



OPEN
Compute Project

OCP NIC 3.0 Design Specification

Version 0.74

Author: OCP Server Workgroup, OCP NIC subgroup

Table of Contents

1 Overview	10
1.1 License	10
1.2 Acknowledgements	11
1.3 Background	12
1.4 Overview	14
1.4.1 Mechanical Form factor overview	14
1.4.2 Electrical overview	16
1.5 Non-NIC Use Cases	17
1.6 References	18
1.6.1 Trademarks	19
2 Mechanical Card Form Factor	20
2.1 Form Factor Options	20
2.1.1 Small Form Factor (SFF) Faceplate Configurations	22
2.1.2 Large Form Factor (LFF) Faceplate Configurations	26
2.2 Line Side I/O Implementations	30
2.3 Top Level Assembly (SFF and LFF)	31
2.4 Faceplate Subassembly (SFF and LFF)	32
2.4.1 Faceplate Subassembly – Exploded View	32
2.4.2 Faceplate Subassembly – Bill of Materials (BOM)	32
2.4.3 SFF Generic I/O Faceplate	34
2.4.4 LFF Generic I/O Faceplate	35
2.4.5 Ejector Lever (SFF)	36
2.4.6 Ejector Levers (LFF)	36
2.4.7 Ejector Lock (SFF and LFF)	37
2.5 Card Keep Out Zones	38
2.5.1 Small Card Form Factor Keep Out Zones	38
2.5.2 Large Card Form Factor Keep Out Zones	40
2.6 Baseboard Keep Out Zones	42
2.7 Insulation Requirements	43
2.7.1 Small Card Insulator	43
2.7.2 Large Card Insulator	44
2.8 Critical-to-Function (CTF) Dimensions (SFF and LFF)	45
2.8.1 CTF Tolerances	45
2.8.2 SFF Pull Tab CTF Dimensions	45
2.8.3 SFF OCP NIC 3.0 Card with Ejector Latch CTF Dimensions	47
2.8.4 SFF OCP NIC 3.0 Baseboard CTF Dimensions	48
2.8.5 LFF OCP NIC 3.0 Card CTF Dimensions	50
2.8.6 LFF OCP NIC 3.0 Baseboard CTF Dimensions	51
2.9 Labeling Requirements	53
2.9.1 General Guidelines for Label Contents	53
2.9.2 MAC Address Labeling Requirements	54
2.10 Mechanical CAD Package Examples	58
3 Electrical Interface Definition – Card Edge and Baseboard	59
3.1 Card Edge Gold Finger Requirements	59
3.1.1 Gold Finger Mating Sequence	60
3.2 Baseboard Connector Requirements	62
3.2.1 Right Angle Connector	62
3.2.2 Right Angle Offset	63
3.2.3 Straddle Mount Connector	63
3.2.4 Straddle Mount Offset and PCB Thickness Options	65
3.2.5 Large Card Connector Locations	66

3.3	Pin definition	66
3.4	Signal Descriptions	70
3.4.1	PCIe Interface Pins.....	70
3.4.2	PCIe Present and Bifurcation Control Pins	75
3.4.3	SMBus Interface Pins.....	78
3.4.4	NC-SI Over RBT Interface Pins	79
3.4.5	Scan Chain Pins.....	85
3.4.6	Power Supply Pins	9391
3.4.7	Miscellaneous Pins	9695
3.5	PCIe Bifurcation Mechanism	9796
3.5.1	PCIe OCP NIC 3.0 Card to Baseboard Bifurcation Configuration (PRSNTA#, PRSNTB[3:0]#).....	9796
3.5.2	PCIe Baseboard to OCP NIC 3.0 Card Bifurcation Configuration (BIF[2:0]#)	9796
3.5.3	PCIe Bifurcation Decoder	9897
3.5.4	Bifurcation Detection Flow.....	10099
3.5.5	PCIe Bifurcation Examples.....	101490
3.6	PCIe Clocking Topology.....	106405
3.7	PCIe Bifurcation Results and REFCLK Mapping.....	107406
3.8	Port Numbering and LED Implementations.....	117446
3.8.1	OCP NIC 3.0 Port Naming and Port Numbering.....	117446
3.8.2	OCP NIC 3.0 Card LED Configuration	117446
3.8.3	OCP NIC 3.0 Card LED Ordering.....	119448
3.8.4	Baseboard LEDs Configuration over the Scan Chain	120449
3.9	Power Capacity and Power Delivery.....	120449
3.9.1	NIC Power Off.....	121420
3.9.2	ID Mode.....	121420
3.9.3	Aux Power Mode (S5).....	121420
3.9.4	Main Power Mode (S0).....	121420
3.10	Power Supply Rail Requirements and Slot Power Envelopes	121420
3.11	Hot Swap Considerations for +12V_EDGE and +3.3V_EDGE Rails.....	122424
3.12	Power Sequence Timing Requirements.....	123422
3.13	Digital I/O Specifications	125424
4	Management and Pre-OS Requirements	126425
4.1	Sideband Management Interface and Transport	126425
4.2	NC-SI Traffic	127426
4.3	Management Controller (MC) MAC Address Provisioning	127426
4.4	Temperature Reporting	129428
4.5	Power Consumption Reporting	130429
4.6	Pluggable Transceiver Module Status and Temperature Reporting	131430
4.7	Management and Pre-OS Firmware Inventory and Update.....	131430
4.7.1	Secure Firmware.....	131430
4.7.2	Firmware Inventory.....	132431
4.7.3	Firmware Inventory and Update in Multi-Host Environments	132431
4.8	NC-SI Package Addressing and Hardware Arbitration Requirements	132431
4.8.1	NC-SI over RBT Package Addressing	132431
4.8.2	Arbitration Ring Connections	133432
4.9	SMBus 2.0 Addressing Requirements.....	133432
4.9.1	SMBus Address Map.....	133432
4.10	FRU EEPROM	134433
4.10.1	FRU EEPROM Address, Size and Availability.....	134433
4.10.2	FRU EEPROM Content Requirements.....	134433
4.10.3	FRU Template	137436
5	Routing Guidelines and Signal Integrity Considerations	138437
5.1	NC-SI Over RBT	138437

