



OPEN
Compute Project

OCP NIC 3.0 Design Specification

Version 0.99

Author: OCP Server Workgroup, OCP NIC subgroup

Table of Contents

1 Overview	10
1.1 License	10
1.2 Acknowledgements	11
1.3 References	12
1.3.1 Trademarks	13
1.4 Acronyms	14
1.5 Background	16
1.6 Overview	18
1.6.1 Mechanical Form Factor Overview	18
1.6.2 Electrical Overview	20
1.6.2.1 Primary Connector	20
1.6.2.2 Secondary Connector	21
1.7 Non-NIC Use Cases	21
2 Mechanical Card Form Factor	22
2.1 Form Factor Options	22
2.1.1 SFF Faceplate Configurations	24
2.1.2 LFF Faceplate Configurations	28
2.2 Line Side I/O Implementations	32
2.3 Top Level Assembly (SFF and LFF)	33
2.4 Faceplate Subassembly (SFF and LFF)	34
2.4.1 Faceplate Subassembly – Exploded View	34
2.4.2 Faceplate Subassembly – Bill of Materials (BOM)	34
2.4.3 SFF Generic I/O Faceplate	37
2.4.4 LFF Generic I/O Faceplate	38
2.4.5 Ejector Lever (SFF)	39
2.4.6 Ejector Levers (LFF)	40
2.4.7 Ejector Lock (SFF and LFF)	41
2.4.8 Clinch Nut (SFF and LFF)	42
2.5 Card Keep Out Zones	43
2.5.1 SFF Keep Out Zones	43
2.5.2 LFF Keep Out Zones	46
2.6 Baseboard Keep Out Zones	49
2.7 Insulation Requirements	50
2.7.1 SFF Insulator	50
2.7.2 LFF Insulator	52
2.8 Critical-to-Function (CTF) Dimensions (SFF and LFF)	55
2.8.1 CTF Tolerances	55
2.8.2 SFF Pull Tab CTF Dimensions	55
2.8.3 SFF Ejector Latch CTF Dimensions	57
2.8.4 SFF Internal Lock CTF Dimensions	58
2.8.5 SFF Baseboard CTF Dimensions	59
2.8.6 LFF Ejector Latch CTF Dimensions	62
2.8.7 LFF Baseboard CTF Dimensions	63
2.9 Labeling Requirements	66
2.9.1 General Guidelines for Label Contents	66
2.9.2 MAC Address Labeling Requirements	67
2.9.2.1 MAC Address Label Example 1 – Quad Port with Single Host, Single Managed Controller	68
2.9.2.2 MAC Address Label Example 2 – Octal Port with Single Host, Dual Managed Controllers	68
2.9.2.3 MAC Address Label Example 3 – Quad Port with Dual Hosts, Dual Managed Controllers	69
2.9.2.4 MAC Address Label Example 4 – Single Port with Quad Host, Single Managed Controller	69
2.10 Mechanical CAD Package Examples	71
3 Electrical Interface Definition – Card Edge and Baseboard	72
3.1 Card Edge Gold Finger Requirements	72
3.1.1 Gold Finger Mating Sequence	74
3.2 Baseboard Connector Requirements	78
3.2.1 Right Angle Connector	78
3.2.2 Right Angle Offset	79

3.2.3	Straddle Mount Connector	79
3.2.4	Straddle Mount Offset and PCB Thickness Options	81
3.2.5	LFF Connector Locations	82
3.3	Pin Definition	82
3.3.1	Primary Connector	83
3.3.2	Secondary Connector	85
3.4	Signal Descriptions	86
3.4.1	PCIe Interface Pins	86
3.4.2	PCIe Present and Bifurcation Control Pins	92
3.4.3	SMBus Interface Pins	95
3.4.4	NC-SI over RBT Interface Pins	96
3.4.5	Scan Chain Pins	104
3.4.6	Power Supply Pins	111
3.4.7	USB 2.0 (A68/A69) – Primary Connector Only	117
3.4.8	UART (A68/A69) – Secondary Connector Only	119
3.4.9	RFU[1:4] Pins	121
3.5	PCIe Bifurcation Mechanism	122
3.5.1	PCIe OCP NIC 3.0 Card to Baseboard Bifurcation Configuration (PRSNTA#, PRSNTB[3:0]#)	122
3.5.2	PCIe Baseboard to OCP NIC 3.0 Card Bifurcation Configuration (BIF[2:0]#)	122
3.5.3	PCIe Bifurcation Decoder	123
3.5.4	Bifurcation Detection Flow	125
3.5.5	PCIe Bifurcation Examples	126
3.5.5.1	Single Host (1 x16) Baseboard with a 1 x16 OCP NIC 3.0 Card (Single Controller)	126
3.5.5.2	Single Host (2 x8) Baseboard with a 2 x8 OCP NIC 3.0 Card (Dual Controllers)	127
3.5.5.3	Quad Host (4 x4) Baseboard with a 4 x4 OCP NIC 3.0 Card (Single Controller)	128
3.5.5.4	Quad Host (4 x4) Baseboard with a 4 x4 OCP NIC 3.0 Card (Quad Controllers)	129
3.5.5.5	Single Host (1 x16, no Bifurcation) Baseboard with a 2 x8 OCP NIC 3.0 Card (Dual Controller)	130
3.6	PCIe REFCLK and PERST# Mapping	131
3.6.1	SFF PCIe REFCLK and PERST# Mapping	132
3.6.2	LFF PCIe REFCLK and PERST# Mapping	135
3.6.3	REFCLK and PERST# Mapping Expansion	137
3.7	Port Numbering and LED Implementations	138
3.7.1	OCP NIC 3.0 Port Naming and Port Numbering	138
3.7.2	OCP NIC 3.0 Card LED Configuration	138
3.7.3	OCP NIC 3.0 Card LED Ordering	140
3.7.4	Baseboard LEDs Configuration over the Scan Chain	141
3.8	Power State Machine	142
3.8.1	NIC Power Off	143
3.8.2	ID Mode	143
3.8.3	Aux Power Mode (S5)	143
3.8.4	Main Power Mode (S0)	144
3.8.5	Programming Mode	144
3.9	Power Supply Rail Requirements and Slot Power Envelopes	145
3.10	Hot Swap Considerations for +12V_EDGE and +3.3V_EDGE Rails	146
3.11	Power Sequence Timing Requirements	147
3.12	Digital I/O Specifications	151
4	Management and Pre-OS Requirements	152
4.1	Sideband Management Interface and Transport	152
4.2	NC-SI Traffic	153
4.3	Management Controller (MC) MAC Address Provisioning	153
4.4	ASIC Die Temperature Reporting	155
4.5	Power Consumption Reporting	157
4.6	Pluggable Transceiver Module Status and Temperature Reporting	158
4.7	Management and Pre-OS Firmware Inventory and Update	158
4.7.1	Secure Firmware	159
4.7.2	Firmware Inventory	159
4.7.3	Firmware Inventory and Update in Multi-Host Environments	159
4.8	NC-SI Package Addressing and Hardware Arbitration Requirements	160

4.8.1	NC-SI over RBT Package Addressing	160
4.8.2	Arbitration Ring Connections	160
4.9	SMBus 2.0 Addressing Requirements	160
4.9.1	SMBus Address Map	161
4.10	FRU EEPROM	161
4.10.1	FRU EEPROM Addressing and Size	161
4.10.2	FRU EEPROM Write Protection	163
4.10.3	FRU EEPROM Content Requirements	163
4.10.4	FRU Template	169
5	Routing Guidelines and Signal Integrity Considerations	170
5.1	NC-SI over RBT	170
5.1.1	SFF Baseboard Requirements	171
5.1.2	LFF Baseboard Requirements	172
5.1.3	SFF OCP NIC 3.0 Card Requirements	172
5.1.4	LFF OCP NIC 3.0 Card Requirements	173
5.2	SMBus 2.0	173
5.3	PCIe	174
5.3.1	Channel Requirements	174
5.3.1.1	REFCLK requirements	174
5.3.1.2	Add-in Card Electrical Budgets	174
5.3.1.3	Baseboard Channel Budget	174
5.3.1.4	SFF-TA-1002 Connector Channel Budget	174
5.3.1.5	Differential Impedance (Informative)	174
5.3.2	Test Fixtures	175
5.3.2.1	Compliance Load Board (CLB)	175
5.3.2.2	Compliance Baseboard (CBB)	176
5.3.3	Test Methodology	176
5.3.3.1	Test Setup	176
6	Thermal and Environmental	177
6.1	Airflow Direction	177
6.1.1	Hot Aisle Cooling	177
6.1.2	Cold Aisle Cooling	178
6.2	Thermal Design Guidelines	179
6.2.1	SFF Card ASIC Cooling – Hot Aisle	179
6.2.2	LFF Card ASIC Cooling – Hot Aisle	183
6.2.3	SFF Card ASIC Cooling – Cold Aisle	185
6.2.4	LFF Card ASIC Cooling – Cold Aisle	188
6.3	Thermal Simulation (CFD) Modeling	190
6.4	Thermal Test Fixture	190
6.4.1	Test Fixture for SFF Card	191
6.4.2	Test Fixture for LFF Card	193
6.4.3	Test Fixture Airflow Direction	195
6.4.4	Thermal Test Fixture Candlestick Sensors	195
6.5	Card Sensor Requirements	197
6.6	Card Cooling Tiers	197
6.7	Non-Operational Shock & Vibration Testing	199
6.7.1	Shock & Vibe Test Fixture	199
6.7.2	Test Procedure	200
6.8	Dye and Pull Test Method	202
6.9	Gold Finger Plating Requirements	204
6.9.1	Host Side Gold Finger Plating Requirements	204
6.9.2	Line Side Gold Finger Durability Requirements	204
7	Regulatory	205
7.1	Required Compliance	205
7.1.1	Required Environmental Compliance	205
7.1.2	Required EMC Compliance	205
7.1.3	Required Product Safety Compliance	206
7.1.4	Required Immunity (ESD) Compliance	206

7.2	Recommended Compliance	207
7.2.1	Recommended Environmental Compliance	207
7.2.2	Recommended EMC Compliance	207
8	Revision History	208
8.1	Document Revision History	208
8.2	FRU Content Revision History	213

List of Figures

Figure 1:	Representative SFF OCP NIC 3.0 Card with Dual QSFP Ports	16
Figure 2:	Representative LFF OCP NIC 3.0 Card with Dual QSFP Ports and on-board DRAM	17
Figure 3:	SFF and LFF Block Diagrams (not to scale)	18
Figure 4:	Primary Connector (4C+) and Secondary Connector (4C) (LFF) OCP NIC 3.0 Cards	22
Figure 5:	Primary Connector (4C+) Only (LFF) OCP NIC 3.0 Cards	23
Figure 6:	Primary Connector (4C+) with 4C and 2C (SFF) OCP NIC 3.0 Cards	23
Figure 7:	SFF NIC Configuration Views	25
Figure 8:	SFF NIC Line Side 3D Views	26
Figure 9:	SFF NIC Chassis Mounted 3D Views	27
Figure 10:	LFF NIC Configuration Views	29
Figure 11:	LFF NIC Line Side 3D Views	30
Figure 12:	LFF NIC Chassis Mounted 3D Views	31
Figure 13:	PBA Exploded Views (SFF and LFF)	33
Figure 14:	Faceplate Assembly Exploded Views (SFF and LFF)	34
Figure 15:	SFF Generic I/O Faceplate with Pulltab Version (2D View)	37
Figure 16:	SFF Generic I/O Faceplate – Ejector Version (2D View)	37
Figure 17:	SFF Generic I/O Faceplate – Internal Lock Version (2D View)	38
Figure 18:	LFF Generic I/O Faceplate – Ejector Version (2D View)	38
Figure 19:	SFF I/O Faceplate – Ejector Lever (2D View)	39
Figure 20:	LFF I/O Faceplate – Ejector Lever (2D View)	40
Figure 21:	Ejector Lock	41
Figure 22:	Clinch Nut Option A	42
Figure 23:	Clinch Nut Option B	42
Figure 24:	SFF Keep Out Zone – Top View	43
Figure 25:	SFF Keep Out Zone – Top View – Detail A	44
Figure 26:	SFF Keep Out Zone – Bottom View	44
Figure 27:	SFF Keep Out Zone – Side View	45
Figure 28:	SFF Keep Out Zone – Side View – Detail D	45
Figure 29:	LFF Keep Out Zone – Top View	46
Figure 30:	LFF Keep Out Zone – Top View – Detail A	47
Figure 31:	LFF Keep Out Zone – Bottom View	48
Figure 32:	LFF Keep Out Zone – Side View	48
Figure 33:	LFF Keep Out Zone – Side View – Detail D	49
Figure 34:	SFF Bottom Side Insulator (3D View)	50
Figure 35:	SFF Bottom Side Insulator (Top and Side View)	51
Figure 36:	SFF Bottom Side Insulator (alternate) (Top and Side View)	52
Figure 37:	LFF Bottom Side Insulator (3D View)	52
Figure 38:	LFF Bottom Side Insulator (Top and Side View)	53
Figure 39:	LFF Bottom Side Insulator (alternate) (Top and Side View)	54
Figure 40:	SFF OCP NIC 3.0 Card with Pull Tab CTF Dimensions (Top View)	55
Figure 41:	SFF OCP NIC 3.0 Card with Pull Tab CTF Dimensions (Front View)	56
Figure 42:	SFF OCP NIC 3.0 Card with Pull Tab CTF Dimensions (Side View)	56
Figure 43:	SFF OCP NIC 3.0 Card with Ejector CTF Dimensions (Top View)	57
Figure 44:	SFF OCP NIC 3.0 Card with Ejector CTF Dimensions (Front View)	57
Figure 45:	SFF OCP NIC 3.0 Card with Ejector CTF Dimensions (Side View)	58
Figure 46:	SFF OCP NIC 3.0 Card with Internal Lock CTF Dimensions (Top View)	58
Figure 47:	SFF OCP NIC 3.0 Card with Internal Lock CTF Dimensions (Front View)	59

Figure 48: SFF OCP NIC 3.0 Card with Internal Lock CTF Dimensions (Side View).....	59
Figure 49: SFF Baseboard Chassis CTF Dimensions (Rear View)	59
Figure 50: SFF Baseboard Chassis to Card Thumb Screw CTF Dimensions (Side View).....	60
Figure 51: SFF Baseboard Chassis to Ejector lever Card CTF Dimensions (Side View)	60
Figure 52: SFF Baseboard Chassis CTF Dimensions (Rear Rail Guide View)	60
Figure 53: SFF Baseboard Chassis CTF Dimensions (Rail Guide Detail) – Detail C	61
Figure 54: LFF OCP NIC 3.0 Card with Ejector CTF Dimensions (Top View).....	62
Figure 55: LFF OCP NIC 3.0 Card with Ejector CTF Dimensions (Front View)	62
Figure 56: LFF OCP NIC 3.0 Card with Ejector CTF Dimensions (Side View)	63
Figure 57: LFF Baseboard Chassis CTF Dimensions (Rear View)	63
Figure 58: LFF Baseboard Chassis CTF Dimensions (Side View)	64
Figure 59: LFF Baseboard Chassis CTF Dimensions (Rail Guide View)	64
Figure 60: LFF Baseboard Chassis CTF Dimensions (Rail Guide – Detail C)	64
Figure 61: SFF Label Area Example.....	66
Figure 62: MAC Address Label Example 1 – Quad Port with Single Host, Single Managed Controller	68
Figure 63: MAC Address Label Example 2 – Octal Port with Single Host, Dual Managed Controller	69
Figure 64: MAC Address Label Example 3 – Quad Port with Dual Hosts, Dual Managed Controllers	69
Figure 65: MAC Address Label Example 4 – Single Port with Quad Host, Single Managed Controller	70
Figure 66: MAC Address Label Example 5 – Octal Port with Single Host, Octal Managed Controller	70
Figure 67: SFF Primary Connector Gold Finger Dimensions – x16 – Top Side (“B” Pins)	72
Figure 68: SFF Primary Connector Card Profile Dimensions	73
Figure 69: SFF Primary Connector Gold Finger - Detail D	73
Figure 70: LFF Gold Finger Dimensions – x32 – Top Side (“B” Pins)	74
Figure 71: LFF Gold Finger Dimensions – x32 – Bottom Side (“A” Pins).....	74
Figure 72: 168-pin Base Board Primary Connector – Right Angle	78
Figure 73: 140-pin Base Board Secondary Connector – Right Angle	79
Figure 74: OCP NIC 3.0 Card and Host Offset for Right Angle Connectors.....	79
Figure 75: 168-pin Base Board Primary Connector – Straddle Mount	80
Figure 76: 140-pin Base Board Secondary Connector – Straddle Mount.....	80
Figure 77: OCP NIC 3.0 Card and Baseboard PCB Thickness Options for Straddle Mount Connectors	81
Figure 78: 0 mm Offset (Coplanar) for 0.062” Thick Baseboards	81
Figure 79: 0.3 mm Offset for 0.076” Thick Baseboards.....	82
Figure 80: Primary and Secondary Connector Locations for LFF Support with Right Angle Connectors	82
Figure 81: Primary and Secondary Connector Locations for LFF Support with Straddle Mount Connectors	82
Figure 82: PCIe Present and Bifurcation Control Pins (Baseboard Controlled BIF[0:2]#)	94
Figure 83: PCIe Present and Bifurcation Control Pins (Static BIF[0:2]#).....	94
Figure 84: Example SMBus Connections.....	96
Figure 85: NC-SI over RBT Connection Example – Single Primary Connector	102
Figure 86: NC-SI over RBT Connection Example – Dual Primary Connectors	103
Figure 87: Example Scan Chain Timing Diagram	105
Figure 88: Scan Chain Connection Example	111
Figure 89: Example Power Supply Topology	116
Figure 90: USB 2.0 Connection Example – Basic Connectivity	118
Figure 91: USB 2.0 Connection Example – USB-Serial / USB-JTAG Connectivity	118
Figure 92: UART Connection Example	120
Figure 93: Single Host (1 x16) and 1 x16 OCP NIC 3.0 Card (Single Controller)	126
Figure 94: Single Host (2 x8) and 2 x8 OCP NIC 3.0 Card (Dual Controllers)	127
Figure 95: Quad Hosts (4 x4) and 4 x4 OCP NIC 3.0 Card (Single Controller).....	128
Figure 96: Quad Hosts (4 x4) and 4 x4 OCP NIC 3.0 Card (Quad Controllers).....	129
Figure 97: Single Host with no Bifurcation (1 x16) and 2 x8 OCP NIC 3.0 Card (Dual Controllers).....	130
Figure 98: SFF PCIe REFCLK Mapping – Single Host – 1, 2 and 4 links	132
Figure 99: SFF PCIe REFCLK Mapping – Dual Host – 2 and 4 links.....	133
Figure 100: SFF PCIe REFCLK Mapping – Quad Host – 4 Links.....	134
Figure 101: LFF PCIe REFCLK Mapping – Single Host – 1, 2 and 4 links	135
Figure 102: LFF PCIe REFCLK Mapping – Dual Host – 2 and 4 links.....	136
Figure 103: LFF PCIe REFCLK Mapping – Quad Host – 4 Links.....	137
Figure 104: Port and LED Ordering – Example SFF Link/Activity and Speed LED Placement.....	141
Figure 105: Baseboard Power States	142

Figure 106: Power-Up Sequencing – Normal Operation	147
Figure 107: Power-Down Sequencing – Normal Operation	148
Figure 108: Programming Mode Sequencing	149
Figure 109: FRU EEPROM Writes with Double Byte Addressing	162
Figure 110: FRU EEPROM Reads with Double Byte Addressing	162
Figure 111: FRU Update Flow	163
Figure 112: NC-SI over RBT Timing Budget Topology	171
Figure 113: NC-SI over RBT Propagation Delay Matching for Two Target ASICs – No Clock Buffer	173
Figure 114: NC-SI over RBT Propagation Delay Matching for Two Target ASICs – Clock Buffer	173
Figure 115: PCIe Load Board Test Fixture for OCP NIC 3.0 SFF	175
Figure 116: PCIe Base Board Test Fixture for OCP NIC 3.0 SFF	176
Figure 117: Airflow Direction for Hot Aisle Cooling (SFF and LFF)	177
Figure 118: Airflow Direction for Cold Aisle Cooling (SFF and LFF)	178
Figure 119: ASIC Supportable Power for Hot Aisle Cooling – SFF	179
Figure 120: OCP NIC 3.0 SFF Reference Design and CFD Geometry	180
Figure 121: Server System Airflow Capability – SFF Card Hot Aisle Cooling	181
Figure 122: Server System Airflow Capability – SFF Card Hot Aisle Cooling in standby (S5) mode	182
Figure 123: ASIC Supportable Power for Hot Aisle Cooling – LFF Card	183
Figure 124: OCP NIC 3.0 LFF Reference Design and CFD Geometry	183
Figure 125: Server System Airflow Capability – LFF Card Hot Aisle Cooling	185
Figure 126: ASIC Supportable Power for Cold Aisle Cooling – SFF Card	186
Figure 127: Server System Airflow Capability – SFF Cold Aisle Cooling	187
Figure 128: ASIC Supportable Power Comparison – SFF Card	187
Figure 129: ASIC Supportable Power for Cold Aisle Cooling – LFF Card	188
Figure 130: Server System Airflow Capability – LFF Cold Aisle Cooling	189
Figure 131: ASIC Supportable Power Comparison – LFF Card	189
Figure 132: SFF Thermal Test Fixture Preliminary Design	191
Figure 133: SFF Thermal Test Fixture Preliminary Design – Cover Removed	192
Figure 134: SFF Card Thermal Test Fixture PCB	192
Figure 135: LFF Card Thermal Test Fixture Design	193
Figure 136: LFF Card Thermal Test Fixture Design – Cover Removed	193
Figure 137: LFF Card Thermal Test Fixture PCB	194
Figure 138: Thermal Test Fixture Airflow Direction	195
Figure 139: SFF Fixture, Hot Aisle Flow – Candlestick Air Velocity vs. Volume Flow	196
Figure 140: LFF Fixture, Hot Aisle Flow – Candlestick Air Velocity vs. Volume Flow	196
Figure 141: Graphical View of Card Cooling Tiers	198
Figure 142: Typical Operating Range for Hot Aisle Configurations	198
Figure 143: Typical Operating Range for Cold Aisle Configurations	199
Figure 144: SFF Shock and Vibe Fixture	200
Figure 145: LFF Shock and Vibe Fixture	200
Figure 146: Dye and Pull Type Locations	203
Figure 147: Dye Coverage Percentage	203

List of Tables

Table 1: Acknowledgements – By Company	11
Table 2: Acronyms	14
Table 3: OCP 3.0 Form Factor Dimensions	19
Table 4: Baseboard to OCP NIC Form Factor Compatibility Chart	19
Table 5: Example Non-NIC Use Cases	21
Table 6: OCP NIC 3.0 Card Definitions	24
Table 7: OCP NIC 3.0 Line Side I/O Implementations	32
Table 8: Line Side I/O Cross Reference to Industry Standards	32
Table 9: Bill of Materials for the SFF and LFF Faceplate Assemblies	35
Table 10: CTF Default Tolerances (SFF and LFF OCP NIC 3.0)	55
Table 11: MAC Address Label Example 1 – Quad Port with Single Host, Single Managed Controller	68
Table 12: MAC Address Label Example 2 – Octal Port with Single Host, Dual Managed Controller	68
Table 13: MAC Address Label Example 3 – Quad Port with Dual Hosts, Dual Managed Controller	69
Table 14: MAC Address Label Example 4 – Single Port with Quad Host, Single Managed Controller	70
Table 15: MAC Address Label Example 5 – Octal Port with Single Host, Octal Managed Controller	70
Table 16: NIC Implementation Examples and 3D CAD	71
Table 17: Contact Mating Positions for the Primary Connector	74
Table 18: Contact Mating Positions for the Secondary Connector	76
Table 19: Right Angle Connector Options	78
Table 20: Straddle Mount Connector Options	79
Table 21: Primary Connector Pin Definition (x16) (4C+)	83
Table 22: Secondary Connector Pin Definition (x16) (4C)	85
Table 23: Pin Descriptions – PCIe	86
Table 24: Pin Descriptions – PCIe Present and Bifurcation Control Pins	92
Table 25: Pin Descriptions – SMBus	95
Table 26: Pin Descriptions – NC-SI over RBT	96
Table 27: Pin Descriptions – Scan Chain	104
Table 28: Pin Descriptions – Scan Chain DATA_OUT Bit Definition	106
Table 29: Pin Descriptions – Scan Chain DATA_IN Bit Definition	106
Table 30: Pin Descriptions – Power	112
Table 31: Pin Descriptions – USB 2.0 – Primary Connector only	117
Table 32: Pin Descriptions – UART – Secondary Connector Only	119
Table 33: Pin Descriptions – RFU[1:4]	121
Table 34: PCIe Bifurcation Decoder for x32, x16, x8, x4, x2 and x1 Card Widths	124
Table 35: PCIe REFCLK and PERST Associations	131
Table 36: SFF PCIe Link / REFCLKn / PERSTn mapping for 1, 2 and 4 Links	131
Table 37: LFF PCIe Link / REFCLKn / PERSTn mapping for 1, 2, 4 and 8 Links	131
Table 38: OCP NIC 3.0 Card LED Configuration with Two Physical LEDs per Port	139
Table 39: Available Card Functions per Power State	143
Table 40: Baseboard Power Supply Rail Requirements – Slot Power Envelopes	145
Table 41: Power Sequencing Parameters	149
Table 42: Digital I/O DC specifications	151
Table 43: Digital I/O AC specifications	151
Table 44: OCP NIC 3.0 Management Implementation Definitions	152
Table 45: Sideband Management Interface and Transport Requirements	152
Table 46: NC-SI Traffic Requirements	153
Table 47: MC MAC Address Provisioning Requirements	153
Table 48: Temperature Reporting Requirements	156
Table 49: Power Consumption Reporting Requirements	157
Table 50: Pluggable Module Status Reporting Requirements	158
Table 51: Management and Pre-OS Firmware Inventory and Update Requirements	158
Table 52: Slot_ID[1:0] to Package ID[2:0] Mapping	160
Table 53: FRU EEPROM Address Map	161
Table 54: FRU EEPROM Record – OEM Record 0xC0, Offset 0x00	164
Table 55: NC-SI over RBT Timing Parameters	170
Table 56: PCIe Electrical Budgets	174
Table 57: PCIe Test Fixtures for OCP NIC 3.0	175

Table 58: Hot Aisle Air Temperature Boundary Conditions178

Table 59: Hot Aisle Airflow Boundary Conditions178

Table 60: Cold Aisle Air Temperature Boundary Conditions178

Table 61: Cold Aisle Airflow Boundary Conditions179

Table 62: Reference OCP NIC 3.0 SFF Card Geometry180

Table 63: Reference OCP NIC 3.0 LFF Card Geometry184

Table 64: Card Cooling Tier Definitions (LFM)197

Table 65: Random Vibration Testing 1.88 G_{RMS} Profile201

Table 66: FCC Class A Radiated and Conducted Emissions Requirements Based on Geographical Location205

Table 67: Safety Requirements206

Table 68: Immunity (ESD) Requirements206

DRAFT FOR REVIEW ONLY

1 Overview

1.1 License

As of January 23rd, 2018, the following persons or entities have made this Specification available under the Open Compute Project Hardware License (Permissive) Version 1.0 (OCPHL-P)

- OCP NIC Subgroup

An electronic copy of the OCPHL-P is available at:

<https://www.opencompute.org/documents/ocphl-permissive-v10>

Your use of this Specification may be subject to other third party rights. THIS SPECIFICATION IS PROVIDED "AS IS." The contributors expressly disclaim any warranties (express, implied, or otherwise), including implied warranties of merchantability, non-infringement, fitness for a particular purpose, or title, related to the Specification. The Specification implementer and user assume the entire risk as to implementing or otherwise using the Specification. IN NO EVENT WILL ANY PARTY BE LIABLE TO ANY OTHER PARTY FOR LOST PROFITS OR ANY FORM OF INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER FROM ANY CAUSES OF ACTION OF ANY KIND WITH RESPECT TO THIS SPECIFICATION OR ITS GOVERNING AGREEMENT, WHETHER BASED ON BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), OR OTHERWISE, AND WHETHER OR NOT THE OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

1.2 Acknowledgements

The OCP NIC 3.0 specification was created under a collaboration from many OCP member companies, and facilitated by the OCP NIC Subgroup under the OCP Server Workgroup.

The OCP NIC Subgroup would like to acknowledge the following member companies for their contributions to the OCP NIC 3.0 specification:

Table 1: Acknowledgements – By Company

Amphenol Corporation	Lenovo Group Ltd
Broadcom Limited	Marvell Semiconductor, Inc.
Dell, Inc.	Mellanox Technologies, Ltd
Facebook, Inc.	Netronome Systems, Inc.
Hewlett Packard Enterprise Company	Quanta Computer Inc.
Intel Corporation	TE Connectivity Corporation
Keysight Technologies	University of New Hampshire InterOperability Lab

1.3 References

- DMTF Standard. *DSP0222, Network Controller Sideband Interface (NC-SI) Specification*. Distributed Management Task Force (DMTF), Rev 1.1.0, September 23rd, 2015.
- DMTF Standard. *DSP0222, Network Controller Sideband Interface (NC-SI) Specification*. Distributed Management Task Force (DMTF), Rev 1.2.0, Work-In-Progress.
- DMTF Standard. *DSP0236, Management Component Transport Protocol (MCTP) Base Specification*. Distributed Management Task Force (DMTF), Rev 1.3.0, November 24th, 2016.
- DMTF Standard. *DSP0237, Management Component Transport Protocol (MCTP) SMBus/I2C Transport Binding Specification*. Distributed Management Task Force (DMTF), Rev 1.1.0, May 21st, 2017.
- DMTF Standard. *DSP0238, Management Component Transport Protocol (MCTP) PCIe VDM Transport Binding Specification*. Distributed Management Task Force (DMTF), Rev 1.0.2, December 7th, 2014.
- DMTF Standard. *DSP0239, MCTP IDs and Codes Specification*. Distributed Management Task Force (DMTF), Rev 1.5.0, December 17th, 2017.
- DMTF Standard. *DSP0240, Platform Level Data Model (PLDM) Base Specification*. Distributed Management Task Force (DMTF), Rev 1.0.0, April 23rd, 2009.
- DMTF Standard. *DSP0240, Platform Level Data Model (PLDM) over MCTP Binding Specification*. Distributed Management Task Force (DMTF), Rev 1.0.0, April 23rd, 2009.
- DMTF Standard. *DSP0245, Platform Level Data Model (PLDM) IDs and Codes Specification*. Distributed Management Task Force (DMTF), Rev 1.2.0, November 24th, 2016.
- DMTF Standard. *DSP0248, Platform Level Data Model (PLDM) for Platform Monitoring and Control Specification*. Distributed Management Task Force (DMTF), Rev 1.1.1, January 10th, 2017.
- DMTF Standard. *DSP0249, Platform Level Data Model (PLDM) State Sets Specification*. Distributed Management Task Force (DMTF), Rev 1.0.0, March 16th, 2009.
- DMTF Standard. *DSP0261, NC-SI over MCTP Binding Specification*. Distributed Management Task Force (DMTF), Rev 1.2.0, August 26th, 2017.
- DMTF Standard. *DSP0267, Platform Level Data Model (PLDM) for Firmware Update Specification*. Distributed Management Task Force (DMTF), Rev 1.0.1, March 28th, 2018.
- EDSFF. *Enterprise and Datacenter SSD Form Factor Connector Specification*. Enterprise and Datacenter SSD Form Factor Working Group, Rev 0.9 (draft), August 2nd 2017.
- IPC. *IPC-TM-650 Test Methods Manual number 2.4.53. Dye and Pull Test Method (Formerly Known as Dye and Pry)*, Association Connecting Electronics Industries, August 2017.
- IPMI Platform Management FRU Information Storage Definition, v1.0 Document Revision 1.3, March 24th, 2015.
- National Institute of Standards and Technology (NIST). *Special Publication 800-193, Platform Firmware Resiliency Guidelines*, May 2018.
- NXP Semiconductors. *I²C-bus specification and user manual*. NXP Semiconductors, Rev 6, April 4th, 2014.
- Open Compute Project. *OCP NIC Subgroup*. Online. <http://www.opencompute.org/wiki/Server/Mezz>
- PCI-SIG®. *PCI Express® Base Specification*, Revision 3.0 December 7th, 2015.

- PCI-SIG®. *PCI Express® Base Specification*, Revision 4.0 Version 1.0, October 5th, 2017.
- PCI-SIG®. *PCI Express® Card Electromechanical Specification*, Revision 3.0, July 21st, 2013.
- PCI-SIG®. *PCI Express® Card Electromechanical Specification*, Revision 4.0 (draft).
- SMBus Management Interface Forum. *System Management Bus (SMBus) Specification*. System Management Interface Forum, Inc, Version 2.0, August 3rd, 2000.
- SNIA. SFF-TA-1002, Specification for Protocol Agnostic Multi-Lane High Speed Connector. SNIA SFF TWG Technology Affiliate, Rev 1.1 draft, January 18th, 2018.
- UEFI Specification Version 2.5, <http://www.uefi.org/sites/default/files/resources/UEFI%202.5.pdf>, April 2015.
- USB Implementers Forum. *Universal Serial Bus Specification*, Revision 2.0, April 27th, 2000.

1.3.1 Trademarks

Names and brands may be claimed as trademarks by their respective companies.

I²C® is a trademark of NXP Semiconductor.

PCIe® and PCI Express® are the registered trademarks of PCI-SIG.

1.4 Acronyms

For the purposes of the OCP NIC 3.0 specification, the following acronyms apply:

Table 2: Acronyms

Acronym	Definition
AIC	Add-in Card
ASIC	Application Specific Integrated Circuit
BGA	Ball Grid Array
BMC	Baseboard Management Controller
BOM	Bill of Materials
CAD	Computer Aided Design
CBB	Compliance Base Board
CEM	Card Electromechanical
CFD	Computational Fluid Dynamics
CFM	Cubic Feet per Minute
CLB	Compliance Load Board
CTD	Chain of Trust for Detection
CTF	Critical to Function
CTU	Chain of Trust for Update
DMTF	Distributed Management Task Force
DRAM	Dynamic Random Access Memory
EDSFF	Enterprise and Datacenter SSD Form Factor
EMI	Electro Magnetic Interference
ESD	Electrostatic Discharge
EU	European Union
FCC	Federal Communications Commission
FRU	Field Replaceable Unit
I/O	Input / Output
I2C	Inter-Integrated Circuit - two wire serial protocol
IEC	International Electrotechnical Commission
IPC	Institute for Printed Circuits
IPMI	Intelligent Platform Management Interface
ISO	International Organization for Standardization
LED	Light Emitting Diode
LFF	Large Form Factor
LFM	Linear Feet per Minute
MAC	Media Access Control
MC	Management Controller
MCTP	Management Component Transport Protocol
ME	Management Entity
MSA	Multi-source Agreement
NC	No Connect
NC-SI	Network Controller Sideband Interface
NEBS	Network Equipment Building-System
NIC	Network Interface Card
OCP	Open Compute Project

ODM	Original Design Manufacturer
OEM	Original Equipment Manufacturer
PBA	Printed Board Assembly
PCB	Printed Circuit Board
PCI™	Peripheral Component Interconnect
PCIe®	PCI Express®
PDR	Platform Descriptor Record
PLDM	Platform Level Data Model
QSFP	Quad Small Form Factor Pluggable
QZ	Quiet Zone
RA	Right Angle
RBT	RMII Based Transport
REACH	Registration, Evaluation, Authorization and Restriction of Chemicals
RFU	Reserved Future Use
RJ45	Registered Jack 45 (IEC 60603-7 8P8C connector)
RoHS	Restriction of Hazardous Substances Directive
RSVD	Reserved
RTU	Root of Trust for Update
SFF	Small Form Factor
SFP	Small Form Factor Pluggable
SMBus	System Management Bus
SMT	Surface Mount Technology
TBD	To be Determined
TDP	Thermal Design Power
UART	Universal Asynchronous Receiver-Transmitter
UDID	Unique Device Identifier for each SMBus device. Refer to the SMBus 2.0 specification for the field definition.
UEFI	Unified Extensible Firmware Interface
USB	Universal Serial Bus
VDM	Vendor Defined Messages
WEEE	Waste Electrical and Electronic Equipment

1.5 Background

The OCP NIC 3.0 specification is a follow-on to the OCP Mezz 2.0 rev 1.00 design specification. The OCP NIC 3.0 specification supports two basic card sizes: Small Form Factor (SFF), and Large Form Factor (LFF). The SFF allows for up to 16 PCIe® lanes on the card edge while the LFF supports up to 32 PCIe lanes. Compared to the OCP Mezz Card 2.0 Design Specification, the updated OCP NIC 3.0 specification provides a broader solution space for the NIC and system vendors to support the following use case scenarios:

- NICs with a higher Thermal Design Power (TDP)
- Power delivery supports up to 80 W to a single connector (SFF) card, and up to 150 W to a dual connector (LFF) card
 - Note: Baseboard vendors need to evaluate if there is sufficient airflow to thermally cool the OCP NIC 3.0 card. Refer to Section 6 for additional details.
- Supports up to PCIe Gen 4 (16 GT/s) on the baseboard and OCP NIC 3.0 card
 - Connector is electrically compatible with PCIe Gen 5 (32 GT/s)
- Support for up to 32 lanes of PCIe per OCP NIC 3.0 card
- Support for single host, multi-root complex, and multi-host environments
- Supports a greater board area for more complex OCP NIC 3.0 card designs
- Support for Smart NIC implementations with on-board DRAM and accelerators
- Simplification of FRU installation and removal while reducing overall down time

A representative SFF OCP NIC 3.0 card is shown in Figure 1 and a representative LFF is shown in Figure 2.

Figure 1: Representative SFF OCP NIC 3.0 Card with Dual QSFP Ports





Figure 2: Representative LFF OCP NIC 3.0 Card with Dual QSFP Ports and on-board DRAM



In order to achieve the features outlined in this specification, OCP NIC 3.0 compliant cards are not backwards compatible with OCP Mezz 2.0 cards.

This specification is created under OCP Server workgroup – OCP NIC subgroup. An electronic copy of this specification can be found on the Open Compute Project and the OCP Marketplace websites:

http://www.opencompute.org/wiki/Server/Mezz#Specifications_and_Designs

<https://www.opencompute.org/contributions?query=OCP%20NIC%203.0>

1.6 Overview


1.6.1 Mechanical Form Factor Overview

The OCP NIC 3.0 specification defines a third generation mechanical form factor that allows for interoperability between compliant baseboards and OCP NIC 3.0 cards.

OCP NIC 3.0 cards have two form factors – SFF and LFF. These cards are shown in Figure 3 below. The components shown in the figures are for illustrative purposes. The SFF uses one connector (Primary Connector) on the baseboard. The LFF uses one or two connectors (Primary Connector only or both the Primary and Secondary Connectors) on the baseboard.

Both the Primary and Secondary Connectors and card edge gold fingers are defined in and compliant to SFF-TA-1002. The Primary Connector is the “4C+” variant, the Secondary Connector is the “4C” version. On the OCP NIC 3.0 card side, the card edge is implemented with gold fingers. The SFF gold finger area only occupies the Primary Connector area for up to 16 PCIe lanes. The LFF gold finger area may occupy both the Primary and Secondary Connectors for up to 32 PCIe lanes, or optionally just the Primary Connector for up to 16 PCIe lane implementations.

Figure 3: SFF and LFF Block Diagrams (not to scale)



The two form factor dimensions are shown in Table 3.

Table 3: OCP 3.0 Form Factor Dimensions

Form Factor	Width	Depth	Primary Connector	Secondary Connector	Typical Use Case
SFF	W1 = 76 mm	L = 115 mm	“4C+” 168 pins	N/A	Low profile and NIC with a similar profile as an OCP NIC 2.0 card; up to 16 PCIe lanes.
LFF	W2 = 139 mm	L = 115 mm	“4C+” 168 pins	“4C” 140 pins	Larger PCB width to support additional NICs; up to 32 PCIe lanes.

The OCP NIC 3.0 design allows downward compatibility between the two card sizes. Table 4 shows the compatibility between the baseboard and NIC combinations. A SFF baseboard slot may only accept a SFF sized NIC. A LFF baseboard slot may accept a SFF or LFF NIC.

Table 4: Baseboard to OCP NIC Form Factor Compatibility Chart

Baseboard Slot Size	NIC Size / Supported PCIe Width	
	SFF	LFF
SFF	Up to 16 PCIe lanes	Not Supported
LFF	Up to 16 PCIe lanes	Up to 32 PCIe lanes

There are two baseboard connector mounting options available for system designers: straddle mount and right angle (RA). The straddle mount connector option allows the OCP NIC and baseboard to exist in a co-planer position. To achieve this, a cutout exists on the baseboard and is defined in this specification. Alternatively, the right angle option allows the OCP NIC to be installed on top of the baseboard. A baseboard cutout is not required for the right angle connector. The right angle option allows the baseboard to use this area for additional routing or backside component placement. The straddle mount and right angle connectors are shown in Section 3.2.

For both the baseboard and OCP NIC 3.0 card, this specification defines the component and routing keep out areas. Refer to Section 2.5 for details.

Both the straddle mount and right angle implementations shall accept the same OCP NIC 3.0 card and shall be supported in the baseboard chassis regardless of the baseboard connector selection (right angle or straddle mount) so long as the baseboard slot and OCP NIC 3.0 card sizes are a supported combination as shown in Table 4.

This specification defines the form factor at the OCP NIC 3.0 card level, including the front panel, latching mechanism and card guide features.

More details about the card form factor is shown in Section 2.

1.6.2 Electrical Overview

This specification defines the electrical interface between baseboard and the OCP NIC 3.0 card. The electrical interface is implemented with a right angle or straddle mount connector on baseboard and gold finger on the OCP NIC 3.0 card. As previously noted in the mechanical overview, each card may implement a Primary Connector or Primary + Secondary Connector. Cards using only the Primary Connector are suitable for both the SFF and LFF and may support up to 16 lanes of PCIe. The Secondary Connector, when used in conjunction with the Primary Connector, allows LFF implementations and may support up to 32 lanes of PCIe.

1.6.2.1 Primary Connector

The Primary Connector provides all OCP specific management functions as well as up to 16 lanes of PCIe between the OCP NIC and the system motherboard.

Management Function Overview (OCP Bay):

- DMTF DSP0222 compliant Network Controller Sideband Interface (NC-SI) RMIII Based Transport (RBT) Physical Interface
- Power management and status reporting
 - Power Brake for emergency power reduction
 - State change control
- Control / status serial bus
 - NIC-to-Host status
 - Port LED Link/Activity
 - Environmental Indicators
 - Host-to-NIC configuration Information
- Multi-host PCIe support signals (2x PCIe resets, 2x reference clocks)
 - The OCP bay provides PERST2#, PERST3#, REFCLK2 and REFCLK3. This enables support for up to four hosts when used in conjunction with PERST0#, PERST1#, REFCLK0 and REFCLK1 in the Primary 4C region.
- PCIe Wake signal

See Section 3.4 for a complete list of pin and function descriptions for the OCP Bay portion of the Primary Connector. The OCP Bay pins are prefixed with “OCP_” in the pin location column.

Interface Overview (4C Connector):

- 16x differential transmit/receive pairs
 - Up to PCIe Gen 4 (16 GT/s) support
 - Connector is electrically compatible with PCIe Gen 5 (32 GT/s)
- 2x 100 MHz differential reference clocks
- Control signals
 - 2x PCIe Resets
 - Link Bifurcation Control
 - Card power disable/enable
- SMBus 2.0
- USB 2.0 interface

- Power
 - +12V_EDGE
 - +3.3V_EDGE
 - Power distribution between the aux and main power domains is up to the baseboard vendor

See Section 3.4 for a complete list of pin and function descriptions for the 4C+ connector.

1.6.2.2 Secondary Connector

The Secondary Connector provides an additional 16 lanes of PCIe and their respective control signals.

Interface Overview (4C Connector):

- 16x differential transmit/receive pairs
 - Up to PCIe Gen 4 (16 GT/s) support
 - Connector is electrically compatible with PCIe Gen 5 (32 GT/s)
- 2x 100 MHz differential reference clocks
- Control signals
 - 2x PCIe Resets
 - Link Bifurcation Control
 - Card power disable/enable
- SMBus 2.0
- UART (transmit and receive)
- Power
 - +12V_EDGE
 - +3.3V_EDGE
 - Power distribution between the aux and main power domains is up to the baseboard vendor

See Section 3.4 for a complete list of pin and function descriptions for the 4C connector.

1.7 Non-NIC Use Cases

The OCP NIC 3.0 specification is mainly targeted for Network Interface Card applications. It is possible to use the same OCP NIC 3.0 card form factor, baseboard interface and mechanical design to enable non-NIC use cases. These non-NIC use cases use the same baseboard/OCP NIC 3.0 card interface as defined in Section 3. The non-NIC use cases are not covered in the current revision of the OCP NIC 3.0 specification. Example non-NIC use cases implement various external I/O interfaces and are shown in Table 5.

Table 5: Example Non-NIC Use Cases

Example Use Case	Card External I/O Interface(s)
PCIe Retimer Card	PCIe
Accelerator Card	N/A
NVMe Card	N/A
Storage HBA / RAID Card	I/O Interface to be determined in future specification release.

2 Mechanical Card Form Factor

2.1 Form Factor Options

OCP NIC 3.0 provides two fundamental form factor options: a SFF (76 mm x 115 mm) and a LFF (139 mm x 115 mm).

