REF_CLKs after receiving TOKEN to when packet data is driven onto the RXD lines Measured from the falling edge of the last bit of the opcode received on ARB_IN to the rising edge of the first clock of the next packet on RXD
Max XOFF Renewal Interval	T12	50,331,648 REF_CLK max	Maximum time period (3 XOFF Frame timer cycles) during which a channel within a package is allowed to request and renew a single XOFF condition after requesting the initial XOFF
IPG to TOKEN Opcode Overlap	T13	6 REF_CLK max	Maximum number of REF_CLKs that the beginning of TOKEN transmission can precede the end of the Inter Packet Gap. For more information, see 7.3.8.
Delayed Execution Interval	T14	4 s, max	Maximum time interval from when a controller receives a command to when it delivers a response to that command, including all responses with "Delayed Response" code 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 "Delayed Response Ready" AEN if enabled or to the moment the NC is internally ready with a response for a polling command.
NOTE If hardware arbitration is in effect, the hardware arbitration output buffer enable/disable timing specifications take precedence.			

10 RBT Electrical specification

This clause provides background information about the NC-SI RBT specification, describes the RBT topology, and defines the electrical, timing, signal behavior, and power-up characteristics for the RBT physical interface.

10.1 Topologies

The electrical specification defines the RBT electrical characteristics for one management processor and one to four Network Controller packages in a bussed “multi-drop” arrangement. The actual number of devices that can be supported may differ based on the trace characteristics and routing used to interconnect devices in an implementation.

Figure 16 shows an example topology.