5.1.1	Channel Budget Requirements	138137
5.2	SMBus 2.0	138137
5.3	PCIe.....	138137
5.3.1	Background.....	138137
5.3.2	Channel Requirements	138137
5.3.3	Test Fixtures	139138
5.3.4	Test Methodology	139138
6	Thermal and Environmental	141140
6.1	Airflow Direction	141140
6.1.1	Hot Aisle Cooling	141140
6.1.2	Cold Aisle Cooling	142141
6.2	Thermal Design Guidelines	143142
6.2.1	SFF Card ASIC Cooling – Hot Aisle.....	143142
6.2.2	LFF Card ASIC Cooling – Hot Aisle.....	146145
6.2.3	SFF Card ASIC Cooling – Cold Aisle	148147
6.2.4	LFF Card ASIC Cooling – Cold Aisle	151150
6.3	Thermal Simulation (CFD) Modeling	153152
6.4	Thermal Test Fixture	153152
6.4.1	Test Fixture for SFF Card.....	154153
6.4.2	Test Fixture for LFF Card.....	155154
6.4.3	Test Fixture Airflow Direction.....	157156
6.4.4	Thermal Test Fixture Candlestick Sensors	157156
6.5	Card Sensor Requirements	159158
6.6	Card Cooling Tiers.....	159158
6.6.1	Hot Aisle Cooling Tiers.....	159158
6.6.2	Cold Aisle Cooling Tiers	159158
6.7	Non-Operational Shock & Vibration Testing	160159
6.7.1	Shock & Vibe Test Fixture.....	160159
6.7.2	Test Procedure	160159
6.8	Dye and Pull Test Method	161160
6.9	Gold Finger Plating Requirements.....	163162
6.9.1	Host Side Gold Finger Plating Requirements.....	163162
6.9.2	Line Side Gold Finger Durability Requirements.....	163162
7	Regulatory.....	164163
7.1	Required Compliance	164163
7.1.1	Required Environmental Compliance.....	164163
7.1.2	Required EMC Compliance	164163
7.1.3	Required Product Safety Compliance.....	165164
7.1.4	Required Immunity (ESD) Compliance	165164
7.2	Recommended Compliance.....	165164
7.2.1	Recommended Environmental Compliance.....	165164
7.2.2	Recommended EMC Compliance	166165
8	Revision History	167166