These form factors support a Primary Connector and optionally, a Secondary Connector. The Primary Connector is defined to be a SFF-TA-1002 compliant 4C+ connector. The 4C+ connector is a 4C compliant implementation plus a 28-pin “OCP bay” for OCP NIC 3.0 specific pins. The Secondary Connector is the 4C connector as defined in SFF-TA-1002. The 4C specification supports up to 32 differential pairs for a x16 PCIe connection per connector. For host platforms, the 28-pin OCP bay is required for all Primary Connector implementations.

The SFF uses the Primary 4C+ connector to provide up to a x16 PCIe interface to the host. The additional 28-pin OCP bay carries sideband management interfaces as well as OCP NIC 3.0 specific control signals for multi-host PCIe support. The SFF card provides sufficient faceplate area to accommodate up to 2x QSFP modules, 4x SFP modules, or 4x RJ45 for BASE-T operation. The SFF supports up to 80 W of delivered power to the card edge. An example SFF is shown in Figure 1.

The LFF uses the Primary 4C+ connector to provide the same functionality as the SFF along with an additional Secondary 4C connector to provide up to a x32 PCIe interface. The LFF Card may utilize both the Primary and Secondary Connectors, or just the Primary Connector for lower PCIe lane count applications. Table 6 summarizes the LFF permutations. The LFF supports higher power envelopes and provides additional board area for more complex designs. The LFF supports up to 150 W of delivered power to the card edge across the two connectors. An example LFF is shown in Figure 2.

For LFF Cards, implementations may use both the Primary and Secondary Connector (as shown in Figure 4), or may use the Primary Connector only (as shown in Figure 5) for the card edge gold fingers.

Figure 4: Primary Connector (4C+) and Secondary Connector (4C) (LFF) OCP NIC 3.0 Cards





Figure 5: Primary Connector (4C+) Only (LFF) OCP NIC 3.0 Cards



For both form factors, an OCP NIC 3.0 card may optionally implement a subset of pins to support less than a x16 PCIe connection. This may be implemented using a 2C+ card edge per SFF-TA-1002. The baseboard Primary Connector shall use a 4C+ in all cases. Figure 6 illustrates the supported 4C+ and 2C+ card edge configurations on a 4C+ Primary Connector.

Figure 6: Primary Connector (4C+) with 4C and 2C (SFF) OCP NIC 3.0 Cards




Table 6 summarizes the supported card form factors. SFF cards support the Primary Connector and up to 16 PCIe lanes. LFF cards support implementations with both the Primary and Secondary Connectors and up to 32 PCIe lanes, or a Primary Connector only implementation with up to 16 PCIe lanes.

Table 6: OCP NIC 3.0 Card Definitions

OCP NIC 3.0 Card Size and PCIe Lane Count	Baseboard Secondary Connector (4C)	Baseboard Primary Connector (4C+)	
	x16 PCIe	x16 PCIe	OCP Bay
SFF (x8)	Not used with SFF 2C+ Card Edge	x8 (Lanes 7:0) PCIe	OCP Bay
SFF (x16)	Not used with SFF 4C+ Card Edge	x16 (Lanes 15:0) PCIe	OCP Bay
LFF (x8)	Not used with LFF 2C+ Card Edge	x8 (Lanes 7:0) PCIe	OCP Bay
LFF (x16)	Not used with LFF 4C+ Card Edge	x16 (Lanes 15:0) PCIe	OCP Bay
LFF (x32)	x16 (Lanes 31:16) PCIe	x16 (Lanes 15:0) PCIe	OCP Bay

All mechanical board assemblies shall meet the safety requirements described in Section 7.1.3.

2.1.1 SFF Faceplate Configurations

The SFF configuration views are shown below. Three different faceplates are available for the SFF – a pull tab, ejector latch and an internal lock version are available. The same SFF OCP NIC 3.0 PBA assembly accepts all three faceplates types and may be interchanged depending on the end application. The drawings shown in Figure 7 below illustrate a representative front, side and top views of the SFF.

Where space is permitted on the faceplate, square vents sized to a maximum of 3.0 mm x 3.0 mm must be added to help optimize airflow while maintaining the integrity of the faceplate structure. EMI considerations should also be taken into account during the design process. Refer to the images shown in Figure 8 for example square vent configurations depending on the line side I/O connectors.

Depending on the OCP NIC 3.0 card implementation, I/O connectors may be placed anywhere within the allowable connector keep in regions as defined by the SFF PBA mechanical drawings and faceplate drawings of Section 2.5.1.

The OCP NIC 3.0 outline provides an optional feature to lock the card into the chassis. This is accomplished with two notches – one on each side of the card guide rail. A baseboard may choose to use one or both notches for the internal locking mechanism. The OCP NIC 3.0 outline provides a notch location on both guide rails to provide flexible configurations to baseboard vendors. If a locking feature is implemented on the baseboard, the OCP NIC 3.0 card may only be inserted or removed after actuating the internal locking mechanism. These retention notches are compatible with all chassis implementations. Please refer to the SFF dimensions in Section 2.5.1 for details. The internal locking mechanism is not available on LFF cards.

Note: The OCP NIC 3.0 card supplier shall add port identification on the faceplate assembly that meet their manufacturing and customer requirements.

All of the OCP NIC 3.0 CAD files are available for download and use on the OCP NIC 3.0 Wiki site: <http://www.opencompute.org/wiki/Server/Mezz>

Figure 7: SFF NIC Configuration Views




Figure 8 illustrates example SFF 3D views for the supported line side I/O implementations. The line side I/O implementations are discussed in Section 2.2.





Figure 9 illustrates example SFF 3D views of the pull tab and ejector latch assemblies mounted in a chassis utilizing a straddle mount connector and a right-angle connector. The baseboard connector options are discussed in Section 3.2. The SFF OCP NIC 3.0 card is identical for both chassis connector options.

The OCP NIC 3.0 card provides a notch on each side of the card rail edge. This feature is used in conjunction with a baseboard's internal locking mechanism to prevent card insertion and removal. The internal locking mechanism is a baseboard vendor's optional feature and is not shown in the views below.


On the internal lock variation only, a 5.4 Kg force with a 1 second ramp may be applied to the heatsink during NIC ejection. The mechanical and thermal solution shall be implemented such that the thermal performance is maintained for 10 force cycles (i.e., the NIC still passes all functional tests). If there exists a minimum of (10 mm x 10 mm x 10 mm) of space behind the line side connectors, then each connector shall be able to sustain the same 5.4 Kg force with a 1 second ramp.

Figure 9: SFF NIC Chassis Mounted 3D Views



SFF with Pull Tab





SFF with Ejector Latch





Right Angle vs Straddle Mount Chassis Configuration



Right Angle Baseboard Connector



Straddle Mount Baseboard Connector



NIC Insertion / Removal (Shown with a Straddle Mount Connector)

2.1.2 LFF Faceplate Configurations

The LFF configuration views are shown below. A single faceplate implementation is available for the LFF – with a single ejector latch. The long ejector is the default configuration, however, a short ejector version is available for non-shadowed front I/O configurations and is being considered for future development. Similar to the SFF, if additional LFF faceplate implementations become available, the same LFF OCP NIC 3.0 PBA assembly shall be able to accept new faceplate types and may be interchanged depending on the end application. The drawings shown in Figure 10 below illustrate a representative front, side and top views of the LFF.

Where space is permitted on the faceplate, square vents sized to a maximum of 3.0 mm x 3.0 mm must be added to help optimize airflow while maintaining the integrity of the faceplate structure. EMI considerations should also be taken into account during the design process. Refer to the images shown in Figure 11 for example square vent configurations depending on the line side I/O connectors.


Depending on the OCP NIC 3.0 card implementation, I/O connectors may be placed anywhere within the allowable connector keep in regions as defined by the PBA mechanical drawings and faceplate drawings of Section 2.5

Note: The OCP NIC 3.0 card supplier shall add port identification on the faceplate assembly that meet their manufacturing and customer requirements.


All of the OCP NIC 3.0 CAD files are available for download and use on the OCP NIC 3.0 Wiki site:
<http://www.opencompute.org/wiki/Server/Mezz>

Figure 10: LFF NIC Configuration Views


LFF with Ejector Latch



Front View



Side View



Top View




Figure 11 illustrates example LFF 3D views for the supported line side I/O implementations. The line side I/O implementations are discussed in Section 2.2.

Figure 11: LFF NIC Line Side 3D Views





Figure 12 illustrates example LFF 3D views of the ejector latch assembly mounted in a chassis utilizing a straddle mount connector and a right angle connector. The baseboard connector options are discussed in Section 3.2. The LFF OCP NIC 3.0 card is identical for both chassis connector options.

Figure 12: LFF NIC Chassis Mounted 3D Views



2.2 Line Side I/O Implementations

At the time of this writing, the SFF and LFF implementations have been optimized to support the standard line side I/O implementations shown in Table 7. OCP NIC 3.0 cards may implement a subset of line side connectors and shall stay within the allowed I/O area as depicted in Section 2.4.3 for SFF and Section 2.4.4 for LFF.

Table 7: OCP NIC 3.0 Line Side I/O Implementations

Form Factor	Max Topology Connector Count
SFF	2x QSFP+/QSFP28
SFF	4x SFP+/SFP28
SFF	4x RJ45
LFF	2x QSFP+/QSFP28
LFF	4x SFP+/SFP28
LFF	4x RJ45

Note: For brevity, references to QSFP+, and QSFP28 shall be referred to as QSFP for the remainder of this document. Similarly, references to SFP+, and SFP28 shall be referred to as SFP.

The following specifications may be used to cross reference the associated line side I/O form-factor and module management specifications. This is not an exhaustive list. Refer to Table 3-1 in SFF-8024 – SFF Cross Reference to Industry Products for additional details.

Table 8: Line Side I/O Cross Reference to Industry Standards

Form Factor	Pluggable Form-Factor Specification	Common Management Specification
SFP	INF-8074 (MSA)	SFF-8472
SFP+	SFF-8084	SFF-8472
SFP28	SFF-8402	SFF-8472
QSFP	INF-8438 (MSA)	SFF-8636
QSFP+	SFF-8436	SFF-8636
QSFP28	SFF-8665	SFF-8636

SFP and QSFP cables come in passive and active variants. An active cable contains equalization or optical transceiver components. For cards that support active cables, appropriate thermal considerations shall be made per Section 6. A passive direct-attach, or RJ45 cable are purely passive elements.

Additional combinations and connector types (such as SFP-DD and QSFP-DD) are permissible as I/O form factor technologies and thermal capabilities evolve.

2.3 Top Level Assembly (SFF and LFF)

The images in Figure 13 illustrate the exploded top level assemblies for both the SFF and the LFF.

Figure 13: PBA Exploded Views (SFF and LFF)




Diagram callouts #8, and #11 through #15 are installed at the NIC assembly level:

Item #8 and #11 – Wave washer and bushing are part of the ejector latch mechanism.

Item #12 & #13 – Screws used to attach the faceplate assembly to the OCP NIC 3.0 PBA.

Item #14 – 2x SMT nuts installed on to the PBA assembly using the reflow process.

Item #15 – Insulator is located on the secondary side and is installed on the PBA prior to the faceplate.


2.4 Faceplate Subassembly (SFF and LFF)

The following section define the generic SFF and LFF faceplates.

2.4.1 Faceplate Subassembly – Exploded View

The images in Figure 14 illustrate the three faceplates subassemblies as exploded views. The bill of materials is shown in Section 2.4.2.

Figure 14: Faceplate Assembly Exploded Views (SFF and LFF)



2.4.2 Faceplate Subassembly – Bill of Materials (BOM)

Table 9 shows the bill of materials for the SFF and LFF assemblies. Item number call outs align with the SFF and LFF numbering of Figure 14.

Note: Dimensionally identical equivalent parts and equivalent materials may be substituted in the assembly. Substituted parts and materials shall meet or exceed the tolerances and requirements specified by the supplier part numbers of Table 9. Refer to the 3D CAD files for hardware specifics not covered by this table.

Table 9: Bill of Materials for the SFF and LFF Faceplate Assemblies

Item #	Item description	Part Number / Drawing	Supplier
1A 1B 1C 1D	Faceplate	See Section 2.4.3: <u>1A NIC_OCPv3_SFF_Faceplate_Pulltab_20190303.pdf</u> <u>1B NIC_OCPv3_SFF_Faceplate_Latch_20190719.pdf</u> 1C NIC_OCPv3_SFF_Faceplate_IntLatch_20190303.pdf See Section 2.4.4: 1D NIC_OCPv3_LFF_Faceplate_Latch_20190719.pdf	Custom
2A 2B 2C 2D 2E	Top and Bottom EMI Fingers	2A LT18CJ1921 – 13 fingers (Laird) TF187VE32F11-2.41-08 (Tech-Etch) 2B LT18CJ1920 – 11 fingers (Laird) TF187VE32F11-2.04-08 (Tech-Etch) 2C LT18CJ1923 – 27 fingers (Laird) TF187VE32F11-5.03-08 (Tech-Etch) 2D LT18CJ1922 – 25 fingers (Laird) TF187VE32F11-4.66-08 (Tech-Etch) 2E LT18CJ4430 – 9 fingers (Laird) TF187VE32F11-1.67-08 (Tech-Etch)	Laird, Tech-ETCH
3	Rivet	1-AC-2424-01	Dong Guan KSETT Hardware Technology
4	Side EMI Fingers	LT18DP1911	Laird
5	Thumbscrew	4C-99-343-K081	Southco, Inc.
6A 6B	Pull tab w/2x screws	CN-99-459	Southco, Inc.
8	Clinch Nut	See Section 2.4.8 and drawing NIC_OCPv3_ClinchNut_20190719.pdf	Custom
9A 9B	Ejector Handle	SFF Ejector: See Section 2.4.5 and drawing 9A NIC_OCPv3_EjectorHandle_Short_20190719.pdf Note: The SFF ejector is also used on the LFF non-shadowed I/O faceplate configuration. LFF Ejector: See Section 2.4.6 & Drawing 9B NIC_OCPv3_EjectorHandle_Long_20190719.pdf	Custom
10	Ejector Lock	See Section 2.4.7 and drawing NIC_OCPv3_EjectorLock_20190719.pdf	Custom
11	N/A	Item no longer used for drawing call out.	N/A
12*	Screw for securing faceplate to NIC	Phillips screw: ICMMS200403N Torx screw: ITMMAE200402X	WUJIANG Screw Tech Precision Industry
13*	Screw for attaching faceplate and ejector to NIC	Phillips screw: FCMMQ200503N Torx screw: FTMMC200502X	WUJIANG Screw Tech Precision Industry

14	SMT nut (on NIC)	82-950-22-010-05-RL	Fivetech Technology Inc.
15A 15B	Insulator	Refer to Section 2.7 for the SFF (15A) and LFF (15B) insulator mechanical requirements	Custom

Note: * Phillips and Torx screws are allowed. Head types should not be mixed on the same assembly.

DRAFT FOR REVIEW ONLY

2.4.3 SFF Generic I/O Faceplate

Figure 15 shows the standard SFF I/O bracket with a thumbscrew and pull tab assembly.

Figure 15: SFF Generic I/O Faceplate with Pulltab Version (2D View)




Figure 16: SFF Generic I/O Faceplate – Ejector Version (2D View)






Figure 17: SFF Generic I/O Faceplate – Internal Lock Version (2D View)



1. DIMENSION ARE IN MILLIMETERS
2. MATERIAL: 0.8mm CRS
3. TOLERANCE UNLESS OTHERWISE SPECIFIED: $\pm 0.25\text{mm}$
4. TOLERANCE SUPPORTS ASSEMBLY CTF'S. OTHER COMBINATIONS OF TOLERANCES ARE ALSO POSSIBLE TO MEET CTF'S.
5. BEND RELIEFS AS REQUIRED. ALL CRITICAL DIMENSIONS MUST BE MAINTAINED.
6. PARTS MUST COMPLY WITH ROHS DIRECTIVE 2002/95/EC. FURTHERMORE, THE USE OF HEXAVALENT CHROMIUM IN THE FABRICATION PROCESS IS NOT ALLOWED.

2.4.4 LFF Generic I/O Faceplate

Figure 18: LFF Generic I/O Faceplate – Ejector Version (2D View)




1. DIMENSION ARE IN MILLIMETER
2. MATERIAL: 0.8mm CRS
3. TOLERANCE UNLESS OTHERWISE SPECIFIED: $\pm 0.25\text{mm}$, MAX BEND RADIUS 0.30mm
4. TOLERANCE SUPPORTS ASSEMBLY CTF'S. OTHER COMBINATIONS OF TOLERANCES ARE ALSO POSSIBLE TO MEET CTF'S.
5. BEND RELIEFS ALLOWED. FACEPLATE GEOMETRY MUST BE MAINTAINED. ANY DEVIATIONS MUST BE APPROVED.
6. PARTS MUST COMPLY WITH ROHS DIRECTIVE 2002/95/EC. FURTHERMORE, THE USE OF HEXAVALENT CHROMIUM IN THE FABRICATION PROCESS IS NOT ALLOWED.

2.4.5 Ejector Lever (SFF)

This section defines the SFF lever dimensions. Note: this SFF ejector lever is also used on the non-shadowed LFF faceplate configuration.

Figure 19: SFF I/O Faceplate – Ejector Lever (2D View)




1. DIMENSION ARE IN MILLIMETERS
2. MATERIAL: 0.8mm 301 SS 1/4 HARD
3. TOLERANCE UNLESS OTHERWISE SPECIFIED: $\pm 0.25\text{mm}$, $\pm 1.0^\circ$
4. PARTS MUST COMPLY WITH RoHS DIRECTIVE 2002/95/EC. FURTHERMORE, THE USE OF HEXAVALENT CHROMIUM IN THE FABRICATION PROCESS IS NOT ALLOWED.

DRAFT

2.4.6 Ejector Levers (LFF)

This section defines the LFF ejector lever dimensions.

Figure 20: LFF I/O Faceplate – Ejector Lever (2D View)




1. DIMENSION ARE IN MILLIMETERS
2. MATERIAL: 0.8mm 301 SS 1/4 HARD
3. TOLERANCE UNLESS OTHERWISE SPECIFIED: $\pm 0.25\text{mm}$, $\pm 1.0^\circ$
4. PARTS MUST COMPLY WITH RoHS DIRECTIVE 2002/95/EC. FURTHERMORE, THE USE OF HEXAVALENT CHROMIUM IN THE FABRICATION PROCESS IS NOT ALLOWED.

DRAFT FOR REVIEW

2.4.7 Ejector Lock (SFF and LFF)

The SFF and LFF ejector uses a locking mechanism at the end of the handle to retain the lever position. This is shown in Figure 21.

Figure 21: Ejector Lock



1. DIMENSION ARE IN MILLIMETERS
2. MATERIAL: 0.3mm 301 SS 1/2 HARD
3. TOLERANCE UNLESS OTHERWISE SPECIFIED: $\pm 0.25\text{mm}$, $\pm 1.0^\circ$
4. PARTS MUST COMPLY WITH RoHS DIRECTIVE 2002/95/EC. FURTHERMORE, THE USE OF HEXAVALENT CHROMIUM IN THE FABRICATION PROCESS IS NOT ALLOWED.

DRAFT FOR

2.4.8 Clinch Nut (SFF and LFF)

The SFF and LFF card ejector handle uses a clinch nut as a spacer and rotation anchor. The clinch nut binds the ejector handle to the faceplate. Two clinch nut options are available to accommodate supplier manufacturing processes. These are shown in Figure 22 and Figure 23.

Figure 22: Clinch Nut Option A






Figure 23: Clinch Nut Option B



2.5 Card Keep Out Zones

2.5.1 SFF Keep Out Zones

Figure 24: SFF Keep Out Zone – Top View



- NOTES:
- UNLESS OTHERWISE SPECIFIED, DIMENSIONS ARE IN MILLIMETERS
 - TOLERANCE UNLESS OTHERWISE SPECIFIED: ±0.13, ±1°
 - ANY BREAK OFF FEATURES MUST ADHERE TO SPECIFIED DIMENSION AND TOLERANCE.
 - TOLERANCE SUPPORTS ASSEMBLY CTF'S. OTHER COMBINATIONS OF TOLERANCES ARE ALSO POSSIBLE.

DRAFT

Figure 25: SFF Keep Out Zone – Top View – Detail A






Figure 26: SFF Keep Out Zone – Bottom View




- NOTES:
1. UNLESS OTHERWISE SPECIFIED, DIMENSIONS ARE IN MILLIMETERS
 2. TOLERANCE UNLESS OTHERWISE SPECIFIED: $\pm 0.13, \pm 1.0^\circ$

Figure 27: SFF Keep Out Zone – Side View



Note: The area defined by DIMENSION "A" is between the faceplate and the primary side of the board. Grounded through-hole mounting pins are permitted but should avoid making contact with the latching mechanism and EMI fingers. Signal pins are not permitted in this area.

Figure 28: SFF Keep Out Zone – Side View – Detail D



2.5.2 LFF Keep Out Zones

Figure 29: LFF Keep Out Zone – Top View






Figure 30: LFF Keep Out Zone – Top View – Detail A




DRAFT FOR REVIEW ONLY

Figure 31: LFF Keep Out Zone – Bottom View




- NOTES:
1. UNLESS OTHERWISE SPECIFIED, DIMENSIONS ARE IN MILLIMETERS
 2. TOLERANCE UNLESS OTHERWISE SPECIFIED: $\pm 0.13, \pm 1.0^\circ$

Figure 32: LFF Keep Out Zone – Side View



- NOTES:
1. UNLESS OTHERWISE SPECIFIED, DIMENSIONS ARE IN MILLIMETERS
 2. TOLERANCE UNLESS OTHERWISE SPECIFIED: $\pm 0.13, \pm 1.0^\circ$
 3. PCBA SHALL ADHERE TO IPC-A-600G, CH. 2-11 FLATNESS: FOR PRINT BOARDS USING SURFACE MOUNT COMPONENTS THE BOW AND TWIST SHALL BE 0.75% OR LESS. THE FLATNESS MEASUREMENT IS DEFINED ACCORDING TO IPC-TM-650 2.4.22

Figure 33: LFF Keep Out Zone – Side View – Detail D



2.6 Baseboard Keep Out Zones

Refer to the 3D CAD files for the baseboard keep out zones for both the SFF and LFF designs. The 3D CAD files are available for download on the OCP NIC 3.0 Wiki:

<http://www.opencompute.org/wiki/Server/Mezz>

DRAFT FOR REVIEW ONLY

2.7 Insulation Requirements

All OCP NIC 3.0 cards shall implement an insulator to prevent the bottom side card components from shorting out to the baseboard chassis. The recommended insulator thickness is 0.25 mm and shall reside within the following mechanical envelope for the SFF and LFF. An alternate insulator thickness of 0.127 mm is permitted. The total stack up height of the secondary side components, insulator and the labels shall not exceed the 2 mm keep-in dimension as shown in Section 2.7.1 and 2.7.2.

A maximum of four circular holes with a 4 mm maximum diameter are permitted for access to mechanical retention components. Any components exposed by these holes must be non-conductive.

2.7.1 SFF Insulator

Figure 34: SFF Bottom Side Insulator (3D View)





Figure 35: SFF Bottom Side Insulator (Top and Side View)




1. DIMENSIONS ARE IN MILLIMETERS
2. MATERIAL:
 1. FORMEX GK-10BK or FORMEX N3 (HALOGEN FREE, BLACK OR NATURAL), 0.25mm THICKNESS
 2. FORMEX GK-5BK (HALOGEN FREE, BLACK), 0.127mm THICKNESS
3. MATERIAL SHALL CONFORM TO 94VTM-0
4. ADHESIVE 3M 467MP 0.05mm THICKNESS
5. FRONT EDGE OF INSULATOR AND ADHESIVE MAY EXTEND TO 107.60MM IF NECESSARY TO PROTECT LEDS ON SECONDARY SIDE OF PCBA.
6. TOLERANCE UNLESS OTHERWISE SPECIFIED: $\pm 0.30\text{mm}$, $\pm 1.0^\circ$

DRAFT

Figure 36: SFF Bottom Side Insulator (alternate) (Top and Side View)

ALTERNATE OPTION USING BUMPER & ADHESIVE FOR ATTACHING THE INSULATOR TO PCB INSTEAD OF USING FRONT ADHESIVE SURFACE. ENSURE CORRESPONDING SURFACE ON THE PCB IS FLAT.



2.7.2 LFF Insulator

Figure 37: LFF Bottom Side Insulator (3D View)





Figure 38: LFF Bottom Side Insulator (Top and Side View)




1. DIMENSION ARE IN MILLIMETERS
2. MATERIAL: FORMEX GK-10BK, FORMEX N3 (HALOGEN FREE, BLACK OR NATURAL), 0.25mm THICKNESS
3. MATERIAL SHALL CONFORM TO 94VTM-0
4. ADHESIVE 3M 467MP 0.05mm THICKNESS
5. FRONT EDGE OF INSULATOR AND ADHESIVE MAY EXTEND TO 107.60MM IF NECESSARY TO PROTECT LEDS ON SECONDARY SIDE OF PCBA.
6. TOLERANCE UNLESS OTHERWISE SPECIFIED: $\pm 0.30\text{mm}$, $\pm 1.0^\circ$

DRAI

Figure 39: LFF Bottom Side Insulator (alternate) (Top and Side View)

ALTERNATE OPTION USING BUMPER & ADHESIVE FOR ATTACHING THE INSULATOR TO PCB INSTEAD OF USING FRONT ADHESIVE SURFACE.
ENSURE CORRESPONDING SURFACE ON PCB IS FLAT



DRAFT

2.8 Critical-to-Function (CTF) Dimensions (SFF and LFF)

2.8.1 CTF Tolerances

The following CTF tolerances are used in this section and are the same for both the SFF and LFF.

Table 10: CTF Default Tolerances (SFF and LFF OCP NIC 3.0)

CTF DEFAULT TOLERANCES	
DIMENSION RANGE	TOLERANCE
	TWO PLACE DECIMALS: X.XX
LINEAR:	± 0.30
ANGULAR:	± 1.00 DEGREES
HOLE DIAMETER:	± 0.13

2.8.2 SFF Pull Tab CTF Dimensions

The following dimensions are considered critical-to-function (CTF) for each SFF OCP NIC 3.0 card with a pull tab and thumbscrew. The CTF default tolerances are shown in Section 2.8.1.

Figure 40: SFF OCP NIC 3.0 Card with Pull Tab CTF Dimensions (Top View)




Figure 41: SFF OCP NIC 3.0 Card with Pull Tab CTF Dimensions (Front View)





Figure 42: SFF OCP NIC 3.0 Card with Pull Tab CTF Dimensions (Side View)



DRAFT FOR REVIEW

2.8.3 SFF Ejector Latch CTF Dimensions

The following dimensions are considered critical-to-function (CTF) for each SFF OCP NIC 3.0 card with an ejector latch. The CTF default tolerances are shown in Section 2.8.1.

Figure 43: SFF OCP NIC 3.0 Card with Ejector CTF Dimensions (Top View)




Figure 44: SFF OCP NIC 3.0 Card with Ejector CTF Dimensions (Front View)





Figure 45: SFF OCP NIC 3.0 Card with Ejector CTF Dimensions (Side View)



2.8.4 SFF Internal Lock CTF Dimensions

The following dimensions are considered critical-to-function (CTF) for each SFF OCP NIC 3.0 card with an internal lock. The CTF default tolerances are shown in Section 2.8.1.

Figure 46: SFF OCP NIC 3.0 Card with Internal Lock CTF Dimensions (Top View)




Figure 47: SFF OCP NIC 3.0 Card with Internal Lock CTF Dimensions (Front View)





Figure 48: SFF OCP NIC 3.0 Card with Internal Lock CTF Dimensions (Side View)



2.8.5 SFF Baseboard CTF Dimensions

The following dimensions are considered critical-to-function (CTF) for each SFF baseboard chassis. The CTF default tolerances are shown in Section 2.8.1.

Note: The SFF baseboard CTF dimensions are applicable to both the right angle and straddle mount connector configurations. The faceplate opening relative to the baseboard changes due to the connector vertical offset, but all CTF dimensions remain identical.

Figure 49: SFF Baseboard Chassis CTF Dimensions (Rear View)




Figure 50: SFF Baseboard Chassis to Card Thumb Screw CTF Dimensions (Side View)




Figure 51: SFF Baseboard Chassis to Ejector lever Card CTF Dimensions (Side View)




Figure 52: SFF Baseboard Chassis CTF Dimensions (Rear Rail Guide View)





Figure 53: SFF Baseboard Chassis CTF Dimensions (Rail Guide Detail) – Detail C



DETAIL C

The right angle and straddle mount card guides are identical between the SFF and LFF cards. The card guide model is included in the 3D CAD packages and may be downloaded from the OCP NIC 3.0 Wiki site: <http://www.opencompute.org/wiki/Server/Mezz>.

2.8.6 LFF Ejector Latch CTF Dimensions

The following dimensions are considered critical-to-function (CTF) for each LFF OCP NIC 3.0 card. The CTF default tolerances are shown in Section 2.8.1.

Figure 54: LFF OCP NIC 3.0 Card with Ejector CTF Dimensions (Top View)




Figure 55: LFF OCP NIC 3.0 Card with Ejector CTF Dimensions (Front View)




Figure 56: LFF OCP NIC 3.0 Card with Ejector CTF Dimensions (Side View)



2.8.7 LFF Baseboard CTF Dimensions

The following dimensions are considered critical-to-function (CTF) for each LFF baseboard chassis. The CTF default tolerances are shown in Section 2.8.1.

Note: The LFF baseboard CTF dimensions are applicable to both the right angle and straddle mount connector configurations. The faceplate opening relative to the baseboard changes due to the connector vertical offset, but all CTF dimensions remain identical.

Figure 57: LFF Baseboard Chassis CTF Dimensions (Rear View)




Figure 58: LFF Baseboard Chassis CTF Dimensions (Side View)






Figure 59: LFF Baseboard Chassis CTF Dimensions (Rail Guide View)



SECTION B-B

Figure 60: LFF Baseboard Chassis CTF Dimensions (Rail Guide – Detail C)

DRAFT



DETAIL C

The right angle and straddle mount card guides are identical between the SFF and LFF. The card guide models are included in the 3D CAD packages and may be downloaded from the OCP NIC 3.0 Wiki site: <http://www.opencompute.org/wiki/Server/Mezz>.

DRAFT FOR REVIEW ONLY

2.9 Labeling Requirements

OCP NIC 3.0 cards shall implement all (or a subset of) label items listed below as required by each customer. All labels shall be placed on the exposed face of the insulator and within their designated zones. All labels shall be placed within the insulator edge and insulator bend lines to prevent labels from peeling or interfering with the faceplate, chassis card guides and card gold finger edge.


The insulator shall be divided into three different zones:

- **Regulatory Zone** – Used for all regulatory markings and filing numbers
- **Customer Zone** – Used for manufacturer markings or any ODM specific labels
- **OCP NIC 3.0 Zone** – Used for MAC addresses, part number labels and optionally the board serial number label if there are no manufacturer requirements to place it on the primary side

Notes:

- Some NIC vendor(s) may require serial number labels to be placed on the primary side of the PBA. This is permitted but it is up to the NIC vendor(s) to find the appropriate location(s) to affix the label. If a label is to be adhered to the PCB, then the label must be ESD safe as defined by ANSI/ESD S541-2008 (between 10^4 and 10^{11} Ohms).
- Regulatory marks may be printed on the insulator or affixed via a label
- Each zone size shall be adjustable to accommodate each vendor's labeling requirements
- All labels shall be oriented and readable in the same direction. The readable direction should be with the line side I/O interfaces facing "up"
- Additional labels may be placed on the primary side or on the PCB itself. This is up to the NIC vendor(s) to find the appropriate location(s)

Figure 61: SFF Label Area Example



2.9.1 General Guidelines for Label Contents

Each board shall have a unique label for identification. The label information shall be both in human readable and machine readable formats (linear or 2D data matrix). The labels may include:

- Serial number
- Part Number
- MAC Address
- Manufacturing Date
- Manufacturing Site Information

Barcode Requirements

- Linear Barcodes
- Code 93, Code 128 Auto or Code 128 Subset B
- Minimum narrow bar width $X \geq 5$ mil (0.127 mm)
- 2D data matrix
- Data matrix shall use ECC200 error correction
- Minimum cell size $X \geq 10$ mil (0.254 mm)
- All linear barcode and data matrix labels shall meet the contrast and print growth requirements per ISO/IEC 16022
- All linear barcode and data matrix labels shall have a quality level C or higher per ISO/IEC 15415
- All linear barcode and data matrix labels shall define a minimum Quiet Zone (QZ) to ensure the label is correctly registered by the scanner per ISO/IEC 15415
- Linear barcode labels shall use a QZ that is 10 times the width of the narrowest bar or 1/8th inch, whichever is greater
- Data matrix labels shall have a Quiet Zone (QZ) that is at least one module (in the X-dimension) around the perimeter of the data matrix
- Multiple Serial Numbers, MAC address may exist in one 2D data matrix, each separated by a comma

Human Readable Font

- Arial or printer font equivalent
- Minimum 5 point font size. 3 point font is acceptable when using 600 DPI printers
- Text must be easily legible under normal lighting 6-to-8 inches away

The label size and typeface may vary based on each vendor and/or customer's label content and requirements.

2.9.2 MAC Address Labeling Requirements

For an OCP NIC 3.0 card with m line side interfaces and n RBT management interfaces, the MAC address label shall list the MAC addresses in sequential order starting with line side port 1 to port m followed by the controller #0 MAC address to controller n . For cards that support multi-host configurations, the label shall associate each MAC address with a host number. The examples below show the MAC addresses presented as a single column, for labels with many MAC addresses, the label may also be formatted in multiple columns for greater readability. Labels may optionally use ellipses for the human readable text to individually indicate a range of Port and ME MAC addresses. In all cases, the 2D data matrix shall include all the card assigned MAC addresses.

2.9.2.1 MAC Address Label Example 1 – Quad Port with Single Host, Single Managed Controller

As an example, the label content of a quad SFP OCP NIC 3.0 card with a single management MAC address shall be constructed to show human readable data per the Label Data column of Table 11. The constructed label is shown in Figure 62. For each human readable line, there is a MAC prefix “Px:” for a line side Port, or “MEx:” for a managed controller instance, followed by the MAC address. The port/controller association for each row is shown in the far right column.

Table 11: MAC Address Label Example 1 – Quad Port with Single Host, Single Managed Controller

Label Data	MAC Prefix	MAC Address	Association
P1: AA:BB:CC:DD:EE:F0	P1:	AA:BB:CC:DD:EE:F0	Port 1
P2: AA:BB:CC:DD:EE:F1	P2:	AA:BB:CC:DD:EE:F1	Port 2
P3: AA:BB:CC:DD:EE:F2	P3:	AA:BB:CC:DD:EE:F2	Port 3
P4: AA:BB:CC:DD:EE:F3	P4:	AA:BB:CC:DD:EE:F3	Port 4
ME1: AA:BB:CC:DD:EE:F4	ME1:	AA:BB:CC:DD:EE:F4	Controller #0

Figure 62: MAC Address Label Example 1 – Quad Port with Single Host, Single Managed Controller




2.9.2.2 MAC Address Label Example 2 – Octal Port with Single Host, Dual Managed Controllers

As a second example, the label content of an octal port (2xQSFP with “breakout” support) OCP NIC 3.0 card with two managed silicon instances is constructed per Table 12. The constructed label is shown in Figure 63. The MAC address label shall also list the four MAC addresses associated with QSFP lanes [1:4] for QSFP connectors that allow “breakout” modes. The Host-MAC address presentation may also be formatted horizontally for easier readability.

Table 12: MAC Address Label Example 2 – Octal Port with Single Host, Dual Managed Controller

Label Data	MAC Prefix	MAC Address	Association
P1: AA:BB:CC:DD:EE:F0	P1:	AA:BB:CC:DD:EE:F0	QSFP1, Port 1
P2: AA:BB:CC:DD:EE:F1	P2:	AA:BB:CC:DD:EE:F1	QSFP1, Port 2
P3: AA:BB:CC:DD:EE:F2	P3:	AA:BB:CC:DD:EE:F2	QSFP1, Port 3
P4: AA:BB:CC:DD:EE:F3	P4:	AA:BB:CC:DD:EE:F3	QSFP1, Port 4
P5: AA:BB:CC:DD:EE:F4	P5:	AA:BB:CC:DD:EE:F4	QSFP2, Port 5
P6: AA:BB:CC:DD:EE:F5	P6:	AA:BB:CC:DD:EE:F5	QSFP2, Port 6
P7: AA:BB:CC:DD:EE:F6	P7:	AA:BB:CC:DD:EE:F6	QSFP2, Port 7
P8: AA:BB:CC:DD:EE:F7	P8:	AA:BB:CC:DD:EE:F7	QSFP2, Port 8
ME1: AA:BB:CC:DD:EE:F8	ME1:	AA:BB:CC:DD:EE:F8	Controller #0
ME2: AA:BB:CC:DD:EE:F9	ME2:	AA:BB:CC:DD:EE:F9	Controller #1

Figure 63: MAC Address Label Example 2 – Octal Port with Single Host, Dual Managed Controller



2.9.2.3 MAC Address Label Example 3 – Quad Port with Dual Hosts, Dual Managed Controllers

For multi-host implementations, each MAC address shall be prefixed with the host association “Hx” prior to the port number, where x represents the host number. An example of this is shown in Table 13 and Figure 64.

Table 13: MAC Address Label Example 3 – Quad Port with Dual Hosts, Dual Managed Controller

Label Data	Host	MAC Prefix	MAC Address	Association
P1: AA:BB:CC:DD:EE:F0	H1	P1:	AA:BB:CC:DD:EE:F0	Port 1
P2: AA:BB:CC:DD:EE:F1	H1	P2:	AA:BB:CC:DD:EE:F1	Port 2
P3: AA:BB:CC:DD:EE:F2	H2	P3:	AA:BB:CC:DD:EE:F2	Port 3
P4: AA:BB:CC:DD:EE:F3	H2	P4:	AA:BB:CC:DD:EE:F3	Port 4
ME1: AA:BB:CC:DD:EE:F4	n/a	ME1:	AA:BB:CC:DD:EE:F4	Controller #0
ME2: AA:BB:CC:DD:EE:F5	n/a	ME2:	AA:BB:CC:DD:EE:F5	Controller #1

Figure 64: MAC Address Label Example 3 – Quad Port with Dual Hosts, Dual Managed Controllers



2.9.2.4 MAC Address Label Example 4 – Single Port with Quad Host, Single Managed Controller

The following example shows a single port device with quad hosts. To conserve space on the MAC address label, this example only shows the MAC addresses for Port 1 through Port 4. The MAC address for each managed host is Px+1. This is shown in Table 14 and Figure 65.

Table 14: MAC Address Label Example 4 – Single Port with Quad Host, Single Managed Controller

Label Data	Host	MAC Prefix	MAC Address	Association
P1: AA:BB:CC:DD:EE:F0	H1	P1:	AA:BB:CC:DD:EE:F0	Port 1
ME1: AA:BB:CC:DD:EE:F1	ME1	P1:	AA:BB:CC:DD:EE:F1	Port 1
P2: AA:BB:CC:DD:EE:F2	H2	P1:	AA:BB:CC:DD:EE:F2	Port 1
ME2: AA:BB:CC:DD:EE:F3	ME2	P1:	AA:BB:CC:DD:EE:F3	Port 1
P3: AA:BB:CC:DD:EE:F4	H3	P1:	AA:BB:CC:DD:EE:F4	Port 1
ME3: AA:BB:CC:DD:EE:F5	ME3	P1:	AA:BB:CC:DD:EE:F5	Port 1
P4: AA:BB:CC:DD:EE:F6	H4	P1:	AA:BB:CC:DD:EE:F6	Port 1
ME4: AA:BB:CC:DD:EE:F7	ME4	P1:	AA:BB:CC:DD:EE:F7	Port 1

Figure 65: MAC Address Label Example 4 – Single Port with Quad Host, Single Managed Controller




Table 15: MAC Address Label Example 5 – Octal Port with Single Host, Octal Managed Controller

Label Data	MAC Prefix	MAC Address	Association
P1: AA:BB:CC:DD:EE:F0	P1:	AA:BB:CC:DD:EE:F0	Port 1
P8: AA:BB:CC:DD:EE:F7	P8:	AA:BB:CC:DD:EE:F7	Port 8
ME1: AA:BB:CC:DD:EE:F8	ME1:	AA:BB:CC:DD:EE:F8	P1
ME8: AA:BB:CC:DD:EE:FF	ME8:	AA:BB:CC:DD:EE:FF	P8

Figure 66: MAC Address Label Example 5 – Octal Port with Single Host, Octal Managed Controller




Figure 66 illustrates the use of ellipses in the human readable text to individually indicate a range of Port and Management MAC addresses. In this example, the 2D data matrix shall contain all 16 MAC addresses.

2.10 Mechanical CAD Package Examples

Typical OCP NIC 3.0 implementation examples are included in the 3D CAD package. The purpose of these examples is to demonstrate the implementation feasibility. Additional use cases beyond the implementation examples are possible as long they adhere to the OCP NIC 3.0 specification.

Note: For brevity, references to QSFP+, and QSFP28 shall be referred to as QSFP in this document. Similarly, references to SFP+, and SFP28 shall be referred to as SFP.

The 3D CAD files may be obtained from the OCP NIC 3.0 Wiki:
<http://www.opencompute.org/wiki/Server/Mezz>

Table 16: NIC Implementation Examples and 3D CAD

Implementation Example	3D CAD File name
SFF Single/Dual QSFP ports	01_nic_v3_sff2q_1tab_asm.stp 01_nic_v3_sff2q_latch_asm.stp
SFF Single/Dual SFP ports	N/A
SFF Quad SFP ports	01_nic_v3_sff4s_1tab_asm.stp 01_nic_v3_sff4s_latch_asm.stp
SFF Quad 10GBASE-T ports	01_nic_v3_sff4r_1tab_asm.stp 01_nic_v3_sff4r_latch_asm.stp
LFF Single/Dual QSFP ports	01_nic_v3_lff2q_asm.stp
LFF Single/Dual SFP ports	N/A
LFF Quad SFP ports	01_nic_v3_lff4s_asm.stp
LFF Quad 10GBASE-T ports	01_nic_v3_lff4r_asm.stp

3 Electrical Interface Definition – Card Edge and Baseboard

3.1 Card Edge Gold Finger Requirements

The OCP NIC 3.0 cards are compliant to the SFF-TA-1002 specification with respect to the gold fingers and connectors.

SFF cards fit in the Primary Connector. Primary Connector compliant cards are 76 mm x 115 mm and may implement the full 168-pins. The Primary Connector cards may optionally implement a subset of gold finger pins if there is a reduced PCIe width requirement (such as 1 x8 and below). In this case, the card edge gold finger may implement a 2C+ design. The overall board thickness is 1.57 mm. The gold finger dimensions for the Primary Connector compliant cards are shown below.

LFF cards support up to a x32 PCIe implementation and may use both the Primary 4C+ and Secondary 4C Connectors. LFF cards may implement a reduced PCIe lane count and optionally implement only the Primary Connector.

Note: The “B” pins on the connector are associated with the top side of the OCP NIC 3.0 card. The “A” pins on the connector are associated with the bottom side of the OCP NIC 3.0 card. The A and B side pins are physically on top of each other with zero x-axis offset.

Figure 67: SFF Primary Connector Gold Finger Dimensions – x16 – Top Side (“B” Pins)

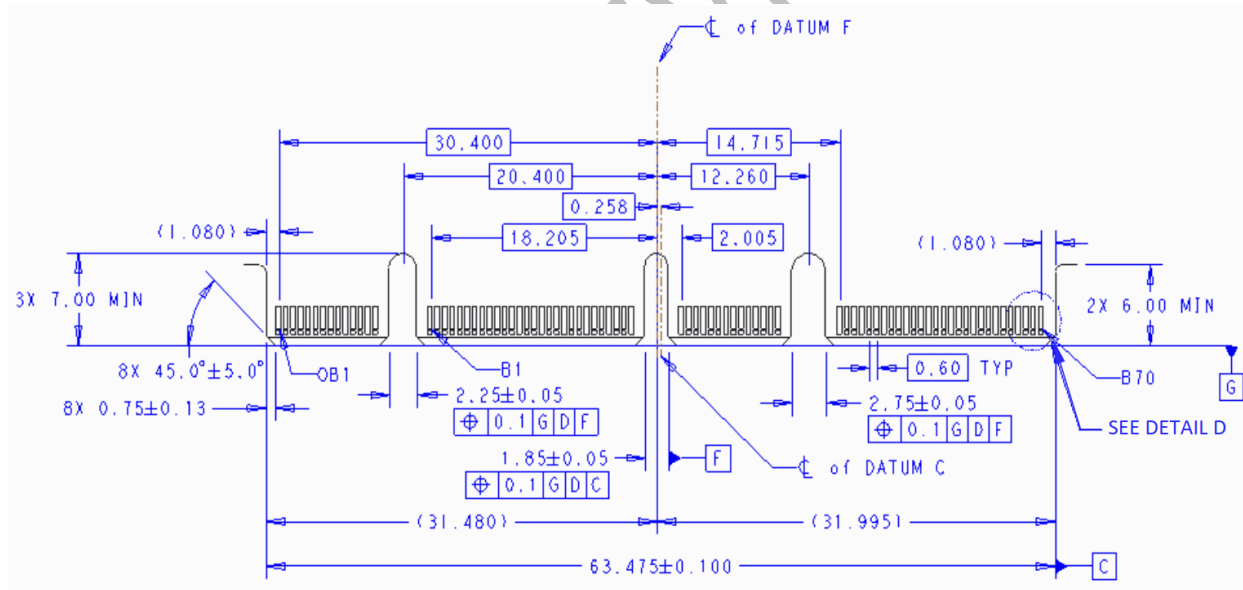


Figure 68: SFF Primary Connector Card Profile Dimensions




Figure 69: SFF Primary Connector Gold Finger - Detail D




Figure 70: LFF Gold Finger Dimensions – x32 – Top Side (“B” Pins)





Figure 71: LFF Gold Finger Dimensions – x32 – Bottom Side (“A” Pins)



3.1.1 Gold Finger Mating Sequence

Per the SFF-TA-1002 specification, the Primary and Secondary Connectors are protocol agnostic and are optimized for high speed differential pairs. For use in the OCP NIC 3.0 application, some pin locations are used for single ended control nets or power and would benefit from a shorter pin length for staggering. As such, the required OCP NIC 3.0 card gold finger staging is shown in Table 17 and Table 18 for a two stage, first-mate, last-break functionality. The two-stage finger length is a normative requirement for the OCP NIC 3.0 card. The host connectors have a single stage mating and do not implement different pin lengths.