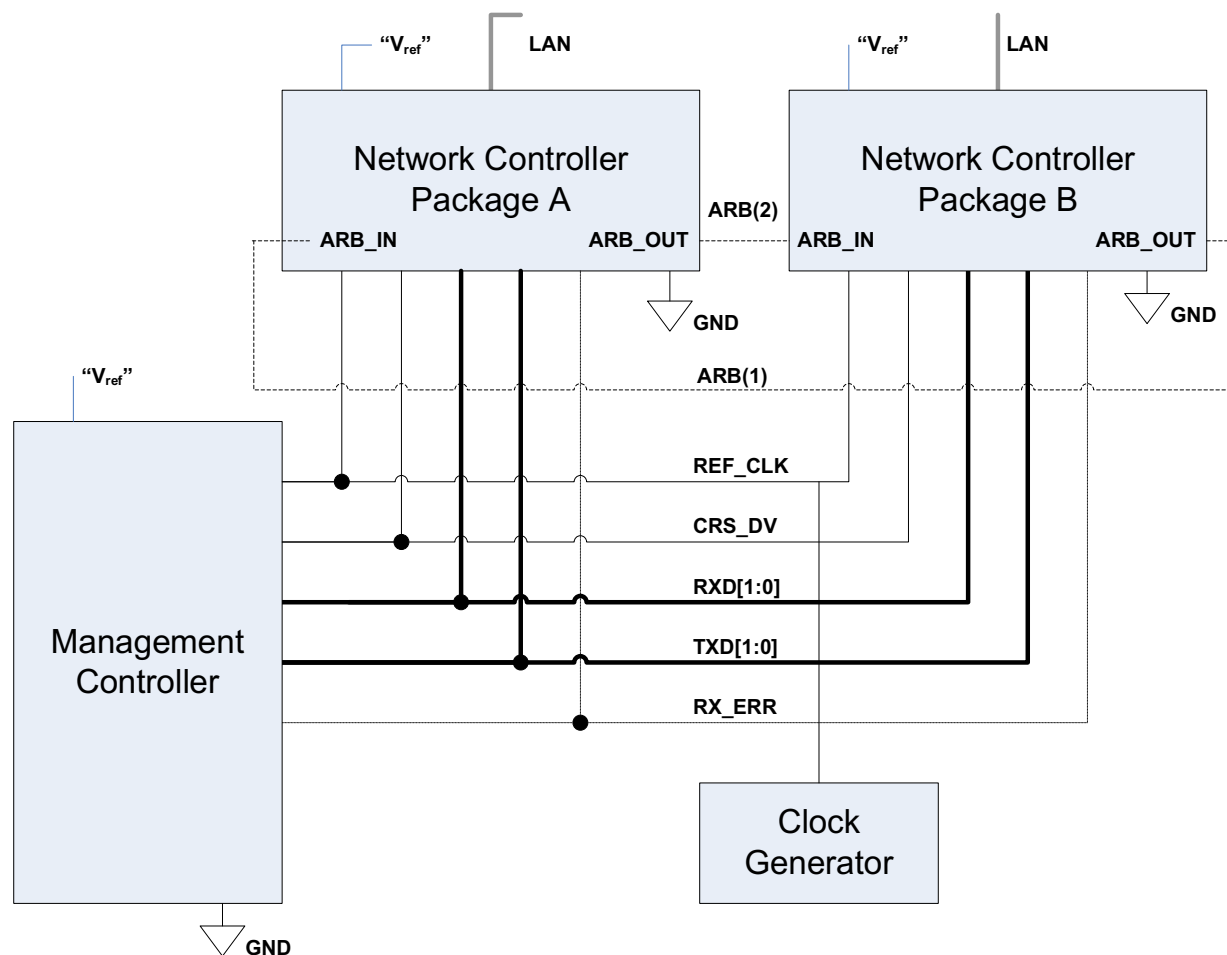


Figure 16 – Example NC-SI RBT signal interconnect topology

10.2 Electrical and signal characteristics and requirements

This clause defines the electrical, timing, signal behavior, and power-up characteristics for the NC-SI RBT physical interface.

10.2.1 Companion specifications

Implementations of the physical interface and signaling for RBT shall meet the specifications in [RMII](#) and [IEEE 802.3](#), except where those requirements differ or are extended with specifications provided in this document, in which case the specifications in this document shall take precedence.

10.2.2 Full-duplex operation

RBT is specified only for full-duplex operation. Half-duplex operation is not covered by this specification.

10.2.3 Signals

Table 263 lists the signals that make up the RBT physical interface.

Unless otherwise specified, the high level of a RBT signal corresponds to its asserted state, and the low level represents the de-asserted state. For data bits, the high level represents a binary '1' and the low level a binary '0'.

Table 263 – Physical RBT signals

Signal Name	Direction (with respect to the Network Controller)	Direction (with respect to the Management Controller MAC)	Use	Mandatory or Optional
REF_CLK ^[a]	Input	Input	Clock reference for receive, transmit, and control interface	M
CRS_DV ^[b]	Output	Input	Carrier Sense/Receive Data Valid	M
RXD[1:0]	Output	Input	Receive data	M
TX_EN	Input	Output	Transmit enable	M
TXD[1:0]	Input	Output	Transmit data	M
RX_ER	Output	Input	Receive error	O
ARB_IN	Input ^[c]	N/A	Network Controller hardware arbitration Input	O ^[c]
ARB_OUT	Output ^[c]	N/A	Network Controller hardware arbitration Output	O ^[c]
<p>^[a] A device can provide an additional option to allow it to be configured as the source of REF_CLK, in which case the device is not required to provide a separate REF_CLK input line, but it can use REF_CLK input pin as an output. The selected configuration shall be in effect at NC power up and remain in effect while the NC is powered up.</p> <p>^[b] In the RMII Specification, the MII Carrier Sense signal, CRS, was combined with RX_DV to form the CRS_DV signal. When RBT is using its specified full-duplex operation, the CRS aspect of the signal is not required; therefore, the signal shall provide only the functionality of RX_DV as defined in IEEE 802.3. (This is equivalent to the CRS_DV signal states in RMII Specification when a carrier is constantly present.) The Carrier Sense aspect of the CRS_DV signal is not typically applicable to RBT because it does not typically detect an actual carrier (unlike an actual PHY). However, the Network</p>				

Controller should emulate a carrier-present status on CRS_DV per [IEEE 802.3](#) in order to support Management Controller MACs that may require a carrier-present status for operation.