List of Figures

Figure 1: Representative Small OCP NIC 3.0 Card with Dual QSFP Ports	12
Figure 2: Representative Large OCP NIC 3.0 Card with Dual QSFP Ports and on-board DRAM	13
Figure 3: Small and Large Card Form-Factors (not to scale)	14
Figure 4: Primary Connector (4C+) and Secondary Connector (4C) (Large) OCP NIC 3.0 Cards	20
Figure 5: Primary Connector (4C+) Only (Large) OCP NIC 3.0 Cards	21
Figure 6: Primary Connector (4C+) with 4C and 2C (Small) OCP NIC 3.0 Cards	21
Figure 7: Small Form Factor NIC Configuration Views	23
Figure 8: Small Form Factor NIC Line Side 3D Views	24
Figure 9: Small Form Factor NIC Chassis Mounted 3D Views	25
Figure 10: Large Form Factor NIC Configuration Views	27
Figure 11: Large Form Factor NIC Line Side 3D Views	28
Figure 12: Large Form Factor NIC Chassis Mounted 3D Views	29
Figure 13: PBA Exploded Views (SFF and LFF)	31
Figure 14: Faceplate Assembly Exploded Views (SFF and LFF)	32
Figure 15: Small Card Generic I/O Faceplate with Pulltab Version (2D View)	34
Figure 16: Small Card Generic I/O Faceplate – Ejector Version (2D View)	34
Figure 17: Small Card Generic I/O Faceplate – Internal Lock Version (2D View)	35
Figure 18: Large Card Generic I/O Faceplate – Dual Ejector Version (2D View)	35
Figure 19: Small Card I/O Faceplate – Ejector Lever (2D View)	36
Figure 20: Large Card I/O Faceplate – Ejector Lever (2D View)	36
Figure 21: Ejector Lock	37
Figure 22: Small Form Factor Keep Out Zone – Top View	38
Figure 23: Small Form Factor Keep Out Zone – Top View – Detail A	38
Figure 24: Small Form Factor Keep Out Zone – Bottom View	39
Figure 25: Small Form Factor Keep Out Zone – Side View	39
Figure 26: Small Form Factor Keep Out Zone – Side View – Detail D	40
Figure 27: Large Form Factor Keep Out Zone – Top View	40
Figure 28: Large Form Factor Keep Out Zone – Top View – Detail A	41
Figure 29: Large Form Factor Keep Out Zone – Bottom View	41
Figure 30: Large Form Factor Keep Out Zone – Side View	42
Figure 31: Large Form Factor Keep Out Zone – Side View – Detail D	42
Figure 32: Small Card Bottom Side Insulator (3D View)	43
Figure 33: Small Card Bottom Side Insulator (Top and Side View)	43
Figure 34: Large Card Bottom Side Insulator (3D View)	44
Figure 35: Large Card Bottom Side Insulator (Top and Side View)	44
Figure 36: Small Form Factor OCP NIC 3.0 Card with Pull Tab CTF Dimensions (Top View)	45
Figure 37: Small Form Factor OCP NIC 3.0 Card with Pull Tab CTF Dimensions (Front View)	46
Figure 38: Small Form Factor OCP NIC 3.0 Card with Pull Tab CTF Dimensions (Side View)	46
Figure 39: Small Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Top View)	47
Figure 40: Small Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Front View)	47
Figure 41: Small Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Side View)	47
Figure 42: Small Form Factor Baseboard Chassis CTF Dimensions (Rear View)	48
Figure 43: Small Form Factor Baseboard Chassis to Card Thumb Screw CTF Dimensions (Side View)	48
Figure 44: Small Form Factor Baseboard Chassis to Ejector lever Card CTF Dimensions (Side View)	49
Figure 45: Small Form Factor Baseboard Chassis CTF Dimensions (Rear Rail Guide View)	49
Figure 46: Small Form Factor Baseboard Chassis CTF Dimensions (Rail Guide Detail) – Detail C	49
Figure 47: Large Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Top View)	50
Figure 48: Large Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Front View)	50
Figure 49: Large Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Side View)	51
Figure 50: Large Form Factor Baseboard Chassis CTF Dimensions (Rear View)	51
Figure 51: Large Form Factor Baseboard Chassis CTF Dimensions (Side View)	52
Figure 52: Large Form Factor Baseboard Chassis CTF Dimensions (Rail Guide View)	52
Figure 53: Large Form Factor Baseboard Chassis CTF Dimensions (Rail Guide – Detail C)	52
Figure 54: Small Card Label Area Example	53
Figure 55: MAC Address Label Example 1 – Quad Port with Single Host, Single Managed Controller	55
Figure 56: MAC Address Label Example 2 – Octal Port with Single Host, Dual Managed Controller	55
Figure 57: MAC Address Label Example 3 – Quad Port with Dual Hosts, Dual Managed Controllers	56