The AIC Plug (Free) side refers to the OCP NIC 3.0 card gold fingers; the receptacle (Fixed) side refers to the physical connector on the host platform. This table is based on the SFF-TA-1002 Table A-1 with modifications for OCP NIC 3.0. Refer to the mechanical drawings for pin the first-mate and second-mate lengths.

Note: Pin names in Table 17 and Table 18 are used for first mate/second mate reference only. Full pin definitions are described in Sections 3.3 and 3.4.

Table 17: Contact Mating Positions for the Primary Connector

Side B			Side A			
Gold Finger Side (Free)		Receptacle (Fixed)	Gold Finger Side (Free)		Receptacle (Fixed)	
2 nd Mate	1 st Mate		2 nd Mate	1 st Mate		
OCP B1	NIC_PWR_GOOD		OCP A1	PERST2#		
OCP B2	MAIN_PWR_EN		OCP A2	PERST3#		
OCP B3	LD#		OCP A3	WAKE#		
OCP B4	DATA_IN		OCP A4	RBT_ARB_IN		
OCP B5	DATA_OUT		OCP A5	RBT_ARB_OUT		
OCP B6	CLK		OCP A6	SLOT_ID1		
OCP B7	SLOT_ID0		OCP A7	RBT_TX_EN		
OCP B8	RBT_RXD1		OCP A8	RBT_TXD1		
OCP B9	RBT_RXD0		OCP A9	RBT_TXD0		
OCP B10	GND		OCP A10	GND		
OCP B11	REFCLKn2		OCP A11	REFCLKn3		
OCP B12	REFCLKp2		OCP A12	REFCLKp3		

OCP B13	GND			OCP A13	GND
OCP B14	RBT_CRS_DV			OCP A14	RBT_CLK_IN
Mechanical Key					
B1	+12V_EDGE			A1	GND
B2	+12V_EDGE			A2	GND
B3	+12V_EDGE			A3	GND
B4	+12V_EDGE			A4	GND
B5	+12V_EDGE			A5	GND
B6	+12V_EDGE			A6	GND
B7	BIF0#			A7	SMCLK
B8	BIF1#			A8	SMDAT
B9	BIF2#			A9	SMRST#
B10	PERST0#			A10	PRSNTA#
B11	+3.3V_EDGE			A11	PERST1#
B12	AUX_PWR_EN			A12	PRSNTB2#
B13	GND			A13	GND
B14	REFCLKn0			A14	REFCLKn1
B15	REFCLKp0			A15	REFCLKp1
B16	GND			A16	GND
B17	PETn0			A17	PERn0
B18	PETp0			A18	PERp0
B19	GND			A19	GND
B20	PETn1			A20	PERn1
B21	PETp1			A21	PERp1
B22	GND			A22	GND
B23	PETn2			A23	PERn2
B24	PETp2			A24	PERp2
B25	GND			A25	GND
B26	PETn3			A26	PERn3
B27	PETp3			A27	PERp3
B28	GND			A28	GND
Mechanical Key					
B29	GND			A29	GND
B30	PETn4			A30	PERn4
B31	PETp4			A31	PERp4
B32	GND			A32	GND
B33	PETn5			A33	PERn5
B34	PETp5			A34	PERp5
B35	GND			A35	GND
B36	PETn6			A36	PERn6
B37	PETp6			A37	PERp6
B38	GND			A38	GND
B39	PETn7			A39	PERn7
B40	PETp7			A40	PERp7
B41	GND			A41	GND
B42	PRSNTB0#			A42	PRSNTB1#
Mechanical Key					
B43	GND			A43	GND
B44	PETn8			A44	PERn8
B45	PETp8			A45	PERp8
B46	GND			A46	GND
B47	PETn9			A47	PERn9
B48	PETp9			A48	PERp9
B49	GND			A49	GND
B50	PETn10			A50	PERn10
B51	PETp10			A51	PERp10
B52	GND			A52	GND
B53	PETn11			A53	PERn11
B54	PETp11			A54	PERp11
B55	GND			A55	GND
B56	PETn12			A56	PERn12
B57	PETp12			A57	PERp12
B58	GND			A58	GND
B59	PETn13			A59	PERn13
B60	PETp13			A60	PERp13
B61	GND			A61	GND
B62	PETn14			A62	PERn14
B63	PETp14			A63	PERp14
B64	GND			A64	GND
B65	PETn15			A65	PERn15
B66	PETp15			A66	PERp15

B67	GND		A67	GND
B68	RFU1, N/C		A68	USB_DATn
B69	RFU2, N/C		A69	USB_DATp
B70	PRSNB3#		A70	PWRBRK0#

Table 18: Contact Mating Positions for the Secondary Connector

Side B				Side A			
Gold Finger Side (Free)		Receptacle (Fixed)		Gold Finger Side (Free)		Receptacle (Fixed)	
2 nd Mate	1 st Mate			2 nd Mate	1 st Mate		
B1	+12V_EDGE			A1	GND		
B2	+12V_EDGE			A2	GND		
B3	+12V_EDGE			A3	GND		
B4	+12V_EDGE			A4	GND		
B5	+12V_EDGE			A5	GND		
B6	+12V_EDGE			A6	GND		
B7	BIF0#			A7	SMCLK		
B8	BIF1#			A8	SMDAT		
B9	BIF2#			A9	SMRST#		
B10	PERST4#			A10	PRSNBTA#		
B11	+3.3V_EDGE			A11	PERST5#		
B12	AUX_PWR_EN			A12	PRSNB2#		
B13	GND			A13	GND		
B14	REFCLKn4			A14	REFCLKn5		
B15	REFCLKp4			A15	REFCLKp5		
B16	GND			A16	GND		
B17	PETn16			A17	PERn16		
B18	PETp16			A18	PERp16		
B19	GND			A19	GND		
B20	PETn17			A20	PERn17		
B21	PETp17			A21	PERp17		
B22	GND			A22	GND		
B23	PETn18			A23	PERn18		
B24	PETp18			A24	PERp18		
B25	GND			A25	GND		
B26	PETn19			A26	PERn19		
B27	PETp19			A27	PERp19		
B28	GND			A28	GND		
Mechanical Key							
B29	GND			A29	GND		
B30	PETn20			A30	PERn20		
B31	PETp20			A31	PERp20		
B32	GND			A32	GND		
B33	PETn21			A33	PERn21		
B34	PETp21			A34	PERp21		
B35	GND			A35	GND		
B36	PETn22			A36	PERn22		
B37	PETp22			A37	PERp22		
B38	GND			A38	GND		
B39	PETn23			A39	PERn23		
B40	PETp23			A40	PERp23		
B41	GND			A41	GND		
B42	PRSNB0#			A42	PRSNB1#		
Mechanical Key							
B43	GND			A43	GND		
B44	PETn24			A44	PERn24		
B45	PETp24			A45	PERp24		
B46	GND			A46	GND		
B47	PETn25			A47	PERn25		
B48	PETp25			A48	PERp25		
B49	GND			A49	GND		
B50	PETn26			A50	PERn26		
B51	PETp26			A51	PERp26		
B52	GND			A52	GND		
B53	PETn27			A53	PERn27		
B54	PETp27			A54	PERp27		
B55	GND			A55	GND		
B56	PETn28			A56	PERn28		
B57	PETp28			A57	PERp28		
B58	GND			A58	GND		

B59	PETn29		A59	PERn29
B60	PETp29		A60	PERp29
B61	GND		A61	GND
B62	PETn30		A62	PERn30
B63	PETp30		A63	PERp30
B64	GND		A64	GND
B65	PETn31		A65	PERn31
B66	PETp31		A66	PERp31
B67	GND		A67	GND
B68	RFU3, N/C		A68	UART_RX
B69	RFU4, N/C		A69	UART_TX
B70	PRSNB3#		A70	PWRBRK1#

DRAFT FOR REVIEW ONLY

3.2 Baseboard Connector Requirements

The OCP NIC 3.0 connectors are compliant to the 4C+ and 4C connectors as defined in the SFF-TA-1002 specification for a right angle or straddle mount form factor. The Primary Connector is a 4C+ implementation with 168-pins. The Secondary Connector is a 4C implementation with 140-pins. Both the Primary and Secondary Connectors includes support for up to 32 differential pairs to support a x16 PCIe connection. Each connector also provides 6 pins of +12V_EDGE, and 1 pin of +3.3V_EDGE for power. This implementation is common between both the Primary and Secondary Connectors. In addition, the 4C+ implementation of the Primary Connector has a 28-pin OCP Bay used for management and support for up to a 4 x2 and 4 x4 multi-host configuration on the Primary Connector. The Primary and Secondary Connector drawings are shown below.

All diagram units are in mm unless otherwise noted.

3.2.1 Right Angle Connector

The following offset and height options are available for the right angle Primary and Secondary Connectors.

Table 19: Right Angle Connector Options

Name	Pins	Style and Baseboard Thickness	Offset (mm)
Primary Connector – 4C+	168 pins	Right Angle	4.05 mm
Secondary Connector – 4C	140 pins	Right Angle	4.05 mm

Figure 72: 168-pin Base Board Primary Connector – Right Angle






Figure 73: 140-pin Base Board Secondary Connector – Right Angle



3.2.2 Right Angle Offset

The OCP NIC 3.0 right angle connectors have a 4.05 mm offset from the baseboard. This is shown in Figure 74.

Figure 74: OCP NIC 3.0 Card and Host Offset for Right Angle Connectors



3.2.3 Straddle Mount Connector

The following offset and height options are available for the straddle mount Primary and Secondary Connectors.

Table 20: Straddle Mount Connector Options

Name	Pins	Style and Baseboard Thickness	Offset (mm)
Primary Connector – 4C+	168 pins	Straddle Mount for 0.062"	Coplanar (0 mm)
Primary Connector – 4C+	168 pins	Straddle Mount for 0.076"	-0.3 mm
Primary Connector – 4C+	168 pins	Straddle Mount for 0.093"	Coplanar (0 mm)
Secondary Connector – 4C	140 pins	Straddle Mount for 0.062"	Coplanar (0 mm)
Secondary Connector – 4C	140 pins	Straddle Mount for 0.076"	-0.3 mm
Secondary Connector – 4C	140 pins	Straddle Mount for 0.093"	Coplanar (0 mm)

Figure 75: 168-pin Base Board Primary Connector – Straddle Mount






Figure 76: 140-pin Base Board Secondary Connector – Straddle Mount



3.2.4 Straddle Mount Offset and PCB Thickness Options

The OCP NIC 3.0 straddle mount connectors have three baseboard PCB thicknesses they can accept. The available options are shown in Figure 77. The thicknesses are 0.062", 0.076", and 0.093". These PCBs must be controlled to a thickness of $\pm 10\%$. These are available for both the Primary and Secondary Connector locations. At the time of this writing, the most commonly used part is expected to be the 0.076" baseboard thickness.

Figure 77: OCP NIC 3.0 Card and Baseboard PCB Thickness Options for Straddle Mount Connectors



The connectors are capable of being used coplanar as shown in Figure 78. Additionally, the connectors are also capable of having a 0.3 mm offset from the centerline of the host board as shown in Figure 79.

Figure 78: 0 mm Offset (Coplanar) for 0.062" Thick Baseboards





Figure 79: 0.3 mm Offset for 0.076" Thick Baseboards



3.2.5 LFF Connector Locations

In order to support the LFF, systems must locate the Primary and Secondary Connectors per the mechanical drawing shown in Figure 80 and Figure 81.

Figure 80: Primary and Secondary Connector Locations for LFF Support with Right Angle Connectors





Figure 81: Primary and Secondary Connector Locations for LFF Support with Straddle Mount Connectors



3.3 Pin Definition

The pin definitions of an OCP NIC 3.0 card with up to a x32 PCIe interface are shown in Table 21 and Table 22. All signal directions are shown from the perspective of the baseboard.

A baseboard system may provide a combination of Primary Connectors only, or Primary and Secondary Connectors to support multiple sizes of OCP NIC 3.0 cards. Both connectors share common functionality with power, SMBus 2.0, x16 PCIe and bifurcation control. The Primary Connector 4C+ definition has an additional OCP Bay (pins OCP_A[1:14], OCP_B[1:14]) with additional REFCLKs for supporting up to four PCIe hosts, NC-SI over RBT connectivity and a Scan Chain for information exchange between the host and card. The NIC is required to implement the Scan Chain, while the baseboard may choose to optionally implement it. Depending on the baseboard form factor, multiple OCP NIC 3.0 compliant cards may be designed into the system.

The Primary and Secondary Connectors pins are shown in Section 3.4. The OCP Bay pins on the Primary Connector only are explicitly called out with the “OCP_” prefix in the pin location column.

Cards or systems that do not require the use of a PCIe x16 connection may optionally implement a subset of electrical connections as applicable to the design. For example, a x8 (or smaller) card using the first 8 PCIe lanes that is compliant with the Primary Connector pinout. Refer to Sections 3.1 and 3.2 for mechanical details. For these cases, the Primary Connector matches the 2C dimensions as defined in SFF-TA-1002.

In all cases, the physical baseboard connectors shall support x16 PCIe widths and must be implemented with the Primary (4C+) and Secondary (4C) connectors.

3.3.1 Primary Connector

Table 21: Primary Connector Pin Definition (x16) (4C+)

Side B		Side A		Primary Connector (4C+, x16, 168-pin OCP NIC 3.0 card with OCP Bay)	Primary Connector (2C+, x8, 112-pin OCP NIC 3.0 card with OCP bay)
OCP_B1	NIC_PWR_GOOD	PERST2#	OCP_A1		
OCP_B2	MAIN_PWR_EN	PERST3#	OCP_A2		
OCP_B3	LD#	WAKE#	OCP_A3		
OCP_B4	DATA_IN	RBT_ARB_IN	OCP_A4		
OCP_B5	DATA_OUT	RBT_ARB_OUT	OCP_A5		
OCP_B6	CLK	SLOT_ID1	OCP_A6		
OCP_B7	SLOT_ID0	RBT_TX_EN	OCP_A7		
OCP_B8	RBT_RXD1	RBT_TXD1	OCP_A8		
OCP_B9	RBT_RXD0	RBT_TXD0	OCP_A9		
OCP_B10	GND	GND	OCP_A10		
OCP_B11	REFCLKn2	REFCLKn3	OCP_A11		
OCP_B12	REFCLKp2	REFCLKp3	OCP_A12		
OCP_B13	GND	GND	OCP_A13		
OCP_B14	RBT_CRS_DV	RBT_CLK_IN	OCP_A14		
Mechanical Key					
B1	+12V_EDGE	GND	A1		
B2	+12V_EDGE	GND	A2		
B3	+12V_EDGE	GND	A3		
B4	+12V_EDGE	GND	A4		
B5	+12V_EDGE	GND	A5		
B6	+12V_EDGE	GND	A6		
B7	BIF0#	SMCLK	A7		
B8	BIF1#	SMDAT	A8		
B9	BIF2#	SMRST#	A9		
B10	PERSTO#	PRSNTA#	A10		
B11	+3.3V_EDGE	PERST1#	A11		
B12	AUX_PWR_EN	PRSNB2#	A12		
B13	GND	GND	A13		
B14	REFCLKn0	REFCLKn1	A14		
B15	REFCLKp0	REFCLKp1	A15		
B16	GND	GND	A16		
B17	PETn0	PERn0	A17		
B18	PETp0	PERp0	A18		
B19	GND	GND	A19		
B20	PETn1	PERn1	A20		
B21	PETp1	PERp1	A21		
B22	GND	GND	A22		

B23	PETn2	PERn2	A23
B24	PETp2	PERp2	A24
B25	GND	GND	A25
B26	PETn3	PERn3	A26
B27	PETp3	PERp3	A27
B28	GND	GND	A28
Mechanical Key			
B29	GND	GND	A29
B30	PETn4	PERn4	A30
B31	PETp4	PERp4	A31
B32	GND	GND	A32
B33	PETn5	PERn5	A33
B34	PETp5	PERp5	A34
B35	GND	GND	A35
B36	PETn6	PERn6	A36
B37	PETp6	PERp6	A37
B38	GND	GND	A38
B39	PETn7	PERn7	A39
B40	PETp7	PERp7	A40
B41	GND	GND	A41
B42	PRSNB0#	PRSNB1#	A42
Mechanical Key			
B43	GND	GND	A43
B44	PETn8	PERn8	A44
B45	PETp8	PERp8	A45
B46	GND	GND	A46
B47	PETn9	PERn9	A47
B48	PETp9	PERp9	A48
B49	GND	GND	A49
B50	PETn10	PERn10	A50
B51	PETp10	PERp10	A51
B52	GND	GND	A52
B53	PETn11	PERn11	A53
B54	PETp11	PERp11	A54
B55	GND	GND	A55
B56	PETn12	PERn12	A56
B57	PETp12	PERp12	A57
B58	GND	GND	A58
B59	PETn13	PERn13	A59
B60	PETp13	PERp13	A60
B61	GND	GND	A61
B62	PETn14	PERn14	A62
B63	PETp14	PERp14	A63
B64	GND	GND	A64
B65	PETn15	PERn15	A65
B66	PETp15	PERp15	A66
B67	GND	GND	A67
B68	RFU1, N/C	USB_DATn	A68
B69	RFU2, N/C	USB_DATp	A69
B70	PRSNB3#	PWRBRK0#	A70

3.3.2 Secondary Connector

Table 22: Secondary Connector Pin Definition (x16) (4C)

Side B		Side A	
B1	+12V_EDGE	GND	A1
B2	+12V_EDGE	GND	A2
B3	+12V_EDGE	GND	A3
B4	+12V_EDGE	GND	A4
B5	+12V_EDGE	GND	A5
B6	+12V_EDGE	GND	A6
B7	BIF0#	SMCLK	A7
B8	BIF1#	SMDAT	A8
B9	BIF2#	SMRST#	A9
B10	PERST4#	PRSNTA#	A10
B11	+3.3V_EDGE	PERST5#	A11
B12	AUX_PWR_EN	PRSNTB2#	A12
B13	GND	GND	A13
B14	REFCLKn4	REFCLKn5	A14
B15	REFCLKp4	REFCLKp5	A15
B16	GND	GND	A16
B17	PETn16	PERn16	A17
B18	PETp16	PERp16	A18
B19	GND	GND	A19
B20	PETn17	PERn17	A20
B21	PETp17	PERp17	A21
B22	GND	GND	A22
B23	PETn18	PERn18	A23
B24	PETp18	PERp18	A24
B25	GND	GND	A25
B26	PETn19	PERn19	A26
B27	PETp19	PERp19	A27
B28	GND	GND	A28
Mechanical Key			
B29	GND	GND	A29
B30	PETn20	PERn20	A30
B31	PETp20	PERp20	A31
B32	GND	GND	A32
B33	PETn21	PERn21	A33
B34	PETp21	PERp21	A34
B35	GND	GND	A35
B36	PETn22	PERn22	A36
B37	PETp22	PERp22	A37
B38	GND	GND	A38
B39	PETn23	PERn23	A39
B40	PETp23	PERp23	A40
B41	GND	GND	A41
B42	PRSNTB0#	PRSNTB1#	A42
Mechanical Key			
B43	GND	GND	A43
B44	PETn24	PERn24	A44
B45	PETp24	PERp24	A45
B46	GND	GND	A46
B47	PETn25	PERn25	A47
B48	PETp25	PERp25	A48
B49	GND	GND	A49

Secondary Connector (4C, x16, 140-pin OCP NIC 3.0 card)

B50	PETn26	PERn26	A50
B51	PETp26	PERp26	A51
B52	GND	GND	A52
B53	PETn27	PERn27	A53
B54	PETp27	PERp27	A54
B55	GND	GND	A55
B56	PETn28	PERn28	A56
B57	PETp28	PERp28	A57
B58	GND	GND	A58
B59	PETn29	PERn29	A59
B60	PETp29	PERp29	A60
B61	GND	GND	A61
B62	PETn30	PERn30	A62
B63	PETp30	PERp30	A63
B64	GND	GND	A64
B65	PETn31	PERn31	A65
B66	PETp31	PERp31	A66
B67	GND	GND	A67
B68	RFU3, N/C	UART_RX	A68
B69	RFU4, N/C	UART_TX	A69
B70	PRSNB3#	PWRBRK1#	A70

3.4 Signal Descriptions

The pins shown in this section are common for both the Primary and Secondary Connectors unless otherwise noted. Pins that exist only for the Primary Connector OCP Bay are explicitly called out in the pin location column with the prefix “OCP_xxx”. USB is only defined on the Primary Connector. UART is only defined on the secondary connector. All pin directions are from the perspective of the baseboard.

Note: The OCP NIC 3.0 card shall implement protection methods to prevent leakage or low impedance paths between the V_{AUX} and V_{MAIN} power domains in the event that a powered-down NIC is physically present in a powered-up baseboard. This specification provides example isolation implementations in the signal description text and appropriate figures. OCP NIC 3.0 implementers may choose to do a different implementation as long as the isolation requirements are met and the same result is achieved.

3.4.1 PCIe Interface Pins

This section provides the pin assignments for the PCIe interface signals. The PCIe signals have unique names on the Primary and Secondary connector. The Primary Connector uses the REFCLK[0:3], TX/RX[0:15], PERST[0:3] indices. The Secondary Connector uses the REFCLK[4:5], TX/RX[16:31] and PERST[4:5] indices. Where applicable, the Primary/Secondary connector naming convention is shown as a pair. The AC/DC specifications are defined in the PCIe CEM Specification, Rev 4.0. Example connection diagrams for are shown in Section 3.6.

Table 23: Pin Descriptions – PCIe

Signal Name (Primary / Secondary)	Pin #	Baseboard Direction	Signal Description
REFCLKn0/REFCLKn4	B14	Output	
REFCLKp0/REFCLKp4	B15		

REFCLKn1/REFCLKn5 REFCLKp1/REFCLKp5	A14 A15	Output	PCIe compliant differential reference clocks. 100MHz reference clocks are used for the OCP NIC 3.0 card PCIe core logic.
REFCLKn2 REFCLKp2	OCP_B11 OCP_B12	Output	<p>REFCLK0 is always available to all OCP NIC 3.0 cards. The card should not assume REFCLK1, REFCLK2 or REFCLK3 are available until the bifurcation negotiation process is complete.</p> <p>For baseboards, the REFCLK0, REFCLK1, REFCLK2 and REFCLK3 signals shall be available at the Primary Connector for supported designs. REFCLK2 and REFCLK3 are only available on the Primary connector in the OCP Bay. REFCLK4 and REFCLK5 are available on the Secondary connector.</p> <ul style="list-style-type: none"> REFCLK0 is required for all designs. REFCLK1, REFCLK2 and REFCLK3 are required for designs that support 2 xn, and 4 xn bifurcation implementations. <p>For baseboard implementations that use REFCLK[1:3], the baseboard should disable the appropriate REFCLKs not used by the OCP NIC 3.0 card.</p> <p>The baseboard shall not advertise the corresponding bifurcation modes if REFCLK[1:3] are not implemented.</p> <p>REFCLK4 and REFCLK5 are only available on the Secondary Connector and are not defined for use this specification release.</p> <p>For OCP NIC 3.0 cards, the required REFCLKs shall be connected per the endpoint datasheet. Unused REFCLKs on the OCP NIC 3.0 card shall be left as a no connect.</p> <p>Note: For cards that only support 1 x16, REFCLK0 is used. For cards that support 2 x8, REFCLK0 is used for the first eight PCIe lanes, and REFCLK1 is used for the second eight PCIe lanes. REFCLK2 and REFCLK3 are only used for cards that only support a four link PCIe bifurcation mode.</p>
REFCLKn3 REFCLKp3	OCP_A11 OCP_A12	Output	

			Refer to Section 2.1 in the PCIe CEM Specification, Rev 4.0 for electrical details.
PETn0 / PETn16 PETp0 / PETp16	B17 B18	Output	<p>Transmitter differential pairs [0:15] (Primary Connector), and differential pairs [16:31] (Secondary Connector). These pins are connected from the baseboard transmitter differential pairs to the receiver differential pairs on the OCP NIC 3.0 card.</p> <p>The PCIe transmit pins shall be AC coupled on the baseboard with capacitors. The AC coupling capacitor value shall use the C_{TX} parameter value specified in the PCIe Base Specification Rev 4.0 Section 8.3.9.</p> <p>For baseboards, the PET[0:15] signals are required at the Primary Connector for a SFF slot. PET[0:15] and PET[16:31] are required for a LFF slot.</p> <p>For SFF OCP NIC 3.0 cards, the required PET[0:15] signals shall be connected to the endpoint silicon. For silicon that uses less than a x16 connection, the appropriate PET[0:15] signals shall be connected per the endpoint datasheet.</p> <p>For LFF implementations, PET[0:15] are assigned to the Primary Connector, and PET[16:31] are assigned to the Secondary Connector.</p> <p>Refer to Section 6.1 in the PCIe CEM Specification, Rev 4.0 for details.</p>
PETn1 / PETn17 PETp1 / PETp17	B20 B21	Output	
PETn2 / PETn18 PETp2 / PETp18	B23 B24	Output	
PETn3 / PETn19 PETp3 / PETp19	B26 B27	Output	
PETn4 / PETn20 PETp4 / PETp20	B30 B31	Output	
PETn5 / PETn21 PETp5 / PETp21	B33 B34	Output	
PETn6 / PETn22 PETp6 / PETp22	B36 B37	Output	
PETn7 / PETn23 PETp7 / PETp23	B39 B40	Output	
PETn8 / PETn24 PETp8 / PETp24	B44 B45	Output	
PETn9 / PETn25 PETp9 / PETp25	B47 B48	Output	
PETn10 / PETn26 PETp10 / PETp26	B50 B51	Output	
PETn11 / PETn27 PETp11 / PETp27	B53 B54	Output	
PETn12 / PETn28 PETp12 / PETp28	B56 B57	Output	
PETn13 / PETn29 PETp13 / PETp29	B59 B60	Output	
PETn14 / PETn30 PETp14 / PETp30	B62 B63	Output	
PETn15 / PETn31 PETp15 / PETp31	B65 B66	Output	
PERn0 / PERn16 PERp0 / PERp16	A17 A18	Input	Receiver differential pairs [0:15] (Primary Connector), and differential pairs [16:31] (Secondary Connector). These pins are connected from the OCP NIC 3.0 card transmitter differential pairs to the receiver differential pairs on the baseboard.
PERn1 / PERn17 PERp1 / PERp17	A20 A21	Input	
PERn2 / PERn18 PERp2 / PERp18	A23 A24	Input	

PERn3 / PERn19 PERp3 / PERp19	A26 A27	Input	<p>The PCIe receive pins shall be AC coupled on the OCP NIC 3.0 card with capacitors. The AC coupling capacitor value shall use the C_{TX} parameter value specified in the PCIe Base Specification Rev 4.0 Section 8.3.9.</p> <p>For baseboards, the PER[0:15] signals are required at the Primary Connector for a SFF slot. PER[0:15] and PER[16:31] are required for a LFF slot.</p> <p>For SFF OCP NIC 3.0 cards, the required PER[0:15] signals shall be connected to the endpoint silicon. For silicon that uses less than a x16 connection, the appropriate PER[0:15] signals shall be connected per the endpoint datasheet.</p> <p>For LFF implementations, PER[0:15] are assigned to the Primary Connector, and PER[16:31] are assigned to the Secondary Connector.</p> <p>Refer to Section 6.1 in the PCIe CEM Specification, Rev 4.0 for details.</p>	
PERn4 / PERn20 PERp4 / PERp20	A30 A31	Input		
PERn5 / PERn21 PERp5 / PERp21	A33 A34	Input		
PERn6 / PERn22 PERp6 / PERp22	A36 A37	Input		
PERn7 / PERn23 PERp7 / PERp23	A39 A40	Input		
PERn8 / PERn24 PERp8 / PERp24	A44 A45	Input		
PERn9 / PERn25 PERp9 / PERp25	A47 A48	Input		
PERn10 / PERn26 PERp10 / PERp26	A50 A51	Input		
PERn11 / PERn27 PERp11 / PERp27	A53 A54	Input		
PERn12 / PERn28 PERp12 / PERp28	A56 A57	Input		
PERn13 / PERn29 PERp13 / PERp29	A59 A60	Input		
PERn14 / PERn30 PERp14 / PERp30	A62 A63	Input		
PERn15 / PERn31 PERp15 / PERp31	A65 A66	Input		
PERST0# / PERST4# PERST1# / PERST5# PERST2# PERST3#	B10 A11 OCP_A1 OCP_A2	Output		<p>PCIe Reset #[0:5]. Active low.</p> <p>When PERSTn# is deasserted, the signal shall indicate the power state is already in Main Power Mode and is within tolerance and stable for the OCP NIC 3.0 card to bring up the PCIe link.</p> <p>PERST# shall be deasserted at least 1s after the NIC_PWR_GOOD assertion to Main Power Mode. This ensures the card power rails are within the operating limits. This value is longer than the minimum value specified in the PCIe CEM Specification. The PCIe REFCLKs shall also become stable within this period of time.</p> <p>PERST[0:5]# shall be asserted low on the baseboard until the platform is ready to deassert reset.</p> <p>For baseboards that support bifurcation, the PERST[0:3]# signals are required at the Primary</p>

			<p>Connector, PERST[4:5]# are required at the Secondary Connector.</p> <p>For OCP NIC 3.0 cards, the required PERST[0:5]# signals shall be connected to the endpoint silicon. Unused PERST[0:5]# signals shall be left as a no connect.</p> <p>Note: For cards that only support 1 x16, PERST0# is used. For cards that support 2 x8, PERST0# is used for the first eight PCIe lanes, and PERST1# is used for the second eight PCIe lanes. PERST2# and PERST3# are only used for cards that support a four link PCIe bifurcation mode.</p> <p>PERST0# is always available to all OCP NIC 3.0 cards. The card should not assume PERST[1:5]# are available until the bifurcation negotiation process is complete.</p> <p>Refer to Section 2.2 in the PCIe CEM Specification, Rev 4.0 for details.</p>
WAKE#	OCP_A3	Input, OD	<p>WAKE#. Open drain. Active low.</p> <p>This signal shall be driven by the OCP NIC 3.0 card to notify the baseboard to restore PCIe link. For OCP NIC 3.0 cards that support multiple WAKE# signals, their respective WAKE# pins may be tied together as the signal is open-drain to form a wired-OR. For multi-homed host configurations, the WAKE# signal assertion shall wake all nodes.</p> <p>For baseboards, this signal shall be pulled up to +3.3V_EDGE on the baseboard with a 10 kOhm resistor. This signals shall be connected to the system WAKE# signal.</p> <p>For OCP NIC 3.0 cards, this signal shall be connected between the endpoint silicon WAKE# pin(s) and the card edge through an isolation buffer. The WAKE# signal shall not assert until the PCIe card is in the D3 state according to the PCIe CEM specification to prevent false WAKE# events. For OCP NIC 3.0, the WAKE# pin shall be buffered or otherwise isolated from the host until the aux voltage source is present. Examples of this are shown in Section 3.5.5 by gating via an on-board “AUX_PWR_GOOD” signal to indicate all the NIC</p>

			<p>AUX power rails are stable. The PCIe CEM specification also shows an example in the WAKE# signal section.</p> <p>This pin shall be left as a no connect if WAKE# is not supported by the silicon.</p> <p>Refer to Section 2.3 in the PCIe CEM Specification, Rev 4.0 for details.</p>
PWRBRK0# / PWRBRK1#	A70	Output, OD	<p>Power Brake. Active low, open drain.</p> <p>This signal shall be pulled up to +3.3V_EDGE on the OCP NIC 3.0 card with 95 kOhm or larger resistance. A baseboard that supports this function must provide a stronger pull up on PWRBRK#. A baseboard pull up value between 4.7 kOhm and 10 kOhm is recommended. The pull up shall meet the T_{PWRBRK} timing parameter as shown in the PCIe CEM Specification.</p> <p>When this signal is driven low by the baseboard, the Emergency Power Reduction State is requested. The OCP NIC 3.0 card shall move to a lower power consumption state.</p> <p>For baseboards, the PWRBRK0# pin shall be implemented and available on the Primary Connector for SFF slots. In addition, the PWRBRK1# pin shall be implemented on the Secondary connector for LFF slots.</p> <p>For OCP NIC 3.0 cards, the PWRBRK[0:1]# pin usage is optional. If used, the PWRBRK0# on the Primary Connector should be connected to the network silicon to enable reduced power state. If not used, the PWRBRK0# signals shall be left as a no connect. PWRBRK1# on the Secondary Connector is reserved for future use cases and shall be left as a no connect.</p> <p>Note: The PWRBRK[0:1]# pins are only available for OCP NIC 3.0 cards that implement a SFF 4C+ edge connector or a LFF. For SFF cards that implement at 2C+ edge connection, the PWRBRK[0:1]# functionality is not available.</p>

3.4.2 PCIe Present and Bifurcation Control Pins

This section provides the pin assignments for the PCIe present and bifurcation control signals. The AC/DC specifications are defined in Section 3.11. Example connection diagrams are shown in Figure 82 and Figure 83.

The PRSNTA#/PRSNTB[0:3]# state shall be used to determine if a card has been physically plugged in. The BIF[0:2]# pins shall be asserted by the baseboard along with the rising edge of AUX_PWR_EN. The BIF[0:2]# pins shall be latched by the OCP NIC 3.0 card when AUX_PWR_EN=1 and NIC_PWR_GOOD=1 to ensure the correct values are detected by the OCP NIC 3.0 card. Changing the pin states after this timing window is not allowed. Refer to the AC timing diagram in Section 3.11 for details.

PRSNTB[0:3]# pins are available to each connector and are independent of each other. For the SFF, the baseboard shall only read the Primary Connector PRSNTB[0:3]# to determine the card type. For the LFF, the baseboard shall read both the Primary and Secondary connector PRSNTB[0:3]# pins to determine the card type. The card type matrix is discussed in Section 3.5.

Table 24: Pin Descriptions – PCIe Present and Bifurcation Control Pins

Signal Name	Pin #	Baseboard Direction	Signal Description
PRSNTA#	A10	Output	<p>Present A is used for OCP NIC 3.0 card presence and PCIe capabilities detection.</p> <p>For baseboards, this pin shall be directly connected to GND.</p> <p>For OCP NIC 3.0 cards, this pin shall be directly connected to the PRSNTB[3:0]# pins.</p>
PRSNTB0# PRSNTB1# PRSNTB2# PRSNTB3#	B42 A42 A12 B70	Input	<p>Present B [0:3]# are used for OCP NIC 3.0 card presence and PCIe capabilities detection.</p> <p>For baseboards, these pins shall be connected to the I/O hub and pulled up to +3.3V_EDGE using 1 kOhm resistors.</p> <p>For OCP NIC 3.0 cards, these pins shall be strapped to PRSNTA# per the encoding definitions described in Section 3.5.</p> <p>Note: PRSNTB3# is located at the bottom of the 4C connector and is only applicable for OCP NIC 3.0 cards with a PCIe width of x16 (or greater). OCP NIC 3.0 cards that implement a 2C card edge do not use the PRSNTB3# pin for capabilities or present detection.</p>
BIF0# BIF1# BIF2#	B7 B8 B9	Output	<p>Bifurcation [0:2]# pins allow the baseboard to force configure the OCP NIC 3.0 card bifurcation.</p>

		<p>For baseboards, the BIF[0:2]# these pins shall be driven from the baseboard I/O hub on the rising edge of AUX_PWR_EN. This allows the baseboard to force the OCP NIC 3.0 card bifurcation. The baseboard may optionally pull the BIF[0:2]# signals to AUX_PWR_EN or to ground per the definitions described in Section 3.5 if no dynamic bifurcation configuration is required. The BIF[0:2]# pins shall be low until AUX_PWR_EN is asserted.</p> <p>For baseboards that allow dynamic bifurcation, the BIF[0:2] pins are driven low prior to AUX_PWR_EN. The state of the BIF[0:2] pins are driven with the rising edge of AUX_PWR_EN when bifurcation is requested. Refer to Figure 82 for an example configuration.</p> <p>For baseboards with static bifurcation, the BIF pins that are intended to be a logical '1' shall be connected to a pull up to AUX_PWR_EN. BIF pins that are a logical '0' may be directly tied to ground. Refer to Figure 83 for an example configuration.</p> <p>For OCP NIC 3.0 cards, these signals shall connect to the endpoint bifurcation pins if it is supported. The BIF[0:2]# signals shall be left as no connects if end point bifurcation is not supported. The value of the BIF[2:0]# pins are latched by the OCP NIC 3.0 card upon entering the AUX power mode state (when AUX_PWR_EN=1 and NIC_PWR_GOOD=1).</p> <p>Note: the required combinatorial logic output for endpoint bifurcation is dependent on the specific silicon and is not defined in this specification.</p>
--	--	--

Figure 82: PCIe Present and Bifurcation Control Pins (Baseboard Controlled BIF[0:2]#)

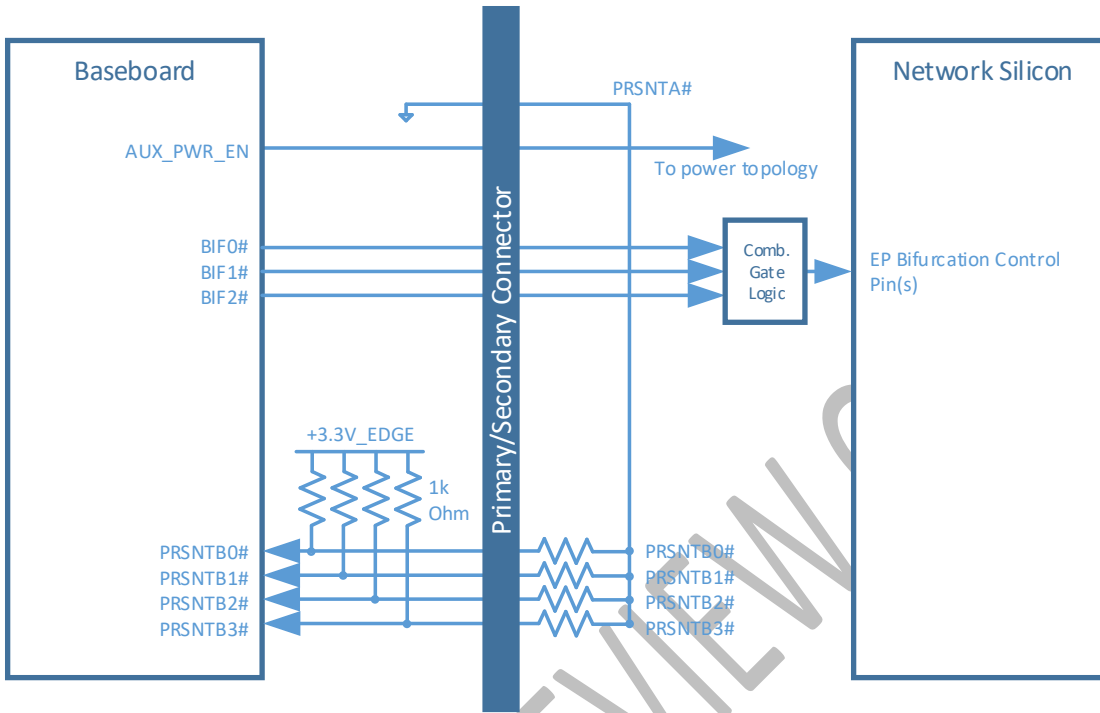
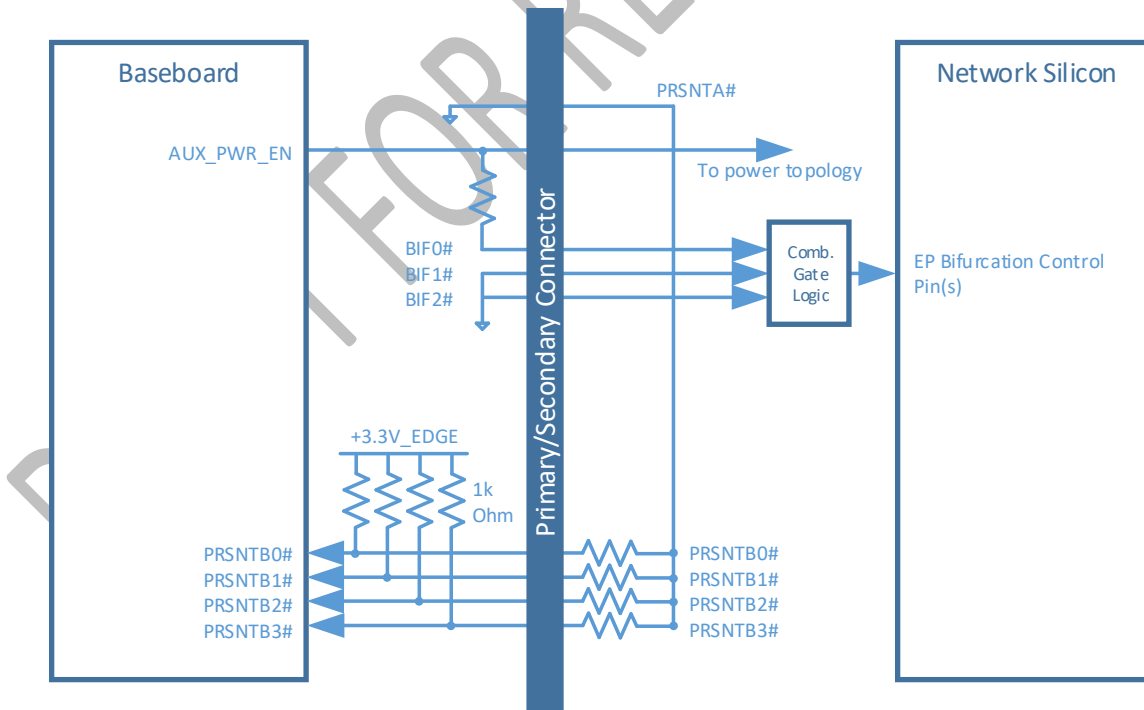


Figure 83: PCIe Present and Bifurcation Control Pins (Static BIF[0:2]#)




3.4.3 SMBus Interface Pins

This section provides the pin assignments for the SMBus interface signals. The AC/DC specifications are defined in the SMBus 2.0 specification. The SMBus interface is pinned out on the Primary and Secondary Connectors. For SFF and LFF OCP NIC 3.0 card implementations, FRU and MCTP over SMBus transactions shall use the Primary Connector only. SMBus on the Secondary Connector may be a separate bus and is reserved for a future use case. An example connection diagram is shown in Figure 84.

Table 25: Pin Descriptions – SMBus

Signal Name	Pin #	Baseboard Direction	Signal Description
SMCLK	A7	Output, OD	<p>SMBus clock. Open drain, pulled up to +3.3V_EDGE on the baseboard.</p> <p>For baseboards, the SMCLK from the platform SMBus master shall be connected to the connector.</p> <p>For OCP NIC 3.0 cards, the SMCLK from the endpoint silicon shall be connected to the card edge gold fingers.</p>
SMDAT	A8	Input / Output, OD	<p>SMBus Data. Open drain, pulled up to +3.3V_EDGE on the baseboard.</p> <p>For baseboards, the SMDAT from the platform SMBus master shall be connected to the connector.</p> <p>For OCP NIC 3.0 cards, the SMDAT from the endpoint silicon shall be connected to the card edge gold fingers.</p>
SMRST#	A9	Output, OD	<p>SMBus reset. Open drain.</p> <p>For baseboards, this pin shall be pulled up to +3.3V_EDGE. The SMRST pin may be used to reset optional downstream SMBus devices (such as temperature sensors). The SMRST# implementation shall be mandatory for baseboard implementations.</p> <p>For OCP NIC 3.0 cards, SMRST# is optional and is dependent on the OCP NIC 3.0 card implementation. If used, the SMRST# is on the +3.3V_EDGE power domain. Isolation logic may be required if the target device(s) exist on a different power domain to prevent a leakage path. The SMRST# signal shall be left as a no connect if it is not used on the OCP NIC 3.0 card.</p>

Figure 84: Example SMBus Connections



3.4.4 NC-SI over RBT Interface Pins

This section provides the pin assignments for the NC-SI over RBT interface signals on the Primary Connector OCP bay. The AC/DC specifications for NC-SI over RBT are defined in the DMTF DSP0222 NC-SI specification. Example connection diagrams are shown in Figure 85 and Figure 86.

Note: The RBT pins must provide the ability to be isolated on the baseboard side when AUX_PWR_EN=0 or when (AUX_PWR_EN=1 and NIC_PWR_GOOD=0). The RBT pins shall remain isolated until the power state machine has transitioned to AUX power mode or to Main Power Mode along with a valid indication of NIC_PWR_GOOD. This prevents a leakage path through unpowered silicon. The RBT REF_CLK must also be disabled until AUX_PWR_EN=1 and NIC_PWR_GOOD=1. Example buffering implementations are shown in Figure 85 and Figure 86. The isolator shall be controlled on the baseboard with a signal called RBT_ISOLATE#.