^[c] If hardware arbitration is implemented, the Network Controller package shall provide both ARB_IN and ARB_OUT connections. In some implementations, ARB_IN may be required to be tied to a logic high or low level if it is not used.

4214 **10.2.4 High-impedance control**

4215 Shared RBT operation requires Network Controller devices to be able to set their outputs (RXD[1:0],
4216 CRS_DV, and, if implemented, RX_ER) into a high-impedance state either upon receipt of a command
4217 being received, or, if hardware-based arbitration is enabled as a result of hardware-based arbitration. A
4218 pull-down resistor should be provided on high impedance signals to prevent them from floating and keep
4219 their C_{load} value when not driven.

4220 Network Controllers shall leave their RBT outputs in the high-impedance state on interface power up and
4221 shall not drive them until the package is selected. For additional information about Network Controller
4222 packages, see 8.4.5.

4223 For RBT output signals in this specification, unless otherwise specified, the high-impedance state is
4224 defined as the state in which the signal leakage meets the I_z specification provided in 10.2.5.

4225 **10.2.5 DC characteristics**

4226 This clause defines the DC characteristics of the RBT physical interface.

4227 **10.2.5.1 Signal levels**

4228 CMOS 3.3 V signal levels are used for this specification.

4229 The following characteristics apply to DC signals:

- 4230 • Unless otherwise specified, DC signal levels and V_{ref} are measured relative to Ground (GND) at
4231 the respective device providing the interface, as shown in Figure 17.
- 4232 • Input specifications refer to the signals that a device shall accept for its input signals, as
4233 measured at the device.
- 4234 • Output specifications refer to signal specifications that a device shall emit for its output signals,
4235 as measured at the device.

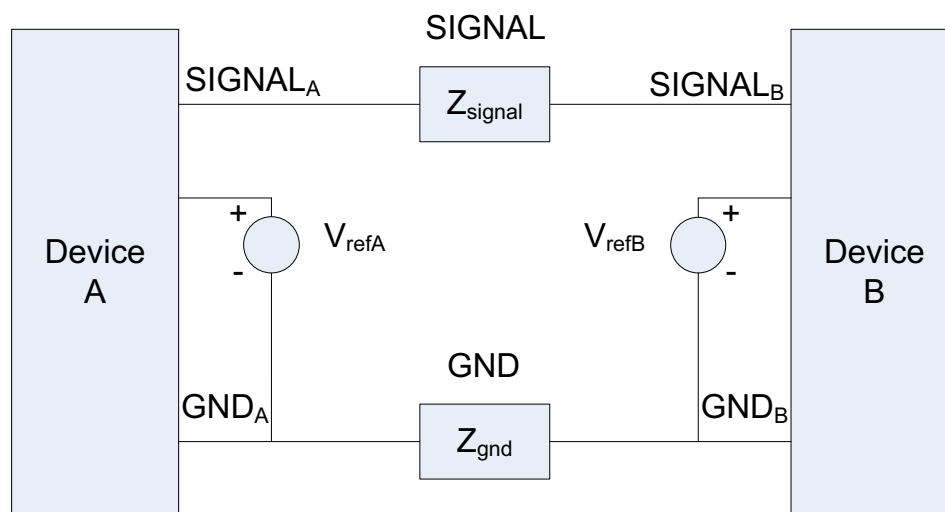


Figure 17 – DC measurements

Table 264 provides DC specifications.