Figure 58: MAC Address Label Example 4 – Single Port with Quad Host, Single Managed Controller	57
Figure 59: Small Size Primary Connector Gold Finger Dimensions – x16 – Top Side (“B” Pins)	59
Figure 60: Large Size Card Gold Finger Dimensions – x32 – Top Side (“B” Pins)	60
Figure 61: Large Size Card Gold Finger Dimensions – x32 – Bottom Side (“A” Pins)	60
Figure 62: 168-pin Base Board Primary Connector – Right Angle	62
Figure 63: 140-pin Base Board Secondary Connector – Right Angle	63
Figure 64: OCP NIC 3.0 Card and Host Offset for Right Angle Connectors	63
Figure 65: 168-pin Base Board Primary Connector – Straddle Mount	64
Figure 66: 140-pin Base Board Secondary Connector – Straddle Mount	64
Figure 67: OCP NIC 3.0 Card and Baseboard PCB Thickness Options for Straddle Mount Connectors	65
Figure 68: 0mm Offset (Coplanar) for 0.062” Thick Baseboards	65
Figure 69: 0.3mm Offset for 0.076” Thick Baseboards	66
Figure 70: Primary and Secondary Connector Locations for Large Card Support with Right Angle Connectors	66
Figure 71: Primary and Secondary Connector Locations for Large Card Support with Straddle Mount Connectors	66
Figure 72: PCIe Present and Bifurcation Control Pins (Baseboard Controlled BIF[0:2]#)	77
Figure 73: PCIe Present and Bifurcation Control Pins (Static BIF[0:2]#)	77
Figure 74: Example SMBus Connections	79
Figure 75: NC-SI Over RBT Connection Example – Single Primary Connector	83
Figure 76: NC-SI Over RBT Connection Example – Dual Primary Connectors	84
Figure 77: Example Scan Chain Timing Diagram	86
Figure 78: Scan Chain Connection Example	90
Figure 79: Example Power Supply Topology	9695
Figure 80: Single Host (1 x16) and 1 x16 OCP NIC 3.0 Card (Single Controller)	101400
Figure 81: Single Host (2 x8) and 2 x8 OCP NIC 3.0 Card (Dual Controllers)	102401
Figure 82: Quad Hosts (4 x4) and 4 x4 OCP NIC 3.0 Card (Single Controller)	103402
Figure 83: Quad Hosts (4 x4) and 4 x4 OCP NIC 3.0 Card (Quad Controllers)	104403
Figure 84: Single Host with no Bifurcation (1 x16) and 2 x8 OCP NIC 3.0 Card (Dual Controllers)	105404
Figure 85: PCIe Interface Connections for 1 x16 and 2 x8 OCP NIC 3.0 Cards	106405
Figure 86: PCIe Interface Connections for a 4 x4 OCP NIC 3.0 Card	107406
Figure 87: Port and LED Ordering – Example Small Card Link/Activity and Speed LED Placement	119418
Figure 88: Baseboard Power States	120419
Figure 89: Power-Up Sequencing	123422
Figure 90: Power-Down Sequencing	124423
Figure 91: Airflow Direction for Hot Aisle Cooling (SFF and LFF)	141440
Figure 92: Airflow Direction for Cold Aisle Cooling (SFF and LFF)	142441
Figure 93: ASIC Supportable Power for Hot Aisle Cooling – Small Card Form Factor	143442
Figure 94: OCP NIC 3.0 SFF Reference Design and CFD Geometry	144443
Figure 95: Server System Airflow Capability – SFF Card Hot Aisle Cooling	145444
Figure 96: ASIC Supportable Power for Hot Aisle Cooling – LFF Card	146445
Figure 97: OCP NIC 3.0 LFF Reference Design and CFD Geometry	146445
Figure 98: Server System Airflow Capability – LFF Card Hot Aisle Cooling	148447
Figure 99: ASIC Supportable Power for Cold Aisle Cooling – SFF Card	149448
Figure 100: Server System Airflow Capability – SFF Cold Aisle Cooling	149448
Figure 101: ASIC Supportable Power Comparison – SFF Card	150449
Figure 102: ASIC Supportable Power for Cold Aisle Cooling – LFF Card	151450
Figure 103: Server System Airflow Capability – LFF Cold Aisle Cooling	152451
Figure 104: ASIC Supportable Power Comparison – LFF Card	152451
Figure 105: Small Card Thermal Test Fixture Preliminary Design	154453
Figure 106: SFF Card Thermal Test Fixture Preliminary Design – Cover Removed	154453
Figure 107: SFF Card Thermal Test Fixture PCB	155454
Figure 108: LFF Card Thermal Test Fixture Design	155454
Figure 109: LFF Card Thermal Test Fixture Design – Cover Removed	156455
Figure 110: LFF Card Thermal Test Fixture PCB	156455
Figure 111: Thermal Test Fixture Airflow Direction	157456
Figure 112: SFF Fixture, Hot Aisle Flow - Candlestick Air Velocity vs. Volume Flow	158457
Figure 113: LFF Fixture, Hot Aisle Flow - Candlestick Air Velocity vs. Volume Flow	158457
Figure 114: Dye and Pull Type Locations	162461
Figure 115: Dye Coverage Percentage	162461