RBT reference clock buffers are permitted on the OCP NIC for multi-endpoint implementations if the NIC timing budget is not violated. Refer to the signal integrity requirements in Section 5.1 for timing budget details.

Command-based (software) arbitration is also permissible per DSP0222. Baseboards may choose to use the Select Package and Deselect Package commands to allow a network controller to transmit on the shared bus when more than one device is present.

Table 26: Pin Descriptions – NC-SI over RBT

Signal Name	Pin #	Baseboard Direction	Signal Description
-------------	-------	---------------------	--------------------

RBT_REF_CLK	OCP_A14	Output	<p>Reference clock input. Synchronous clock reference for receive, transmit and control interface. The clock shall have a typical frequency of 50MHz \pm50 ppm.</p> <p>For baseboards, this pin shall be connected between the baseboard NC-SI over RBT PHY and the Primary Connector OCP bay. The RBT_REF_CLK shall not be driven until the card has transitioned into AUX Power Mode. The RBT_REF_CLK shall be continuous once it has started.</p> <p>If the baseboard does not support NC-SI over RBT, then this pin shall be terminated to ground through a 100 kOhm pull down resistor.</p> <p>For OCP NIC 3.0 cards, this pin shall be connected between the card gold finger and the endpoint silicon. This pin shall be left as a no connect if NC-SI over RBT is not supported.</p>
RBT_CRS_DV	OCP_B14	Input	<p>Carrier sense/receive data valid. This signal is used to indicate to the baseboard that the carrier sense/receive data is valid.</p> <p>For baseboards, this pin shall be connected between the baseboard NC-SI over RBT PHY and the connector. This signal requires a 100 kOhm pull down resistor on the baseboard between the BMC and the RBT isolator to prevent the signal from floating when no card is installed.</p> <p>If the baseboard does not support NC-SI over RBT, then this pin shall be terminated to ground through a 100 kOhm pull down resistor.</p> <p>For OCP NIC 3.0 cards, this pin shall be connected between the card gold finger and the endpoint silicon. This pin shall be left as a no connect if NC-SI over RBT is not supported.</p>
RBT_RXD0 RBT_RXD1	OCP_B9 OCP_B8	Input	<p>Receive data. Data signals from the network controller to the BMC.</p> <p>For baseboards, this pin shall be connected between the baseboard NC-SI over RBT PHY and the connector. This signal requires a 100 kOhm pull down resistor to GND on the baseboard between the BMC and the RBT isolator to prevent the signal from floating when no card is installed.</p>

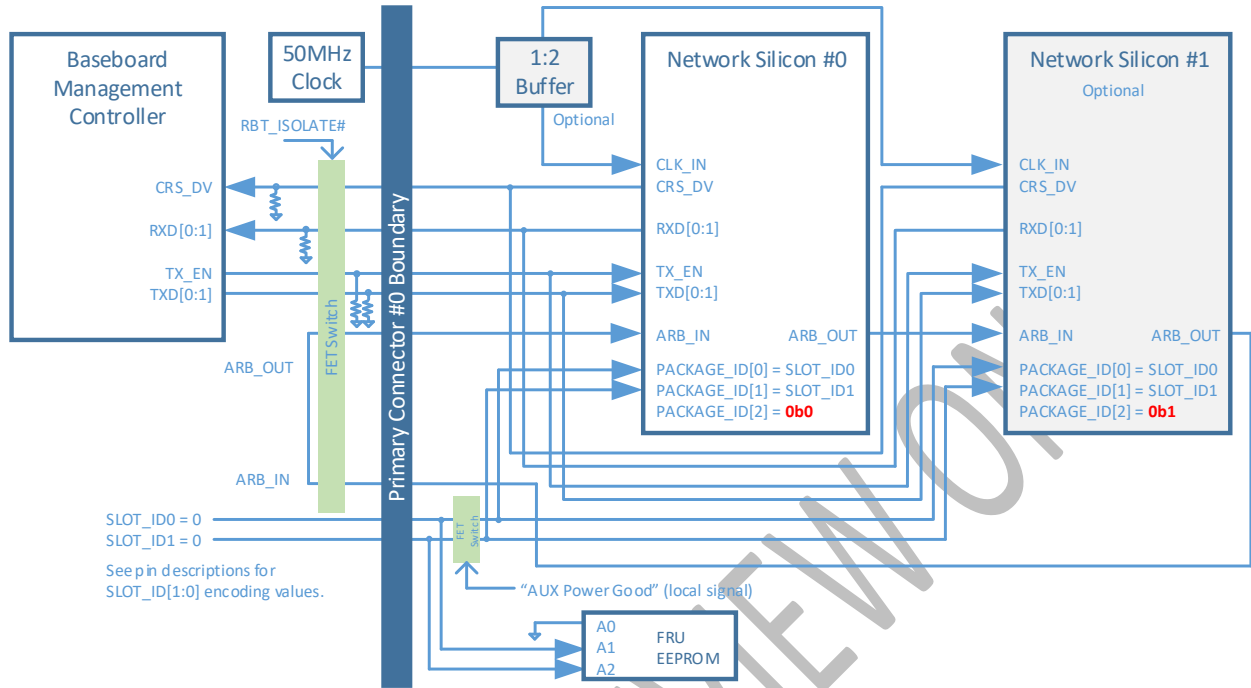
			<p>If the baseboard does not support NC-SI over RBT, then this pin shall be terminated to GND through a 100 kOhm pull down.</p> <p>For OCP NIC 3.0 cards, this pin shall be connected between the card gold fingers and the RBT_RXD[0:1] pins on endpoint silicon. This pin shall be left as a no connect if NC-SI over RBT is not supported.</p>
RBT_TX_EN	OCP_A7	Output	<p>Transmit enable.</p> <p>For baseboards, this pin shall be connected between the baseboard NC-SI over RBT PHY and the connector. This signal requires a 100 kOhm pull down resistor to ground on the baseboard between the RBT isolator and the OCP connector to prevent the card-side signals from floating when the RBT signals are isolated.</p> <p>If the baseboard does not support NC-SI over RBT, then this pin shall be terminated to ground through a 100 kOhm pull down.</p> <p>For OCP NIC 3.0 cards, this pin shall be connected between the card gold finger and the endpoint silicon. This pin shall be left as a no connect if NC-SI over RBT is not supported.</p>
RBT_TXD0 RBT_TXD1	OCP_A9 OCP_A8	Output	<p>Transmit data. Data signals from the BMC to the network controller.</p> <p>For baseboards, this pin shall be connected between the baseboard NC-SI over RBT PHY and the connector. This signal requires a 100 kOhm pull down resistor to GND on the baseboard between the RBT isolator and the OCP connector to prevent the card-side signals from floating when the RBT signals are isolated.</p> <p>If the baseboard does not support NC-SI over RBT, then this pin shall be terminated to GND through a 100 kOhm pull down.</p> <p>Note: Some BMC vendors use the RBT_TXD[0:1] pins as hardware configuration straps. A 4.7 kOhm to 10 kOhm pull up/pull down resistor is permitted between the BMC and the RBT isolator for this purpose. The resulting network will not violate the NC-SI V_{IHMIN} of 2.0V.</p>

			For OCP NIC 3.0 cards, this pin shall be connected between the card gold fingers and the RBT_TXD[0:1] pins on the endpoint silicon. This pin shall be left as a no connect if NC-SI over RBT is not supported.
RBT_ARB_OUT	OCP_A5	Output	<p>NC-SI hardware arbitration output.</p> <p>If the baseboard supports multiple OCP NIC 3.0 cards connected to the same RBT interface, it shall implement logic that connects the RBT_ARB_OUT pin of the first populated OCP NIC 3.0 card to its RBT_ARB_IN pin if it is the only card present or to the RBT_ARB_IN pin of the next populated card and so on sequentially for all cards on the specified RBT bus to ensure the arbitration ring is complete. This logic shall bypass slots that are not populated, powered off, or in ID mode. A two OCP NIC 3.0 card example using an analog mux is shown in Figure 86.</p> <p>If the baseboard does not support NC-SI over RBT or implements only one OCP NIC 3.0 interface, this signal shall be directly connected to the RBT_ARB_IN pin to complete the hardware arbitration ring on the OCP NIC 3.0 card.</p> <p>For OCP NIC 3.0 cards that support hardware arbitration, this pin shall be connected between the card gold finger and the RBT_ARB_IN pin on the endpoint silicon. If the card implements two controllers, both must be connected internally to complete the ring, see Figure 86. If hardware arbitration is not supported, then this pin shall be directly connected to the card edge RBT_ARB_IN pin. This allows the hardware arbitration signals to pass through in a multi-Primary Connector baseboard.</p>
RBT_ARB_IN	OCP_A4	Input	<p>NC-SI hardware arbitration input.</p> <p>If the baseboard supports multiple OCP NIC 3.0 cards connected to the same RBT interface, it shall implement logic that connects the RBT_ARB_IN pin of the first populated OCP NIC 3.0 card to its RBT_ARB_OUT pin if it is the only card present or to the RBT_ARB_OUT pin of the next populated card and so on sequentially for all cards on the specified RBT bus to ensure the arbitration ring is complete. This logic shall bypass slots that are not populated, powered off, or in ID mode. A two OCP NIC 3.0 card example using an analog mux is shown in Figure 86.</p>

			<p>If the baseboard does not support NC-SI over RBT or implements only one OCP NIC 3.0 interface, this signal shall be directly connected to the RBT_ARB_OUT pin to complete the hardware arbitration ring on the OCP NIC 3.0 card.</p> <p>For OCP NIC 3.0 cards that support hardware arbitration, this pin shall be connected between the card gold finger and the RBT_ARB_OUT pin on the endpoint silicon. If the card implements two controllers, both must be connected internally to complete the ring, see Figure 86. If hardware arbitration is not supported, then this pin shall be directly connected to the card edge RBT_ARB_OUT pin. This allows the hardware arbitration signals to pass through in a multi-Primary Connector baseboard.</p>															
<p>SLOT_ID0 SLOT_ID1</p>	<p>OCP_B7 OCP_A6</p>	<p>Output</p>	<p>NC-SI / FRU EEPROM Address 0/1.</p> <p>For baseboards, the SLOT_ID[1:0] pins shall be connected to GND through a 100 Ohm pull down or to +3.3V_EDGE through a 4.7 kOhm pull up. The SLOT[1:0] values are based on the following mapping on a per slot basis:</p> <table border="1" data-bbox="781 1083 1414 1304"> <thead> <tr> <th>Physical Slot (Decimal)</th> <th>SLOT_ID1 OCP_A6</th> <th>SLOT_ID0 OCP_B7</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>0</td> </tr> <tr> <td>1</td> <td>0</td> <td>1</td> </tr> <tr> <td>2</td> <td>1</td> <td>0</td> </tr> <tr> <td>3</td> <td>1</td> <td>1</td> </tr> </tbody> </table> <p>For OCP NIC 3.0 cards, the SLOT_ID[1:0] pins shall be used to set the RBT Package ID and the FRU EEPROM address on the OCP NIC 3.0 card. The OCP NIC 3.0 card may optionally implement weak pull up or pull down resistors (>47 kOhm) to prevent the silicon pins from floating prior to the local silicon “Aux power good.”</p> <p>For OCP NIC 3.0 cards, SLOT_ID0 shall be connected to the endpoint device GPIO associated with Package ID[0]. SLOT_ID1 shall be associated with Package ID[1]. Refer to Section 4.8.1 and the device datasheet for details.</p>	Physical Slot (Decimal)	SLOT_ID1 OCP_A6	SLOT_ID0 OCP_B7	0	0	0	1	0	1	2	1	0	3	1	1
Physical Slot (Decimal)	SLOT_ID1 OCP_A6	SLOT_ID0 OCP_B7																
0	0	0																
1	0	1																
2	1	0																
3	1	1																


		<p>For OCP NIC 3.0 cards with multiple endpoint devices, Package ID[2] shall be used to identify a second physical RBT capable controller on the same physical card.</p> <p>For Package ID addressing, the SLOT_ID[1:0] pins shall be buffered on NIC side with a FET switch (or a similar implementation) to prevent a leakage path when the OCP NIC 3.0 card is in ID mode. The SLOT_ID[1:0] buffers shall isolate the signals to the network silicon until an “Aux Power Good” is generated locally from the NIC. This indication shall be generated from an on-board voltage monitor or similar logic. OCP NIC 3.0 designers may omit isolation logic for the Package ID addressing if the target silicon properly isolates the signals when it is unpowered.</p> <p>For FRU EEPROM addressing, the SLOT_ID0 pin shall be directly connected to the EEPROM A1 address pin; SLOT_ID1 shall be connected to the EEPROM A2 address pin. No isolation shall be used for the FRU EEPROM connections.</p> <p>For endpoint devices without NC-SI over RBT support, these pins shall only be connected to the FRU EEPROM as previously described.</p>
--	--	---

Figure 85: NC-SI over RBT Connection Example – Single Primary Connector



DRAFT FOR REVIEW ONLY

Figure 86: NC-SI over RBT Connection Example – Dual Primary Connectors



Note 1: For baseboard designs with a single Primary Connector, connect ARB_IN to ARB_OUT to complete the NC-SI hardware arbitration ring. For designs with multiple Primary Connectors, connect ARB_IN and ARB_OUT to an analog mux to complete the NC-SI arbitration ring based on the number of cards installed in the system. An example dual Primary Connector implementation is shown in Figure 86.

Note 2: For baseboard implementations having two or more RBT busses, the baseboard hardware arbitration rings shall remain within their respective bus and shall not cross RBT bus domains.

Note 3: The baseboard implementation shall maintain the arbitration ring integrity when there exists one or more cards that are not present, plugged in but are powered off, or in ID Mode.

Note 4: For OCP NIC 3.0 cards with two discrete endpoint silicon, the Package ID[2] bit shall be statically set based on the silicon instance. For example, the figure above shows Network Silicon #0 and Network Silicon #1. Network Silicon #0 has Package ID[2] = 0b0, Network Silicon #1 has Package ID[2] = 0b1.

Note 5: Designs that implement a clock fan out buffer will affect the RBT timing budget. Careful analysis of the timing budget is required. Refer to Section 5.1 for RBT signal integrity and timing budget considerations.

3.4.5 Scan Chain Pins

This section provides the pin assignments for the Scan Chain interface signals on the Primary Connector OCP Bay. The scan chain is a point-to-point bus on a per OCP slot basis. The scan chain consists of two unidirectional busses, a common clock and a common load signal. The DATA_OUT signal serially shifts control signals from the baseboard to the OCP NIC 3.0 card. The DATA_IN signal serially shifts bits from the OCP NIC 3.0 card to the baseboard. The DATA_OUT and DATA_IN chains are independent of each other. The scan chain CLK is driven from the baseboard. The LD pin, when asserted by the baseboard, allows loading of the data on to the shift registers. An example timing diagram is shown in Figure 87. An example connection diagram is shown in Figure 88.


Note: The DATA_OUT chain is provisioned, but is not used on OCP NIC 3.0 cards for this revision of the specification.

Table 27: Pin Descriptions – Scan Chain

Signal Name	Pin #	Baseboard Direction	Signal Description
CLK	OCP_B6	Output	<p>Scan clock. The CLK is an output pin from the baseboard to the OCP NIC 3.0 card. The CLK may run up to 12.5MHz.</p> <p>For baseboard implementations, the CLK pin shall be connected to the Primary Connector. The CLK pin shall be tied directly to GND if the scan chain is not used.</p> <p>For NIC implementations, the CLK pin shall be connected to Shift Registers 0 & 1, and optionally connected to Shift Registers 2 & 3 (if implemented) as defined in the text and Figure 88, below. The CLK pin shall be pulled up to +3.3V_EDGE through a 1 kOhm resistor.</p>
DATA_OUT	OCP_B5	Output	<p>Scan data output from the baseboard to the OCP NIC 3.0 card. This bit stream is used to shift configuration data out to the NIC.</p>

			<p>For baseboard implementations, the DATA_OUT pin shall be connected to the Primary Connector. The DATA_OUT pin shall be pulled down to GND through a 1 kOhm resistor if the scan chain is not used.</p> <p>For NIC implementations, the DATA_OUT pin shall be pulled down to GND on the OCP NIC 3.0 card through a 10 kOhm resistor.</p>
DATA_IN	OCP_B4	Input	<p>Scan data input to the baseboard. This bit stream is used to shift out NIC status bits to the baseboard.</p> <p>For baseboard implementations, the DATA_IN pin shall be pulled up to +3.3V_EDGE through a 10 kOhm resistor to prevent the input signal from floating if a card is not installed. This pin may be left as a no connect if the scan chain is not used.</p> <p>For NIC implementations, the DATA_IN scan chain is required. The DATA_IN pin shall be connected to Shift Register 0, as defined in the text and Figure 88.</p>
LD#	OCP_B3	Output	<p>Scan shift register load. Used to latch configuration data on the OCP NIC 3.0 card.</p> <p>For baseboard implementations, the LD# pin shall be pulled up to +3.3V_EDGE through a 1 kOhm resistor if the scan chain is not used to prevent the OCP NIC 3.0 card from erroneous data latching.</p> <p>For NIC implementations, the LD# pin implementation is required. The LD# pin shall be connected to Shift Registers 0 & 1, and optionally connected to Shift Registers 2 & 3 (if implemented) as defined in the text and Figure 88. The LD# pin shall be pulled up to +3.3V_EDGE through a 10 kOhm resistor.</p>

Figure 87: Example Scan Chain Timing Diagram



The scan chain provides sideband status indication between the OCP NIC 3.0 card and the baseboard. The scan chain bit definition is defined in the two tables below. The scan chain data stream is 32-bits in length for both the DATA_OUT and the DATA_IN streams. The scan chain implementation is optional on

the host, but its implementation is mandatory per Table 28 and Table 29 on all OCP NIC 3.0 cards. The scan chain components operates on the +3.3V_EDGE power domain.

The DATA_OUT bus is an output from the host. The DATA_OUT bus provides initial configuration options to the OCP NIC 3.0 card. At the time of this writing, the DATA_OUT bus is not used. All baseboard systems that implement the Scan Chain shall connect DATA_OUT between the platform and the Primary Connector for subsequent revisions of this specification. The DATA_OUT data stream shall shift out all 0's prior to AUX_PWR_EN assertion to prevent leakage paths into unpowered silicon.

Table 28: Pin Descriptions – Scan Chain DATA_OUT Bit Definition

Byte.bit	DATA_OUT Field Name	Default Value	Description
0.[0..7]	Reserved	0h00	Reserved. Byte 0 value is 0h00.
1.[0..7]	Reserved	0h00	Reserved. Byte 1 value is 0h00.
2.[0..7]	Reserved	0h00	Reserved. Byte 2 value is 0h00.
3.[0..7]	Reserved	0h00	Reserved. Byte 3 value is 0h00.

The DATA_IN bus is an input to the host and provides NIC status indication. The default implementation is completed with two 8-bit 74LV165 parallel in to serial out shift registers in a cascaded implementation. Up to four shift registers may be implemented to provide additional NIC status indication to the host platform. Alternatively, an OCP NIC 3.0 card vendor may choose to implement this chain using an active device (such as a microcontroller or CPLD). For active device implementations, there is an associated device start-up time. Refer to Section 3.11 for details on the +3.3V_EDGE stable to the first data valid read in ID Mode.

DATA_IN shift register 0 shall be mandatory for scan chain implementations for the card present, WAKE_N and thermal threshold features. DATA_IN shift registers 1, 2 & 3 are optional depending on the line side I/O and LED fields being reported to the host. Dual port LED applications require shift register 1. Quad port LED applications require shift registers 1 & 2. Octal port applications require shift registers 1, 2 & 3.

The host should read the DATA_IN bus multiple times to qualify the incoming data stream. The number of data qualification reads is dependent on the baseboard implementation.

On the OCP NIC 3.0 card, a 1 kOhm pull up resistor shall be connected to the SER input of the last DATA_IN shift register. Doing so ensures the default bit value of 0b1 for implementations using less than four shift registers.

Table 29: Pin Descriptions – Scan Chain DATA_IN Bit Definition

Byte.bit	DATA_IN Field Name	Default Value	Description
0.0	PRSNTB[0]_P#	0bX	PRSNTB[3:0]# bits shall reflect the same state as the signals on the Primary Connector. Connect these scan chain signals directly to the OCP NIC 3.0 card edge PRSNTB[3:0]# pins. The OCP NIC 3.0 implementer may alternatively choose to locally
0.1	PRSNTB[1]_P#	0bX	
0.2	PRSNTB[2]_P#	0bX	
0.3	PRSNTB[3]_P#	0bX	

			populate pull up and pull down resistors to these scan chain inputs as long as the PRSNTB[3:0]# values are the same on the scan chain and card edge.
0.4	WAKE_N / PRSNTB[0]_S#	0bX / 0bX	<p>PCIe WAKE_N signal shall reflect the same state as the signal on the Primary Connector.</p> <p>For LFF designs, this bit shall also serve as the PRSNTB[0]# signal from the Secondary Connector when the card is in ID Mode (AUX_PWR_EN==0). Multiplexing between the two functions shall be controlled via AUX_PWR_EN. Refer to Figure 88 for details.</p>
0.5	TEMP_WARN_N / PRSNTB[1]_S#	0b1 / 0bX /	<p>Temperature monitoring pin from the on-card thermal solution. This pin shall be asserted low when the network silicon or transceiver module temperature sensors exceed the temperature warning threshold.</p> <p>For cards that do not require temperature reporting, the TEMP_WARN_N value shall be statically set to 0b1. Refer to Section 4.4 for details on temperature reporting.</p> <p>For LFF designs, this bit shall also serve as the PRSNTB[1]# signal from the Secondary Connector when the card is in ID Mode (AUX_PWR_EN==0). Multiplexing between the two functions shall be controlled via AUX_PWR_EN. Refer to Figure 88 for details.</p>
0.6	TEMP_CRIT_N / PRSNTB[2]_S#	0b1 / 0bX	<p>Temperature monitoring pin from the on-card thermal solution. This pin shall be asserted low when the network silicon or transceiver module temperature sensors exceed the temperature critical threshold.</p> <p>For cards that do not require temperature reporting, the TEMP_CRIT_N value shall be statically set to 0b1. Refer to Section 4.4 for details on temperature reporting.</p> <p>For LFF designs, this bit shall also serve as the PRSNTB[2]# signal from the Secondary Connector when the card is in ID Mode (AUX_PWR_EN==0). Multiplexing between the two functions shall be controlled via AUX_PWR_EN. Refer to Figure 88 for details.</p>

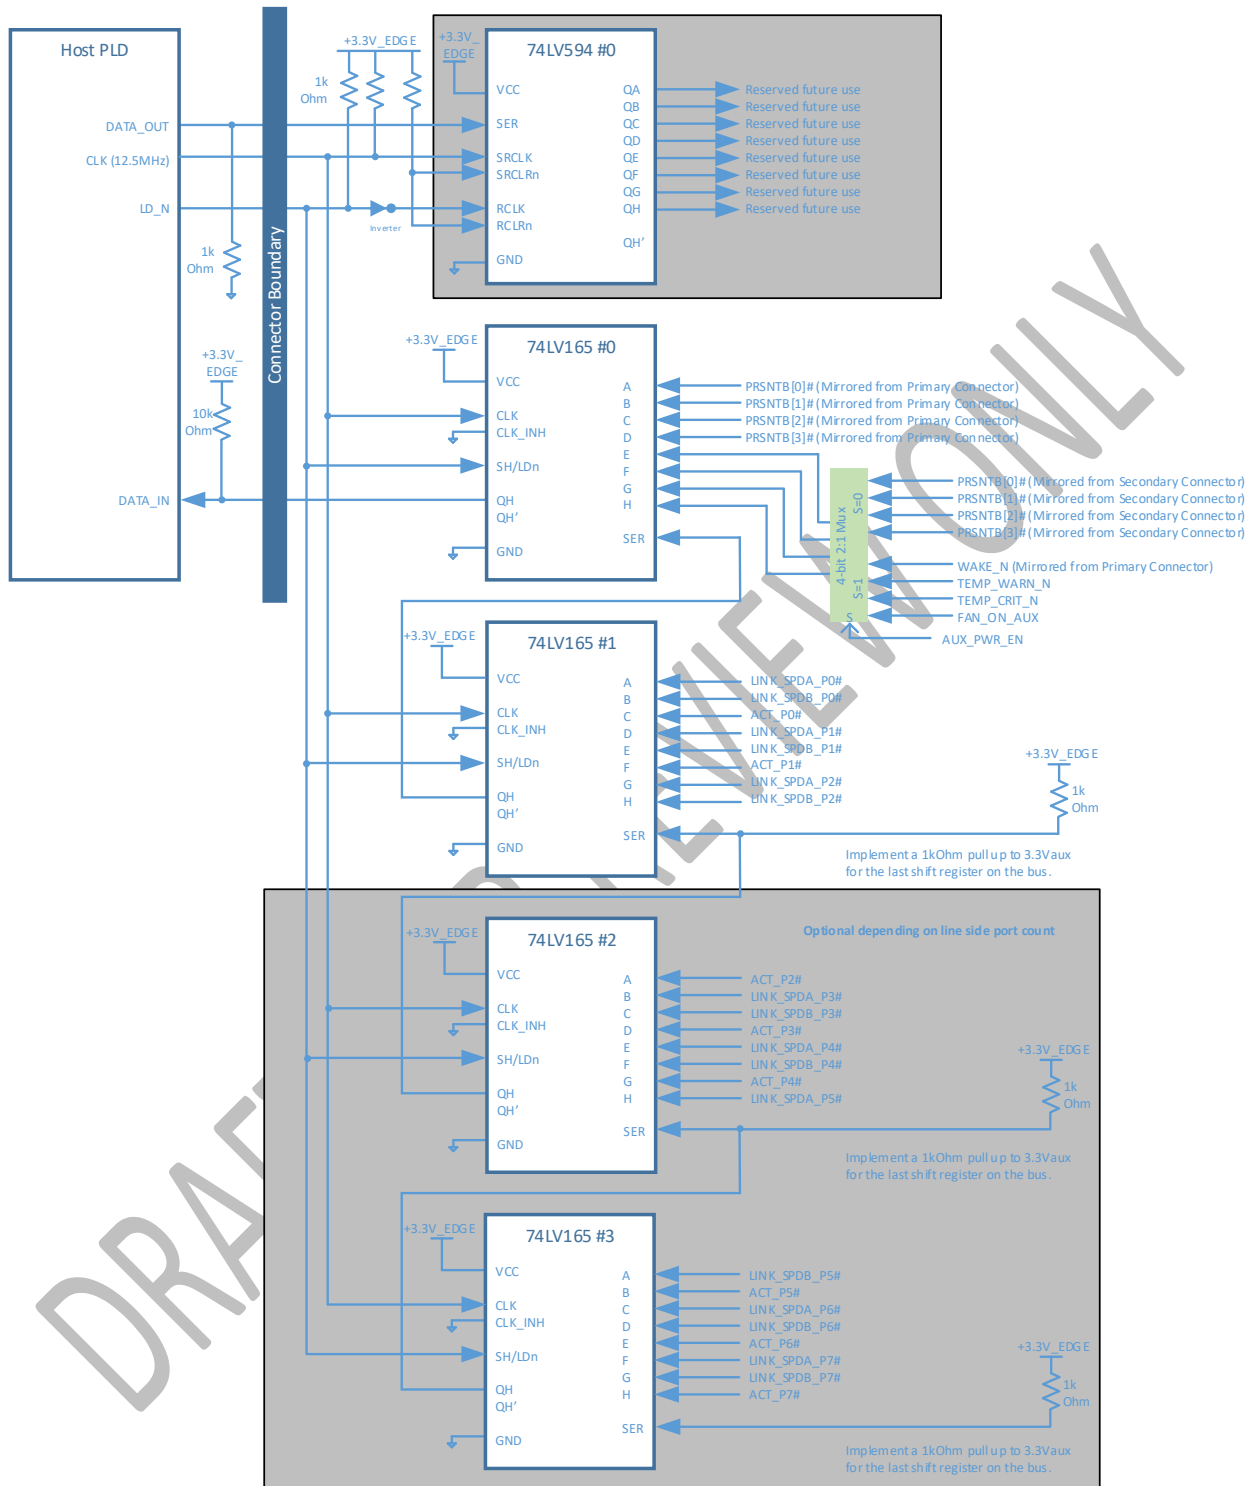
0.7	FAN_ON_AUX / PRSNTB[3]_S# /	0b0 / 0bX	<p>When high, FAN_ON_AUX shall request the system fan to be enabled for extra cooling in the S5 state.</p> <p>The FAN_ON_AUX bit shall be asserted when the network silicon or transceiver module cooling requested threshold has been exceeded. The temperature at which this assertion occurs is device dependent.</p> <p>The FAN_ON_AUX bit shall deassert when the network silicon or transceiver module temperature is at least 5°C below the assertion threshold.</p> <p>0b0 – The system fan is not requested/off in S5. 0b1 – The system fan is requested/on in S5.</p> <p>For cards that do not require temperature reporting, the FAN_ON_AUX value shall be statically set to 0b0. Refer to Section 4.4 for details on temperature reporting.</p> <p>For LFF designs, this bit shall also serve as the PRSNTB[3]# signal from the Secondary Connector when the card is in ID Mode (AUX_PWR_EN==0). Multiplexing between the two functions shall be controlled via AUX_PWR_EN. Refer to Figure 88 for details.</p>
1.0	LINK_SPDA_P0#	0b1	<p>Port 0 link and speed A indication (max speed). Active low.</p> <p>0b0 – Link LED is illuminated on the host platform. 0b1 – Link LED is not illuminated on the host platform.</p> <p>On = link is detected on the port and is at the maximum speed. Off = the physical link is down, the link is not operating at the maximum speed or the port is disabled.</p> <p>Note: The link and speed A LED may also be blinked for use as port identification.</p>
1.1	LINK_SPDB_P0#	0b1	<p>Port 0 link and speed B indication (not max speed). Active low.</p> <p>0b0 – Link LED is illuminated on the host.</p>

			<p>0b1 – Link LED is not illuminated on the host.</p> <p>On = link is detected on the port and is not at the maximum speed. Off = the physical link is down, the link is operating at the maximum speed, or the port is disabled.</p> <p>Note: The link and speed B LED may also be blinked for use as port identification.</p>
1.2	ACT_P0#	0b1	<p>Port 0 activity indication. Active low.</p> <p>0b0 – ACT LED is illuminated on the host. 0b1 – ACT LED is not illuminated on the host.</p> <p>Blinking = activity is detected on the port. The LED should blink at the rate of ½ Hz to 5 Hz. Off = no activity is detected on the port.</p>
1.3	LINK_SPDA_P1#	0b1	Port 1 link and speed A indication (max speed). Active low.
1.4	LINK_SPDB_P1#	0b1	Port 1 link and speed B indication (not max speed). Active low.
1.5	ACT_P1#	0b1	Port 1 activity indication. Active low.
1.6	LINK_SPDA_P2#	0b1	Port 2 link and speed A indication (max speed). Active low.
1.7	LINK_SPDB_P2#	0b1	Port 2 link and speed B indication (not max speed). Active low.
2.0	ACT_P2#	0b1	Port 2 activity indication. Active low.
2.1	LINK_SPDA_P3#	0b1	Port 3 link and speed A indication (max speed). Active low.
2.2	LINK_SPDB_P3#	0b1	Port 3 link and speed B indication (not max speed). Active low.
2.3	ACT_P3#	0b1	Port 3 activity indication. Active low.
2.4	LINK_SPDA_P4#	0b1	Port 4 link and speed A indication (max speed). Active low.
2.5	LINK_SPDB_P4#	0b1	Port 4 link and speed B indication (not max speed). Active low.
2.6	ACT_P4#	0b1	Port 4 activity indication. Active low.
2.7	LINK_SPDA_P5#	0b1	Port 5 link and speed A indication (max speed). Active low.
3.0	LINK_SPDB_P5#	0b1	Port 5 link and speed B indication (not max speed). Active low.
3.1	ACT_P5#	0b1	Port 5 activity indication. Active low.
3.2	LINK_SPDA_P6#	0b1	Port 6 link and speed A indication (max speed). Active low.
3.3	LINK_SPDB_P6#	0b1	Port 6 link and speed B indication (not max speed). Active low.

3.4	ACT_P6#	0b1	Port 6 activity indication. Active low.
3.5	LINK_SPDA_P7#	0b1	Port 7 link and speed A indication (max speed). Active low.
3.6	LINK_SPDB_P7#	0b1	Port 7 link and speed B indication (not max speed). Active low.
3.7	ACT_P7#	0b1	Port 7 activity indication. Active low.

DRAFT FOR REVIEW ONLY

Figure 88: Scan Chain Connection Example



3.4.6 Power Supply Pins

This section provides the pin assignments for the power supply interface signals. The AC/DC specifications are defined in the PCIe CEM Specification, Rev 4.0 and amended in Section 3.9. An example connection diagram is shown in Figure 89.

Table 30: Pin Descriptions – Power

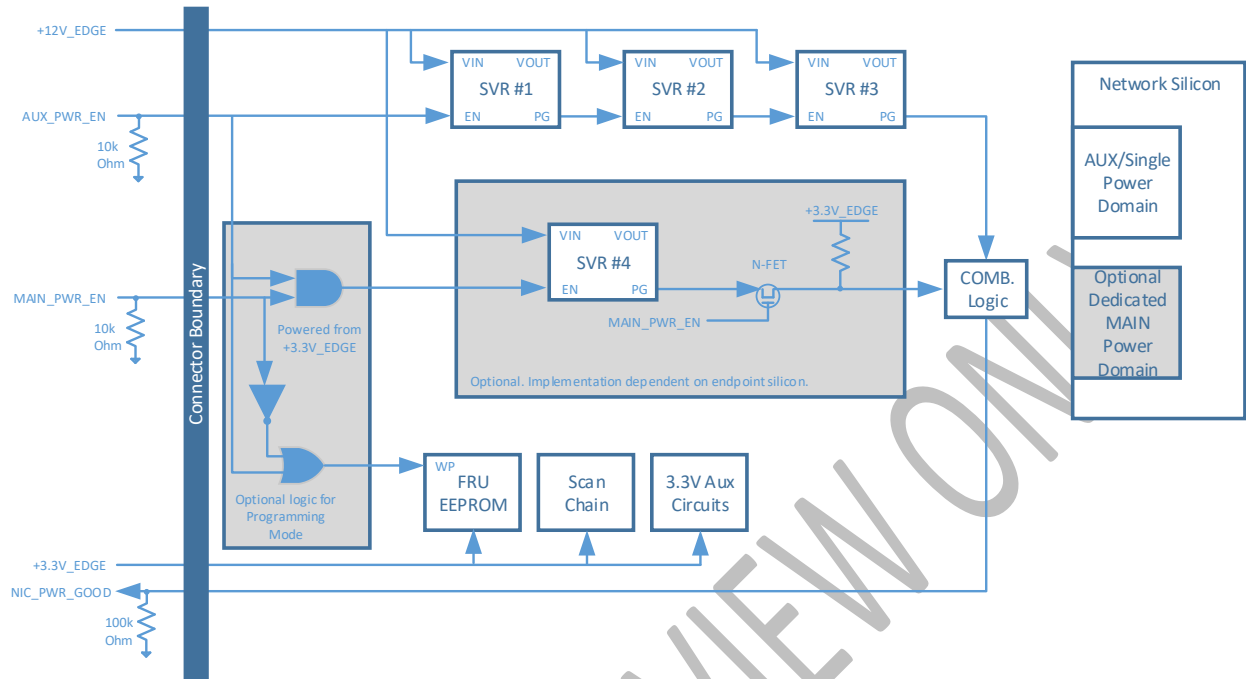
Signal Name	Pin #	Baseboard Direction	Signal Description
GND	Various	GND	Ground return; a total of 46 ground pins are on the main 140-pin connector area. Additionally, a total of 4 ground pins are in the OCP bay area. Refer to Section 3.3 for details.
+12V_EDGE	B1, B2, B3, B4, B5, B6	Power	<p>+12 V main or +12 V aux power; total of 6 pins per connector. The +12V_EDGE pins shall be rated to 1.1 A per pin with a maximum derated power delivery of 80 W.</p> <p>The +12V_EDGE power pins shall be within the rail tolerances as defined in Section 3.9 when the PWR_EN pin is driven high by the baseboard.</p> <p>The OCP NIC 3.0 card may optionally implement a fuse on +12V_EDGE to protect against electrical faults.</p>
+3.3V_EDGE	B11	Power	<p>+3.3 V main or +3.3 V aux power; total of 1 pin per connector. The +3.3V_EDGE pin shall be rated to 1.1 A for a maximum derated power delivery of 3.63 W.</p> <p>The +3.3V_EDGE power pin shall be within the rail tolerances as defined in Section 3.9 when the PWR_EN pin is driven high by the baseboard.</p> <p>The OCP NIC 3.0 card may optionally implement a fuse on +3.3V_EDGE to protect against electrical faults.</p>
AUX_PWR_EN	B12	Output	<p>Aux Power enable. Active high.</p> <p>This pin indicates that the baseboard +12V_EDGE and +3.3V_EDGE power are supplied per the Aux Power Mode requirements described in Section 3.9. Additionally, this signal notifies the OCP NIC 3.0 card to enable any power supplies that run only in the Aux Power Mode.</p> <p>AUX_PWR_EN is pinned out on both the Primary and Secondary Connector. For SFF and LFF OCP NIC 3.0 cards, the AUX_PWR_EN connection shall be implemented on the Primary Connector only. The AUX_PWR_EN connection on the Secondary Connector is reserved for a future use case.</p>

			<p>This signal shall be pulled down to GND through a 10 kOhm resistor on the baseboard. This ensures the OCP NIC 3.0 card power is disabled until instructed to turn on by the baseboard.</p> <p>When low, the OCP NIC 3.0 card supplies running on aux power shall be disabled.</p> <p>When high, the OCP NIC 3.0 card supplies running on aux power shall be enabled.</p> <p>For OCP NIC 3.0 cards that do not use a separate “main power” domain circuitry (or can operate in a single power domain), the AUX_PWR_EN signal serves as the primary method to enable all the card power supplies.</p> <p>For OCP NIC 3.0 cards that support the Programming Mode power state, the condition $AUX_PWR_EN == 0$ and $MAIN_PWR_EN == 1$ shall prevent the aux power supplies from being enabled. An example of this logic is shown in Figure 89.</p> <p>It is expected that a baseboard will not drive signals other than SMBus and the Scan Chain to the OCP NIC 3.0 card when this signal is low.</p>
MAIN_PWR_EN	OCP_B2	Output	<p>Main Power Enable. Active high.</p> <p>This pin indicates that the baseboard +12V_EDGE and +3.3V_EDGE power are supplied per the Main Power Mode requirements described in Section 3.9. Additionally, this signal notifies the OCP NIC 3.0 card to enable any power supplies that run only in the Main Power Mode.</p> <p>The MAIN_PWR_EN pin is driven by the baseboard and may only be asserted when AUX_PWR_EN is already asserted. The MAIN_PWR_EN pin must be implemented on baseboard systems, but may optionally be used to control main power rail power supplies by the OCP NIC 3.0 card depending on the end point silicon implementation. Depending on the silicon vendor, end point devices may be able to operate in a single power domain, or may require separate power domains to function.</p> <p>For OCP NIC 3.0 cards that support the Programming Mode power state, the condition $AUX_PWR_EN == 0$</p>

			<p>and MAIN_PWR_EN == 1 shall prevent the main power supplies from being enabled. An example of this gating logic is shown in Figure 89.</p> <p>For baseboard implementations, this signal shall be pulled down to GND through a 10 kOhm resistor on the baseboard. This ensures the OCP NIC 3.0 card power is disabled until instructed to turn on by the baseboard.</p> <p>When low, the OCP NIC 3.0 card supplies running on main power shall be disabled.</p> <p>When high, the OCP NIC 3.0 card supplies running on main power shall be enabled.</p> <p>This pin may be left as a no connect for OCP NIC 3.0 cards that do not use a separate “main power” domain SVR circuitry.</p>																				
NIC_PWR_GOOD	OCP_B1	Input	<p>NIC Power Good. Active high. This signal is driven by the OCP NIC 3.0 card.</p> <p>The NIC_PWR_GOOD signal is used to indicate when the aux power domain, and main power domain rails are within operational tolerances.</p> <p>The truth table shows the expected NIC_PWR_GOOD state for power up sequencing depending on the values of AUX_PWR_EN and MAIN_PWR_EN.</p> <table border="1" data-bbox="760 1226 1409 1478"> <thead> <tr> <th>AUX_PWR_EN</th> <th>MAIN_PWR_EN</th> <th>NIC_PWR_GOOD Nominal Steady State Value</th> <th>State Name</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>0</td> <td>0</td> <td>ID Mode</td> </tr> <tr> <td>1</td> <td>0</td> <td>1</td> <td>Aux Power Mode</td> </tr> <tr> <td>0</td> <td>1</td> <td>0</td> <td>Programming Mode</td> </tr> <tr> <td>1</td> <td>1</td> <td>1</td> <td>Main Power Mode</td> </tr> </tbody> </table> <p>Refer to the power up and power down sequencing diagrams (Figure 106 and Figure 107) for timing details.</p> <p>Where appropriate, designs that have a separate Main Power domain should also connect to the main power good indication to the NIC_PWR_GOOD signal via a FET to isolate the domains. Refer to Figure 89 for an example implementation.</p> <p>When low, this signal shall indicate that the OCP NIC 3.0 card power supplies are not yet within nominal</p>	AUX_PWR_EN	MAIN_PWR_EN	NIC_PWR_GOOD Nominal Steady State Value	State Name	0	0	0	ID Mode	1	0	1	Aux Power Mode	0	1	0	Programming Mode	1	1	1	Main Power Mode
AUX_PWR_EN	MAIN_PWR_EN	NIC_PWR_GOOD Nominal Steady State Value	State Name																				
0	0	0	ID Mode																				
1	0	1	Aux Power Mode																				
0	1	0	Programming Mode																				
1	1	1	Main Power Mode																				

		<p>tolerances or are in a fault condition after the power ramp times (T_{APL} and T_{MPL}) have expired.</p> <p>For cards that support the Programming Mode power state, the NIC_PWR_GOOD signal shall remain low when the OCP NIC 3.0 card is in Programming Mode as both the Aux Power Mode and Main Power Mode power supplies are disabled. Programming Mode shall only be used on OCP NIC 3.0 cards that advertise that the Programming Mode power state is supported per Section 4.10.3.</p> <p>For baseboards, this pin may be connected to the platform I/O hub as a NIC power health status indication. This signal shall be pulled down to ground with a 100 kOhm resistor on the baseboard to prevent a false power good indication if no OCP NIC 3.0 card is present.</p> <p>For OCP NIC 3.0 cards this signal shall indicate the OCP NIC 3.0 card power is “good” for the given power mode. This signal may be implemented by combinatorial logic, a cascaded power good tree or a discrete power good monitor output.</p> <p>When high, this signal should be treated as V_{REF} is available for NC-SI communications. Refer to timing parameter T4 in the DMTF DSP0222 specification for details.</p>
--	--	---

Figure 89: Example Power Supply Topology



DRAFT FOR REVIEW ONLY

3.4.7 USB 2.0 (A68/A69) – Primary Connector Only

This section provides the pin assignments for the USB 2.0 interface signals. USB 2.0 is only defined for operation on the Primary Connector. USB 2.0 may be used for applications with end point silicon that requires a USB connection to the baseboard. Implementations may also allow for a USB-Serial or USB-JTAG translator for serial or JTAG applications. If multiple USB devices are required, an optional USB hub may be implemented on the OCP NIC 3.0 card. Downstream device discovery is completed as part of the bus enumeration per the USB 2.0 specification. A basic example connection diagram is shown in Figure 90. An example depicting USB-Serial and USB-JTAG connectivity with an USB hub is shown in Figure 91.

Table 31: Pin Descriptions – USB 2.0 – Primary Connector only

Signal Name	Pin #	Baseboard Direction	Signal Description
USB_DATn USB_DATp	A68 A69	Bi-directional	<p>USB 2.0 Differential Pair – Primary Connector Only.</p> <p>A baseboard implementation shall provide a USB connection to the OCP NIC 3.0 primary connector.</p> <p>NIC implementations that require USB shall connect the bus to the end point silicon. This pin shall be left as a no connect if it is not used on the OCP NIC 3.0 card.</p> <p>The USB pins shall be directly connected between the end point silicon or USB device and the card gold fingers.</p> <p>The USB interface shall be based on a $V_{BUS} = 5\text{ V}$ per the USB specification. Both the baseboard and NIC device shall be capable of driving a differential signal using 3.3 V logic. Differential termination (V_{TERM}) pull-up values on the D+ line, if implemented, shall also be 3.3V compliant. The OCP NIC 3.0 card may implement protection diodes and is up to the adapter vendor for placement.</p> <p>To prevent leakage paths, a baseboard shall not use USB pull up resistors on the USB_DATp/n lines to indicate the bus data transmission rate. If used, pull up resistors shall only exist on the NIC side.</p> <p>The AUX_PWR_EN signal may be used for downstream USB devices that require a V_{BUS} detection indication. Designers would have to ensure the V_{BUS} detection threshold supports 3.3V signaling. Examples of this may include USB-serial converting devices.</p>

Figure 90: USB 2.0 Connection Example – Basic Connectivity





Figure 91: USB 2.0 Connection Example – USB-Serial / USB-JTAG Connectivity




3.4.8 UART (A68/A69) – Secondary Connector Only

This section provides the pin assignments for the UART interface signals. UART is only defined for operation on the Secondary Connector. The UART pins may be used with end point silicon that require console redirection over the baseboard – such as LFF SmartNICs. An example connection diagram is shown in Figure 92.