Table 264 – DC specifications

Parameter	Symbol	Conditions	Minimum	Typical	Maximum	Units
IO reference voltage	$V_{\text{ref}}^{[a]}$		3.0	3.3	3.6	V
Signal voltage range	V_{abs}		-0.300		3.765	V
Input low voltage	V_{il}				0.8	V
Input high voltage	V_{ih}		2.0			V
Input high current	I_{ih}	$V_{\text{in}} = V_{\text{ref}} = V_{\text{ref,max}}$	0		200	μA
Input low current	I_{il}	$V_{\text{in}} = 0 \text{ V}$	-20		0	μA
Output low voltage	V_{ol}	$I_{\text{ol}} = 4 \text{ mA}$, $V_{\text{ref}} = \text{min}$	0		400	mV
Output high voltage	V_{oh}	$I_{\text{oh}} = -4 \text{ mA}$, $V_{\text{ref}} = \text{min}$	2.4		V_{ref}	V
Clock midpoint reference level	V_{ckm}				1.4	V
Leakage current for output signals in high-impedance state	I_{z}	$0 \leq V_{\text{in}} \leq V_{\text{ref}}$ at $V_{\text{ref}} = V_{\text{ref,max}}$	-20		20	μA

^[a] V_{ref} = Bus high reference level (typically the NC-SI logic supply voltage). This parameter replaces the term *supply voltage* because actual devices may have internal mechanisms that determine the operating reference for RBT that are different from the devices' overall power supply inputs.

V_{ref} is a reference point that is used for measuring parameters (such as overshoot and undershoot) and for determining limits on signal levels that are generated by a device. To facilitate system implementations, a device shall provide a mechanism (for example, a power supply pin, internal programmable reference, or reference level pin) to allow V_{ref} to be set to within 20 mV of any point in the specified V_{ref} range. This approach enables a system integrator to establish an interoperable V_{ref} level for devices on RBT.

10.2.6 AC characteristics

This clause defines the AC characteristics of the RBT physical interface.

10.2.6.1 Rise and fall time measurement

Rise and fall time are measured between points that cross 10% and 90% of V_{ref} (see Table 264). The middle points (50% of V_{ref}) are marked as V_{ckm} and V_m for clock and data, respectively.

10.2.6.2 REF_CLK measuring points

In Figure 18, REF_CLK duty cycle measurements are made from V_{ckm} to V_{ckm} . Clock skew T_{skew} is measured from V_{ckm} to V_{ckm} of two RBT devices and represents the maximum clock skew between any two devices in the system.

10.2.6.3 Data, control, and status signal measuring points

In Figure 18, all timing measurements are made between V_{ckm} and V_m . T_{co} is measured with a capacitive load between 10 pF and 50 pF. Propagation delay T_{prop} is measured from V_m on the transmitter to V_m on the receiver.




Figure 18 – AC measurements

Table 265 provides AC specifications.

Table 265 – AC specifications

Parameter	Symbol	Minimum	Typical	Maximum	Units
REF_CLK Frequency			50	50+100 ppm	MHz
REF_CLK Duty Cycle		35		65	%
Clock-to-out ^[a] (10 pF ≤ C_{load} ≤ 50 pF)	T_{co}	2.5		12.5	ns
Skew between clocks	T_{skew}			1.5	ns

TXD[1:0], TX_EN, RXD[1:0], CRS_DV, RX_ER, and ARB_IN data setup to REF_CLK rising edge	T_{su}	3			ns
TXD[1:0], TX_EN, RXD[1:0], CRS_DV, RX_ER, and ARB_OUT data hold from REF_CLK rising edge	T_{hd}	1			ns
Signal Rise/Fall Time	T_r/T_f	0.5		6	ns
REF_CLK Rise/Fall Time	T_{ckr}/T_{ckf}	0.5		3.5	ns
Interface Power-Up High-Impedance Interval	T_{pwrz}	2			μ s
Power Up Transient Interval (recommendation)	T_{pwrt}			100	ns
Power Up Transient Level (recommendation)	V_{pwrt}	-200		200	mV
REF_CLK Startup Interval	$T_{clkstrt}$			100	ms
[a] This timing relates to the output pins, while T_{su} and T_{hd} relate to timing at the input pins.					