List of Tables

Table 1: Acknowledgements – By Company	11
Table 2: OCP 3.0 Form Factor Dimensions	15
Table 3: Baseboard to OCP NIC Form factor Compatibility Chart.....	15
Table 4: Example Non-NIC Use Cases	17
Table 5: OCP NIC 3.0 Card Definitions	22
Table 6: OCP NIC 3.0 Line Side I/O Implementations	30
Table 7: Bill of Materials for the SFF and LFF Faceplates for the Large Card Assembly.....	33
Table 8: CTF Default Tolerances (SFF and LFF OCP NIC 3.0).....	45
Table 9: MAC Address Label Example 1 – Quad Port with Single Host, Single Managed Controller	55
Table 10: MAC Address Label Example 2 – Octal Port with Single Host, Dual Managed Controller	55
Table 11: MAC Address Label Example 3 – Quad Port with Dual Hosts, Dual Managed Controller	56
Table 12: MAC Address Label Example 4 – Single Port with Quad Host, Single Managed Controller	56
Table 13: NIC Implementation Examples and 3D CAD	58
Table 14: Contact Mating Positions for the Primary and Secondary Connectors	60
Table 15: Right Angle Connector Options.....	62
Table 16: Straddle Mount Connector Options	63
Table 17: Primary Connector Pin Definition (x16) (4C+).....	67
Table 18: Secondary Connector Pin Definition (x16) (4C)	69
Table 19: Pin Descriptions – PCIe	70
Table 20: Pin Descriptions – PCIe Present and Bifurcation Control Pins	75
Table 21: Pin Descriptions – SMBus.....	78
Table 22: Pin Descriptions – NC-SI Over RBT.....	79
Table 23: Pin Descriptions – Scan Chain	85
Table 24: Pin Descriptions – Scan Chain DATA_OUT Bit Definition	87
Table 25: Pin Descriptions – Scan Chain DATA_IN Bit Definition	87
Table 26: Pin Descriptions – Power.....	9391
Table 27: Pin Descriptions – Miscellaneous	9695
Table 28: PCIe Bifurcation Decoder for x16 and x8 Card Widths.....	9998
Table 29: PCIe Clock Associations	106405
Table 30: Bifurcation for Single Host, Single Socket and Single Upstream Link (BIF[2:0]#=0b000)	108407
Table 31: Bifurcation for Single Host, Single Socket and Single/Dual Upstream Links (BIF[2:0]#=0b000)	109408
Table 32: Bifurcation for Single Host, Single Socket and Single/Dual/Quad Upstream Links (BIF[2:0]#=0b000)	110409
Table 33: Bifurcation for Single Host, Dual Sockets and Dual Upstream Links (BIF[2:0]#=0b001).....	111410
Table 34: Bifurcation for Single Host, Quad Sockets and Quad Upstream Links (BIF[2:0]#=0b010).....	112411
Table 35: Bifurcation for Single Host, Quad Sockets and Quad Upstream Links – First 8 PCIe Lanes (BIF[2:0]#=0b011)	113412
Table 36: Bifurcation for Dual Host, Dual Sockets and Dual Upstream Links (BIF[2:0]#=0b101)	114413
Table 37: Bifurcation for Quad Host, Quad Sockets and Quad Upstream Links (BIF[2:0]#=0b110)	115414
Table 38: Bifurcation for Quad Host, Quad Sockets and Quad Upstream Links – First 8 lanes (BIF[2:0]#=0b111)	116415
Table 39: OCP NIC 3.0 Card LED Configuration with Two Physical LEDs per Port	118417
Table 40: Power States.....	120419
Table 41: Baseboard Power Supply Rail Requirements – Slot Power Envelopes	121420
Table 42: Power Sequencing Parameters	124423
Table 43: Digital I/O DC specifications	125424
Table 44: Digital I/O AC specifications	125424
Table 45: OCP NIC 3.0 Management Implementation Definitions	126425
Table 46: Sideband Management Interface and Transport Requirements.....	126425
Table 47: NC-SI Traffic Requirements	127426
Table 48: MC MAC Address Provisioning Requirements	127426
Table 49: Temperature Reporting Requirements	129428
Table 50: Power Consumption Reporting Requirements.....	130429
Table 51: Pluggable Module Status Reporting Requirements	131430
Table 52: Management and Pre-OS Firmware Inventory and Update Requirements	131430
Table 53: Slot_ID[1:0] to Package ID[2:0] Mapping	132431
Table 54: FRU EEPROM Address Map	133432
Table 55: FRU EEPROM Record – OEM Record 0xC0, Offset 0x00	134433
Table 56: Hot Aisle Air Temperature Boundary Conditions	142441
Table 57: Hot Aisle Airflow Boundary Conditions	142441

Table 58: Cold Aisle Air Temperature Boundary Conditions	142 141
Table 59: Cold Aisle Airflow Boundary Conditions	143 142
Table 60: Reference OCP NIC 3.0 SFF Card Geometry	144 143
Table 61: Reference OCP NIC 3.0 LFF Card Geometry	147 146
Table 62: Hot Aisle Card Cooling Tier Definitions (LFM)	159 158
Table 63: Cold Aisle Card Cooling Tier Definitions (LFM)	160 159
Table 64: Random Vibration Testing 1.88G _{RMS} Profile	161 160
Table 65: FCC Class A Radiated and Conducted Emissions Requirements Based on Geographical Location	164 163
Table 66: Safety Requirements	165 164
Table 67: Immunity (ESD) Requirements	165 164

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:

<http://www.opencompute.org/assets/download/01-Contribution-Licenses/OCPHL-Permissive-v1.0.pdf>

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	Mellanox Technologies, Ltd
Dell, Inc.	Netronome Systems, Inc.
Facebook, Inc.	Quanta Computer Inc.
Hewlett Packard Enterprise Company	TE Connectivity Corporation
Intel Corporation	

1.3 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 Card, and Large Card. The Small Card allows for up to 16 PCIe lanes on the card edge while the Large Card 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 80W to a single connector (Small) card, and up to 150W to a dual connector (Large) 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 Small Card OCP NIC 3.0 card is shown in Figure 1 and a representative Large Card is shown in Figure 2.

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





Figure 2: Representative Large 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

<http://opencompute.org/products/specsanddesign?keyword=SPEC%2C+NIC>

1.4 Overview


1.4.1 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 – Small and Large. These cards are shown in Figure 3 below. The components shown in the figures are for illustrative purposes. The Small form factor card has one connector (Primary Connector) on the baseboard. The Large form factor card has 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 Small Card gold finger area only occupies the Primary Connector area for up to 16 PCIe lanes. The Large Card 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: Small and Large Card Form-Factors (not to scale)



The two form factor dimensions are shown in Table 2.