Table 32: Pin Descriptions – UART – Secondary Connector Only

Signal Name	Pin #	Baseboard Direction	Signal Description
UART_RX	A68	Input	<p>UART Receive. +3.3 V signaling levels. Secondary Connector Only.</p> <p>A baseboard implementation shall provide a UART receive connection from the OCP NIC 3.0 connector. The UART_RX pin shall be pulled up to +3.3 V_{AUX} on the baseboard to prevent erroneous data reception when the OCP NIC 3.0 card is powered off or not present.</p> <p>NIC implementations that require a UART shall connect the network silicon UART_RX pin to the UART_TX pin on the OCP NIC 3.0 connector. This pin shall be left as a no connect if it is not used on the OCP NIC 3.0 card.</p> <p>The UART_RX pin shall be buffered on the NIC to prevent a leakage path into unpowered silicon when the card is in ID Mode. The buffer may be controlled via a local “Power Good” indicator.</p>
UART_TX	A69	Output	<p>UART Transmit. +3.3 V signaling levels. Secondary Connector Only.</p> <p>A baseboard implementation shall provide a UART transmit connection to the OCP NIC 3.0 connector.</p> <p>NIC implementations that require a UART shall connect the UART_TX pin from the OCP NIC 3.0 connector to the target silicon UART_RX pin. This pin shall be left as a no connect if it is not used on the OCP NIC 3.0 card.</p> <p>The UART_TX pin shall be buffered on the NIC to prevent a leakage path into unpowered silicon when the card is in ID Mode. The buffer may be controlled via a local “Power Good” indicator.</p>

Figure 92: UART Connection Example



DRAFT FOR REVIEW ONLY

3.4.9 RFU[1:4] Pins

This section provides the pin assignments for the RFU[1:4] interface signals.

Table 33: Pin Descriptions – RFU[1:4]

Signal Name (Primary / Secondary)	Pin #	Baseboard Direction	Signal Description
RFU1 / RFU3, N/C	B68	Input / Output	Reserved future use pins. These pins shall be left as no connect. These pins may also be used as a differential pair for future implementations. In this release of the OCP NIC 3.0 specification, the RFU[1:2] pins are defined on the Primary Connector. RFU[3:4] are defined on the Secondary Connector. A total of two reserved pins are available for the SFF; a total of four reserved pins are available for the LFF.
RFU2 / RFU4, N/C	B69		

3.5 PCIe Bifurcation Mechanism

OCP NIC 3.0 baseboards and OCP NIC 3.0 cards support multiple bifurcation combinations. Single socket baseboards with a single or multiple root ports, as well as multi-socket baseboards with a single or multiple root ports are supported. The bifurcation mechanism also supports OCP NIC 3.0 cards with a single or multiple end points. These features are accomplished via I/O pins on the Primary and Secondary Connector:

- PRSNTA#, PRSNTB[3:0]#. The PRSNTA# pin shall connect to the PRSNTB# pins as a hard coded value on the OCP NIC 3.0 card. The encoding of the PRSNTB[3:0]# pins allows the baseboard to determine the PCIe Links available on the OCP NIC 3.0 card. PRSNTA# and PRSNTB[3:0]# pins exist for each connector. For the SFF, a baseboard shall read the pins associated with the Primary Connector to determine the card type. For the LFF, a baseboard shall read the pins associated with both the Primary and Secondary Connector to determine the card type.
- BIF[3:0]#. The BIF# pin states shall be controlled by the baseboard to allow the baseboard to override the default end point bifurcation for silicon that support bifurcation. Additional combinatorial logic is required and is specific to the card silicon. The combinatorial logic is not covered in this specification. The BIF[3:0]# pins may optionally be hardcoded for baseboards that do not require a dynamic bifurcation override. BIF[2:0]# pins exist on each connector.

A high level bifurcation connection diagram is shown in Figure 82.

3.5.1 PCIe OCP NIC 3.0 Card to Baseboard Bifurcation Configuration (PRSNTA#, PRSNTB[3:0]#)

The OCP NIC 3.0 card to baseboard configuration mechanism consists of four dual use pins (PRSNTB[3:0]#) on the OCP NIC 3.0 card and a grounded PRSNTA# pin on the baseboard per connector. For the SFF, a baseboard shall read the pins associated with the Primary Connector to determine the card type. For the LFF, a baseboard shall read the pins associated with both the Primary and Secondary Connector to determine the card type. These pins provide card presence detection as well as mechanism to notify the baseboard of the pre-defined PCIe lane width capabilities. The PRSNTB[3:0]# pins are pulled up to +3.3V_EDGE on the baseboard and are active low signals. A state of 0b1111 indicates that no card is present in the connector(s). Depending on the capabilities of the OCP NIC 3.0 card, a selection of PRSNTB[3:0]# signals may be strapped to the PRSNTA# signal and is pulled low by the baseboard. The encoding of the PRSTNB[3:0]# bits is shown in Table 34 for x32, x16 and x8 PCIe cards. While SFF and LFF cards are allowed in an LFF compliant slot, the condition where the Primary Connector PRSNTB[3:0]# equals 0b1111 and the Secondary Connector PRSNTB[3:0]# pins is not equal to 0b1111 is invalid.

3.5.2 PCIe Baseboard to OCP NIC 3.0 Card Bifurcation Configuration (BIF[2:0]#)

Three signals (BIF[2:0]#) are driven by the baseboard to notify requested bifurcation on the OCP NIC 3.0 card silicon. This allows the baseboard to set the lane configuration on the OCP NIC 3.0 card that supports multiple bifurcation options. BIF[2:0]# pins exist on each connector. For the SFF, the BIF[2:0]# pins associated with the Primary Connector are used. For the LFF, the BIF[2:0]# pins associated with both the Primary and Secondary Connector are used to determine the requested bifurcation.

For example, a baseboard that has four separate hosts that support a 4 x4 connection, should appropriately drive the BIF[2:0]# pins per Table 34 and indicate to the SFF OCP NIC 3.0 card silicon to setup a 4 x4 configuration.

As previously noted, the BIF[2:0]# signals require additional combinatorial logic to decode the BIF[2:0]# value and appropriately apply it to the end-point silicon. The combinatorial logic is not covered in the specification as its implementation is specific to the vendor silicon used.

3.5.3 PCIe Bifurcation Decoder

The combination of the PRSNTB[3:0]# and BIF[2:0]# pins deterministically sets the PCIe lane width for a given combination of baseboard and OCP NIC 3.0 cards. Table 34 shows the resulting number of PCIe links and its width for known combinations of baseboards and OCP NIC 3.0 cards. A copy of this bifurcation decoder, along with a detailed PCIe lane, PERST# and REFCLK assignments is available on the OCP NIC 3.0 Wiki site. Please refer to: <https://www.opencompute.org/wiki/Server/Mezz>.

Note 1: Baseboard vendors do not have to support the full complement of cases enumerated in the bifurcation table. Instead, baseboard vendors may choose to support only a subset of cases that is applicable to their hardware topology.

Note 2: The baseboard must disable PCIe lanes during the initialization phase if the number of detected PCIe links are greater than what is supported on the baseboard to prevent a nondeterministic solution. For example, if the baseboard only supports a 1 x16 connection, and the OCP NIC 3.0 card only supports a 2 x8 connection, the baseboard must disable PCIe lanes 8-15 to prevent any potential LTSSM issues during the discovery phase.

Note 3: Due to separate PCIe REFCLKs and power state timing differences in multi-host configurations, Table 34 shows the expected resulting links for a given baseboard and OCP NIC 3.0 card combination.

Table 34: PCIe Bifurcation Decoder for x32, x16, x8, x4, x2 and x1 Card Widths

Line Card Width	Supported PCIe Configurations	Network Card Card Short Not Present	Primary Connector Secondary Connector Bifurcation (LFF only)	1 Upstream Socket 1.2, 4 or 8 Links	2 Upstream Sockets 4 Links	4 Upstream Sockets 8 Links	8 Upstream Sockets 16 Links	16 Upstream Sockets 32 Links	32 Upstream Sockets 64 Links	4 Sockets First 8 PCIe lanes 4.2 Links	8 Sockets First 8 PCIe lanes per connector (Note 1) 8.2 Links	16 Sockets First 8 PCIe lanes per connector (Note 1) 16.2 Links	32 Sockets First 8 PCIe lanes per connector (Note 1) 32.2 Links	64 Sockets First 8 PCIe lanes per connector (Note 1) 64.2 Links
SFF 20C	1x8, 1x4, 1x2, 1x1	0b110	0b111 (not used)	1x8	1x8 (Socket 0 only)	1x4 (Socket 0 only)	1x4 (Socket 0 only)	1x4 (Socket 0 only)	1x4 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)
SFF 20C	1x4, 1x2, 1x1	0b110	0b111 (not used)	1x4	1x4 (Socket 0 only)	1x4 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)
SFF 20C	1x2	0b110	0b111 (not used)	1x2	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)
SFF 20C	1x1	0b101	0b111 (not used)	1x1	1x1 (Socket 0 only)	1x1 (Socket 0 only)	1x1 (Socket 0 only)	1x1 (Socket 0 only)	1x1 (Socket 0 only)	1x1 (Socket 0 only)	1x1 (Socket 0 only)	1x1 (Socket 0 only)	1x1 (Socket 0 only)	1x1 (Socket 0 only)
SFF 20C	1x8 Option B 4x4, 4x2, 4x1	0b101	0b111 (not used)	1x8	1x8 (Socket 0 only)	4x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)
SFF 40C	2x8 Option B 4x4, 4x2, 4x1	0b100	0b111 (not used)	2x8	2x8 (Socket 0 only)	4x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)
SFF 20C	1x8 Option D 2x4, 2x2, 2x1	0b100	0b111 (not used)	1x8	1x8 (Socket 0 only)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)
SFF 40C	1x16 Option D 2x8, 2x4, 2x2, 2x1	0b100	0b111 (not used)	1x16	1x16 (Socket 0 only)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)
SFF 20C	2x4	0b010	0b111 (not used)	2x4	2x4 (Socket 0 only)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)	2x4 (Socket 0 & 1)
SFF 20C	4x2	0b001	0b111 (not used)	4x2	4x2 (Socket 0 only)	4x2 (Socket 0 & 1)	4x2 (Socket 0 & 1)	4x2 (Socket 0 & 1)	4x2 (Socket 0 & 1)	4x2 (Socket 0 & 1)	4x2 (Socket 0 & 1)	4x2 (Socket 0 & 1)	4x2 (Socket 0 & 1)	4x2 (Socket 0 & 1)
SFF 40C	RSVD	0b000	0b111 (not used)	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD
SFF 40C	1x16 Option A 2x8, 2x4, 2x2, 2x1	0b010	0b111 (not used)	1x16	1x16 (Socket 0 only)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)
SFF 40C	2x8 Option A 2x8, 2x4, 2x2, 2x1	0b010	0b111 (not used)	2x8	2x8 (Socket 0 only)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)	2x8 (Socket 0 & 1)
SFF 40C	1x16 Option B 4x4, 4x2, 4x1	0b010	0b111 (not used)	1x16	1x16 (Socket 0 only)	4x4 (Socket 0 & 1)	4x4 (Socket 0 & 1)	4x4 (Socket 0 & 1)	4x4 (Socket 0 & 1)	4x4 (Socket 0 & 1)	4x4 (Socket 0 & 1)	4x4 (Socket 0 & 1)	4x4 (Socket 0 & 1)	4x4 (Socket 0 & 1)
SFF 40C	4x4	0b011	0b111 (not used)	4x4	4x4 (Socket 0 only)	4x4 (Socket 0 & 1)	4x4 (Socket 0 & 1)	4x4 (Socket 0 & 1)	4x4 (Socket 0 & 1)	4x4 (Socket 0 & 1)	4x4 (Socket 0 & 1)	4x4 (Socket 0 & 1)	4x4 (Socket 0 & 1)	4x4 (Socket 0 & 1)
LFF	2x16 Option A 4x8, 4x4, 4x2, 4x1	0b010	0b011	2x16	2x16 (Socket 0 & 2 only)	4x8 (Socket 0 & 4)	4x8 (Socket 0 & 4)	4x8 (Socket 0 & 4)	4x8 (Socket 0 & 4)	4x8 (Socket 0 & 4)	4x8 (Socket 0 & 4)	4x8 (Socket 0 & 4)	4x8 (Socket 0 & 4)	4x8 (Socket 0 & 4)
LFF	4x8 Option A	0b010	0b010	4x8	4x8 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)
LFF	2x16 Option B 4x8, 4x4, 4x2, 4x1	0b010	0b010	2x16	2x16 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)	4x8 (Socket 0 & 2, 4 & 6)
LFF	2x16 Option C 8x4, 8x2, 8x1	0b010	0b010	2x16	2x16 (Socket 0 & 2, 4 & 6)	8x4 (Socket 0 & 2, 4 & 6)	8x4 (Socket 0 & 2, 4 & 6)	8x4 (Socket 0 & 2, 4 & 6)	8x4 (Socket 0 & 2, 4 & 6)	8x4 (Socket 0 & 2, 4 & 6)	8x4 (Socket 0 & 2, 4 & 6)	8x4 (Socket 0 & 2, 4 & 6)	8x4 (Socket 0 & 2, 4 & 6)	8x4 (Socket 0 & 2, 4 & 6)
RSVD	1x32 Option A 1x16, 1x8, 1x4, 1x2, 1x1	0b000	0b000	1x32	1x32 (Socket 0 only)	1x32 (Socket 0 only)	1x32 (Socket 0 only)	1x32 (Socket 0 only)	1x32 (Socket 0 only)	1x32 (Socket 0 only)	1x32 (Socket 0 only)	1x32 (Socket 0 only)	1x32 (Socket 0 only)	1x32 (Socket 0 only)
LFF	1x32 Option B 2x16, 2x8, 2x4, 2x2, 2x1	0b000	0b001	1x32	1x32 (Socket 0 only)	2x16 (Socket 0 & 2)	2x16 (Socket 0 & 2)	2x16 (Socket 0 & 2)	2x16 (Socket 0 & 2)	2x16 (Socket 0 & 2)	2x16 (Socket 0 & 2)	2x16 (Socket 0 & 2)	2x16 (Socket 0 & 2)	2x16 (Socket 0 & 2)
LFF	1x32 Option B 4x8, 4x4, 4x2, 4x1	0b000	0b001	1x32	1x32 (Socket 0 only)	4x8 (Socket 0 & 2)	4x8 (Socket 0 & 2)	4x8 (Socket 0 & 2)	4x8 (Socket 0 & 2)	4x8 (Socket 0 & 2)	4x8 (Socket 0 & 2)	4x8 (Socket 0 & 2)	4x8 (Socket 0 & 2)	4x8 (Socket 0 & 2)
LFF	1x32 Option B 1x16, 1x8, 1x4, 1x2, 1x1	0b000	0b001	1x32	1x32 (Socket 0 only)	1x16 (Socket 0 & 2)	1x16 (Socket 0 & 2)	1x16 (Socket 0 & 2)	1x16 (Socket 0 & 2)	1x16 (Socket 0 & 2)	1x16 (Socket 0 & 2)	1x16 (Socket 0 & 2)	1x16 (Socket 0 & 2)	1x16 (Socket 0 & 2)
LFF	1x32 Option B 9x4, 9x2, 9x1	0b000	0b001	1x32	1x32 (Socket 0 only)	9x4 (Socket 0 & 2)	9x4 (Socket 0 & 2)	9x4 (Socket 0 & 2)	9x4 (Socket 0 & 2)	9x4 (Socket 0 & 2)	9x4 (Socket 0 & 2)	9x4 (Socket 0 & 2)	9x4 (Socket 0 & 2)	9x4 (Socket 0 & 2)

3.5.4 Bifurcation Detection Flow

The following detection flow shall be used to determine the resulting link count and lane width based on the baseboard and OCP NIC 3.0 card configurations.

1. The baseboard shall read the state of the PRSNTB[3:0]# pins for the Primary Connector and Secondary Connector (if applicable). An OCP NIC 3.0 card is present in the system if the resulting value is not 0b1111 on the Primary Connector.
2. Firmware determines the OCP NIC 3.0 card PCIe lane width capabilities per Table 34 by reading the PRSNTB[3:0]# pins.
3. The baseboard reconfigures the PCIe bifurcation on its ports to match the highest common lane width and lowest common link count on the card.
4. For cases where the baseboard request a link count override (such as requesting a 4-host baseboard requesting 4 x4 operation on a supported card that would otherwise default to a 2 x8 case), the BIF[0:2]# pins shall be asserted as appropriate. Asserting the BIF[0:2]# pins assumes the OCP NIC 3.0 card supports the requested link override.
 - Note: For cards that are already powered up, BIF[0:2]# reconfiguration requires a transition back to ID Mode. During this transition, the card power rails are inactive and manageability links may be briefly lost due to the RBT isolation state.
5. The BIF[0:2]# pins must be in their valid states upon the assertion of AUX_PWR_EN.
6. AUX_PWR_EN is asserted. An OCP NIC 3.0 card is allowed a max ramp time T_{APL} between AUX_PWR_EN assertion and NIC_PWR_GOOD assertion.
7. MAIN_PWR_EN is asserted. An OCP NIC 3.0 card is allowed a max ramp time T_{MPL} between MAIN_PWR_EN assertion and NIC_PWR_GOOD reassertion. For cards that do not have a separate AUX and MAIN power domain, this state is an unconditional transition to NIC_PWR_GOOD.
8. The PCIe REFCLK shall become valid a minimum of 100 μ s before the deassertion of PERST#.
9. PERST# shall be deasserted >1 s after NIC_PWR_GOOD assertion as defined in Figure 106. Refer to Section 3.11 for timing details.


3.5.5 PCIe Bifurcation Examples

For illustrative purposes, the following figures show several common bifurcation permutations using a SFF.

3.5.5.1 Single Host (1 x16) Baseboard with a 1 x16 OCP NIC 3.0 Card (Single Controller)

Figure 93 illustrates a single host baseboard that supports x16 with a single controller OCP NIC 3.0 card that also supports x16. The PRSTNB[3:0]# state is 0b0111. The BIF[2:0]# state is 0b000 to set the card as a 1x16 for bifurcation capable controllers. For controllers without bifurcation support, the BIF[2:0] pin connections are not required on the card. The PRSNTB encoding notifies the baseboard that this card is only capable of 1 x16. The single host baseboard determines that it is also capable of supporting 1 x16. The resulting link width is 1 x16.

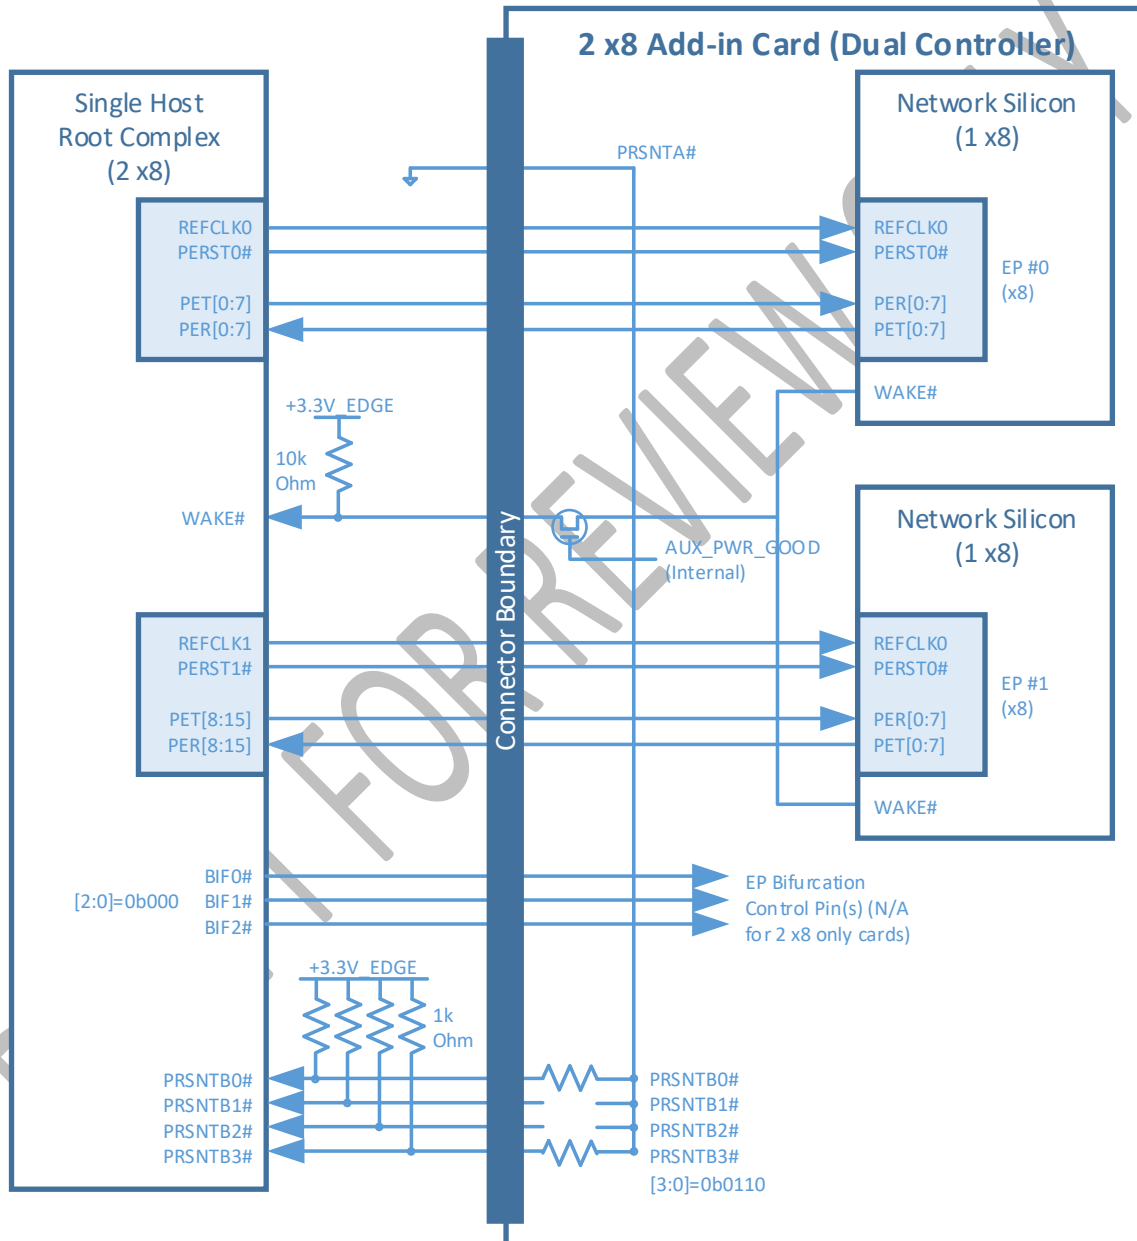
Figure 93: Single Host (1 x16) and 1 x16 OCP NIC 3.0 Card (Single Controller)



3.5.5.2 Single Host (2 x8) Baseboard with a 2 x8 OCP NIC 3.0 Card (Dual Controllers)

Figure 94 illustrates a single host baseboard that supports 2 x8 with a dual controller OCP NIC 3.0 card that also supports 2 x8. The PRSTNB[3:0]# state is 0b0110. The BIF[2:0]# state is 0b000 in this example because the network card only supports a 2x8. The PRSNTB encoding notifies the baseboard that this card is only capable of 2 x8. The single host baseboard determines that it is also capable of supporting 2 x8. The resulting link width is 2 x8.


Figure 94: Single Host (2 x8) and 2 x8 OCP NIC 3.0 Card (Dual Controllers)



3.5.5.3 Quad Host (4 x4) Baseboard with a 4 x4 OCP NIC 3.0 Card (Single Controller)

Figure 95 illustrates a quad host baseboard that supports 4 x4 with a single controller OCP NIC 3.0 card that supports 1 x16, 2 x8 and 4 x4. The PRSTNB[3:0]# state is 0b0100. The BIF[2:0]# state in this example is 0b110 as the end point network controller is forced to bifurcate to 4 x4. The PRSNTB encoding notifies the baseboard that this card is only capable of 1 x16, 2 x8 and 4 x4. The quad host baseboard determines that it is also capable of supporting 4 x4. The resulting link width is 4 x4.


Figure 95: Quad Hosts (4 x4) and 4 x4 OCP NIC 3.0 Card (Single Controller)



3.5.5.4 Quad Host (4 x4) Baseboard with a 4 x4 OCP NIC 3.0 Card (Quad Controllers)

Figure 96 illustrates a quad host baseboard that supports 4 x4 with a quad controller OCP NIC 3.0 card that supports 4 x4. The PRSTNB[3:0]# state is 0b0011. The BIF[2:0]# state is a don't care value as there is no need to instruct the end-point network controllers to a specific bifurcation (each controller only supports 1x4 in this example). The PRSNTB encoding notifies the baseboard that this card is only capable of 4 x4. The quad host baseboard determines that it is also capable of supporting 4 x4. The resulting link width is 4 x4.

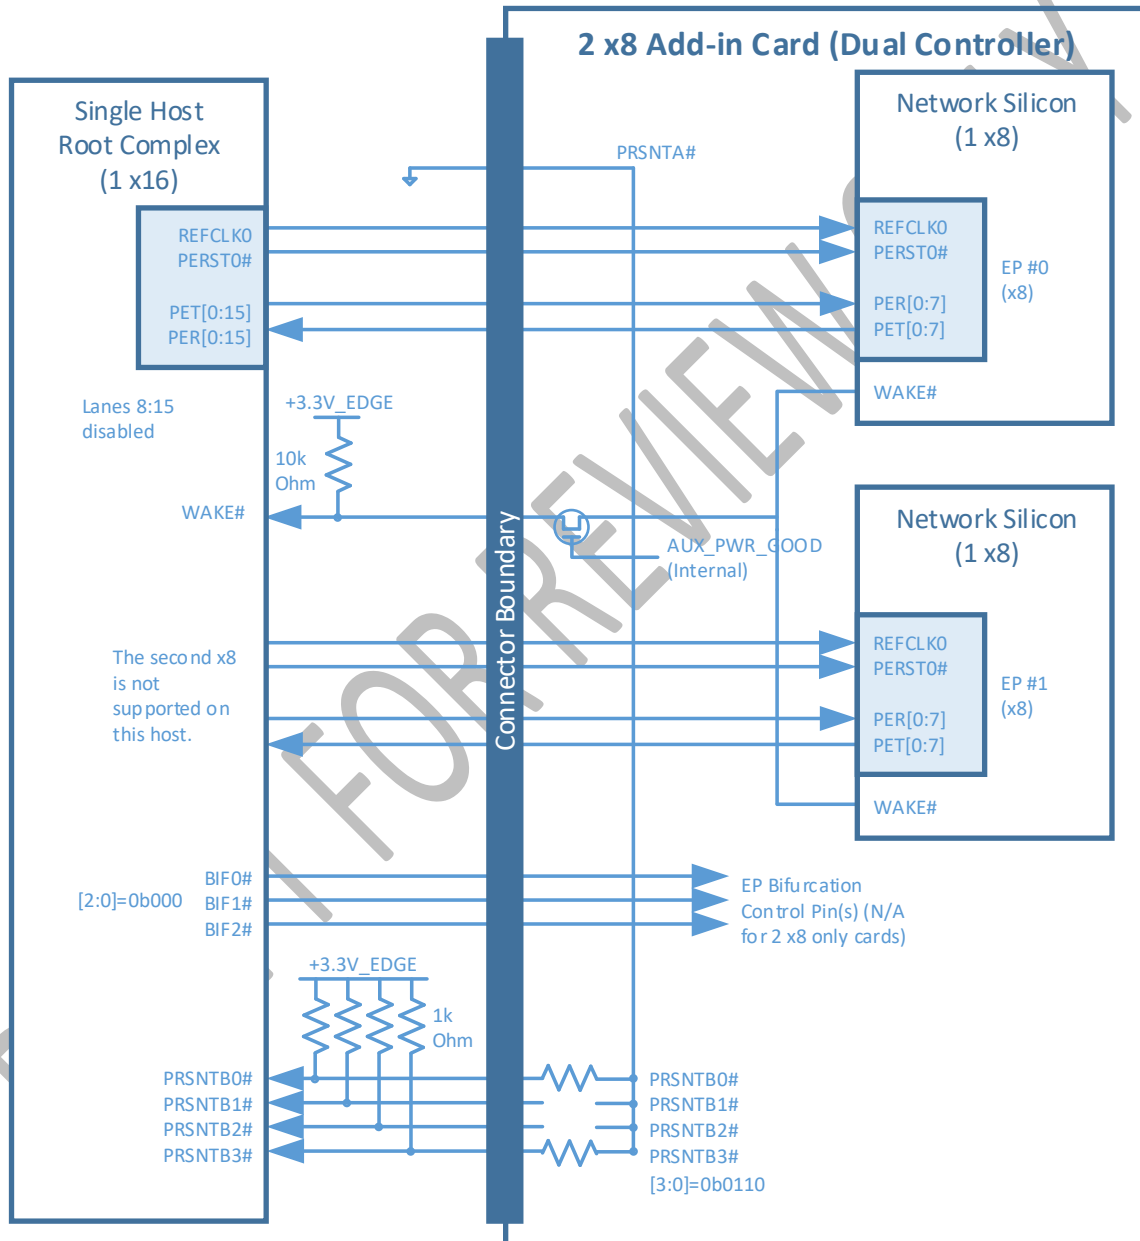
Figure 96: Quad Hosts (4 x4) and 4 x4 OCP NIC 3.0 Card (Quad Controllers)



3.5.5.5 Single Host (1 x16, no Bifurcation) Baseboard with a 2 x8 OCP NIC 3.0 Card (Dual Controller)

Figure 97 illustrates a single host baseboard that supports 1 x16 with a dual controller OCP NIC 3.0 card that supports 2 x8. The PRSTNB[3:0]# state is 0b0110. The BIF[2:0]# state is 0b000 as each silicon instance only supports 1x8. The PRSNTB encoding notifies the baseboard that this card is only capable of 2 x8. The quad host baseboard determines that it is capable of 1x 16, but down shifts to 1 x8. The resulting link width is 1 x8 and only on endpoint 0.

Figure 97: Single Host with no Bifurcation (1 x16) and 2 x8 OCP NIC 3.0 Card (Dual Controllers)



3.6 PCIe REFCLK and PERST# Mapping

The OCP NIC 3.0 specification allows for up to four PCIe REFCLKs and PERST# signals on the Primary Connector and up to two PCIe REFCLKs and PERST# signals on the Secondary Connector. The association of each REFCLK and PERST# is based on the card PCIe Link number and is shown in Table 35. Cards that implement both the Primary and Secondary Connectors have a total of up to 6 available REFCLKs and 6 PERST# signals.

REFCLK[0:3] and PERST[0:3]# are defined for use in this release of the specification. REFCLK[4:5] and PERST[4:5]# are not currently defined for use. The following tables enumerate the REFCLK and PERST# mapping for SFF cards for 1, 2 and 4 links; LFF cards for 1, 2, 4 and 8 links. For a LFF 8 link scenario, the lower x4 “link-a” and upper x4 “link-b” of each x8 lanes are expected to use the same REFCLK and PERST (see Table 37). A 1:2 clock driver circuit is expected on the OCP NIC 3.0 card in this case.

For multi-host use cases, the baseboard may require a multiplexer circuit to direct the Host 1, Host 2 PCIe reference clock to the connector REFCLK1 signal to maintain proper REFCLK associations for a card with two links. Refer to the diagrams in Sections 3.6.1 and 3.6.2.

Table 35: PCIe REFCLK and PERST Associations

REFCLK #	PERST #	Description	Availability (Connector)
REFCLK0	PERST0#	Associated with Link 0.	Primary Connector only.
REFCLK1	PERST1#	Associated with Link 1.	Primary Connector only.
REFCLK2	PERST2#	Associated with Link 2.	Primary Connector only.
REFCLK3	PERST3#	Associated with Link 3.	Primary Connector only.
REFCLK4	PERST4#	Not used.	Secondary Connector only.
REFCLK5	PERST5#	Not used.	Secondary Connector only.

Table 36: SFF PCIe Link / REFCLKn / PERSTn mapping for 1, 2 and 4 Links

Primary Connector			
Lanes [0:3]	Lanes [4:7]	Lanes [8:11]	Lanes [12:15]
Link 0 – x16, REFCLK0, PERST0#			
Link 0 – x8, REFCLK0, PERST0#		Link 1 – x8, REFCLK1, PERST1#	
Link 0 – x4, REFCLK0, PERST0#	Link 1 – x4, REFCLK1, PERST1#	Link 2 – x4, REFCLK2, PERST2#	Link 3 – x4, REFCLK3, PERST3#

Table 37: LFF PCIe Link / REFCLKn / PERSTn mapping for 1, 2, 4 and 8 Links

Primary Connector				Secondary Connector			
Lanes [0:3]	Lanes [4:7]	Lanes [8:11]	Lanes [12:15]	Lanes [16:19]	Lanes [20:23]	Lanes [24:27]	Lanes [28:31]
Link 0 – x32, REFCLK0, PERST0#							
Link 0 – x16, REFCLK0, PERST0#				Link 2 – x16, REFCLK2, PERST2#			
Link 0 – x8, REFCLK0, PERST0#		Link 1 – x8, REFCLK1, PERST1#		Link 2 – x8, REFCLK2, PERST2#		Link 3 – x8 REFCLK3, PERST3#	
Link 0a – x4, REFCLK0, PERST0#	Link 0b – x4, REFCLK0, PERST0#	Link 1a – x4, REFCLK1, PERST1#	Link 1b – x4, REFCLK1, PERST1#	Link 2a – x4, REFCLK2, PERST2#	Link 2b – x4, REFCLK2, PERST2#	Link 3a – x4, REFCLK3, PERST3#	Link 3b – x4, REFCLK3, PERST3#

3.6.1 SFF PCIe REFCLK and PERST# Mapping

The following figures show the Link n, REFCLKn, PERSTn mapping for the SFF with 1, 2 and 4 links as single, dual and quad host configurations. For clarity, the PCIe sideband signals are not illustrated in this section. Please refer to the signal descriptions and associated diagrams for connectivity requirements.

Figure 98: SFF PCIe REFCLK Mapping – Single Host – 1, 2 and 4 links


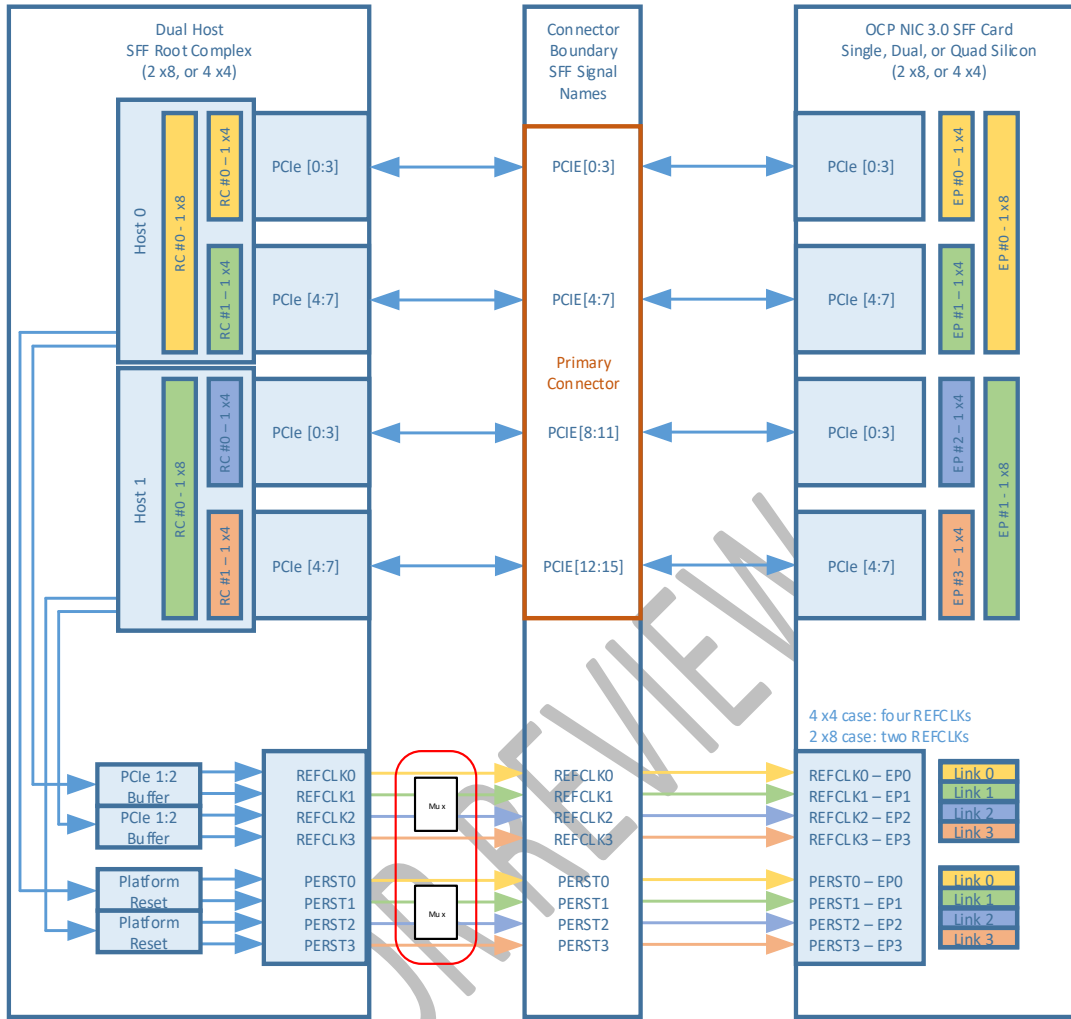



Figure 99: SFF PCIe REFCLK Mapping – Dual Host – 2 and 4 links



Note: For dual host applications that connect to a two link endpoint, the baseboard Host 1 REFCLK0 and PERST0 signal needs to be multiplexed to the REFCLK1 and PERST1 pins of the OCP NIC 3.0 card edge. This ensures the mandated Link n, REFCLKn and PERSTn mappings are maintained.

Figure 100: SFF PCIe REFCLK Mapping – Quad Host – 4 Links



Note: For quad host applications that connect to a two link endpoint, the baseboard Host 2 REFCLK and PERST signal needs to be multiplexed to the REFCLK1 and PERST1 pins of the OCP NIC 3.0 card edge. This ensures the mandated Link n, REFCLKn and PERSTn mappings are maintained.

3.6.2 LFF PCIe REFCLK and PERST# Mapping

The following figures show the Link n, REFCLKn, PERSTn mapping for the LFF with 1, 2 and 4 links as single, dual and quad host configurations. For clarity, the PCIe sideband signals are not illustrated this section. Please refer to the signal descriptions and associated diagrams for connectivity requirements.

Figure 101: LFF PCIe REFCLK Mapping – Single Host – 1, 2 and 4 links



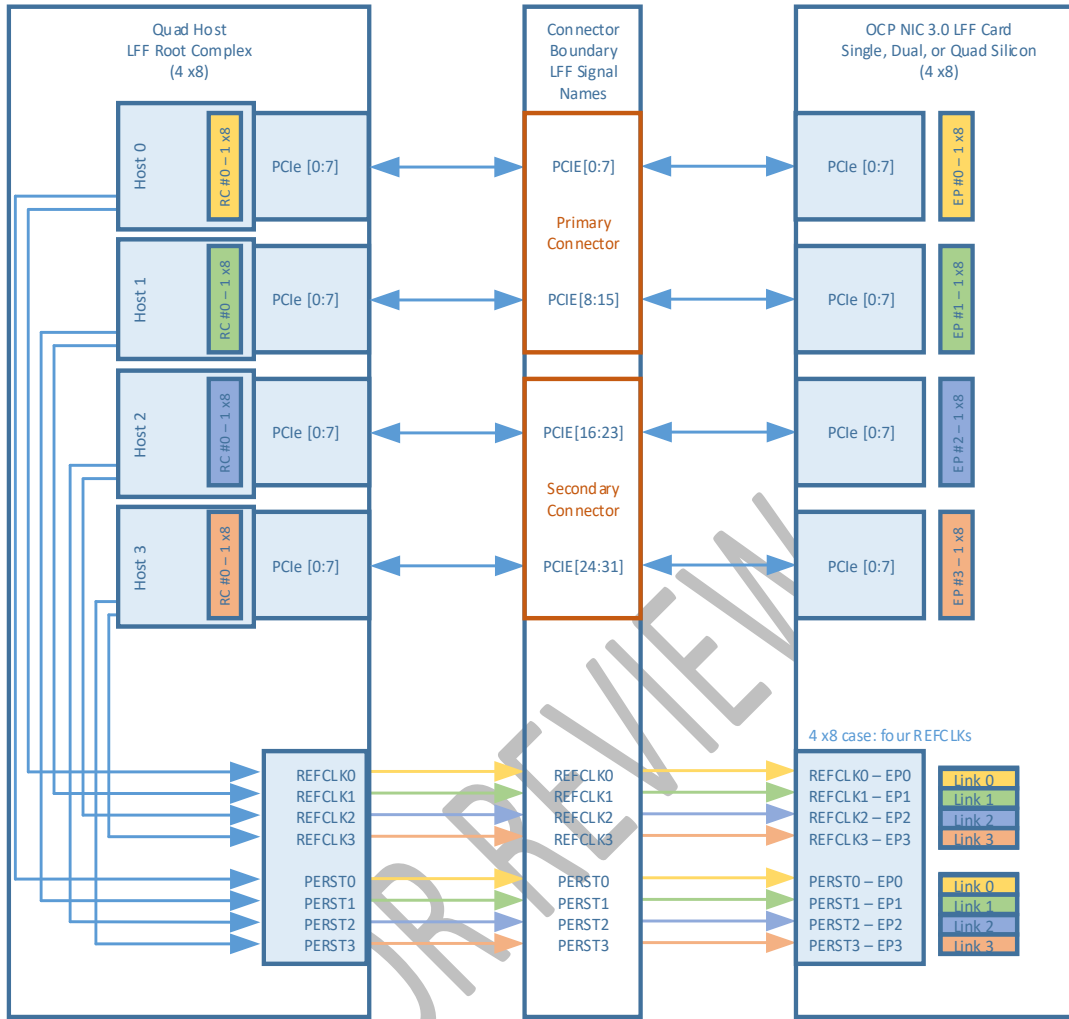


Figure 102: LFF PCIe REFCLK Mapping – Dual Host – 2 and 4 links



Note: For dual host applications that connect to a two link endpoint, the baseboard Host 1 REFCLK0 and PERST0 signal needs to be multiplexed to the REFCLK1 and PERST1 pins of the OCP NIC 3.0 card edge. This ensures the mandated Link n, REFCLKn and PERSTn mappings are maintained.

Figure 103: LFF PCIe REFCLK Mapping – Quad Host – 4 Links



Note: For quad host applications that connect to a two link endpoint, the baseboard Host 2 REFCLK and PERST signal needs to be multiplexed to the REFCLK1 and PERST1 pins of the OCP NIC 3.0 card edge. This ensures the mandated Link n, REFCLKn and PERSTn mappings are maintained.

3.6.3 REFCLK and PERST# Mapping Expansion

For cases where bifurcation are permissible on the baseboard and OCP NIC 3.0 card, an expanded PCIe Bifurcation spreadsheet is available on the OCP Wiki site:

<https://www.opencompute.org/wiki/Server/Mezz>.

Implementers shall use the spreadsheet version that is aligned with this version of the OCP NIC 3.0 specification.

The spreadsheet enumerates all of the supported PCIe link, lane and REFCLK mappings for each supported configuration. The bifurcation decoder is shown in Section 3.5.3.

3.7 Port Numbering and LED Implementations

The OCP NIC 3.0 I/O bracket shall provide port labeling for user identification.

LEDs shall be implemented on the OCP NIC 3.0 I/O bracket when there is sufficient space for local indication. LEDs are typically placed on the primary side. LEDs may be optionally implemented on the secondary side of the card for space constrained implementations. LEDs may be remotely implemented on the card Scan Chain (as defined in Section 3.4.5) for link/activity indication on the baseboard. LED configurations for the local and remote cases are described in the sections below. In all cases, the actual link rate may be directly queried through the management interface.

3.7.1 OCP NIC 3.0 Port Naming and Port Numbering

The numbering of all OCP NIC 3.0 external ports shall start from Port 1. When oriented with the primary side components facing up and viewing directly into the port, Port 1 shall be located on the left hand side. The port numbers shall sequentially increase to the right. Refer to Figure 104 as an example implementation.

3.7.2 OCP NIC 3.0 Card LED Configuration

For low I/O count SFF cards without built in light pipes (such as 1x QSFP, 2x QSFP, 2x SFP, or 2x RJ45), or LFF cards, where additional I/O bracket area is available, the card shall locally implement on-board link/activity indications. The card may additionally implement LEDs on the optional Scan Chain data stream.

For 4x SFP, a permissible LED implementation may include right angle SMT mount LEDs placed on the secondary side of the OCP NIC 3.0 card. The LEDs shall be located below the line side I/O cages.

A QSFP typically operates in a single port mode, but may be optionally configured to operate in dual or quad port modes. The LED definition extends to these modes to indicate the overall operating status of the cage.

Note: Depending on the end faceplate implementation (e.g., with an ejector latch), the secondary side LED implementation may be obstructed and biased to the left to prevent interference with the ejector cam mechanism.

The recommended local (on-card) LED implementation uses two physical LEDs (a bicolored Speed A/Speed B Link LED and a discrete Activity LED). Table 38 describes the OCP NIC 3.0 card LED implementations.