4257 10.2.6.4 Timing calculation (informative)

4258 10.2.6.4.1 Setup time calculation

4259
$$T_{su} \leq T_{clk} - (T_{skew} + T_{co} + T_{prop})$$

4260 10.2.6.4.2 Hold time calculation

4261
$$T_{hd} \leq T_{co} - T_{skew} + T_{prop}$$

4262 10.2.6.5 Overshoot specification

4263 Devices shall accept signal overshoot within the ranges specified in Figure 19, measured at the device,
 4264 without malfunctioning.




Figure 19 – Overshoot measurement

The signal may overshoot up to the specified V_{max} for the first 3 ns following the transition above V_{ih} . Following that interval is an exponential decay envelope equal to the following:

$$V_{ref} + V_{os} * e^{[-K * (t - 3 \text{ ns}) / T_d]}$$

Where, for $t = 3$ to 10 ns:

$t = 0$ corresponds to the leading crossing of V_{ih} , going high.

V_{ref} is the bus high reference voltage (see 10.2.5).

$V_{abs,max}$ is the maximum allowed signal voltage level (see 10.2.5).

$V_{os} = V_{max} - V_{ref}$

$K = I_n(25 \text{ mV} / V_{os})$

$T_d = 7 \text{ ns}$

For $t > 10 \text{ ns}$, the $V_{ref} + 25 \text{ mV}$ limit holds flat until the conclusion of T_{high} .

10.2.6.6 Undershoot specification

Devices are required to accept signal undershoot within the ranges specified in Figure 20, measured at the device, without malfunctioning.




Figure 20 – Undershoot measurement

The signal is allowed to undershoot up to the specified $V_{abs,min}$ for the first 3 ns following the transition above V_{il} . Following that interval is an exponential envelope equal to the following:

$$* ([t - 3 \text{ ns}] / T_d)$$

Where, for $t = 3$ to 10 ns:

$t = 0$ corresponds to the leading crossing of V_{il} , going low.

$V_{abs,min}$ is the minimum allowed signal voltage level (see 10.2.5).

$$K = \ln(25 \text{ mV} / V_{os})$$

$$T_d = 7 \text{ ns}$$

For $t > 7$ ns, the GND – 25 mV limit holds flat until the conclusion of T_{low} .

10.2.7 Interface power-up

To prevent signals from back-powering unpowered devices, it is necessary to specify a time interval during which signals are not to be driven until devices sharing the interface have had time to power up. To facilitate system implementation, the start of this interval shall be synchronized by an external signal across devices.

4297 **10.2.7.1 Power-up control mechanisms**

4298 The device that provides the interface shall provide one or more of the following mechanisms to enable
 4299 the system integrator to synchronize interface power-up among devices on the interface:

4300

- **Device power supply pin**

4301 The device has a power supply pin that the system integrator can use to control power-up of the
 4302 interface. The device shall hold its outputs in a high-impedance state (current $< I_z$) for at least
 4303 T_{pwrz} seconds after the power supply has initially reached its operating level (where the power
 4304 supply operating level is specified by the device manufacturer).

4305

- **Device reset pin or another similar signal**

4306 The device has a reset pin or other signal that the system integrator can use to control the
 4307 power-up of the interface. This signal shall be able to be driven asserted during interface power-
 4308 up and de-asserted afterward. The device shall hold its outputs in a high-impedance state
 4309 (current $< I_z$) for at least T_{pwrz} seconds after the signal has been de-asserted, other than as
 4310 described in 10.2.7.2. It is highly recommended that a single signal be used; however, an
 4311 implementation is allowed to use a combination of signals if required. Logic levels for the signals
 4312 are as specified by the device manufacturer.

4313

- **REF_CLK detection**

4314 The device can elect to detect the presence of an active REF_CLK and use that for determining
 4315 whether NC-SI power up has occurred. It is recommended that the device should count at least
 4316 100 clocks and continue to hold its outputs in a high-impedance state (current $< I_z$) for at least
 4317 T_{pwrz} seconds more (Informational: 100 clocks at 50 MHz is 2 us).