Table 2: OCP 3.0 Form Factor Dimensions

Form Factor	Width	Depth	Primary Connector	Secondary Connector	Typical Use Case
Small	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.
Large	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 3 shows the compatibility between the baseboard and NIC combinations. A Small size baseboard slot may only accept a small sized NIC. A Large size baseboard slot may accept a small or large sized NIC.

Table 3: Baseboard to OCP NIC Form factor Compatibility Chart

Baseboard Slot Size	NIC Size / Supported PCIe Width	
	Small	Large
Small	Up to 16 PCIe lanes	Not Supported
Large	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 3.

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.

1.4.2 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 Small and Large form-factors and may support up to 16 lanes of PCIe. The Secondary Connector, when used in conjunction with the Primary Connector, allows Large form-factor implementations and may support up to 32 lanes of PCIe.

1.4.2.1 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 1.1 compliant Network Controller Sideband Interface (NC-SI) RMIII Based Transport (RBT) Physical Interface
- Power management and status reporting
 - Power break 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.

PCIe 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
- 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.

1.4.2.2 Secondary Connector

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

PCIe 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
- 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.

1.5 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 4.

Table 4: 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	TBD

1.6 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.
- 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*, draft, May 2017.
- 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>
- PCIe Base Specification. *PCI Express Base Specification*, Revision 3.0 December 7th, 2015.
- PCIe Base Specification. *PCI Express Base Specification*, Revision 4.0 Version 1.0, October 5th, 2017.

- PCIe CEM Specification. *PCI Express Card Electromechanical Specification*, Revision 3.0, July 21st, 2013.
- PCIe CEM Specification. *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.

1.6.1 Trademarks

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

2 Mechanical Card Form Factor

2.1 Form Factor Options

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

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 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 the Primary Connector. This is also mandatory for OCP NIC 3.0 cards.

The Small Card 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 small size card provides sufficient faceplate area to accommodate up to 2x QSFP modules, 4x SFP modules, or 4x RJ-45 for BASE-T operation. The Small Card form factor supports up to 80W of delivered power to the card edge. An example Small Card form factor is shown in Figure 1.

The Large Card uses the Primary 4C+ connector to provide the same functionality as the Small Card along with an additional Secondary 4C connector to provide up to a x32 PCIe interface. The Large Card may utilize both the Primary and Secondary Connectors, or just the Primary Connector for lower PCIe lane count applications. Table 5 summarizes the Large Card permutations. The Large Card supports higher power envelopes and provides additional board area for more complex designs. The Large Card form factor supports up to 150W of delivered power to the card edge across the two connectors. An example Large Card form factor is shown in Figure 2.

For Large 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) (Large) OCP NIC 3.0 Cards





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



For both form-factors, an OCP NIC 3.0 card may optionally implement a subset of pins to support up to a x8 PCIe connection. This is implemented using a 2C card edge per SFF-TA-1002. The Primary Connector may support a 2C sized OCP NIC 3.0 card along with the 28 pin OCP bay shown in the 4C+ drawings. The following diagram from the SFF-TA-1002 specification illustrates the supported host Primary and Secondary Connectors and OCP NIC 3.0 card configurations.

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




Table 5 summarizes the supported card form factors. Small form factors cards support the Primary Connector and up to 16 PCIe lanes. Large form factor 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 5: OCP NIC 3.0 Card Definitions

Add in Card Size and max PCIe Lane Count	Secondary Connector		Primary Connector	
	4C Connector, x16 PCIe		4C+ Connector, x16 PCIe	
Small (x8)			2C+	OCP Bay
Small (x16)			4C+	OCP Bay
Large (x8)			2C+	OCP Bay
Large (x16)			4C+	OCP Bay
Large (x32)	4C		4C+	OCP Bay

2.1.1 Small Form Factor (SFF) Faceplate Configurations

The small form factor (SFF) configuration views are shown below. Two different faceplates are available for the SFF – a pull tab version (on the left) and an ejector latch version (on the right). The same SFF OCP NIC 3.0 PBA assembly accepts both type of faceplates 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 3.0mm x 3.0mm 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 guiderail. If the locking feature is implemented on the baseboard, the OCP NIC 3.0 card may only be removed after pressing on an internal detent mechanism. This retention notch is compatible with all chassis implementations. Please refer to the SFF and LFF CTF dimensions in Section 2.8 for details.

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: Small Form Factor 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 8: Small Form Factor NIC Line Side 3D Views




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.



As previously noted, the OCP NIC 3.0 card provides a notch on the rail edge for an internal locking mechanism to prevent card removal. The internal locking mechanism is an optional feature and is not shown in the views below.

Figure 9: Small Form Factor NIC Chassis Mounted 3D Views



Small Form Factor (SFF) with Pull Tab

Small Form Factor (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 Large Form Factor (LFF) Faceplate Configurations

The large form factor (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 3.0mm x 3.0mm 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 2.5.2.