The LEDs shall be uniformly illuminated across the indicator surface. LED surfaces with a diffusion treatment are preferred. For ease of indication within the operating environment, all OCP NIC 3.0 cards shall implement measures to prevent bleed-through between LED indicators and their surrounding chassis components.

Table 38: OCP NIC 3.0 Card LED Configuration with Two Physical LEDs per Port

LED Pin	LED Color	Description
Link	Green	<p>Active low. The LED is illuminated when the signal is low. Bicolor multifunction LED.</p> <p>This LED shall be used to indicate link.</p> <p>When the link is up, then this LED shall be lit and solid. This indicates that the link is established, there are no local or remote faults, and the link is ready for data packet transmission/reception.</p> <p>For all single port cage implementations (SFP, QSFP and BASE-T), the LED indication shall operate as follows:</p> <ul style="list-style-type: none"> • The LED is Green when the port is linked at its maximum speed. • The LED is Amber when the port is linked but not operating at the highest speed. • The LED is off when no link is present. <p>When a multi-lane cage is configured for dual port or quad port operation (e.g., QSFP breakout), the LED indication shall operate as follows:</p> <ul style="list-style-type: none"> • The LED is Green when all ports of the multi-lane cage are linked at its maximum speed. • The LED is Amber when one or more ports of the multi-lane cage are linked, but not operating at the highest speed, or one or more ports are down. • The LED is off when no link is present on all ports of the multi-lane cage. <p>For silicon with limited I/O, the Amber LED may be omitted. In this case, the Green LED shall simply indicate link is up at any configured speed.</p> <p>The illuminated Link LED indicator may be blinked and used for port identification through vendor specific link diagnostic software.</p> <p>The Link LED shall be located on the left hand side or located on the top for each port when the OCP NIC 3.0 card is viewed in the horizontal plane.</p> <p>For serviceability, green LEDs shall emit light at a wavelength between 513 nm and 537 nm while amber LEDs shall emit light at a wavelength between 580 nm and 589 nm.</p>
	Amber	
	Off	

		For uniformity across OCP NIC 3.0 products, all link LEDs shall have their luminance across the total surface area measured in millicandelas (mcd) with an average value between 12 mcd to 18 mcd.
Activity	Green	Active low. The LED is illuminated when the signal is low.
	Off	<p>For all single port cage implementations (SFP, QSFP and BASE-T), the LED indication shall operate as follows:</p> <ul style="list-style-type: none"> • The LED is off when there is no activity. • The LED should blink at the rate of ½ Hz to 5 Hz when there is activity. <p>When a multi-lane cage is configured for dual port or quad port operation (e.g., QSFP breakout), the LED indication shall operate as follows:</p> <ul style="list-style-type: none"> • The LED is off when there is no activity on all ports of the multi-lane cage. • The LED should blink at the rate of ½ Hz to 5 Hz when there is activity on one or more ports of the multi-lane cage. <p>The activity LED shall be located on the right hand side or located on the bottom for each port when the OCP NIC 3.0 card is viewed in the horizontal plane.</p> <p>For serviceability, green LEDs shall emit light at a wavelength between 513 nm and 537 nm.</p> <p>For uniformity across OCP NIC 3.0 products, all activity LEDs shall have their luminance across the total surface area measured in millicandelas (mcd) with an average value between 12 mcd and 18 mcd.</p>

3.7.3 OCP NIC 3.0 Card LED Ordering

For all OCP NIC 3.0 card use cases, each port shall implement the green/amber Link LED and a green activity LED. For I/O limited silicon, the amber LED may be omitted.

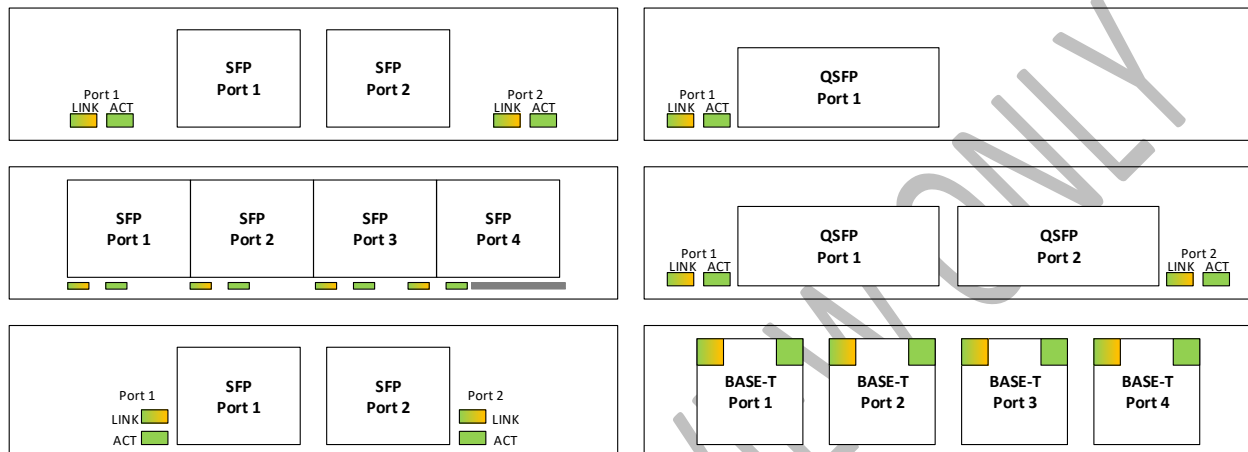
When the OCP NIC 3.0 card is viewed from the horizontal position, and with the primary component side facing up, the Link LED shall be located on the left side and the activity LED shall be located on the right. The LED placement may also make use of a stacked LED assembly, or light pipe in the vertical axis. In this case, the Link Activity LED shall be on the top of the stack, and the Activity LED shall be on the bottom of the stack when viewed from the horizontal position. In all cases, the port ordering shall increase from left to right when viewed from the same horizontal position.

The actual placement of the Link and Activity LEDs on the faceplate may be left up to the discretion of the OCP NIC 3.0 card designer. The LED port association should be clearly labeled on the OCP NIC 3.0

card if the space allows. Similarly, the LED for link and the LED for Activity indication should also be marked on the faceplate. Vendors shall use the largest text permissible for increased readability.

For 4xSFP configurations, the LEDs may be placed on the secondary side of the card using right-angle SMT components. OCP NIC 3.0 designers may opt to use the scan chain LEDs instead or in addition to the on-card indicators.

Figure 104: Port and LED Ordering – Example SFF Link/Activity and Speed LED Placement



Note 1: The example port and LED ordering diagrams shown in Figure 104 are viewed with the card in the horizontal position and the primary side is facing up.

Note 2: The 4xSFP LED implementation is biased to the left to allow clearance for the ejector latch cam.

3.7.4 Baseboard LEDs Configuration over the Scan Chain

A SFF OCP NIC 3.0 card with a fully populated I/O bracket (2x QSFP, 4x SFP or 4x RJ45) does not have sufficient space for primary-side discrete on-board (faceplate) LED indicators. Section 3.7.2 presents an implementation for placing LEDs on the secondary side.

In this scenario, the line side link and activity LED indicators are implemented on the baseboard system via the Scan Chain for remote indication. The Scan Chain bit stream is defined in Section 3.4.5.

The baseboard LED implementation uses two discrete LEDs – a green/amber Link LED and a discrete green Activity. The physical baseboard LED implementation is left up to the baseboard vendor and is not defined in this specification. The LED implementation is optional for baseboards.

For serviceability, green LEDs shall emit light at a wavelength between 513 nm and 537 nm while amber LEDs shall emit light at a wavelength between 580 nm and 589 nm.

At the time of this writing, the Scan Chain definition allows for up to two link and one activity LED per port. A total of up to 8 ports are supported in the Scan Chain. The bit stream defines the LEDs to be active low (on). The Scan Chain LED implementation allows the NIC LED indicators to be remotely located on the OCP NIC 3.0 compliant chassis (e.g., front LED indicators with rear I/O cards).

3.8 Power State Machine

There are four permissible power states for normal operation of the card: NIC Power Off, ID Mode, Aux Power Mode (S5), and Main Power Mode (S0). These four power states are mandatory for each OCP NIC 3.0 card. An optional fifth power state is Programming Mode and allows the FRU EEPROM to be updated in the field under the baseboard control. The normal transition order for these states is shown in Figure 105 and described in detail in the sections below. For simplicity, only the signal transitions resulting in a state change are shown. The available functions per power state are defined in Table 39. The minimum transition time between power states is defined in Section 3.11.

Figure 105: Baseboard Power States




Table 39: Available Card Functions per Power State

Power State	AUX_PWR_EN	MAIN_PWR_EN	PERSTn	FRU	Scan Chain	WAKEn	RBT Link	PCIe Link	+3.3V_EDGE	+12V_EDGE
NIC Power Off	Low	Low	Low							
ID Mode	Low	Low	Low	X	X ¹				X	X ²
Programming Mode	Low	High	Low	X ⁴					X	X ²
Aux Power Mode (S5)	High	Low	Low	X	X	X	X		X	X ³
Main Power Mode (S0)	High	High	High	X	X	X	X	X	X	X

Note 1: Only the PRSNTB[0:3]# scan chain signals are valid in ID mode as the OCP NIC 3.0 card power rails have not yet been enabled via the AUX_PWR_EN/MAIN_PWR_EN signals.

Note 2: The +12V_EDGE rail is on, but the max permissible current draw is up to the ID Mode / Programming Mode current limit defined in Section 3.9.

Note 3: The +12V_EDGE rail is on, but the max permissible current draw is up to the Aux Power Mode current limit defined in Section 3.9.

Note 4: The FRU EEPROM WP pin is low (not write protected) to enable programming using any of the defined methods for controlling the EEPROM WP signal. Refer to the FRU Write Protection Mechanism field in Section 4.10.3.

3.8.1 NIC Power Off

In NIC Power Off mode, all power delivery has been turned off or the card is not physically plugged into the baseboard. Transition to this state can be from any other state.

3.8.2 ID Mode

In the ID Mode, only +3.3V_EDGE is available for powering up the FRU EEPROM and the Scan Chain devices. All OCP NIC 3.0 cards must enter the ID Mode state for FRU EEPROM and scan chain queries immediately following the NIC Power Off state. The baseboard queries the EEPROM and determines the OCP NIC 3.0 device capabilities. The FRU EEPROM content requirements are defined in Section 4.10.3. Only the card PRSNTB[0:3]# bits are valid on the scan chain in this mode as the OCP NIC 3.0 card power rails have not yet been enabled via the AUX_PWR_EN and MAIN_PWR_EN signals. The WAKE#, TEMP_WARN#, TEMP_CRIT#, Link and Activity bits are invalid and should be masked by the baseboard in ID Mode.

The +12V_EDGE rail is not intended to be used in ID Mode, however leakage current may be present. The max leakage is defined in Section 3.9. An OCP NIC 3.0 card shall transition to this mode when AUX_PWR_EN=0 and MAIN_PWR_EN=0.

3.8.3 Aux Power Mode (S5)

The Aux Power Mode provides both +3.3V_EDGE as well as +12V_EDGE is available. +12V_EDGE in Aux mode may be used to deliver power to the OCP NIC 3.0 card, but only up to the Aux mode budget as defined in Table 40. An OCP NIC 3.0 card shall transition to this mode when AUX_PWR_EN=1, MAIN_PWR_EN=0, NIC_PWR_GOOD=1 and the duration (T_{APL}) has passed for the ID-Aux Power Mode

ramp. This guarantees the ID mode to Aux Power Mode transition (as shown in Figure 106) has completed and all Aux Power Mode rails are within operating tolerances. The WAKE#, TEMP_WARN#, and TEMP_CRIT# bits shall not be sampled until these conditions are met.

3.8.4 Main Power Mode (S0)

The Main Power Mode provides both +3.3V_EDGE and +12V_EDGE across the OCP connector. The OCP NIC 3.0 card operates in full capacity. Up to 80 W may be delivered on +12V_EDGE for a SFF Card and up to 150 W for a LFF Card. Additionally, up to 3.63 W is delivered on each +3.3V_EDGE pin. An OCP NIC 3.0 card shall transition to this mode when AUX_PWR_EN=1, MAIN_PWR_EN=1, NIC_PWR_GOOD=1 and the duration (T_{MPL}) has passed for the Aux-Main Power Mode ramp. This guarantees the Aux Power Mode to Main Power Mode transition (as shown in Figure 106) has completed and all Main Power Mode rails are within operating tolerances. The WAKE#, TEMP_WARN#, and TEMP_CRIT# bits shall not be sampled until these conditions are met.

3.8.5 Programming Mode

The Programming Mode only provides +3.3V_EDGE for powering up the FRU EEPROM. This is an optional state that disables the FRU EEPROM write protection mechanism via combinatorial logic and allows the baseboard to reprogram the FRU EEPROM. This state shall only be entered from ID mode if the FRU EEPROM advertises support for this power state as defined in Section 4.10.3.

Additionally, the Aux Power Mode and Main Power Mode SVRs need to remain disabled and NIC_PWR_GOOD shall remain low while the card is in this state. An example discrete logic circuit to accomplish this gating is shown in Figure 89. NIC vendors may implement their own gating circuitry as long as the same result is achieved.

Devices which use GPIO, NC-SI or PLDM for FRU EEPROM WP control may allow programming of the FRU EEPROM when in operating in the Aux Power or Main Power Modes. Usage of the GPIO signal, NC-SI or PLDM to control the WP mechanism are implementation specific and is outside the scope of this specification.

3.9 Power Supply Rail Requirements and Slot Power Envelopes

The baseboard provides +3.3V_EDGE and +12V_EDGE to both the Primary and Secondary Connectors. The rail requirements are leveraged from the PCIe CEM 4.0 specification. For OCP NIC 3.0 implementations, there are five total power envelopes. Four are defined for SFF, and one is defined for LFF. The max current draw is defined in Table 40 for each state and power envelope and is inclusive of the line side transceivers installed on the card.

Table 40: Baseboard Power Supply Rail Requirements – Slot Power Envelopes

Power Rail	15 W Slot SFF Hot Aisle	25 W Slot SFF Hot Aisle	35 W Slot SFF Hot Aisle	80 W Slot SFF Cold Aisle	150 W LFF Cold Aisle
+3.3V_EDGE					
Voltage Tolerance	±9% (max)	±9% (max)	±9% (max)	±9% (max)	±9% (max)
Supply Current					
ID Mode	100 mA (max)	100 mA (max)	100 mA (max)	100 mA (max)	100 mA (max)
Programming Mode	100 mA (max)	100 mA (max)	100 mA (max)	100 mA (max)	100 mA (max)
Aux Mode	1.1 A (max)	1.1 A (max)	1.1 A (max)	1.1 A (max)	2.2 A (max)
Main Mode	1.1 A (max)	1.1 A (max)	1.1 A (max)	1.1 A (max)	2.2 A (max)
Capacitive Load	150 µF (max)	150 µF (max)	150 µF (max)	150 µF (max)	300 µF (max)
+12V_EDGE					
Voltage Tolerance	+8%/-12% (max)	+8/-12% (max)	+8/-12% (max)	+8/-12% (max)	+8/-12% (max)
Supply Current					
ID Mode	50 mA (max)	50 mA (max)	50 mA (max)	50 mA (max)	50 mA (max)
Programming Mode	50 mA (max)	50 mA (max)	50 mA (max)	50 mA (max)	50 mA (max)
Aux Mode	0.7 A (max)	1.1 A (max)	1.5 A (max)	3.3 A (max)	6.3 A (max)
Main Mode	1.25 A (max)	2.1 A (max)	2.9 A (max)	6.6 A (max)	12.5 A (max)
Capacitive Load ³	500 µF (max)	500 µF (max)	500 µF (max)	500 µF (max)	1000 µF (max)

Note 1: While cards may draw up to the published power ratings, the baseboard vendor shall evaluate its cooling capacity for each slot power envelope to determine if a transition to Aux Power Mode is allowed.

Note 2: The maximum slew rate for each OCP NIC 3.0 card shall be no more than 0.1 A/µs per the PCIe CEM specification.

Note 3: Each OCP NIC 3.0 card shall limit the bulk capacitance to the max values published (500 µF for a SFF card, 1000 µF for a LFF card).

Note 4: For systems that implement hot plug, the baseboard shall limit the voltage slew rate such that the instantaneous inrush current shall not exceed the specified max current. The equation is defined in the PCIe CEM specification and is $dV/dt = I/C$; where:

I = max allowed current (A)

C = max allowed bulk capacitance (F)

dV/dt = maximum allowed voltage slew rate (V/s)

The OCP NIC 3.0 FRU definition provides a record for the max power consumption of the card. This value shall be used to aid in determining if the card may be enabled in a given OCP slot. Refer to Section 4.10.3 for the available FRU records.

Additionally, the baseboard shall advertise its slot power limits to aid in the overall board power budget allocation to prevent a high power card from being enabled in a lower power class slot. This is implemented via the Slot Power Limit Control mechanism as defined in the PCIe Base Specification. The end point silicon will power up in a low power state until power is negotiated.

3.10 Hot Swap Considerations for +12V_EDGE and +3.3V_EDGE Rails

Hot plug and hot swap support is optional for baseboard implementers. However, the OCP NIC 3.0 form factor lends itself to potential hot plug and removal events while the baseboard is powered on. These events need to be carefully orchestrated with the operating system and system management entity to prevent a system hang. A surprise extraction may occur in some instances when resources have not been quiesced and the card is removed. Many aspects of the system are involved in processing such an event in both cases. The current scope of this specification does not define an overall hardware or software system architecture to support hot plug. Instead, this specification only highlights the hardware elements that can be utilized to support hot plug for implementations.

The system implementer shall consider the use of hot swap controllers on both the +12V_EDGE and +3.3V_EDGE pins to prevent damage to the baseboard or the OCP NIC 3.0 card. Hot swap controllers limit the in-rush current while also providing overcurrent, undervoltage and overvoltage protection capabilities.

The hot swap controller may gate the +12V_EDGE and +3.3V_EDGE based on the PRSNTB[3:0]# value. Per Section 3.5.3, a card is present in the system when the encoded value is not 0b1111. The PRSNTB[3:0]# may be AND'ed together and connected to the hot swap controller to accomplish this result. Per the OCP NIC 3.0 mechanical definition (Section 3.1.1), the present pins are short pins and engage only when the card is positively seated.

The PRSNTB[3:0]# pins are used to detect an OCP 3.0 NIC card insertion and removal event. The card type detection is described in Section 3.5. Using in-band signaling, the PCIe link may be enabled to periodically train when a card is plugged in. Similarly, the signals may be used to detect a card removal. The card type is determined by querying the FRU data over the SMBus.


At the time of this writing, the DSP0222 Network Controller Sideband Interface (NC-SI) Specification 1.1 and 1.2 does not define a mechanism to discover hot-plug support. Future work is needed for supporting this feature on NCSI over RBT interfaces.

Baseboards that do not support hot insertion, or hot extractions may opt to not implement these features.

3.11 Power Sequence Timing Requirements


The following figure shows the power sequence of PRSNTB[3:0]#, +3.3V_EDGE, +12V_EDGE relative to AUX_PWR_EN, RBT_ISOLATE#, BIF[2:0]#, MAIN_PWR_EN, PERSTn*, and PCIe REFCLK stable on the baseboard. Additionally, the OCP NIC 3.0 card power ramp, and NIC_PWR_GOOD are shown. Please refer to Section 3.4.6 for the NIC_PWR_GOOD definition. Refer to DMTF DSP0222 for details on the NC-SI controller and clock startup requirements.

Figure 106: Power-Up Sequencing – Normal Operation



Note 1: The RBT isolation state is controlled with the baseboard RBT_ISOLATE# signal. This controls the baseboard RBT isolator.
 Note 2: For NC-SI over RBT, the rising edge of NIC_PWR_GOOD shall be treated as V_{REF} is available for NC-SI communication. Refer to timing parameter T4 in the DMTF DSP0222 specification for details.

Figure 107: Power-Down Sequencing – Normal Operation



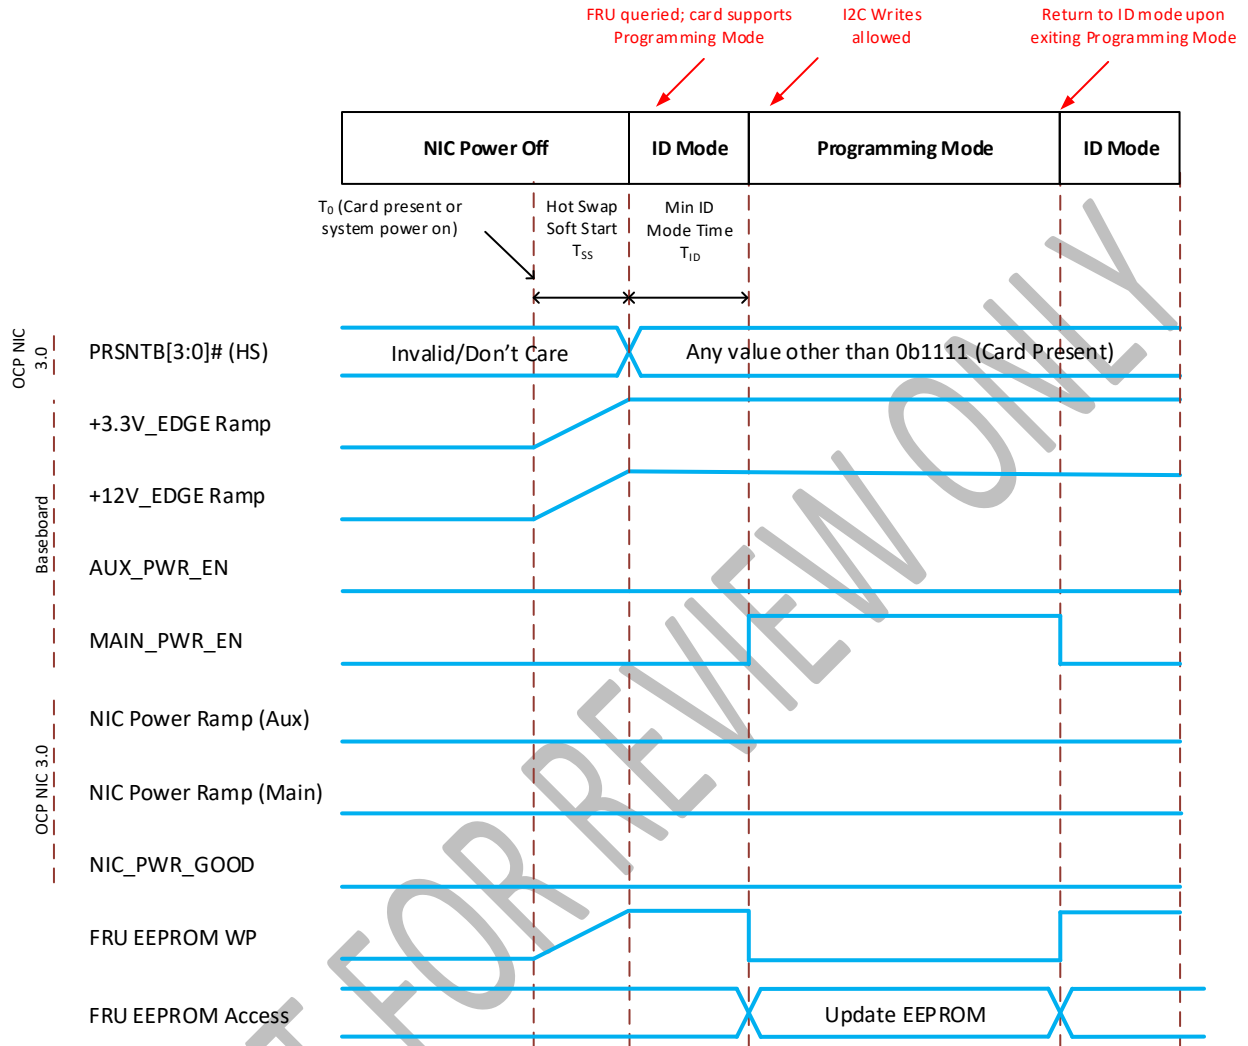
Note 1: REFCLK go inactive after PERST# goes active. (PCIe CEM Section 2.2.3)

Note 2: PERST# goes active before the power on the connector is removed. (PCIe CEM Section 2.2.3)

Note 3: In the case of a surprise power down, PERST# goes active T_{FAIL} after power is no longer stable.

Note 4: The baseboard shall have a minimum delay of T_{CYCLE_SFF} and T_{CYCLE_LFF} for the respective form-factors between AUX_PWR_EN deassertion (power off) and subsequent AUX_PWR_EN assertion (power on) to prevent powering up into a pre-biased condition.

Figure 108: Programming Mode Sequencing



Note 1: All other signals are electrically low and are not shown.

Table 41: Power Sequencing Parameters

Parameter	Value	Units	Description
T _{SS}	20	ms	Maximum time between system +3.3V_EDGE and +12V_EDGE ramp to power stable. This parameter is only applicable to hot swap controller-based implementations. For non-hot swap applications, the +3.3V_EDGE and +12V_EDGE ramp is system dependent.
T _{SCAN_RDY}	1	ms	Minimum time required between +3.3V_EDGE stable to the first scan chain data valid read.
T _{ID}	20	ms	Minimum guaranteed time per spec to spend in ID mode.

T_{APL}	50	ms	Maximum time between AUX_PWR_EN assertion and NIC_PWR_GOOD assertion.
T_{MPL}	50	ms	Maximum time between MAIN_PWR_EN assertion and NIC_PWR_GOOD assertion.
T_{PVPERL}	1	s	Minimum time between NIC_PWR_GOOD assertion in Main Power Mode and PERST# deassertion. For OCP NIC 3.0 applications, this value is >1 second. This is longer than the minimum value specified in the PCIe CEM Specification, Rev 4.0.
T_4	2	s	Maximum time interval from when the network controller NC-SI over RBT interface is able to respond to commands. The parameter T_4 is defined in the DSP0222 specification and is measured from when V_{REF} becomes available. For OCP NIC 3.0, the value T_4 is measured from the deassertion of RBT_ISOLATE#.
$T_{PERST-CLK}$	100	μ s	Minimum Time PCIe REFCLK is stable before PERST# inactive.
T_{FAIL}	500	ns	In the case of a surprise power down, PERST# goes active at minimum T_{FAIL} after power is no longer stable.
T_{AUX-ID}	10	ms	Maximum time from AUX_PWR_EN deassertion to NIC_PWR_GOOD deassertion.
T_{CYCLE_SFF}	30	ms	Minimum time between AUX_PWR_EN deassertion to AUX_PWR_EN reassertion for SFF cards. Delay time allows for OCP NIC 3.0 card capacitors to discharge and prevent reapplying power into a pre-biased condition.
T_{CYCLE_LFF}	60	ms	Minimum time between AUX_PWR_EN deassertion to AUX_PWR_EN reassertion for LFF cards. Delay time allows for OCP NIC 3.0 card capacitors to discharge and prevent reapplying power into a pre-biased condition.

3.12 Digital I/O Specifications

All digital I/O pins on the connector boundary are +3.3 V signaling levels. Table 42 following tables provide the absolute max levels. Refer to the appropriate specifications for the RBT, PCIe and SMBus DC/AC specifications.

Table 42: Digital I/O DC specifications

Symbol	Parameter	Min	Max	Units	Note
V _{OH}	Output voltage		3.6	V	
V _{OL}	Output low voltage		0.8	V	
I _{OH}	Output high current			mA	
I _{OH}	Output low current			mA	
V _{IH}	Input voltage		3.6	V	
V _{IL}	Input low voltage		0.8	V	
I _{OH}	Input current			mA	

Table 43: Digital I/O AC specifications

Symbol	Parameter	Min	Max	Units	Note
T _{OR}	Output rise time			ns	
T _{OF}	Output fall time			ns	

4 Management and Pre-OS Requirements

OCP NIC 3.0 card management is an important aspect to overall system management. This section specifies a common set of management requirements for OCP NIC 3.0 implementations. There are three types of implementations (RBT+MCTP Type, RBT Type, and MCTP Type) depending on the physical sideband management interfaces, transports, and traffic supported over different transports. An OCP NIC 3.0 implementation shall support at least one type of implementation for card management. For a given type of implementation, an OCP NIC 3.0 card shall support type specific requirements described in Sections 4.1 through 4.7.

Table 44: OCP NIC 3.0 Management Implementation Definitions

Management Type	Definition
RBT Type	The RBT Type management interface is exclusive to the Reduced Media Independent Interface (RMII) Based Transport (RBT). The NIC card is required to support the DSP0222 Network Controller Sideband Interface (NC-SI) Specification for this management
RBT+MCTP Type	The RBT+MCTP management interface supports both the RBT and MCTP standards, specifically the DSP0222 Network Controller Sideband Interface (NC-SI) Specification, DSP0236 Management Component Transport Protocol (MCTP) Base Specification, and the associated binding specifications. This is the preferred management implementation for baseboard NIC cards. See MCTP Type below for more details
MCTP Type	The MCTP management interface supports MCTP standards specifically the DSP0236 Management Component Transport Protocol (MCTP) Base Specification and the associated binding specifications.

4.1 Sideband Management Interface and Transport

OCP NIC 3.0 sideband management interfaces are used by a Management Controller (MC) or Baseboard Management Controller (BMC) to communicate with the OCP NIC 3.0 card. Table 45 summarizes the sideband management interface and transport requirements.

Table 45: Sideband Management Interface and Transport Requirements

Requirement	RBT+MCTP Type	RBT Type	MCTP Type
NC-SI compliant RMII Based Transport (RBT) including physical interface defined in Section 10 of DMTF DSP0222	Required	Required	N/A
I ² C compliant physical interface for FRU EEPROM	Required	Required	Required
SMBus 2.0 compliant physical interface	Required	N/A	Required
Management Component Transport Protocol (MCTP) Base (DSP0236 compliant) over MCTP/SMBus Binding (DSP0237 compliant)	Required	N/A	Required
PCIe VDM compliant physical interface	Optional	Optional	Optional
Management Component Transport Protocol (MCTP) Base (DSP0236 compliant) over MCTP/PCIe VDM Binding (DSP0238 compliant)	Optional	Optional	Optional

4.2 NC-SI Traffic

DMTF DSP0222 defines two types of NC-SI traffic: Pass-Through and Control. Table 46 summarizes the NC-SI traffic requirements.

Table 46: NC-SI Traffic Requirements

Requirement	RBT+MCTP Type	RBT Type	MCTP Type
NC-SI Control over RBT (DMTF DSP0222 compliant)	Required	Required	N/A
NC-SI Control over MCTP (DMTF DSP0261 compliant)	Required	N/A	Required
NC-SI Pass-Through over RBT (DMTF DSP0222 compliant)	Required	Required	N/A
NC-SI Pass-Through over MCTP (DMTF DSP0261 compliant)	Optional	N/A	Optional

Note: A Management Controller (MC) is allowed to use NC-SI Control traffic only without enabling NC-SI pass-through.

4.3 Management Controller (MC) MAC Address Provisioning

An OCP NIC 3.0 compliant card that supports NC-SI pass-through shall provision one or more MAC addresses per Package (refer to the Package definition as detailed in the DMTF DSP0222 specification) for Out-Of-Band (OOB) management traffic. The number of MC MAC addresses provisioned is implementation dependent. These MAC addresses are not exposed to the host(s) as available MAC addresses. The MC is not required to use these provisioned MAC addresses. Table 47 summarizes the MC MAC address provisioning requirements. A MAC address algorithm calculator is provided on the OCP NIC 3.0 wiki page and may be downloaded from: <http://www.opencompute.org/wiki/Server/Mezz>

Table 47: MC MAC Address Provisioning Requirements

Requirement	RBT+MCTP Type	RBT Type	MCTP Type
<p>One or more MAC Addresses per package shall be provisioned for the MC.</p> <p>The OCP NIC 3.0 platform may choose to use the NIC vendor allocated MAC addresses for the BMC.</p> <p>The usage of provisioned MAC addresses are BMC implementation specific and is outside the scope of this specification.</p> <p>The recommended MAC address allocation scheme is stated below. This algorithm assumes all of the port MAC addresses are sequentially allocated first, followed by the BMC MAC addresses. For multi-host capable cards, the MAC addresses shall be provisioned for the maximum number of supported hosts.</p> <p>Assumptions:</p>	Required	Required	Optional

<p>1. The number of BMCs or virtual BMCs is the same as the number of hosts (1:1 relationship between each host and the BMC).</p> <p>2. The maximum number of partitions on each port is the same.</p> <p>Variables:</p> <ul style="list-style-type: none"> • <code>num_ports</code> – Number of Ports on the OCP NIC 3.0 card • <code>max_parts</code> – Maximum number of partitions on a port per host • <code>num_hosts</code> – Number of hosts supported by the NIC • <code>first_addr</code> – The MAC address of the first host for the first partition on the first port • <code>host_addr[i]</code> – base MAC address of i^{th} host ($0 \leq i \leq \text{num_hosts}-1$) • <code>bmc_addr[i]</code> – base MAC address of i^{th} BMC ($0 \leq i \leq \text{num_hosts}-1$) <p>Formula:</p> <ul style="list-style-type: none"> • $\text{host_addr}[i] = \text{first_addr} + i * \text{num_ports} * \text{max_parts}$ • The assignment of MAC address used by i^{th} host on port j for the partition k is out of the scope of this specification. • $\text{bmc_addr}[i] = \text{first_addr} + \text{num_ports} * \text{max_parts} * \text{num_hosts} + i * \text{num_ports}$ • The MAC address used by i^{th} BMC on port j, where $0 \leq i \leq \text{num_hosts}-1$ and $0 \leq j \leq \text{num_ports} - 1$ is $\text{bmc_addr}[i] + j$ 			
<p>For multi-host capable OCP NIC 3.0 cards, the MAC addresses shall be provisioned for the maximum number of supported hosts.</p> <p>This algorithm assumes each host and port MAC address combination shall have a corresponding BMC MAC address allocated in sequential order.</p> <p>Variables:</p> <ul style="list-style-type: none"> • <code>num_ports</code> – Number of Ports on the OCP NIC 3.0 card 	<p>Required</p>	<p>Required</p>	<p>Required</p>

<ul style="list-style-type: none"> • <code>max_parts</code> – Maximum number of partitions on a port per host • <code>num_hosts</code> – Number of hosts supported by the NIC • <code>first_addr</code> – The MAC address of the first host for the first partition on the first port • <code>host_addr[i][j]</code> – base MAC address of i^{th} host for the j^{th} port • <code>bmc_addr[i][j]</code> – base MAC address of i^{th} BMC for the j^{th} port • i is the host number ($0 \leq i \leq \text{num_hosts}-1$) • j is the port number ($0 \leq j \leq \text{num_ports}-1$) • k is the partition number ($0 \leq k \leq \text{max_parts}-1$) <p>Formula:</p> <ul style="list-style-type: none"> • $\text{host_addr}[i][j] = \text{first_addr} + i * (\text{max_parts} + 1) * \text{num_ports} + j$ • $\text{bmc_addr}[i][j] = \text{first_addr} + i * (\text{max_parts} + 1) * \text{num_ports} + \text{max_parts} * \text{num_ports} + j$ • The assignment of MAC address used by i^{th} host on port j for the partition k is out of the scope of this specification. 			
<p>Support at least one of the following mechanisms for provisioned MC MAC Address retrieval:</p> <ul style="list-style-type: none"> • NC-SI Control/RBT (DMTF DSP0222 compliant) • NC-SI Control/MCTP (DMTF DSP0261 compliant) <p>Note: This capability is planned to be included in revision 1.2 of the DSP0222 NC-SI specification.</p> <p>For DMTF DSP0222 1.1 compliant OCP NIC 3.0 implementations, MC MAC address retrieval shall be supported using NC-SI OEM commands. An OCP NIC 3.0 implementation, that is compliant with DMTF DSP0222 that defines standard NC-SI commands for MC MAC address retrieval, shall support those NC-SI commands.</p>	Required	Required	Optional

4.4 ASIC Die Temperature Reporting

An OCP NIC 3.0 implementation can have several silicon components including one or more ASICs implementing NIC functions. For the system management, it is important that the die temperatures of these ASIC components can be retrieved over sideband interfaces. The ASIC die temperature reporting

requirements of this section are independent of the transceiver module temperature reporting discussed in Section 4.6.

The temperature reporting interface shall be accessible in Aux Power Mode (S5), and Main Power Mode (S0). Table 48 summarizes temperature reporting requirements. These requirements improve the system thermal management and allow the baseboard management device to access key component temperatures on an OCP NIC 3.0 card. When the temperature reporting function is implemented, it is required that the temperature reporting accuracy is within $\pm 3^{\circ}\text{C}$.

Table 48: Temperature Reporting Requirements

Requirement	RBT+MCTP Type	RBT Type	MCTP Type
ASIC die temperature reporting for a component with TDP ≥ 5 W	Required	Required	Required
ASIC die temperature reporting for a component with TDP < 5 W	Recommended	Recommended	Recommended
When the temperature sensor reporting function is implemented, the OCP NIC 3.0 card shall support PLDM for Platform Monitoring and Control (DSP0248 compliant) for temperature reporting.	Required	Required	Required
When the temperature sensor reporting function is implemented, the OCP NIC 3.0 card shall report upper warning, upper critical, and upper fatal thresholds for PLDM numeric sensors. Note: Refer to DSP0248 for definitions of the upper warning, upper critical, and upper fatal thresholds.	Required	Required	Required
When the temperature reporting function is implemented using PLDM numeric sensors, the temperature tolerance shall be reported as part of the sensor Platform Descriptor Record (PDR) format.	Required	Required	Required
Support for self-shutdown. The purpose of the self-shutdown feature is to “self-protect” the NIC ASIC from permanent damage due to high operating temperatures. The NIC may accomplish this by reducing the power consumed by the ASIC. A BMC may continuously monitor the NIC ASIC temperature and shutdown the NIC ASIC as soon as the temperature reaches a threshold value. There may be scenarios and implementations where the OCP NIC ASIC may be required to self-shutdown without depending on an external entity like the BMC. For those scenarios and implementations, the self-shutdown feature is a final effort in preventing permanent damage to the NIC ASIC at the expense of potential data loss.	Optional	Optional	Optional

<p>If the self-shutdown feature is implemented, the NIC ASIC shall monitor its temperature and shut-down itself as soon as the self-shutdown threshold value is reached. The value of the self-shutdown threshold is implementation specific. It is recommended that the self-shutdown threshold value is higher than the maximum junction temperature of the ASIC implementing the NIC function. It is also recommended that the self-shutdown threshold value is between the critical and fatal temperature thresholds of the ASIC.</p> <p>If the self-shutdown feature is implemented, care shall be taken to ensure that the board power down state is latched and the board does not autonomously resume normal operation.</p> <p>The OCP NIC 3.0 card does not need to know the reason for the NIC ASIC self-shutdown threshold crossing (e.g., fan failure). After the NIC ASIC enters the self-shutdown state, the OCP NIC 3.0 card may not be operational. This might cause the system with the OCP NIC 3.0 card to become unreachable via the NIC.</p> <p>In order to recover the NIC ASIC from the self-shutdown state, the OCP NIC 3.0 card shall go through the NIC ID Mode state as described in Section 3.8.1.</p> <p>If the self-shutdown feature is implemented, the implementation shall provide a mechanism to enable/disable the feature.</p> <p>Note: It is assumed that a system management function will prevent a component from reaching its fatal threshold temperature.</p>			
--	--	--	--

4.5 Power Consumption Reporting

An OCP NIC 3.0 implementation may be able to report the power consumed at the board level. It is important for the system management that the information about the power consumption can be retrieved over sideband interfaces. Table 49 summarizes the power consumption reporting requirements.

Table 49: Power Consumption Reporting Requirements

Requirement	RBT+MCTP Type	RBT Type	MCTP Type
Board Only Estimated Power Consumption Reporting. The value of this field is encoded into the FRU EEPROM contents.	Required	Required	Required

This field reports the board max power consumption value without transceivers plugged into the line side receptacles.			
Pluggable Transceiver Module Power Reporting. The pluggable transceivers plugged into the line side receptacles shall be inventoried (via an EEPROM query) and the Power Class of the module shall be reported.	Required	Required	Required
Board Runtime Power Consumption Reporting. This value shall be optionally reported over the management binding interface. The runtime power value shall report the card edge power.	Optional	Optional	Optional
PLDM for Platform Monitoring and Control (DSP0248 compliant) shall be used for transceiver and board power consumption reporting.	Required	Required	Required

4.6 Pluggable Transceiver Module Status and Temperature Reporting

A pluggable transceiver module is a compact, hot-pluggable transceiver used to connect the OCP 3.0 NIC to an external physical medium. It is important for proper system operation to know the presence and temperature of pluggable transceiver modules. Table 50 summarizes pluggable module status reporting requirements. The transceiver temperature is always reported and is independent of the ASIC die temperature reporting requirements as discussed in Section 4.4. The temperature reporting interface shall be accessible in Aux Power Mode (S5), and Main Power Mode (S0).

Table 50: Pluggable Module Status Reporting Requirements

Requirement	RBT+MCTP Type	RBT Type	MCTP Type
Pluggable Transceiver modules Presence Status and Temperature Reporting	Required	Required	Required
PLDM for Platform Monitoring and Control (DSP0248 compliant) for reporting the pluggable transceiver module presence status and pluggable transceiver module temperature	Required	Required	Required

4.7 Management and Pre-OS Firmware Inventory and Update

An OCP NIC 3.0 implementation can have different types of firmware components for data path, control path, and management path operations. It is desirable that OCP NIC 3.0 implementations support an OS-independent mechanism for the management firmware update. It is desirable that the management firmware update does not require a system reboot for the new firmware image to become active. Table 51 summarizes the firmware inventory and update requirements.

Table 51: Management and Pre-OS Firmware Inventory and Update Requirements

Requirement	RBT+MCTP Type	RBT Type	MCTP Type
Network boot in UEFI driver (supporting both IPv4 and IPv6 addressing for network boot)	Required	Required	Required
UEFI secure boot for UEFI drivers	Required	Required	Required
UEFI Firmware Management Protocol (FMP)	Required	Required	Required
PLDM for Firmware Update (DSP0267 compliant)	Required	Recommended	Required

4.7.1 Secure Firmware

It is highly recommended that an OCP NIC 3.0 card supports a secure firmware feature. In the future versions of the OCP NIC 3.0 specification, the secure firmware feature is intended to be required. When the secure firmware feature is enabled and where export compliance permits, the OCP NIC 3.0 card shall verify firmware components prior to the execution, execute only signed and verified firmware components, and only allow authenticated firmware updates. Where applicable, an OCP NIC 3.0 implementation shall use the guidelines provided in NIST SP 800-193 Platform Resiliency Guidelines for the following secure firmware functions:

- Signed Firmware Updates
- Ensure only valid/authenticated firmware updates can be applied. Refer to: NIST 800-193 Section 3.5 Firmware Update Mechanisms, and 4.1.2 Root of Trust for Update (RTU) and Chain of Trust for Update (CTU)
- Ensure authentication mechanisms cannot be bypassed. Refer to NIST 800-193 Section 4.2 Protection.
- Secure Boot
- Only boot trusted/authenticated firmware: NIST 800-193 4.1.3 Root of Trust for Detection (RTD) and Chain of Trust for Detection (CTD), and Section 4.3 Detection
- Recovery mechanism in case of boot failure: NIST 800-193 Section 4.4 Recovery

4.7.2 Firmware Inventory

The OCP NIC 3.0 card shall allow queries to obtain the firmware component versions, device model, and device ID via in-band and out-of-band interfaces without impacting NIC function and performance of said paths.

4.7.3 Firmware Inventory and Update in Multi-Host Environments

A multi-host capable OCP NIC 3.0 card shall gracefully handle concurrent in-band queries from multiple hosts and out-of-band access from the BMC for firmware component versions, device model, and device ID information.

A multi-host capable OCP NIC 3.0 card shall only permit one entity to perform write accesses to NIC firmware at a time, without creating contention.

A multi-host capable OCP NIC 3.0 card shall gracefully handle exceptions when more than one entity attempts to perform concurrent NIC firmware writes.

4.8 NC-SI Package Addressing and Hardware Arbitration Requirements

NC-SI over RBT is implemented via RMII pins between the MC and the OCP NIC 3.0 card. Protocol and implementation details of NC-SI over RBT can be found in the DMTF DSP0222 standard.

4.8.1 NC-SI over RBT Package Addressing

NC-SI over RBT capable OCP NIC 3.0 cards shall use a unique Package ID per ASIC when multiple ASICs share the single NC-SI physical interconnect to ensure there are no addressing conflicts.

Baseboards use the Slot_ID[1:0] values on the Primary Connector for this identification. The value of Slot_ID[1:0] is determined by the encoding shown in Table 52. SLOT_ID[1:0] is statically set high or low on the baseboard and is available on the OCP Bay portion of the Primary Connector.

Table 52: Slot_ID[1:0] to Package ID[2:0] Mapping

Physical Slot (Dec.)	SLOT_ID[1:0]		Package ID[2:0]		
	Pin OCP_A6	Pin OCP_B7	Package ID[2]	Package ID[1]	Package ID[0]
	SLOT_ID1	SLOT_ID0	PhysDev#	SLOT_ID1	SLOT_ID0
Slot 0	0	0	0/1	0	0
Slot 1	0	1	0/1	0	1
Slot 2	1	0	0/1	1	0
Slot 3	1	1	0/1	1	1

Package ID[2:0] is a 3-bit field and is encoded in the NC-SI Channel ID as bits [7:5]. SLOT_ID1 is associated with Package ID[1]. SLOT_ID0 is associated with Package ID[0]. The Package ID[2] value is based on the silicon instance on the same physical OCP NIC 3.0 card. Package ID[2]==0b0 is assigned for physical controller #0. Package ID[2]==0b1 is assigned for physical controller #1. In this case, physical controller #1 on the same card is at an address offset of +0x4. Refer to the specific endpoint device datasheet for details on the Package ID configuration options.