4318 **10.2.7.2 Power-up transients**

4319 It is possible that a device may briefly drive its outputs while the interface or device is first receiving
 4320 power, due to ramping of the power supply and design of its I/O buffers. It is recommended that devices
 4321 be designed so that such transients, if present, are less than V_{pwrt} and last for no more than T_{pwrt} .

4322 **10.2.8 REF_CLK startup**

4323 REF_CLK shall start up, run, and meet all associated AC and DC specifications within $T_{clkstrt}$ seconds of
 4324 interface power up.

4325 **10.3 RBT Implementation guidance**

4326 This specification does not define implementation requirements due to the wide variation in architectures,
 4327 devices and materials used. Following good engineering practices are a key part of a successful NC-SI
 4328 RBT implementation:

- 4329
 - Care must be taken in placement and layout
- 4330
 - Do a complete signal integrity analysis including determining what, if any, termination is required
- 4331
 - Minimize stubs
- 4332
 - Have uniform clock trace lengths
- 4333
 - Minimize noise on high-impedance0 signals

4334

ANNEX A (normative)

Extending the model

This annex explains how the model can be extended to include vendor-specific content.

Commands extension

A Network Controller vendor can implement extensions and expose them using OEM commands, as described in 8.8.15.

Design considerations

This clause describes certain design considerations for vendors of Management Controllers.

PHY support

Although not a requirement of this specification, a Management Controller vendor can design the RBT interface in such a manner that it could also be configured for use with a conventional RMII PHY. This would enable the vendor's controller to also be used in applications where a direct, non-shared network connection is available or preferred for manageability.

Multiple Management Controllers support

Currently, there is no requirement for Management Controllers to be able to put their TXD output lines and other output lines into a high-impedance state, because the present definition assumes only one Management Controller on the bus. However, component vendors can provide such control capabilities in their devices to support possible future system topologies where more than one Management Controller shares the bus to enable functions such as Management Controller fail-over or to enable topologies where more than one Management Controller can participate in NC-SI communications on the bus. If a vendor elects to make such provision, it is recommended that the TXD line and the remaining output lines be independently and dynamically switched between a high-impedance state and re-enabled under firmware control.

ANNEX B (informative)

Relationship to RMII Specification

Differences with the *RMII Specification*

The following list presents key differences and clarifications between the *NC-SI Specification* and sections in the [RMII Specification](#). (Section numbers refer to the [RMII Specification](#).)

- General: Where specifications from [IEEE 802.3](#) apply, this specification uses the version specified in clause **Error! Reference source not found.**, rather than the earlier IEEE 802.3u version that is referenced by [RMII](#).
- Section 1.0:
 - The *NC-SI Specification* requires 100 Mbps support, but it does not specify a required minimum. (10 Mbps support is not required by NC-SI.)
 - Item 4. (Signals may or may not be considered to be TTL. NC-SI is not 5-V tolerant.)
- Section 2.0:
 - Comment: NC-SI chip-to-chip includes considerations for multi-drop and allows for non-PCB implementations and connectors (that is, not strictly point-to-point).
- Section 3.0:
 - Note/Advisory: The NC-SI clock is provided externally. An implementation can have REF_CLK provided by one of the devices on the bus or by a separate device.
- Section 5.0:
 - For NC-SI, the term *PHY* is replaced by *Network Controller*.
- Table 1:
 - The information in Table 1 in the [RMII Specification](#) is superseded by tables in this specification.
- Section 5.1, paragraph 2:
 - The *NC-SI Specification* allows 100 ppm. This supersedes the [RMII Specification](#), which allows 50 ppm.
- Section 5.1, paragraph 3:
 - The NC-SI inherits the same requirements. The NC-SI MTU is required only to support Ethernet MTU with VLAN, as defined in the [IEEE 802.3](#) version listed in clause **Error! Reference source not found.**
- Section 5.1 paragraph 4:
 - The [RMII Specification](#) states: "During a false carrier event, CRS_DV shall remain asserted for the duration of carrier activity." This statement is not applicable to full-duplex operation of the NC-SI. CRS_DV from the Network Controller is used only as a data valid (DV) signal. Because the Carrier Sense aspect of CRS_DV is not used for full-duplex operation of the NC-SI, the Network Controller would not generate false carrier events for the NC-SI. However, it is recommended that the MAC in the Management Controller be able to