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: Large Form Factor NIC Configuration Views
Large Form Factor (LFF) with Ejector Latch





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: Large Form Factor NIC Line Side 3D Views

Large Form Factor (LFF) with Long Ejector Latch Large Form Factor (LFF) with Short Ejector Latch



Shadowed Configuration




Non-Shadowed Configuration



Dual QSFP




Quad SFP



Under development.

Quad RJ45




Under development.

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: Large Form Factor NIC Chassis Mounted 3D Views

Large Form Factor (LFF) with Ejector Latch



Right Angle Baseboard Connector




Straddle Mount Baseboard Connector



NIC Installed in Chassis



NIC Insertion / Removal (As shown with a Straddle Mount Connector)



2.2 Line Side I/O Implementations

At the time of this writing, the Small and Large form-factor implementations have been optimized to support the following standard line side I/O implementations:

Table 6: OCP NIC 3.0 Line Side I/O Implementations

Form Factor	Max Topology Connector Count
Small	2x QSFP+/QSFP28
Small	4x SFP28+/SFP28
Small	4x RJ-45
Large	2x QSFP+/QSFP28
Large	4x SFP+/SFP28
Large	4x RJ-45

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.

Additional combinations and connector types 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, #11, and #12 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 small form factor and large form factor 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 7 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 7.

Table 7: Bill of Materials for the SFF and LFF Faceplates for the Large Card Assembly

Item #	Item description	Part Number / Drawing	Supplier								
1A 1B 1C 1D	Faceplate	See Section 2.4.3: <table border="1"> <tr> <td>1A</td> <td>NIC_OCPv3_SFF_Faceplate_Pulltab_20180601.pdf</td> </tr> <tr> <td>1B</td> <td>NIC_OCPv3_SFF_Faceplate_Latch_20180601.pdf</td> </tr> <tr> <td>1C</td> <td>NIC_OCPv3_SFF_Faceplate_IntLock_20180601.pdf</td> </tr> </table> See Section 2.4.4: <table border="1"> <tr> <td>1D</td> <td>NIC_OCPv3_LFF_Faceplate_Latch_20180601.pdf</td> </tr> </table>	1A	NIC_OCPv3_SFF_Faceplate_Pulltab_20180601.pdf	1B	NIC_OCPv3_SFF_Faceplate_Latch_20180601.pdf	1C	NIC_OCPv3_SFF_Faceplate_IntLock_20180601.pdf	1D	NIC_OCPv3_LFF_Faceplate_Latch_20180601.pdf	Custom
1A	NIC_OCPv3_SFF_Faceplate_Pulltab_20180601.pdf										
1B	NIC_OCPv3_SFF_Faceplate_Latch_20180601.pdf										
1C	NIC_OCPv3_SFF_Faceplate_IntLock_20180601.pdf										
1D	NIC_OCPv3_LFF_Faceplate_Latch_20180601.pdf										
2A 2B 2C 2D	Top and Bottom EMI Fingers	<table border="1"> <tr> <td>2A</td> <td>LT18CJ1921 – 13 fingers (Laird) TF187VE32F11-2.41-08 (Tech-Etch)</td> </tr> <tr> <td>2B</td> <td>LT18CJ1920 – 11 fingers (Laird) TF187VE32F11-2.04-08 (Tech-Etch)</td> </tr> <tr> <td>2C</td> <td>LT18CJ1923 – 27 fingers (Laird) TF187VE32F11-5.03-08 (Tech-Etch)</td> </tr> <tr> <td>2D</td> <td>LT18CJ1922 – 25 fingers (Laird) TF187VE32F11-4.66-08 (Tech-Etch)</td> </tr> </table>	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)	Laird, Tech-ETCH
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)										
3	Rivet	1-AC-2421-03_2.4x2.1	Dong Guan KSETT Hardware Technology								
4	Side EMI Fingers	LT18DP1911	Laird								
5	Thumbscrew	J-4C-99-343-KEEE_rev06	Southco, Inc.								
6A 6B	Pull tab w/2x screws	J-CN-99-459	Southco, Inc.								
8	Ejector Wave Compression Washer	NIC_OCPv3_EjectorWasher_20180601.pdf	Custom								
9A 9B	Ejector Handle	SFF Ejector: See Section 2.4.5 and drawing <table border="1"> <tr> <td>9A</td> <td>NIC_OCPv3_EjectorHandle_Short_20180601.pdf</td> </tr> </table> Note: The SFF ejector is also used on the LFF non-shadowed I/O faceplate configuration. LFF Ejector: See Section 2.4.6 & Drawing <table border="1"> <tr> <td>9B</td> <td>NIC_OCPv3_EjectorHandle_Long_20180601.pdf</td> </tr> </table>	9A	NIC_OCPv3_EjectorHandle_Short_20180601.pdf	9B	NIC_OCPv3_EjectorHandle_Long_20180601.pdf	Custom				
9A	NIC_OCPv3_EjectorHandle_Short_20180601.pdf										
9B	NIC_OCPv3_EjectorHandle_Long_20180601.pdf										
10	Ejector Lock	See Section 2.4.7 and drawing NIC_OCPv3_EjectorLock_20180601.pdf	Custom								
11	Ejector Bushing	NIC_OCPv3_EjectorBushing_20180601.pdf	Custom								
12	Screw for securing faceplate to NIC	ICMMAJ200403N3	WUJIANG Screw Tech Precision Industry								
13	Screw for attaching faceplate and ejector to NIC	FCMMQ200503N	WUJIANG Screw Tech Precision Industry								
14	SMT nut (on NIC)	82-950-22-010-01-RL	Fivetech Technology Inc.								
15A 15B	Insulator	Refer to Section 2.7 for the SFF (15A) and LFF (15B) insulator mechanical requirements	Custom								

2.4.3 SFF Generic I/O Faceplate

Figure 15 shows the standard Small Card form factor I/O bracket with a thumbscrew and pull tab assembly.