Note: The Package ID[2] field is optionally configurable in the NC-SI specification. If the target silicon hard codes this bit to 0b0, then a card must only implement a single silicon instance to prevent addressing conflicts.

Refer to the DMTF DSP0222 standard for more information on package addressing and Package ID.

4.8.2 Arbitration Ring Connections

For baseboards that implement two or more Primary Connectors, the NC-SI over RBT arbitration ring shall be connected to each other. The arbitration ring shall support operation with one card, or multiple cards installed. Figure 86 shows an example connection with dual Primary Connectors.

4.9 SMBus 2.0 Addressing Requirements

The SMBus provides a low speed management bus for the OCP NIC 3.0 card. The FRU EEPROM is directly connected to the OCP NIC 3.0 card edge on this bus and can be read by the baseboard in the ID Mode, Aux Power Mode and Main Power Mode. Network controllers may utilize the SMBus 2.0 interface for MCTP communications. OCP NIC 3.0 does not support MCTP over I²C due to the use of specific SMBus 2.0 addressing. Proper power domain isolation shall be implemented on the NIC.

4.9.1 SMBus Address Map

OCP NIC 3.0 cards shall support the SMBus Address Resolution Protocol (ARP). This allows for dynamic assignment of slave device addresses. This method automatically resolves address conflicts and eliminates the need for manual address configuration. The SMBus address type can be either a Dynamic and Persistent Address or a Dynamic and Volatile Address. Refer to the SMBus 2.0 specification and Section 6.11 of DSP0237 for details on SMBus address assignment. Due to the prevalent use of SMBus muxes in many baseboard designs, the OCP NIC 3.0 card is discouraged from sending unsolicited messages, which includes the optional “Notify ARP Master” command.

A baseboard implementation may choose to only use fixed addresses for OCP NIC 3.0 cards. The assignment of these fixed addresses is system dependent and is outside the scope of this specification. When fixed addresses are used, then the OCP NIC 3.0 card shall be a “Fixed and Discoverable” SMBus device. Refer to the SMBus 2.0 specification for more details.

All predefined SMBus addresses for OCP NIC 3.0 are shown in Table 53. Baseboard and OCP NIC 3.0 card designers must ensure additional devices do not conflict. The addresses shown are in 8-bit format and represent the read/write address pair.

Table 53: FRU EEPROM Address Map

Physical Slot (Dec.)	SLOT_ID[1:0]		FRU EEPROM Address				
	Pin OCP_A6	Pin OCP_B7	A2	A1	A0	Binary Address	Hex Address
	SLOT_ID1	SLOT_ID0	SLOT_ID1	SLOT_ID0	Fixed		
Slot 0	0	0	0	0	0	0b1010_000X	0xA0/0xA1
Slot 1	0	1	0	1	0	0b1010_010X	0xA4/0xA5
Slot 2	1	0	1	0	0	0b1010_100X	0xA8/0xA9
Slot 3	1	1	1	1	0	0b1010_110X	0xAC/0xAD

4.10 FRU EEPROM

4.10.1 FRU EEPROM Addressing and Size

A FRU EEPROM is implemented on the OCP NIC 3.0 cards and is used by the baseboard to determine the card type and capabilities. The FRU EEPROM is directly connected to the SMBus on the card edge. Only one EEPROM is required for a single physical OCP NIC 3.0 card regardless of the PCIe width or number of physical card edge connectors it occupies. The FRU EEPROM is mandatory and shall be connected to the Primary Connector SMBus.

The FRU EEPROM is readable in all four power states: ID mode, Programming mode, Aux Power mode (S5), and Main Power mode (S0).

The permissible EEPROM addresses are indicated in Table 53. The write/read pair is presented in 8-bit format. The EEPROM shall use double byte addressing and, at minimum, shall be of sufficient size to hold the base FRU contents and any vendor specific information. The double byte write and read accesses are shown in 109 and Figure 110. Refer to the I²C specification for timing details.

Figure 109: FRU EEPROM Writes with Double Byte Addressing

Byte Write



Page Write






Figure 110: FRU EEPROM Reads with Double Byte Addressing


Random Read



Current Address Read



Current Sequential Read




4.10.2 FRU EEPROM Write Protection

The FRU EEPROM should be write protected for production cards. The FRU write protection may optionally be overridden for field updates via the use of a mechanical jumper or switch, a GPIO, controlled via a NC-SI or PLDM based mechanism, or the Programming Mode Power state. An OCP NIC 3.0 card may implement one or more of these write protection implementations. The definitions of these mechanisms are outside the scope of this specification. The NIC vendor write protection mechanism(s) shall be noted in the OEM Record. The FRU shall be writable for Programming mode power state, manufacturing test and during the card development cycle.

The FRU update flow is shown in Figure 111.

Figure 111: FRU Update Flow



4.10.3 FRU EEPROM Content Requirements

The FRU EEPROM shall follow the data format specified in Section 1 of the IPMI Platform Management FRU Information Storage Definition specification. For OCP NIC 3.0, the FRU Information Device shall, at a minimum, contain the Common Header, Board Info Area, Product Info Area and a MultiRecord Area for storing the OEM record. These fields shall be populated in the FRU EEPROM. Where applicable, fields

common to the Board Info and Product Info records shall be populated with the same values so they are consistent.

The OEM record 0xC0 is used to store specific records for the OCP NIC 3.0 and is stored in the MultiRecord area of the FRU layout. For an OCP NIC 3.0 card, the FRU EEPROM OEM record content based on the format defined in Table 54 shall be populated.

Note: Table 54 only shows a portion of the OEM record. The complete record includes a Common Header and valid record checksum as defined in the IPMI Platform Management FRU Information Storage Definition specification.

Table 54: FRU EEPROM Record – OEM Record 0xC0, Offset 0x00

Offset	Length	Description
0	3	<p>Manufacturer ID.</p> <p>For OCP NIC 3.0 compliant cards, the value of this field shall be set to the OCP IANA assigned number. This value is 42623 in decimal or 0x00A67F in hexadecimal. The least significant byte (0x7F) is first at offset 0 and the most significant byte (0x00) is at offset 2.</p>
3	1	<p>OCP NIC 3.0 FRU OEM Record Version.</p> <p>This field indicates the card OEM Record Version. Baseboards shall read this field to determine the OEM Record format. OCP NIC 3.0 cards compliant to this version of the specification shall be set the field to 0x01. Future changes to the OEM Record format will result in an additional record version value added to this list.</p> <p>0x00 – Reserved 0x01 – OCP NIC 3.0 card FRU record released with version 0.90 0x02 – 0xFF – Reserved</p>
4	1	<p>Card Max power (in Watts) in MAIN (S0) mode.</p> <p>The encoded value is the calculated max power of the OCP NIC 3.0 card in the Main Power (S0) mode only and does not include the consumed power by transceivers plugged into the line side receptacle(s).</p> <p>0x00 – 0xFE – Card power rounded up to the nearest Watt for fractional values. 0xFF – Unknown</p>
5	1	<p>Card Max power (in Watts) in AUX (S5) mode.</p> <p>The encoded value is the calculated max power of the OCP NIC 3.0 card in the Aux Power (S5) mode only and does not include the consumed power by transceivers plugged into the line side receptacle(s).</p> <p>0x00 – 0xFE – Card power rounded up to the nearest Watt for fractional values. 0xFF – Unknown</p>
6	1	<p>Hot Aisle Card Cooling Tier with Passive Cables or RJ45.</p> <p>The encoded value reports the OCP NIC 3.0 Card Hot Card Cooling Tier as defined in Section 6.6 for cards with passive pluggable cables or with RJ45 for the line side I/O.</p> <p>0x00 – Reserved 0x01 – Hot Aisle Cooling Tier 1</p>

		<p>0x02 – Hot Aisle Cooling Tier 2 0x03 – Hot Aisle Cooling Tier 3 0x04 – Hot Aisle Cooling Tier 4 0x05 – Hot Aisle Cooling Tier 5 0x06 – Hot Aisle Cooling Tier 6 0x07 – Hot Aisle Cooling Tier 7 0x08 – Hot Aisle Cooling Tier 8 0x09 – Hot Aisle Cooling Tier 9 0x0A – Hot Aisle Cooling Tier 10 0x0B – Hot Aisle Cooling Tier 11 0x0C – Hot Aisle Cooling Tier 12 0x0D – 0xFE – Reserved 0xFF – Unknown</p>
7	1	<p>Cold Aisle Card Cooling Tier with Passive Cables or RJ45.</p> <p>The encoded value reports the OCP NIC 3.0 Card Cold Aisle Cooling Tier as defined in Section 6.6 for cards with passive pluggable cables or with RJ45 for the line side I/O.</p> <p>0x00 – Reserved 0x01 – Cold Aisle Cooling Tier 1 0x02 – Cold Aisle Cooling Tier 2 0x03 – Cold Aisle Cooling Tier 3 0x04 – Cold Aisle Cooling Tier 4 0x05 – Cold Aisle Cooling Tier 5 0x06 – Cold Aisle Cooling Tier 6 0x07 – Cold Aisle Cooling Tier 7 0x08 – Cold Aisle Cooling Tier 8 0x09 – Cold Aisle Cooling Tier 9 0x0A – Cold Aisle Cooling Tier 10 0x0B – Cold Aisle Cooling Tier 11 0x0C – Cold Aisle Cooling Tier 12 0x0D – 0xFE – Reserved 0xFF – Unknown</p>
8	1	<p>Card active/passive cooling.</p> <p>This byte defines if the card has passive cooling (there is no fan on the card) or active cooling (a fan is located on the card).</p> <p>0x00 – Passive Cooling 0x01 – Active Cooling 0x02 – 0xFE – Reserved 0xFF – Unknown</p>
9	2	<p>Hot aisle standby airflow requirement.</p> <p>The encoded value represents the amount of airflow, in LFM, required to cool the card in AUX (S5) mode while operating in a hot aisle environment with an approach air temperature of 45°C. Refer to Section 6 for more information about the thermal and environmental requirements.</p> <p>Byte 9 is the least significant byte, byte 10 is the most significant byte.</p> <p>0x0000-0xFFFFE – LFM required for cooling card in Hot Aisle Operation. 0xFFFF – Unknown.</p>
11	2	<p>Cold aisle standby airflow requirement.</p>

		<p>The encoded value represents the amount of airflow, in LFM, required to cool the card in AUX (S5) mode while operating in a cold aisle environment with an approach air temperature of 35°C. Refer to Section 6 for more information about the thermal and environmental requirements.</p> <p>Byte 11 is the least significant byte, byte 12 is the most significant byte.</p> <p>0x0000-0xFFFE – LFM required for cooling card in Cold Aisle Operation. 0xFFFF – Unknown.</p>
13	1	<p>UART Configuration 1 – Secondary Connector.</p> <p>This byte denotes the UART configuration 1. A value 0x00 means no serial connection is implemented on the Secondary Connector card edge.</p> <p>Bits [2:0] denotes the UART baud rate per the encoding table below. If implemented, the encoded field value defines the default baud rate of the OCP NIC 3.0 card serial port.</p> <p>0b000 – No serial connection 0b001 – 115200 baud 0b010 – 57600 baud 0b011 – 38400 baud 0b100 – 19200 baud 0b101 – 9600 baud 0b110 – 4800 baud 0b111 – 2400 baud</p> <p>Bits [4:3] denotes the number of data bits.</p> <p>0b00 – No serial connection 0b01 – 7 data bits 0b10 – 8 data bits 0b11 – Reserved</p> <p>Bits [7:5] denotes the parity bit character.</p> <p>0b000 – No serial connection 0b001 – None (N) 0b010 – Odd (O) 0b011 – Even (E) 0b100 – Mark (M) 0b101 – Space (S) 0b110 – Reserved 0b111 – Reserved</p>
14	1	<p>UART Configuration 2 – Secondary Connector.</p> <p>This byte denotes the UART configuration 2. A value 0x00 means no serial connection is implemented on the Secondary Connector card edge.</p> <p>Bits [1:0] denotes the number of stop bits.</p> <p>0b00 – No serial connection 0b01 – 1 stop bit 0b10 – 1.5 stop bits 0b11 – 2 stop bits</p> <p>Bits [3:2] denotes the flow control method.</p> <p>0b00 – No serial connection 0b01 – Software handshaking 0b10 – No handshaking 0b11 – Reserved</p> <p>Bits [7:4] are reserved and shall be encoded to a value of 0b0000.</p>

15	1	<p>USB Present – Primary Connector.</p> <p>This byte denotes a USB 2.0 connection is implemented on the Primary Connector card edge.</p> <p>0x00 – No USB 2.0 is present or is not implemented on the card edge. 0x01 – A USB 2.0 connection is implemented on the card edge.</p>
16	1	<p>Manageability Type.</p> <p>This byte denotes the card manageability type and interface used.</p> <p>0x00 – No manageability 0x01 – RBT Type 0x02 – MCTP Type 0x03 – RBT + MCTP Type 0x04-0xFF – Reserved for future use</p>
17	1	<p>FRU Write Protection Mechanism.</p> <p>This byte defines the FRU write protection mechanism implemented on the OCP NIC 3.0 card. Multiple FRU write protection mechanisms may be simultaneously supported on the OCP NIC 3.0 card. Refer to Section 4.10.2 for details.</p> <p>Bit 0 – FRU EEPROM is not write protected. Bit 1 – FRU EEPROM is statically write protected. No field updates permissible. Bit 2 – FRU EEPROM is write protected via a mechanical jumper or switch. Bit 3 – FRU EEPROM write protection is controlled via network silicon GPIO. Bit 4 – FRU EEPROM write protection is controlled via an NC-SI or PLDM based mechanism. Bit 5 – FRU EEPROM write protection is controlled by the Programming Mode power state. Bits [6:7] – Reserved for future use. 0xFF – Reserved.</p>
18	1	<p>Programming Mode Power State supported.</p> <p>This byte defines support for the Programming Mode power state. If supported, the FRU EEPROM write protect mechanism is disabled for field updates. Refer to Section 3.8 for details on the Programming Mode power state.</p> <p>0x00 – Programming Mode power state not supported. 0x01 – Programming Mode power state supported. 0x02 – 0xFE – Reserved for future use. 0xFF – Reserved.</p>
19	1	<p>Hot Aisle Card Cooling Tier with Active Cables.</p> <p>The encoded value reports the OCP NIC 3.0 Card Hot Aisle Card Cooling Tier as defined in Section 6.6 for cards with active pluggable cables for the line side I/O.</p> <p>0x00 – Card only supports passive cables (e.g., RJ45) 0x01 – Hot Aisle Cooling Tier 1 0x02 – Hot Aisle Cooling Tier 2 0x03 – Hot Aisle Cooling Tier 3 0x04 – Hot Aisle Cooling Tier 4 0x05 – Hot Aisle Cooling Tier 5 0x06 – Hot Aisle Cooling Tier 6 0x07 – Hot Aisle Cooling Tier 7 0x08 – Hot Aisle Cooling Tier 8</p>

		<p>0x09 – Hot Aisle Cooling Tier 9 0x0A – Hot Aisle Cooling Tier 10 0x0B – Hot Aisle Cooling Tier 11 0x0C – Hot Aisle Cooling Tier 12 0x0D – 0xFE – Reserved 0xFF – Unknown</p>
20	1	<p>Cold Aisle Card Cooling Tier with Active Cables.</p> <p>The encoded value reports the OCP NIC 3.0 Card Cold Aisle Card Cooling Tier as defined in Section 6.6 for cards with active pluggable cables for the line side I/O.</p> <p>0x00 – Card only supports passive cables (e.g., RJ45) 0x01 – Cold Aisle Cooling Tier 1 0x02 – Cold Aisle Cooling Tier 2 0x03 – Cold Aisle Cooling Tier 3 0x04 – Cold Aisle Cooling Tier 4 0x05 – Cold Aisle Cooling Tier 5 0x06 – Cold Aisle Cooling Tier 6 0x07 – Cold Aisle Cooling Tier 7 0x08 – Cold Aisle Cooling Tier 8 0x09 – Cold Aisle Cooling Tier 9 0x0A – Cold Aisle Cooling Tier 10 0x0B – Cold Aisle Cooling Tier 11 0x0C – Cold Aisle Cooling Tier 12 0x0D – 0xFE – Reserved 0xFF – Unknown</p>
21	1	<p>Transceiver Reference Power Level.</p> <p>The encoded value denotes the reference power envelope for active transceivers that was used to determine the card thermal tier. The active cable power class is defined in the respective module management interface specifications:</p> <p>SFF-8472 for SFP modules – Power level 1 through 3 SFF-8636 for QSFP modules – Power Classes 1 through 7.</p> <p>0x00 – Passive Cable 0x01 – QSFP Active cable Power Class 1 (1.5 W max) / SFP Level 1 (1.0 W max) 0x02 – QSFP Active cable Power Class 2 (2.0 W max) / SFP Level 2 (1.5 W max) 0x03 – QSFP Active cable Power Class 3 (2.5 W max) / SFP Level 3 (2.0 W max) 0x04 – QSFP Active cable Power Class 4 (3.5 W max) 0x05 – QSFP Active cable Power Class 5 (4.0 W max) 0x06 – QSFP Active cable Power Class 6 (4.5 W max) 0x07 – QSFP Active cable Power Class 7 (5.0 W max) 0x08 – 0xFF – Reserved</p>
22	1	<p>Transceiver Reference Temperature Level.</p> <p>The encoded value denotes the reference active transceiver temperature limit, in degrees Celsius, that was used to determine the card thermal tier.</p> <p>0x00 – Passive Cable 0x01 – 0xFE – Transceiver max operating temperature 0xFF – Reserved</p>
23	1	<p>Card Thermal Tier with Local Fan Fail.</p>

		<p>The encoded value denotes the required thermal tier when the OCP NIC 3.0 active cooling fails.</p> <p>0x00 – Card requires the same thermal tier when the active cooling fan fails or the OCP NIC 3.0 card uses passive cooling.</p> <p>0x01 – Card requires Cooling Tier 1 with card fan failure</p> <p>0x02 – Card requires Cooling Tier 2 with card fan failure</p> <p>0x03 – Card requires Cooling Tier 3 with card fan failure</p> <p>0x04 – Card requires Cooling Tier 4 with card fan failure</p> <p>0x05 – Card requires Cooling Tier 5 with card fan failure</p> <p>0x06 – Card requires Cooling Tier 6 with card fan failure</p> <p>0x07 – Card requires Cooling Tier 7 with card fan failure</p> <p>0x08 – Card requires Cooling Tier 8 with card fan failure</p> <p>0x09 – Card requires Cooling Tier 9 with card fan failure</p> <p>0x0A – Card requires Cooling Tier 10 with card fan failure</p> <p>0x0B – Card requires Cooling Tier 11 with card fan failure</p> <p>0x0C – Card requires Cooling Tier 12 with card fan failure</p> <p>0x0D – FF – Reserved</p>
24:30	7	<p>Reserved for future use.</p> <p>Set each byte to 0xFF for this version of the specification.</p>
31	1	<p>Number of Physical Controllers (N).</p> <p>This byte denotes the number of SMBus connected physical controllers on the OCP NIC 3.0 card. If N=0, no controllers exist on this OCP NIC 3.0 card and this is the last byte in the FRU OEM Record.</p> <p>If N≥1, then the controller UDID records below shall be included for each controller N.</p>
32:47	16	<p>Controller 1 UDID (if applicable).</p> <p>This field reports the Controller 1 Unique Device Identifier (UDID) and is used to aid in the dynamic slave address assignment over the SMBus Address Resolution Protocol. The format of the UDID string is defined in the SMBus 2.0 specification.</p> <p>This field shall list the most significant byte first (to align the FRU order to the reported UDID order on the SMBus). This field is populated with the UDID for Controller 1.</p> <p>This field is omitted and is of zero length if controller 1 is not present.</p>
48:63	16	Controller 2 UDID (if applicable). See Controller 1 UDID description above.
64:79	16	Controller 3 UDID (if applicable). See Controller 1 UDID description above.
80:95	16	Controller 4 UDID (if applicable). See Controller 1 UDID description above.
96:111	16	Controller 5 UDID (if applicable). See Controller 1 UDID description above.
112:127	16	Controller 6 UDID (if applicable). See Controller 1 UDID description above.

4.10.4 FRU Template

A FRU template is provided as a baseline implementation example. This FRU template contains the IPMI Platform Management FRU Information Storage Definition Product Info, Board Info records as well as the OEM record for OCP NIC 3.0.

The FRU template file may be downloaded from the OCP NIC 3.0 Wiki site:
<http://www.opencompute.org/wiki/Server/Mezz>.

5 Routing Guidelines and Signal Integrity Considerations

5.1 NC-SI over RBT

The NC-SI over RBT requirements in this section apply to both the SFF and LFF OCP NIC 3.0 cards. Designers shall use the appropriate SFF or LFF timing parameters for the design calculations.

The overall end-to-end timing budget is computed by the formula below. The values of each parameter are shown in Table 55 and in DSP0222. The overall BMC pad to ASIC pad timing budget is 3 ns assuming the values shown. This value is inclusive of the RBT isolation buffer on the baseboard, propagation delay through the OCP connector, clock jitter or any clock buffers that may be implemented on the OCP NIC 3.0 card. The addition of these components subtract from the total available budget.


$$\begin{aligned}
 \text{Total timing budget} &= T_{\text{CLK}} - T_{\text{CO}(\text{max})} - T_{\text{SU}(\text{min})} - T_{\text{SKEW}(\text{max})} \\
 &= 20 \text{ ns} - 12.5 \text{ ns} - 3 \text{ ns} - 1.5 \text{ ns} \\
 &= 3 \text{ ns}
 \end{aligned}$$

Table 55: NC-SI over RBT Timing Parameters

Parameter	Value	Description
T _{PD,BUDGET}	3000 ps	Total propagation delay budget between BMC and the target ASIC.
T _{BASEBOARD,SFF}	2100 ps	Max permissible propagation delay on a SFF baseboard.
T _{BASEBOARD,LFF}	1650 ps	Max permissible propagation delay on a LFF baseboard.
T _{NIC,SFF}	900 ps	Max permissible propagation delay for a SFF OCP NIC 3.0 card.
T _{NIC,LFF}	1350 ps	Max permissible propagation delay for a LFF OCP NIC 3.0 card.
T _{CLK}	20 ns	50 MHz REF_CLK frequency period.
T _{CO(max)}	12.5 ns	Clock-to-out value per DSP0222
T _{SU(min)}	3 ns	RBT single ended signal data setup to REF_CLK rising edge
T _{SKEW(max)}	1.5 ns	Max permissible clock REF_CLK skew between any two devices in the system – BMC and target NIC ASIC(s).
T _{RBT_ISOLATOR(max)}	-	Baseboard RBT isolator propagation delay. Value is dependent on device selection. Baseboard implementers shall include this value in the timing budget calculations.
T _{CONNECTOR,STRADDLE}	110 ps	OCP NIC 3.0 straddle mount connector propagation delay. This value is the assumed worst case value for all straddle mount connector configurations.
T _{CONNECTOR,RA}	130 ps	OCP NIC 3.0 right angle connector propagation delay. This value is the assumed worst case value for all right angle connector configurations.
T _{CLK_BUF(max)}	-	OCP NIC 3.0 clock buffer propagation delay. Value is dependent on device selection. OCP NIC 3.0 implementers shall include this value in the timing budget calculations.

T_{JITTER_REF}	-	Baseboard reference clock generator cycle-to-cycle clock jitter. Value is dependent on device selection. Baseboard implementers shall include this value in the timing budget calculations.
T_{JITTER_NIC}	-	OCP NIC 3.0 clock buffer cycle-to-cycle clock jitter (if implemented on the NIC). Value is dependent on device selection. OCP NIC 3.0 implementers shall include this value in the timing budget calculations.

Figure 112: NC-SI over RBT Timing Budget Topology



The following sections define the portion of the overall propagation delay budget allocated to the baseboard and to the OCP NIC 3.0 card, as well as additional requirements for each. Baseboard and NIC implementers shall analyze their design to ensure the timing budget is not violated.

The traces shall be implemented as 50 Ohm $\pm 15\%$ impedance controlled nets. Baseboard and NIC designers are encouraged to follow the guidelines defined in the RMI and NC-SI specifications for physical routing. Refer to Section 3.4.4 and the DSP0222 specification for example interconnect topologies.

5.1.1 SFF Baseboard Requirements

The SFF baseboard is allocated a maximum propagation time of 2100 ps between the BMC and the connector edge. NC-SI over RBT isolation buffers are required on the baseboard. The requirements for additional add-in card loading are reduced. The available timing budget for the SFF baseboard is computed by the formula below.

$$T_{BASEBOARD,SFF} = 2100 \text{ ps} - T_{RBT_ISOLATOR} - T_{CONNECTOR}$$

The skew requirement defines the max permissible clock skew (T_{SKEW}) between any two system devices. The T_{SKEW} calculation is computed by the formula below. This applies to both the devices on the baseboard and the NIC. L1 is the REF_CLK segment from the baseboard 50 MHz reference clock generator to the BMC. L2 is the REF_CLK segment between the baseboard clock generator to the OCP

NIC 3.0 connector and L3 is the segment between the SFF OCP NIC 3.0 card gold fingers and the target ASIC. Refer to Figure 112 for details. The max permissible value of L3 is $T_{NIC,SFF} = 900$ ps as discussed in Section 5.1.3. Baseboard vendors shall take this value into consideration when analyzing the available timing budget.

$$T_{SKEW(max)} \leq | L1 - (L2 + T_{CONNECTOR} + L3) |$$

5.1.2 LFF Baseboard Requirements

The LFF baseboard is allocated a maximum propagation time of 1650 ps between the BMC and the connector edge. NC-SI over RBT isolation buffers are required on the baseboard. The requirements for additional add-in card loading are reduced. The available timing budget for the LFF baseboard is computed by the formula below.

$$T_{BASEBOARD,LFF} = 1650 \text{ ps} - T_{RBT_ISOLATOR} - T_{CONNECTOR}$$

Similar to SFF, the LFF clock skew parameter $T_{SKEW(max)}$ must not be exceeded. The max permissible value of L3 is $T_{NIC,LFF} = 1350$ ps. Refer to Section 5.1.1 for the skew computation requirements.

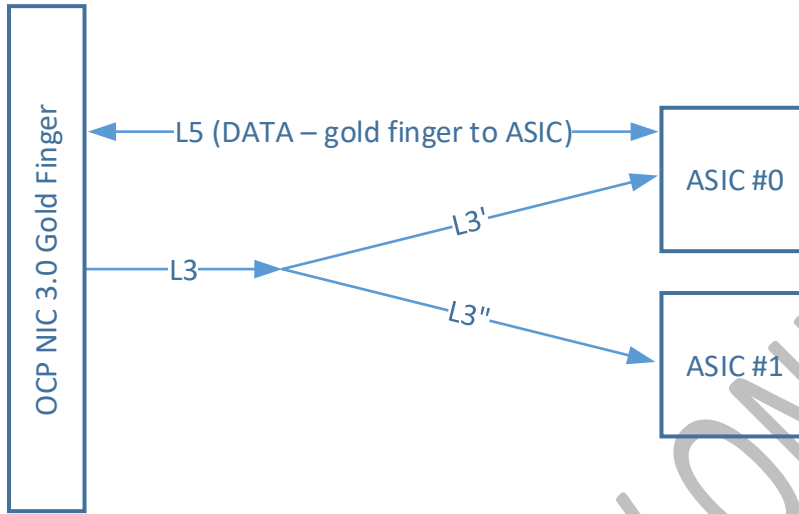
5.1.3 SFF OCP NIC 3.0 Card Requirements

The SFF OCP NIC 3.0 card is allocated a maximum propagation time of $T_{NIC,SFF} = 900$ ps between the card gold finger and the ASIC pad for both the clock and data signals. The total card propagation delay from the RBT clock input towards the card to the RBT outputs from the card shall be less than 14.3 ns ($T_{CO(max)} + 2 \times T_{NIC,SFF}$) when measured at the corresponding SFF OCP NIC 3.0 gold fingers. This can be achieved by using an ASIC with the worst case Clock-to-Out ($T_{CO,MAX}$) value of 12.5 ns specified by DSP0222, and each SFF OCP NIC 3.0 card-side RBT signal not exceeding a max propagation time of 900 ps.

This propagation delay is equivalent to a max length of 5.1 inches assuming standard FR4 material with a propagation delay of 175 ps/in. Additional trace length may be achieved with the use of a higher propagation velocity material (e.g., material with a lower dielectric constant) on the baseboard and OCP NIC 3.0 card or simultaneously using both BMC and ASIC devices with an improved timing from Clock-to-Out ($T_{CO,MAX}$) value compared to the published value of 12.5 ns in DSP0222. For NIC implementations with clock buffers, the propagation delay of the buffer needs to be included in this timing budget (i.e., $L3 + T_{CLK_BUF} + L3' + T_{CO,MAX} + L5$ shall be less than 14.3 ns) as shown in Figure 114.


If multiple ASICs are utilized, the RBT_CLK_IN signal may be routed with a T-topology as shown in Figure 113. The trace length would be calculated as the delay summation of segment $L3 + L3'$ for ASIC #0 and $L3 + L3''$ for ASIC #1. The data path delay to the ASIC is L5.

Figure 113: NC-SI over RBT Propagation Delay Matching for Two Target ASICs – No Clock Buffer



A clock buffer is optionally permitted if the NIC timing budget is not violated. This is shown in Figure 114. In this case, the trace length would be calculated as the delay summation of trace segment L3 + T_{CLK_BUF} + L3' for ASIC #0, and L3 + T_{CLK_BUF} + L3'' for ASIC #1.

Figure 114: NC-SI over RBT Propagation Delay Matching for Two Target ASICs – Clock Buffer



5.1.4 LFF OCP NIC 3.0 Card Requirements

Similar to the SFF, a LFF OCP NIC 3.0 Card is allocated a maximum propagation time of $T_{NIC,LFF} = 1350$ ps between the card gold finger and the ASIC pad for both the clock and data signals. The segment L3 between the LFF OCP NIC 3.0 card gold fingers and the ASIC shall not exceed this propagation delay for a single and multi-ASIC target implementations. Refer to Section 5.1.3 for computation and topology considerations.

5.2 SMBus 2.0

For the SMBus 2.0 interface, both the baseboard and OCP NIC 3.0 designers shall follow the applicable requirements found in the SMBus, PCIe CEM specifications. This applies to the routing guidelines, SI

considerations, bus operational speed range, capacitive loading, and range of pull up resistance values. Doing so allows the baseboard suppliers to design a SMBus interface that is compatible with OCP NIC 3.0 products.

5.3 PCIe

OCP NIC 3.0 card suppliers shall follow the PCIe routing specifications. Refer to the PCIe CEM and PCIe Base specifications for end-to-end channel signal integrity considerations.

5.3.1 Channel Requirements

The OCP NIC 3.0 PCIe channel requirements align with the electrical budget and constraints as detailed in the PCI Express® CEM 4.0 Rev 1.0 and PCI Express Base Specification Rev 4.0. Exceptions or clarifications to the referenced specifications are noted in the sections below.

5.3.1.1 REFCLK requirements

REFCLK requirements are detailed in the PCI Express CEM 4.0 Rev 1.0 Section 2.1.

5.3.1.2 Add-in Card Electrical Budgets

This section defines the OCP NIC 3.0 card channel budget from the gold finger edge to the end point silicon. The values listed below are shown for reference and mirrors that of the PCIe CEM 4.0 specification.

Table 56: PCIe Electrical Budgets

Parameter	PCIe CEM 4.0 Rev 1.0 Specification Section
AC coupling capacitors	Section 4.7.1
Insertion Loss Values (Voltage Transfer Function)	Section 4.7.2 and Appendix A. Section 4.7.10 for 16 GT/s
Jitter Values	Section 4.7.3 for 8 GT/s and 16 GT/s. Also refer to the PCIe Base Specification Section 8.3.5
Crosstalk	Section 4.7.4
Lane-to-lane skew (S_A) for Add-in cards	Section 4.7.5
Transmitter Equalization	Section 4.7.6 and PCIe Base Spec Chapter 9
Skew within a differential pair	Section 4.7.7
Differential data trace impedance	Section 4.7.8
Differential data trace propagation delay	Section 4.7.9

5.3.1.3 Baseboard Channel Budget

The baseboard channel budget directly follows the PCI Express CEM 4.0 Rev 1.0 specification. Details of the budget are outside of the scope of this specification.

5.3.1.4 SFF-TA-1002 Connector Channel Budget

Reference the SFF-TA-1002 Revision 1.1 or later.

5.3.1.5 Differential Impedance (Informative)

For PCIe transmit and receive differential pairs, the target impedance is 85 Ohms $\pm 10\%$.

For the PCIe REFCLKs, the target impedance is 100 Ohms $\pm 10\%$.

5.3.2 Test Fixtures


Test Fixtures are designed using the PCIe CEM 4.0 CLB and CBB. The fixtures host interface has been modified to the OCP connector standard and there are two version of the fixtures, one for Gen 3 PCIe and one for Gen 4 PCIe. Careful attention has been placed on these fixtures to help insure that standard test equipment automation should work without significant modification.

Table 57: PCIe Test Fixtures for OCP NIC 3.0

Test Fixture	PCIe Generation	PCB Material
Load Board	Gen 3	TU863
	Gen 4	TU883
Base Board	Gen 3	TU863
	Gen 4	TBD (+vISI board)


5.3.2.1 Compliance Load Board (CLB)

Figure 115: PCIe Load Board Test Fixture for OCP NIC 3.0 SFF



5.3.2.2 Compliance Baseboard (CBB)

Figure 116: PCIe Base Board Test Fixture for OCP NIC 3.0 SFF



5.3.3 Test Methodology

OCP NIC 3.0 is compliant to the applicable PCIe specifications. The electrical interface may be tested against the PCI Express® Architecture PHY Test Specification Revision 4.0, providing that the appropriate test fixtures from Section 5.3.2 are used.

5.3.3.1 Test Setup

This section is a work-in-progress by the OCP NIC 3.0 SI Subgroup. The following information will be added in a future document release:

- Description of the OCP NIC 3.0 CLB and CBB test figure for use in the PCIe Architecture PHY Test Specifications.
- A user guide is in development through UNH at the time of publication.
- The test procedure is based on the PCIe procedures. The differences between the procedures for PCIe Adapters and OCP NIC 3.0 cards will be documented in the user guide.

6 Thermal and Environmental

6.1 Airflow Direction


The OCP NIC 3.0 card is designed to operate in either of two different airflow directions which are referred to as Hot Aisle and Cold Aisle. In both Hot Aisle and Cold Aisle configurations all airflow is directed over the topside of the card. Component placement must assume that there will be no airflow on the bottom side of the card. The local approach air temperature and velocity to the card is dependent on the capability of the system adopting OCP NIC 3.0 card. These parameters may be impacted by the operational altitude and relative humidity in Hot Aisle or Cold Aisle configurations. Design boundary conditions for Hot Aisle and Cold Aisle cooling are included below in Sections 6.1.1 and 6.1.2 respectively.

The two airflow directions of the Hot and Cold Aisle cases should not result in multiple thermal solutions to separately satisfy the varying thermal boundary conditions. Ideally, any specific OCP NIC 3.0 card design should function in systems with either Hot Aisle or Cold Aisle cooling. Thermal analysis in support of this specification have shown the Hot Aisle configuration to be more challenging than Cold Aisle but card vendors should make that determination for each card that is developed.

6.1.1 Hot Aisle Cooling

The airflow in typical server systems will approach from the card edge or heatsink side of the card. This airflow direction is referred to as Hot Aisle cooling and is illustrated below in Figure 117. The term Hot Aisle refers to the card being located at the rear of the system where the local inlet airflow is preheated by the upstream system components (e.g., HDD, CPU, DIMM, etc.).

Figure 117: Airflow Direction for Hot Aisle Cooling (SFF and LFF)



The boundary conditions for Hot Aisle cooling are shown below in Table 58 and Table 59. The low temperature is listed at 5°C and assumes fresh air could be ducted to the back of the system from the front. More typically the inlet temperature to the OCP NIC 3.0 card will be in the same range as PCIe cards located at the back of the system (55°C local inlet temperature). Depending on the system design, power density, and airflow the inlet temperature to the OCP NIC 3.0 card may be as high as 60°C or 65°C. The airflow velocities listed in Table 59 represent the airflow velocities typical in mainstream

servers. Higher airflow velocities are available within the Hot Aisle cooling tiers listed in Table 64 but card designers must be sure to understand the system level implications of such high card LFM requirements.

Table 58: Hot Aisle Air Temperature Boundary Conditions

	Low	Typical	High	Max
Local Inlet air temperature	5°C (system inlet)	55°C	60°C	65°C


Table 59: Hot Aisle Airflow Boundary Conditions

	Low	Typical	High	Max
Local inlet air velocity	50 LFM	100-200 LFM	300 LFM	System Dependent

6.1.2 Cold Aisle Cooling

When installed in the front of a server the airflow will approach from the I/O connector (e.g., SFP, QSFP or RJ45) side of the card. This airflow direction is referred to as Cold Aisle cooling and is illustrated below in Figure 118. The term Cold Aisle refers to the card being located at the front of the system where the local inlet airflow is assumed to be the same temperature as the system inlet airflow.

Figure 118: Airflow Direction for Cold Aisle Cooling (SFF and LFF)



The boundary conditions for Cold Aisle cooling are shown below in Table 60 and Table 61. The temperature values listed in Table 60 assume the inlet temperature to the OCP NIC 3.0 card to be the same as the system inlet. The low, typical, high, and max temperatures listed align with the ASHRAE A1, A2, A3, and A4 environmental classes. Depending on the system, the supported ASHRAE class may limit the maximum temperature to the OCP 3.0 NIC card. However, for more broad industry support, cards should be designed to the upper end of the ASHRAE classes (e.g., class A4).

Table 60: Cold Aisle Air Temperature Boundary Conditions

	Low	Typical	High	Max
Local Inlet Air Temperature	5°C	25-35°C ASHRAE A1/A2	40°C ASHRAE A3	45°C ASHRAE A4

Table 61: Cold Aisle Airflow Boundary Conditions

	Low	Typical	High	Max
Local Inlet Air Velocity	50 LFM	100 LFM	200 LFM	System Dependent


6.2 Thermal Design Guidelines

The information in this section is intended to serve as a quick reference guide for OCP NIC 3.0 designers early in the design process. The information should be used as a reference for upfront thermal design and feasibility and should not replace detailed card thermal design analysis. The actual cooling capability of the card shall be defined based on the testing with the OCP NIC 3.0 thermal test fixture as defined in Section 6.4.

6.2.1 SFF Card ASIC Cooling – Hot Aisle

The ASIC or controller chip is typically the highest power component on the card. Thus, as OCP NIC 3.0 cards are developed it is important to understand the ASIC cooling capability. Figure 119 below provides an estimate of the maximum ASIC power that can be supported as a function of the local inlet velocity for the SFF card in a hot aisle cooling configuration. Each curve in Figure 119 represents a different local inlet air temperature from 45°C to 65°C.

Figure 119: ASIC Supportable Power for Hot Aisle Cooling – SFF



The curves shown in Figure 119 were obtained using CFD analysis of a reference OCP NIC 3.0 SFF card. The reference card has a 20 mm x 20 mm ASIC with two QSFP connectors. Figure 120 shows a comparison of the 3D CAD and CFD model geometry for the reference OCP NIC 3.0 card. Additional card geometry parameters and boundary conditions used in the reference CFD analysis are summarized in Table 62. The OCP NIC 3.0 simulation was conducted within a virtual version of the test fixture defined in Section 6.4.

Figure 120: OCP NIC 3.0 SFF Reference Design and CFD Geometry




Table 62: Reference OCP NIC 3.0 SFF Card Geometry

OCP NIC 3.0 Form Factor	SFF Card
Heatsink Width	65 mm
Heatsink Length	45 mm
Heatsink Height	9.24 mm
Heatsink Base Thickness	1.5 mm
Fin Count/Thickness	28/0.5 mm
Heatsink Material	Extruded Aluminum
ASIC Width	20
ASIC Length	20
ASIC Height	2.26
ASIC Theta-JC	0.17 C/W
ASIC Theta-JB	10 C/W
OCP PCB In-Plane Conductivity	34 W/mK
OCP PCB Normal Conductivity	0.33 W/mK
ASIC Max T-case	95°C
OCP NIC 3.0 I/O Connectors	Two QSFP @ 3.5 W each


An increase in the supported ASIC power or a decrease in the required airflow velocity may be achieved through heatsink size and material changes. For example, a larger heatsink or a heatsink made out of copper could improve ASIC cooling and effectively shift up the supportable power curves shown in Figure 119.

It is important to point out that the curves shown in Figure 119 represent only the maximum ASIC power that can be supported vs. the supplied inlet velocity. Other heat loads on the card may require airflow velocities above and beyond that required to cool the ASIC. SFP or QSFP optical transceivers located downstream of the ASIC will in many cases pose a greater cooling challenge than the ASIC cooling.

Cooling the optical transceivers becomes even more difficult as the ASIC power is increased due to additional preheating of the air as it moves through the ASIC heatsink. OCP NIC 3.0 designers must consider all heat sources early in the design process to ensure the card thermal solution is sufficient for the feature set. In addition, OCP NIC 3.0 designers must consider all power modes in the design process – including S0 (Main Power Mode) and S5 (Aux Power Mode). For both modes, the card designer must provide the airflow requirements in the OEM FRU record as described in Section 4.10.3.

Card designers must also consider the airflow capability of the server systems that the cards are targeted for use within. Figure 121 below shows the SFF ASIC supportable power curves with an overlay of three server airflow capability ranges. Designers must ensure that their thermal solutions and resulting card airflow requirements fall within the range of supportable system airflow velocity. Cards that are under-designed (e.g., require airflow greater than the system capability) will have thermal issues when deployed into the server system. Card designers are advised to work closely with system vendors to ensure they target the appropriate airflow and temperature boundary conditions.


Figure 121: Server System Airflow Capability – SFF Card Hot Aisle Cooling



The server airflow capability is typically more restrictive when the OCP card is operating in standby mode (S5). The airflow available from the system cooling solution is often more restrictive due to limited power budget available in an S5 state. The graph below emphasizes that though the local OCP NIC 3.0 air temperature is typically lower than in S0 (since preheating from upstream components should be greatly reduced while in standby), the upper bound of local air velocity is a fraction of what can be generated when the system is in main power mode (S0). Card vendors must test for these conditions, making sure that the provided cooling is sufficient for both the ASIC and any installed transceivers (which will still be receiving preheat from the ASIC in S5 mode).

The NIC vendor shall provide the required LFM during S5 state in the FRU EEPROM (see Section 4.10.3). The cold aisle should be tested at 35°C; the hot aisle should be tested at 45°C.

Figure 122: Server System Airflow Capability – SFF Card Hot Aisle Cooling in standby (S5) mode




DRAFT

6.2.2 LFF Card ASIC Cooling – Hot Aisle

Figure 123 below provides an estimate of the maximum ASIC power that can be supported as a function of the local inlet velocity for the LFF card in a hot aisle cooling configuration. Each curve in Figure 123 represents a different local inlet air temperature from 45°C to 65°C.

Figure 123: ASIC Supportable Power for Hot Aisle Cooling – LFF Card



The curves shown in Figure 123 were obtained using CFD analysis of the reference OCP NIC 3.0 LFF card. The reference card has a 45 mm x 45 mm ASIC with two QSFP connectors. Additional card geometry parameters and boundary conditions used in the reference CFD analysis are summarized in Table 63. Figure 124 shows a comparison of the 3D CAD and CFD model geometry for the reference OCP NIC 3.0 card.

Figure 124: OCP NIC 3.0 LFF Reference Design and CFD Geometry





Table 63: Reference OCP NIC 3.0 LFF Card Geometry

OCP NIC 3.0 Form Factor	LFF Card
Heatsink Width	75 mm
Heatsink Length	85 mm
Heatsink Height	9.3 mm
Heatsink Base Thickness	1.5 mm
Fin Count/Thickness	33/0.5 mm
Heatsink Material	Extruded Aluminum
ASIC Width	45
ASIC Length	45
ASIC Height	2.13
ASIC Theta-JC	0.17 C/W
ASIC Theta-JB	10 C/W
OCP PCB In-Plane Conductivity	34 W/mK
OCP PCB Normal Conductivity	0.33 W/mK
ASIC T-case Max	95°C
OCP NIC 3.0 I/O Connectors	Two QSFP @ 3.5 W each

It is important to note that the supportable power for the LFF card is considerably higher than for the SFF card due to the increased size of the ASIC heatsink. In addition, optics module cooling on the LFF card will also be considerably improved due to the arrangement of the optics in parallel to the ASIC heatsink rather than in series. These thermal advantages are key drivers for the LFF card geometry. The OCP NIC 3.0 simulation was conducted within a virtual version of the LFF card test fixture defined in Section 6.4. In addition, OCP NIC 3.0 designers must consider all power modes in the design process – including S0 (Main Power Mode) and S5 (Aux Power Mode). For both modes, the card designer must provide the airflow requirements in the OEM FRU record as described in Section 4.10.3.

Figure 125 below shows the LFF ASIC supportable power curves with an overlay of three server airflow capability ranges. Designers must ensure that their thermal solutions and resulting card airflow requirements fall within the range of supportable system airflow velocity. Cards that are under-designed (e.g., require airflow greater than the system capability) will have thermal issues when deployed into the server system. Card designers are advised to work closely with system vendors to ensure they target the appropriate airflow and temperature boundary conditions.