- 4400 correctly detect and handle these patterns if they occur, as this would be part of enabling
4401 the Management Controller MAC to also be able to work with an RMII PHY.
- 4402 • Section 5.2:
 - 4403 – The NC-SI does not specify a 10 Mbps mode. The Carrier Sense aspect of CRS_DV is not
4404 used for full-duplex operation of NC-SI.
 - 4405 • Section 5.3.1:
 - 4406 – While the NC-SI does not specify Carrier Sense usage of CRS_DV, it is recommended that
4407 a Management Controller allow for CRS_DV toggling, in which CRS_DV toggles at 1/2
4408 clock frequency, and that Management Controller MACs tolerate this and realign bit
4409 boundaries correctly in order to be able to work with an RMII PHY also.
 - 4410 • Section 5.3.2:
 - 4411 – There is no 10 Mbps mode specified for the NC-SI RBT interface.
 - 4412 • Section 5.3.3:
 - 4413 – Generally, there is no expectation that the Network Controller will generate these error
4414 conditions for the NC-SI; however, the MAC in the Management Controller should be able
4415 to correctly detect and handle these patterns if they occur.
 - 4416 • Section 5.3.3:
 - 4417 – The NC-SI does not specify or require support for RMII Registers.
 - 4418 • Section 5.5.2:
 - 4419 – Ignore (N/A) text regarding 10 Mbps mode. RBT does not specify or require interface
4420 operation in 10 Mbps mode.
 - 4421 • Section 5.6:
 - 4422 – The Network Controller will not generate collision patterns for the specified full-duplex
4423 operation of the NC-SI; however, the MAC in the Management Controller should be able to
4424 detect and handle these patterns if they occur in order to be able to work with an RMII PHY
4425 also.
 - 4426 • Section 5.7:
 - 4427 – NC-SI RBT uses the [IEEE 802.3](#) version listed in clause 2 **Error! Reference source not**
4428 **found.** instead of 802.3u as a reference.
 - 4429 • Section 5.8:
 - 4430 – Loopback operation is not specified for the NC-SI RBT interface.
 - 4431 • Section 7.0:
 - 4432 – The NC-SI RBT electrical specifications (clause 0) take precedence. (For example, section
4433 7.4.1 in the [RMII Specification](#) for capacitance is superseded by *NC-SI Specification* 25 pF
4434 and 50 pF target specifications.)
 - 4435 • Section 8.0:
 - 4436 – NC-SI RBT uses the [IEEE 802.3](#) version listed in clause 2 **Error! Reference source not**
4437 **found.** as a reference, instead of 802.3u.

ANNEX C**(informative)****Change log**

Version	Date	Description
1.0.0	2009-07-21	
1.0.1	2013-01-24	DMTF Standard release
1.1.0	2015-09-23	DMTF Standard release
1.1.1	~2021-04-13	Updated to comply with ISO guidelines
1.2.0a	2019-08-19	DMTF Work in Progress release
1.2WIP90	2022-04-27	DMTF WIP - Addition of configuration and monitoring support for Ethernet, Fibre Channel and InfiniBand controllers, including partitioning of ports and multiple host bus interfaces

4443

Bibliography

- 4444 IANA, Internet Assigned Numbers Authority (www.iana.org). A body that manages and organizes
4445 numbers associated with various Internet protocols.
- 4446 DMTF [DSP4014](#), *DMTF Process for Working Bodies 2.2*, August 2015,
4447 https://www.dmtf.org/sites/default/files/standards/documents/DSP4014_2.12.0.pdf