Figure 125: Server System Airflow Capability – LFF Card Hot Aisle Cooling




6.2.3 SFF Card ASIC Cooling – Cold Aisle

Compared to the Hot Aisle cooling configuration, there are several key differences for Cold Aisle ASIC cooling. With Cold Aisle cooling the airflow is pulled from the I/O connector side of the card. The I/O connectors and faceplate venting may affect the airflow through the ASIC heatsink. The I/O connectors may also preheat the airflow by some amount. In a Cold Aisle cooling configuration, other parallel airflow paths may result in less airflow passing over and through the OCP NIC 3.0 card compared to the Hot Aisle.


The ASIC cooling analysis for the SFF Card in the Cold Aisle configuration was conducted utilizing the same geometry and boundary conditions described in Figure 120 and Table 62 with airflow moving from I/O connector to ASIC (opposite to the Hot Aisle analysis). Figure 126 below shows the results of this analysis for the Cold Aisle cooling configuration. Each curve in Figure 126 represents a different system inlet air temperature from 25°C to 45°C.

Figure 126: ASIC Supportable Power for Cold Aisle Cooling – SFF Card




Similar to Figure 121 for Hot Aisle cooling, Figure 127 below shows the ASIC supportable power curves with an overlay of three Cold Aisle server airflow capability ranges. Designers must ensure that their thermal solutions and resulting card airflow requirements fall within the range of supportable Cold Aisle system airflow velocity. Cards that are under-designed (e.g., require airflow greater than the system capability) will have thermal issues when deployed into the server system. Similar to the Hot Aisle cases, cooling of the optical transceivers need to be taken into consideration even though they are not preheated by the ASIC in the Cold Aisle case. OCP NIC 3.0 designers must consider all power modes in the design process – including S0 (Main Power Mode) and S5 (Aux Power Mode). For both modes, the card designer must provide the airflow requirements in the OEM FRU record as described in Section 4.10.3. Card designers are advised to work closely with system vendors to ensure they target the appropriate airflow and temperature boundary conditions for both Hot and Cold Aisle cooling.

Figure 127: Server System Airflow Capability – SFF Cold Aisle Cooling



A comparison of Hot Aisle (55°C) and Cold Aisle (35°C) SFF ASIC cooling capability curves is shown below in Figure 128. The comparison shows the Hot Aisle ASIC cooling capability at 12 W at 150 LFM while the cold Aisle cooling capability shows support for 19 W at 150 LFM. In general, based on the reference geometry, the Cold Aisle cooling configuration allows for higher supported ASIC power at lower velocities due primarily to the lower inlet temperatures local to the OCP NIC 3.0 card when in the Cold Aisle cooling configuration.


Figure 128: ASIC Supportable Power Comparison – SFF Card



6.2.4 LFF Card ASIC Cooling – Cold Aisle


The ASIC cooling analysis for the LFF card in Cold Aisle configuration was conducted utilizing the same geometry and boundary conditions described in Figure 124 and Table 63 with airflow moving from I/O connector to ASIC (opposite to the Hot Aisle analysis). Figure 129 below shows the results of this analysis for the Cold Aisle cooling configuration. Each curve in Figure 129 represents a different system inlet air temperature from 25°C to 45°C.

Figure 129: ASIC Supportable Power for Cold Aisle Cooling – LFF Card




Similar to Figure 127 for LFF Hot Aisle cooling, Figure 130 below shows the LFF ASIC supportable power curves with an overlay of three Cold Aisle server airflow capability ranges. Designers must ensure that their thermal solutions and resulting card airflow requirements fall within the range of supportable Cold Aisle system airflow velocity. Cards that are under-designed (e.g., require airflow greater than the system capability) will have thermal issues when deployed into the server system. Similar to the Hot Aisle cases, cooling of the optical transceivers need to be taken into consideration even though they are not preheated by the ASIC in the Cold Aisle case. OCP NIC 3.0 designers must consider all power modes in the design process – including S0 (Main Power Mode) and S5 (Aux Power Mode). For both modes, the card designer must provide the airflow requirements in the OEM FRU record as described in Section 4.10.3. Card designers are advised to work closely with system vendors to ensure they target the appropriate airflow and temperature boundary conditions for both Hot and Cold Aisle cooling.

Figure 130: Server System Airflow Capability – LFF Cold Aisle Cooling



A comparison of Hot Aisle (55°C) and Cold Aisle (35°C) LFF ASIC cooling capability curves is shown below in Figure 131. The comparison shows the Hot Aisle ASIC cooling capability at 19 W at 150 LFM while the cold Aisle cooling capability shows support for 42 W at 150 LFM. In general, based on the reference geometry, the Cold Aisle cooling configuration allows for higher supported ASIC power at lower velocities due primarily to the lower inlet temperatures local to the OCP NIC 3.0 card when in the Cold Aisle cooling configuration.

Figure 131: ASIC Supportable Power Comparison – LFF Card



6.3 Thermal Simulation (CFD) Modeling

CFD models of the SFF and LFF cards developed for the analysis detailed in Section 6.2 are available for download on the OCP NIC 3.0 Wiki: <http://www.opencompute.org/wiki/Server/Mezz>

The thermal models available on the wiki site are in Icepak format. CAD step file exports from those models are also available to aid in re-creation of the models in other CFD software tools. Note that the geometry utilized in the CFD models is based on the OCP NIC 3.0 thermal test fixture detailed in Section 6.4.

Thermal simulation of OCP NIC 3.0 cards using the provided CFD models is recommended. Ideally, vendors developing OCP NIC 3.0 cards would perform CFD analysis to validate card thermal solutions using the provided CFD models prior to building card prototypes. Once prototypes are available, vendors would then perform thermal testing on the functional cards using the thermal test fixtures detailed in Section 6.4.

6.4 Thermal Test Fixture

Thermal test fixtures have been developed for SFF and LFF OCP NIC 3.0 cards. The test fixtures are intended to provide a common thermal test platform for card vendors, server vendors, and other industry groups planning to develop or utilize the OCP NIC 3.0 card form factors. Details of the thermal test fixtures are as follows:

- Sheet metal side walls, base, faceplate, and top cover
- Thumbscrew top cover access
- PCB sandwiched between base and side walls
- Intended for attachment to wind tunnel or flow bench such as those available at: <http://www.fantester.com/>
- Allows for thermal testing of functional OCP NIC 3.0 cards in a metered airflow environment
- Input power from external power supplies allows for OCP NIC 3.0 card power measurement
- SFF fixture power connections for 3.3 V, GND, GND, 12 V
- LFF fixture power connections for 3.3 V, GND, GND, GND, 12 V, 12 V
- RJ45 connector for NC-SI pass-through
- USB Type-C connector for microprocessor connectivity and Micro-B to the USB pins on the primary connector
- Functions as a remote PCIe extension with intent to position host server under the fixture for connection to system PCIe slot
- Single x16 connection to server host on bottom side of the fixture PCB (SFF)
- Dual x16 connection to server host on bottom side of the fixture PCB (LFF)
- Predefined locations for fixture airflow/temperature sensors on fixture PCB silkscreen. Quantity 3x per SFF board and quantity 4x for LFF – see Figure 137

- Candlestick style sensors are available at:
<https://www.qats.com/Products/Instruments/Temperature-and-Velocity-Measurement/Sensors/Candlestick-Sensor>
- Candlestick sensors must be procured separately, not integrated with fixture PCB
- Low profile PCIe card serves as blockage above OCP NIC 3.0 card to mimic system geometry and prevent airflow bypass
- LFF fixture uses a sheet metal obstruction built into the top cover

CAD Files for the current revision of the test fixture are available for download on the OCP NIC 3.0 Wiki:
<http://www.opencompute.org/wiki/Server/Mezz.>

6.4.1 Test Fixture for SFF Card

Images of the SFF thermal test fixture are shown in Figure 132 and Figure 133. The SFF fixture PCB is shown in Figure 134. Note the three candlestick sensor locations directly next to the OCP NIC 3.0 connectors.

Figure 132: SFF Thermal Test Fixture Preliminary Design




Figure 133: SFF Thermal Test Fixture Preliminary Design – Cover Removed





Figure 134: SFF Card Thermal Test Fixture PCB



6.4.2 Test Fixture for LFF Card

Images of the LFF thermal test fixture are shown in Figure 135 and Figure 136. The LFF fixture PCB is shown in Figure 137. Note the three candlestick sensor locations directly next to the OCP NIC 3.0 connectors.

Figure 135: LFF Card Thermal Test Fixture Design




Figure 136: LFF Card Thermal Test Fixture Design – Cover Removed






Figure 137: LFF Card Thermal Test Fixture PCB



6.4.3 Test Fixture Airflow Direction

When utilizing the OCP NIC 3.0 thermal test fixture, the wind tunnel or flow bench must be configured to push airflow for Hot Aisle cooling or to pull airflow for Cold Aisle cooling as shown in Figure 138.

Figure 138: Thermal Test Fixture Airflow Direction



6.4.4 Thermal Test Fixture Candlestick Sensors

As noted previously, candlestick sensor locations are included on the fixture PCB silkscreen. These candlestick sensors provide point measurements for both airflow velocity (LFM) and air temperature. The airflow at the inlet to the OCP NIC 3.0 will differ from the fixture mean velocity due to the obstructions above the OCP NIC 3.0 cards within the fixture. Thus, the fixture flow rate and cross-sectional area should not be used to determine the local velocity at the OCP NIC 3.0 card. Instead, the candlestick velocity/temperature sensors should be utilized to directly measure the local inlet velocity to the cards for hot aisle cooling.

Figure 139 and Figure 140 below show the air velocity at each sensor location vs. the total fixture flow rate in CFM. The curves shown in these figures are based on the data collected from the CFD models discussed in Section 6.3. Note the error between the velocities obtained from the sensor locations vs. the velocity based on the duct cross-sectional area.

Figure 139: SFF Fixture, Hot Aisle Flow – Candlestick Air Velocity vs. Volume Flow





Figure 140: LFF Fixture, Hot Aisle Flow – Candlestick Air Velocity vs. Volume Flow



6.5 Card Sensor Requirements

See Sections 4.4 to 4.6 for information relating to temperature sensor and reporting requirements.

6.6 Card Cooling Tiers

Section 4.10.3 defines a number of FRU fields that may be read by the baseboard management controller (BMC). Two of these fields provide the Hot Aisle and Cold Aisle Card Cooling Tiers that may be used for open loop fan speed control. The Card Cooling Tiers relate the card local inlet temperature to the required local inlet velocity in linear feet per minute (LFM) which allows the system to set fan speeds according to the cooling requirements of the card.

The Card Cooling Tier FRU fields are particularly useful for systems that do not implement temperature sensor monitoring (open loop fan control). The FRU fields may also be used as a backup for systems that implement fan control based on temperature sensor monitoring (closed loop control).

Card Cooling Tiers for Hot and Cold Aisle Cooling configurations are defined in Table 64. The values in the table are listed with units shown in LFM. Future releases of this specification will provide more detail to the Card Cooling Tier curve definition.

Table 64: Card Cooling Tier Definitions (LFM)

OCP NIC 3.0 Local Inlet Temperature [°C]	Target Operating Region				Server Airflow High Fan Speed		Non-Typical Server Airflow - Subject to System Capability					
	Tier 1	Tier 2	Tier 3	Tier 4	Tier 5	Tier 6	Tier 7	Tier 8	Tier 9	Tier 10	Tier 11	Tier 12
5	0	100	100	100	100	100	110	150	200	255	320	395
10	0	100	100	100	100	100	118	159	211	271	341	426
15	0	100	100	100	100	105	128	170	226	291	366	461
20	0	100	100	100	100	109	137	181	241	311	391	496
25	0	100	100	100	100	114	147	198	260	337	420	532
30	0	100	100	100	110	130	167	215	283	363	453	578
35	0	100	100	100	120	146	186	236	308	393	488	623
40	0	100	100	111	139	172	212	260	335	430	525	673
45	0	100	100	125	161	200	242	290	370	472	575	736
50	0	100	110	155	198	243	288	335	425	530	645	838
55	0	100	150	200	250	300	350	400	500	625	750	1000
60	0	150	200	255	315	380	450	525	643	796	936	1241
65	0	200	260	330	410	500	600	710	850	1025	1200	1550

A graphical view of the Card Cooling Tiers is shown in Figure 141. The Tiers range from 0 LFM to as high as 1000 LFM at 55°C local inlet temperature. It is important to understand that the cooling tiers extend well beyond the airflow range of most server systems. As noted in Section 6.2, card designers must consider the airflow capability of the systems that the cards are to be used in when designing the card thermal solution and component placement. Figure 142 and Figure 143 below show the range of typical system capability for Hot Aisle and Cold Aisle configurations. Cards designed to these typical ranges (Tiers 1-6) will be low risk to support in most if not all server systems. Alternatively, cards that require Tier 7 or greater may still work in many server systems but may require extra validation testing, specific system slot location requirements, and potentially ambient temperature and hardware restrictions.

Figure 141: Graphical View of Card Cooling Tiers




Figure 142: Typical Operating Range for Hot Aisle Configurations





Figure 143: Typical Operating Range for Cold Aisle Configurations



6.7 Non-Operational Shock & Vibration Testing

OCP NIC 3.0 components are deployed in various environments. As such, all OCP NIC 3.0 cards shall be subjected to shock and vibration testing to ensure products do not sustain damage during normal operational or transportation conditions. While end customer deployments may require an additional final system level test, this section sets the minimum shock and vibration requirements for an OCP NIC 3.0 card that must also be considered.

Shock and vibration testing shall be done in accordance with the procedures listed below. The tests shall be conducted using a vertical shock table. The OCP NIC 3.0 card shall be secured in the standard test fixture as described in Section 6.7.1.

6.7.1 Shock & Vibe Test Fixture

The OCP NIC 3.0 shock and vibe fixture supports simultaneous testing of up to four SFF, or four LFF cards. The SFF fixture accepts card configurations that utilize the pull tab, ejector lever and internal lock mechanisms. The LFF fixture only accepts the single latch faceplate.

The fixture is comprised of a universal baseplate that allows for attaching SFF or LFF rail guides and simulated chassis faceplates. The baseplate includes an industry standard vibration table hole pattern for securing the UUT for test. Figure 144 and Figure 145 show the SFF and LFF fixtures, respectively.

CAD files for the current revision of the test fixture are available for download from the OCP NIC 3.0 wiki: <http://www.opencompute.org/wiki/Server/Mezz>.

Figure 144: SFF Shock and Vibe Fixture





Figure 145: LFF Shock and Vibe Fixture



6.7.2 Test Procedure

The following procedures shall be followed for the shock and vibration testing:

- A minimum sample size of three OCP NIC 3.0 cards shall be subjected to shock and vibration. Each sample shall be subjected to testing in the order listed below.
- All samples shall be verified for functionality prior to test.

- The OCP NIC 3.0 card shall be fixtured to simulate how the card will be mounted within a system. For example, the OCP NIC 3.0 card may be fixtured in the horizontal plane with the primary component side facing up for certain chassis configurations.
- The fixture shall be tested on all 6 sides. Each side shall be clearly labeled as 1-6 for test identification purposes. Testing shall be performed in the vertical axis only. The fixture shall be rotated until all six sides have been tested as the product may be dropped from any orientation during handling. Testing shall not be conducted on a three axis slip table.
- Non-operational vibration testing is performed at 1.88 G_{RMS} for a duration of 15 minutes per side per Table 65.

Table 65: Random Vibration Testing 1.88 G_{RMS} Profile


Frequency (Hz)	G ² /Hz
10	0.13
20	0.13
70	0.004
130	0.004
165	0.0018
500	0.0018

- Non-operational half-sine shock test at 71 G ±10% with a 2 ms duration per surface.
- Non-operational trapezoidal wave shock test at 50 G ±10% at a rate of 170 inches/sec per surface.
- All cards shall be checked for proper operation after the shock and vibration tests have been conducted. All three samples must be in full operating order to consider the product as a pass.

6.8 Dye and Pull Test Method


All Dye and Pull test methods shall be implemented per the IPC-TM-650 method 2.4.53 (Dye and Pull Test Method – formerly known as Dye and Pry). The Dye and Pull test uses a colored dye penetrant to visually indicate cracked solder joints on BGA devices. The test shall only be conducted after the Shock and Vibration testing has been conducted on the test samples. The Dye and Pull Test Method is a destructive test.

- A minimum sample size of three OCP NIC 3.0 cards shall be subjected to the Dye and Pull Test Method.
- All samples shall be first subjected to the Shock and Vibration testing outlined in Section 6.7.
- All samples shall be subjected to the preparation and test procedures of IPC-TM-650 method 2.4.53.
- Following the pull-test operation, the board sample shall be examined for dye indication at the target BGA area. Separation locations are categorized in to the following five areas:
 - Type 1 – Separation between the BGA copper pad and the BGA substrate.
 - Type 2 – Separation between the BGA copper pad and the BGA solder sphere.
 - Type 3 – Separation between the BGA solder sphere and the copper pad on the PCB.
 - Type 4 – Separation between the copper pad on the PCB and the PCB laminate.
 - Type 5 – Separation of the BGA solder sphere.



- Samples shall be subjected to the following failure criteria:
 - Dye coverage of >50% (“D” and “E” in Figure 147) of any Type 2 or Type 3 BGA cracks are present in the test sample.
 - One or more Type 1 or Type 4 BGA cracks are present in the test sample.

Figure 147: Dye Coverage Percentage



The following exceptions are allowed:

- For “via-in-pad” designs, dye is allowed on the laminate surface (under the pad), as long as the dye has not entered the inner-via laminate area, or is found on the separated via-barrel wall.
- Allowances for dye indications on mechanical (non-electrical) BGA corner locations or multiple use locations (grounds, powers) may be determined by the appropriate Engineering Team.

6.9 Gold Finger Plating Requirements

This section defines the minimum plating/quality requirements for the OCP NIC 3.0 gold fingers.

6.9.1 Host Side Gold Finger Plating Requirements

Per Section 6.4 (Environmental Requirements) of the PCIe CEM specification, the minimum host side gold finger plating is 30 microinches of gold over 50 microinches of nickel. OCP NIC 3.0 card vendors shall individually evaluate the minimum plating required.

The recommendation for OCP NIC 3.0 is to 30 microinches of gold over 150 microinches of nickel.

6.9.2 Line Side Gold Finger Durability Requirements

The line side connectors must be designed to support a minimum of 250 error free insertion cycles. In order to accomplish this, it is required that the minimum contact plating be as follows:

- SFP and QSFP connectors: 30 microinches of gold over 50 microinches of nickel
- RJ45 connectors have a minimum of 50 microinches of gold over 50 microinches of nickel

7 Regulatory

7.1 Required Compliance

An OCP NIC 3.0 card shall meet the following Environmental, EMC and safety requirements.

Note: Emissions and immunity tests in Section 7.1.4 are to be completed at the system level. The OCP NIC 3.0 vendors should work with the system vendors to achieve the applicable requirements listed in this section.

7.1.1 Required Environmental Compliance

- China RoHS Directive
- **EU RoHS 2 Directive (2011/65/EU)** aims to reduce the environmental impact of electronic and electrical equipment (EEE) by restricting the use of certain hazardous materials. The substances banned under RoHS are lead, mercury, cadmium, hexavalent chromium, polybrominated biphenyls, polybrominated diphenyl ether, and four phthalates.
- **EU REACH Regulation (EC) No 1907/2006** addresses the production and use of chemical substances and their potential impact on human health and the environment.
- **EU Waste Electrical and Electronic Equipment (“WEEE”) Directive (2012/19/EU)** – mandates the treatment, recovery and recycling of EEE.
- **The Persistent Organic Pollutants Regulation (EC) No. 850/2004** bans production, placing on the market and use of certain persistent organic pollutants.
- **The California Safe Drinking Water and Toxic Enforcement Act of 1986 (“Proposition 65”)** sets forth a list of regulated chemicals that require warnings in the State of California.
- **The Packaging and Packaging Waste Directive 94/62/EC** limits certain hazardous substances in the packaging materials
- **Batteries Directive 2006/66/EC** regulates the manufacture and disposal of all batteries and accumulators, including those included in appliances.

7.1.2 Required EMC Compliance

- Radiated and Conducted Emissions requirements are based on deployed geographical locations. Refer to Table 66 for details.

Table 66: FCC Class A Radiated and Conducted Emissions Requirements Based on Geographical Location

Targeted Geography	Applicable Specifications
USA	FCC, 47 CFR Part 15, Class A digital device (USA)
Canada	ICES-003, class A (CAN)
EU	EN 55032: 2015+AC:2016 Class A Radiated and Conducted Emissions requirements for European Union EN 55035: 2017 Immunity requirements for the European Union (EU) EN 55024: 2010+A1:2015 Immunity requirements for European Union (EU) may alternatively be reported.

	Note: EN55024 is scheduled to be superseded by EN55035. OCP NIC 3.0 implementors are encouraged to test to EN55035 to avoid recertifying their product when EN55024 is withdrawn.
Australia/New Zealand	AS/NZS CISPR 32:2015 Class A CISPR 32:2015 for Radiated and Conducted Emissions requirements
Japan	VCCI 32-1 Class A Radiated and Conducted Emissions requirements
Korea	KN32 – Radiated and Conducted Emissions KN35- Immunity
Taiwan	BSMI CNS13438: 2006 (complete) Class A Radiated and Conducted Emissions requirements

- **CE** – Equipment must pass the CE specification
- All technical requirements covered under **EMC Directive (2014/30/EU)**

7.1.3 Required Product Safety Compliance

- Safety – requirements are listed in Table 67.

Table 67: Safety Requirements

Targeted Category	Applicable Specifications
Safety	UL 60950-1/CSA C22.2 No. 60950-1-07, 2 nd Edition + Amendment 1 + Amendment 2, dated 2011/12/19. The Bi-National Standard for Safety of Information Technology Equipment, EN60950-1: 2006+A11:2009+A1:2010+A12:2010+A2:2013 IEC 60950-1 (Ed 2) + A1 + A2. IEC 62368-1 may also be co-reported depending on region

7.1.4 Required Immunity (ESD) Compliance

The OCP NIC 3.0 card shall meet or exceed the following ESD immunity requirements listed in Table 68.

Table 68: Immunity (ESD) Requirements

Targeted Category	Applicable Specifications
Immunity (ESD)	EN 55035 2017, and IEC 61000-4-2 2008 for ESD. EN 55024 may alternatively be reported. Required ± 4 kV contact charge and ± 8 kV air discharge. Note: EN55024 is scheduled to be superseded by EN55035. OCP NIC 3.0 implementors are encouraged to test to EN55035 to avoid recertifying their product when EN55024 is withdrawn.
NEBS Level 3 (optional)	Optionally test devices to NEBS Level 3 – Required ± 8 kV contact charge and ± 15 kV air discharge with interruptions not greater than 2 seconds. The device shall self-recover without operator intervention. Note: NEBS compliance is part of the system level testing. The OCP NIC 3.0 specification is providing a baseline minimum recommendation for ESD immunity.

7.2 Recommended Compliance

All OCP NIC 3.0 cards are required to meet the requirements specified in Section 7.1. Card vendors should also consider meeting the requirements below.

7.2.1 Recommended Environmental Compliance

- **Halogen Free:** IEC 61249-2-21 Definition of halogen free: 900ppm for Bromine or Chlorine, or 1500ppm combined total halogens.
- **Arsenic:** 1000 ppm (or 0.1% by weight)
- **Emerging: US Conflict Minerals law: section 1502 of the Dodd-Frank Act** requires companies using tin, tantalum, tungsten, and gold (“3TG”) in their products to verify and disclose the mineral source. While this does not apply to products that are used to provide services, such as Infrastructure hardware products, the OCP NIC Subgroup is considering voluntarily reporting of this information.

7.2.2 Recommended EMC Compliance

- FCC, 47 CFR Part 15, Subpart B Class A digital device (USA) with 10dB margin. Refer to the baseline requirements shown in Section 7.1.2 for details.

8 Revision History

8.1 Document Revision History

Author	Description	Revision	Date
OCP NIC 3.0 Subgroup	Initial public review.	0.70	01/25/2018
OCP NIC 3.0 Subgroup	<ul style="list-style-type: none"> - Implemented comments from 0.70 review. - LED implementation updated. - Gold finger lengths updated. All pins are full length except for PCIe TX/RX, REFCLKS and PRSNT pins. 	0.71	02/06/2018
OCP NIC 3.0 Subgroup	- Updates to Section 4.x per the working group session.	0.72	02/21/2018
OCP NIC 3.0 Subgroup	<ul style="list-style-type: none"> - Change NC-SI over RBT RXD/TXD pins to a pull-up instead of a pull down. - Update power sequencing diagram. REFCLK is disabled before silicon transitions to AUX Power Mode. - Merge pinout sections 3.4 and 3.5 together for structural clarity. - Add text to gate WAKE# signal on AUX_PWR_GOOD (internal) assertion; updated diagrams with WAKE# signals to reflect implementation. - Add initial signal integrity outline to document (WIP) - Add Initial draft of the Shock and Vibration, and Dye and Pull test requirements. - Rearrange Section 2 for structure; changed section name to Mechanical Card Form Factor - Move non-NIC use cases to Section 1.5. - Moved Port numbering and LED definitions to Section 3.8. - Add secondary side LED placement for 4x SFP and 2x QSFP implementations in Section 3.8. - Revised labeling section (Section 2.9). - Optimize the scan chain LED bit stream for dual port applications. - Add SLOT_ID[1]. Updated text and diagrams for mapping SLOT_ID[1:0] to Package ID[2:0] and FRU EEPROM A[2:0] fields. - Reduce ID Mode power consumption on +12V_EDGE 	0.73	05/01/2018
OCP NIC 3.0 Subgroup	<ul style="list-style-type: none"> - Text clean up. All minor / generally agreed upon items within the OCP NIC 3.0 Workgroup have been accepted. - Clarify PCIe bifurcation is on a per-slot basis. Add 1x32 and 2x16 implementation examples for a Large Form Factor card. - Removed reference to a x24 PCIe width LFF card from Table 5 – OCP NIC 3.0 Card Definitions. - Move SLOT_ID[1] to OCP_A6 for immediate power on indication of the card physical location for RBT and FRU EEPROM addressing. Updated RBT addressing and Scan Chain definition to match. - Updated diagrams and text in Section 6.x based on feedback from the OCP NIC 3.0 Thermal Workgroup. - Updated diagrams and text in Section 2.0 based on feedback received from the OCP NIC 3.0 Mechanical Workgroup. 	0.74	06/04/2018
OCP NIC 3.0 Subgroup	Ov80 public release	0.80	06/04/2018
OCP NIC 3.0 Subgroup	<p>Ov81 public release. Changes are as follows:</p> <ul style="list-style-type: none"> - Section 1.3 – Update Figure 1 with latest thumbscrew design. - Section 2.4.2 – Mechanical corrections to BOM items 5, 6A/B, 8 & 11. - Section 3.4.3 – Add statement to isolate SMRST# if target device voltage is not powered from +3.3V_EDGE. 	0.81	07/06/2018

	<ul style="list-style-type: none"> - Section 3.4.4 – Clarified the RBT_ARB_IN and RBT_ARB_OUT pin descriptions. - Section 3.4.4 – Clarified SLOT_ID[1:0] description and example diagrams; move SLOT_ID[1:0] isolation to NIC and use direct connection to FRU EEPROM. - Section 3.4.5 – DATA_IN bit PRSNTB[3:0] definition to optionally use pull up/down to match PRSNTB[3:0]# card edge connections. - Section 3.4.7 – Add USB 2.0 definition to the Primary Connector. - Section 3.4.8 – Add UART definition to the Secondary Connector. - Section 3.4.9 – Changed Miscellaneous pins to RFU[1:2] pins. - Section 3.8 – Clarified LED placement. - Section 3.9.x – Clarified ID-Aux and Aux-Main Power Mode transition requirements to prevent sampling health status pins until cards have fully entered into Aux and Main modes to prevent false indication. - Section 3.11 – Updated hot swap consideration text to highlight available hot swap mechanisms. Actual hot swap design is outside the scope of this specification. - Section 4 – Update MCTP Type management description. - Section 4.9 – Clarified the FRU EEPROM is directly connected to the card edge. No isolation is used for the FRU EEPROM. 		
OCP NIC 3.0 Subgroup	<ul style="list-style-type: none"> - Minor editorial changes. - Changed names to “SFF” and “LFF” when referencing the two board form factors for uniformity. - Section 3.4.1 – Changed PERST[3:0]# to be asserted low until the platform is ready to bring cards out of reset. - Section 3.5.3 – Corrected typos in the PCIe Bifurcation Decoder (Table 31) for hosts that implement 4 x2 links on the first 8 lanes when using a 4 x4 OCP NIC 3.0 card. - Section 3.5.5.3 & 3.5.5.4 – Corrected the BIF[2:0] values in the diagrams. - Section 3.7 – Corrected typos in the PCIe Bifurcation result and REFCLK mapping (Table 38 and Table 41) for single host/quad host cases with PCIe on the first 8 lanes. This change was due to propagating corrections from Table 31 from Section 3.5.3. - Section 3.8.3 – Changed faceplate LED placement for 2xQSFP to primary side. - Section 4.10.2 – Added FRU field to identify the card manageability type. 	0.82	08/03/2018
OCP NIC 3.0 Subgroup	<ul style="list-style-type: none"> - Section 5.3.4 – Removed subheadings for the PCIe test methodology. Replaced with reference text to the PCIe test specifications. 	0.83	08/29/2018
OCP NIC 3.0 Subgroup	<ul style="list-style-type: none"> - Editorial changes. - Add appropriate trademarks per entity usage guidelines - Section 1.x – Reorganized section, added list of acronyms. - Section 1.5, 2.x – Mechanical updates from ME team. Updates to Vendor PN, pull tab/thumbscrew color. - Section 3.1.x – Add card edge gold finger Detail D and profile dimensions. - Section 3.1.1 – Change gold finger lengths to align with SFF-TA-1002 and SFF-TA-1009. All ground pins mate first; power, single ended and differential pairs mate second. - Section 3.4.1, Section 3.5.x – Updated for LFF and applications with 32 lanes of PCIe. Included TX/RX lane indices for lanes[16:31], REFCLK[4:5], PERST[4:5]# and PWRBRK1 on the Secondary Connector. - Section 3.4.2 – Clarified BIF[0:2]# assertion timing. - Section 3.4.4 – Clarified RBT isolation state with respect to AUX_PWR_EN and NIC_PWR_GOOD assertion. 	0.84	11/13/2018

- Section 3.4.5 – Clarified that the scan chain shift registers may also be implemented with a CPLD as long as the logic is equivalent. Clarified the FAN_ON_AUX bit.
- Section 3.4.9 – Changed RFU[1:2] to RFU[3:4] for the Secondary Connector.
- Section 3.7 – Removed bifurcation expansion tables. Redirected readers to the pinout/bifurcation spreadsheet on the OCP Mezz Wiki site instead.
- Section 3.6 and Section 3.7 – Merged together as they both discuss PCIE REFCLK mapping. Add new diagrams depicting single, dual and quad host implementations with 1, 2, and 4 links. Diagrams show association of the link to the REFCLK and PERST signals. Update bifurcation table color codes to differentiate SFF (2C+ / 4C+) implementations from LFF implementations.
- Section 3.11 – Add RBT_ISOLATE# to the power up/down sequencing diagrams. Changed BIF[2:0]# pins to 'low' state when AUX_PWR_EN is deasserted. Add timing value "T4" from DSP0222 to power up diagram and sequencing parameters table. Add $T_{CYCLE_SFF} / T_{CYCLE_LFF}$ parameters to power down sequencing to prevent a pre-biased output condition when power cycling cards.
- Section 4.4 – Clarified requirements for self-shutdown if the optional feature is implemented on the card.
- Section 4.9 – Cleaned up SMBus 2.0 Address Map text. Add text to discourage the use of unsolicited SMBus messages including the optional 'Notify ARP Master' command.
- Section 4.10.2 – Clarified FRU data format requirements per the IPMI Platform Management FRU Information Storage Definition. List minimum FRU content requirements.
- Section 5.1.x – Clarified NC-SI over RBT physical routing and length matching requirements.
- Section 7.1.4 – Corrected typo on NEBS air discharge value for ESD testing. Changed from 16 kV to 15 kV.

OCP NIC 3.0 Subgroup	November release with all 0.84 subgroup changes	0.85	11/20/2018
OCP NIC 3.0 Subgroup	<ul style="list-style-type: none"> - December Hot fix release. Minor text corrections and clarifications. - Section 2.4.2 – Update faceplate and ejector handle drawing filenames - Sections 2.4.3, 2.4.4, and 2.7.x – Mechanical drawing updates. Fixed typos and updated critical dimension text. - Section 2.8.x – Mechanical drawing updates. Changed bottom of PCB datum from H to J. - Section 3.4.5 – Clarified the default scan chain pin states for TEMP_WARN_N, TEMP_CRIT_N and FAN_ON_AUX when temperature reporting is not required. - Section 3.4.7 – Clarified USB text on differential signaling voltage, termination voltage and V_{BUS} detection indication. Updated USB figures to include V_{BUS} and V_{BUS} detection inputs. - Section 3.8 – Updated power state machine. +12V_EDGE clarified as being "on, but limited up to the ID Mode budget." This clarification aligns with the existing text. Added clarification notes to Table 37 for permissible +12V_EDGE current draw in the ID Mode and Aux Power Mode states. 	0.85b	12/14/2018
OCP NIC 3.0 Subgroup	<ul style="list-style-type: none"> - Section 1.2 – Change Cavium entity name to Marvell Semiconductor, Inc. - Section 1.3 – Added DSP0267 reference - Sections 1.5, 2.1.x – All 3D renderings have been updated - Section 2.4.2 – Mechanical BOM updated. Added notes for Phillips + Torx screw heads. 	0.86	04/02/2019

- Sections 2.4.3, 2.4.4, 2.5.x, 2.7.x, 2.8.2 – Mechanical 2D drawings updated.
- Section 2.7 – Added insulator note regarding the use of 0.127 mm material for constrained regions; 4x 4mm diameter holes permitted for non-metallic mechanical retention pin cutouts
- Section 3.1 – Gold finger figure updated to clarify dimension from center of pad.
- Section 3.4.4 – Add note that RBT clock buffers are permitted so long as the card timing budget is not violated. Clarify permitted baseboard and OCP NIC 3.0 side pull up/pull down implementations on SLOT_ID[1:0].
- Sections 3.4.5, 3.7.2 – Scan Chain LED definition updated. Activity LED is OFF when “idle”
- Section 3.4.7 – USB figures now shown as a BMC or Platform I/O hub connectivity instead of just the BMC.
- Section 3.4.8 – UART figures now shown as a BMC or Platform I/O hub connectivity instead of just the BMC.
- Section 3.6 – Clarify REFCLK/PERST mapping for LFF (Table 35).
- Section 4.4 – Temperature Reporting clarified as ASIC die temp reporting. Also clarified die temp reporting is independent of the transceiver temp reporting.
- Section 4.6 – Transceiver module temperature reporting is always required and is independent of the network ASIC TDP 5W threshold.
- Section 4.10.1 – Add double byte FRU EEPROM access diagrams for clarity. FRU EEPROM WP mechanism may optionally be over written in the field to allow for field updates.
- Section 5.1 – Re-wrote NC-SI over RBT SI recommendations. Broke up into three sections: common, baseboard, and OCP NIC 3.0 requirements.
- Section 5.2 – Reference SMBus and PCIe CEM specifications for SI, routing and signal requirements.
- Sections 6.2.x – Add notes to consider airflow requirements in Aux and Main power modes for SFF & LFF in Hot & Cold Aisle implementations.
- Section 6.7 – Non-operational shock fixture – CAD files on wiki. Updated test procedure requirements.

OCP NIC 3.0 Subgroup	<ul style="list-style-type: none"> - Section 1.1 – Update OCPHL-P license link due to opencompute.org site restructuring. - Section 1.3 – Update NIST Special Publication 800-193 to May 2018 release - Section 1.3.1 – Add I²C as a trademark of NXP Semiconductor. - Section 2.1 – Add safety requirement text and cross reference Section 7.1.3. - Section 2.1.1 – Add requirement regarding maximum extraction force on heatsink for internal lock implementations. - Section 2.x – Update Mechanical 2D drawings from WG notes - Section 2.7 – Clarify insulator access hole requirements. - Section 3.4.1 – Clarify PWRBRK[0:1]# card pull up resistor value as 95kOhm or larger and must meet the T_{PWRBRK} timing parameter per PCIe CEM. Provided recommended Baseboard value between 4.7 kOhm and 10 kOhm. - Section 3.4.4 – Updated NC-SI over RBT signal list pull down requirements. TX_EN and TXD[1:0] pull downs moved between RBT isolation circuit and OCP NIC 3.0 connector to prevent NIC side signals from floating when RBT_ISOLATE# is asserted. Updated NC-SI diagrams to depict this change. Clarify HW ARB_IN / ARB_OUT ring integrity requirements for baseboards with multiple OCP NIC 3.0 slots. 	0.90	06/20/2019
----------------------	---	------	------------

- Section 3.4.6 – Clarify +12V_EDGE and +3.3V_EDGE requirements when AUX_PWR_EN and MAIN_PWR_EN are asserted.
- Section 3.4.6, 3.8.x, 3.9, 4.10.1 – Add optional Programming Mode power state. Add state of NIC_PWR_GOOD when in Programming Mode. Update Power state diagram with Programming Mode.
- Section 3.5.3 – Bifurcation decoder table updated to align with bifurcation spreadsheet.
- Section 3.8 – Rename section from “Power Capacity and Power Delivery” to “Power State Machine.” Add text about Programming Power State. Clarify normal power state flow.
- Section 3.11 – Add Programming Mode sequencing diagram detailing transition from ID Mode to Programming Mode and back to ID Mode.
- Section 4.10.2 – Move FRU write protection mechanism to a dedicated section heading.
- Section 4.10.3 – Modify OEM Record Offsets for supporting FRU write protection, Programming Mode power state, and thermal requirements.
- Offset 03 – Values now advertise compliance to specification 0v90.
- Offset 06/07 – Changed Hot/Cold Aisle tier for use with passive cables.
- Offset 17 – advertise FRU write protection mechanism.
- Offset 18 – advertise if the Programming Mode Power State is supported.
- Offset 19/20 – Add Hot/Cold Aisle tier for use with active cables.
- Offset 21/22 – Add Reference module power and temperature level.
- Offset 23 – Add reference active cooling fan fail tier requirement.
- Offset 31 – Defined "controller" as a SMBus connected device in the context of the FRU UDID.
- Section 5.1, 5.1.2, 5.1.4 – Add NC-SI over RBT requirements for LFF
- Section 6.4 – Update thermal test fixture feature list.
- Section 6.6 – Update card cooling tiers. Combined Hot Aisle and Cold Aisle airflow definition.
- Section 6.7.2 – Clarified Shock and Vibe testing shall be tested in the order noted in the specification.

OCP NIC 3.0 Subgroup	<ul style="list-style-type: none"> - Section 2.2 – Clarified line side implementations may be a subset of I/O listed in Table 7. - Section 2.x – Update ejector lever and faceplate to support a clinch nut implementation. <p>Updated Figure 13 & Figure 14 with new ejector lever diagrams Updated Table 9 with updated drawing & clinch nut. Removed bushing and wave washer from BOM. Updated Figures 16, 18 generic I/O plate for ejector clinch nut Updated Figures 19, 20 with Ejector Lever clinch nut updates. Updated Figure 21 with Ejector lock design. Renamed Section 2.4.8 (was Ejector Bushing, now Clinch Nut), Removed Section 2.4.9 (was Ejector Wave washer)</p> <ul style="list-style-type: none"> - Section 3.4.1, 3.4.3 and 3.4.6 – Clarified the PWRBRK[0:1]#, SMBus and AUX_PWR_EN of SFF and LFF. These pins are defined on both the Primary and Secondary Connectors. SFF and LFF cards shall use the Primary Connector pin. The Secondary Connector pin is reserved for a future use case. - Section 4.3 - Clarified MC MAC address provisioning requirements for multi-host capable cards as the maximum number of supported hosts. 	0.91	07/24/2019
OCP NIC 3.0 Subgroup	<ul style="list-style-type: none"> - Section 3.2.4 – Figure 77 was inadvertently changed to include an “Option D” straddle mount connector for a host PCB thickness of 0.105”. This was an error and has been removed. 	0.92	11/22/2019

- Section 3.4.3 – Add “AUX Power Good (local signal)” to the isolator figure.
- Section 3.5.3 – Add implementation note to state baseboard vendors only need to implement PCIe bifurcation options applicable to their topology.
- Section 3.10 – Clarified that RBT hot-plug is not supported in DSP0222 version 1.1 and 1.2. This is feature would be dependent on future revisions of the DSP0222 specification.
- Section 3.11 - Change the max T_{APL} / T_{MPL} (AUX_PWR_EN to NIC_PWR_GOOD and MAIN_PWR_EN to NIC_PWR_GOOD) parameters from 25ms to 50ms.
- Section 4.10.3 – Changed offset three to state “OCP NIC 3.0 card FRU record released with version 0.90”. The same FRU Record field may be used with incremental revisions of the specification.
- Section 7.1.2, 7.1.4 – Add text for co-reporting testing for EN 55024 and EN 55035.

OCP NIC 3.0 Subgroup	<ul style="list-style-type: none"> - Section 2.2.1 – Update max permissible heatsink force applied for internal lock variant. - Section 2.4.3 – Update SFF Generic I/O faceplate – Figures 15-17 - Section 2.4.4 – Update LFF Generic I/O faceplate – Figure 18 - Section 2.4.8 – Update Clinch nut drawings – Figures 22 and 23 - Section 2.5.x – Update Keep Out zone drawings – Figures 24-33 - Section 2.7.1 – Update SFF insulator drawings – Figures 35 and 36 - Section 2.7.2 – Update LFF insulator drawings – Figures 38 and 39 - Section 2.8.2 – Update SFF Pull Tab CTF dimensions – Figures 40-42 - Section 2.8.3 – Update SFF Ejector latch CTF dimensions – Figs 43-45 - Section 2.8.4 – Update SFF Internal latch CTF dimensions – Figs 46-48 - Section 2.8.5 – Update SFF Baseboard CTF dimensions – Figs 50-53 - Section 2.8.6 – Update LFF Ejector latch CTF dimensions – Figs 54-56 - Section 2.8.7 – Update LFF Baseboard CTF dimensions – Figs 58-60 - Section 3.1 – Update SFF-TA-1002 card edge drawings – Figs 67-69 - Section 3.4.5 – For LFF, multiplex scan chain thermal pins (TEMP_WARN, TEMP_CRIT, FAN_ON_AUX) and WAKEn for PRSNTB[3:0] from the secondary connector. - Section 4.3 – Add second Management Controller MAC Address Algorithm. Add link to OCP NIC 3.0 wiki to download MAC Address algorithm calculator spreadsheet. - Section 4.10.3 – Offset 3 – Rearranged sentences for the OEM Record Format description. Technical content remains the same. 	0.99	12/05/2019
----------------------	--	------	------------

8.2 FRU Content Revision History

The following table summarizes the FRU content revision history and maps it against specification releases.

Date	Change Description	Released with Spec Version
01/25/2018	- Initial FRU contents.	0.70
06/04/2018	<ul style="list-style-type: none"> - Offset 4, 5 – Clarified that the Card Max Power in Main mode and the Card Max Power in Aux Mode do not include the power consumed by transceivers plugged into the line side receptacle(s). - Offset 8 – Corrected Card active/passive cooling) as a byte value (not bitwise). - Offset 13, 14 – Removed ASIC temperature target; replaced with UART Configuration 1 and UART Configuration 2. 	0.80

	- Offset 15 – Added USB Present. - Offset 32-127 – Explicitly called out UDID offset ranges for controllers 1 through 6.	
07/06/2018	- Offset 13, 14 – Clarified these offsets are only applicable to the Secondary Connector (LFF designs) - Offset 15 – Clarified this offset is only applicable to the Primary Connector (SFF designs)	0.81
11/20/2018	- Offset 3 – Corrected Manufacturer ID as 0x00A67F. - Offset 9, 10 – Clarified hot aisle standby airflow LSB, MSB byte order. - Offset 11, 12 – Clarified cold aisle standby airflow LSB, MSB byte order. - Offset 16 – Add card manageability type.	0.85
04/02/2019	- Offset 32-127 – Add note declaring the controller UDID fields may be omitted and left as zero length if no corresponding controller is present on the card.	0.86
06/20/2019	- Offset 3 – Updated FRU OEM Record Version compliance. Cards compliant to spec version 0v90 shall use the value 0x01. All other values reserved. - Offset 6 – Clarified Hot Aisle Card Cooling Tier offset is applicable to passive cable or RJ-45 implementations. - Offset 7 – Clarified Cold Aisle Card Cooling Tier offset is applicable to passive cable or RJ-45 implementations. - Offset 9, 10 – Clarified hot aisle standby airflow requirement with an approach air temperature of 45°C. - Offset 11, 12 – Clarified cold aisle standby airflow requirement with an approach air temperature of 35°C. - Offset 17 – Add FRU Write Protection Mechanism indication. - Offset 18 – Add Programming Mode Power State indication. - Offset 19 – Add Hot Aisle Cooling Tier with Active Cables indication. - Offset 20 – Add Cold Aisle Card Cooling Tier with Active Cables indication. - Offset 21 – Add Transceiver Reference Power Level indication. - Offset 22 – Add Transceiver Reference Temperature Level indication. - Offset 23 – Add Card Thermal Tier with local Fan Fail indication. - Offset 31 – Clarify Number of Physical Controllers as the number of SMBus connected controllers on the OCP NIC 3.0 card.	0.90
11/08/2019	- Offset 3 – Changed FRU OEM Record Version definition to declare “this FRU record is released with version x.xx” of the specification. This prevents revising the field to match future specification releases when no FRU changes have been made.	0.92