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




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






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



2.4.4 LFF Generic I/O Faceplate


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



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: Small Card I/O Faceplate – Ejector Lever (2D View)




1. DIMENSION ARE IN MILLIMETER
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.

2.4.6 Ejector Levers (LFF)


This section defines the LFF ejector lever dimensions.

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



1. DIMENSION ARE IN MILLIMETER
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.


Figure 24: Small Form Factor Keep Out Zone – Bottom View



NOTES:
1. TOLERANCE UNLESS OTHERWISE SPECIFIED: $\pm 0.13, \pm 1.0^\circ$


Figure 25: Small Form Factor Keep Out Zone – Side View

DIMENSION "A" CONNECTOR HEIGHT	MAX (mm)
WITH LATCH	1.70
W/O LATCH	2.77



NOTES:
1. TOLERANCE UNLESS OTHERWISE SPECIFIED: $\pm 0.13, \pm 1.0^\circ$

Figure 26: Small Form Factor Keep Out Zone – Side View – Detail D



2.5.2 Large Card Form Factor Keep Out Zones

Figure 27: Large Form Factor Keep Out Zone – Top View




Figure 28: Large Form Factor Keep Out Zone – Top View – Detail A




Figure 29: Large Form Factor Keep Out Zone – Bottom View




Figure 30: Large Form Factor Keep Out Zone – Side View





Figure 31: Large Form Factor 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 Small and Large Card form factor designs. The 3D CAD files are available for download on the OCP NIC 3.0 Wiki:

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

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 0.25mm and shall reside within the following mechanical envelope for the Small and Large size cards.

2.7.1 Small Card Insulator

Figure 32: Small Card Bottom Side Insulator (3D View)





Figure 33: Small Card Bottom Side Insulator (Top and Side View)



2.7.2 Large Card Insulator

Figure 34: Large Card Bottom Side Insulator (3D View)





Figure 35: Large Card Bottom Side Insulator (Top and Side View)

1. DIMENSION ARE IN MILIMETER
2. MATERIAL: FORMEX GK-10BK 0.25mm THICKNESS
3. ADHESIVE 3M 467MP 0.05mm THICKNESS
4. TOLERANCE UNLESS OTHERWISE SPECIFIED: $\pm 0.30\text{mm}$, $\pm 1.0^\circ$



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 small form factor and large form factor cards.

Table 8: 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 small form factor OCP NIC 3.0 card with a pull tab and thumbscrew. The CTF default tolerances are shown in Section 2.8.1.

Figure 36: Small Form Factor OCP NIC 3.0 Card with Pull Tab CTF Dimensions (Top View)




Figure 37: Small Form Factor OCP NIC 3.0 Card with Pull Tab CTF Dimensions (Front View)





Figure 38: Small Form Factor OCP NIC 3.0 Card with Pull Tab CTF Dimensions (Side View)



2.8.3 SFF OCP NIC 3.0 Card with Ejector Latch CTF Dimensions

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

Figure 39: Small Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Top View)




Figure 40: Small Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Front View)





Figure 41: Small Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Side View)



2.8.4 SFF OCP NIC 3.0 Baseboard CTF Dimensions

The following dimensions are considered critical-to-function (CTF) for each small form factor 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 42: Small Form Factor Baseboard Chassis CTF Dimensions (Rear View)




Figure 43: Small Form Factor Baseboard Chassis to Card Thumb Screw CTF Dimensions (Side View)




Figure 44: Small Form Factor Baseboard Chassis to Ejector lever Card CTF Dimensions (Side View)




Figure 45: Small Form Factor Baseboard Chassis CTF Dimensions (Rear Rail Guide View)





Figure 46: Small Form Factor Baseboard Chassis CTF Dimensions (Rail Guide Detail) – Detail C



The right angle and straddle mount card guides are identical between the Small and Large form factor 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>.

2.8.5 LFF OCP NIC 3.0 Card CTF Dimensions

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

Figure 47: Large Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Top View)




Figure 48: Large Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Front View)





Figure 49: Large Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Side View)



2.8.6 LFF OCP NIC 3.0 Baseboard CTF Dimensions

The following dimensions are considered critical-to-function (CTF) for each large form factor 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 50: Large Form Factor Baseboard Chassis CTF Dimensions (Rear View)




Figure 51: Large Form Factor Baseboard Chassis CTF Dimensions (Side View)






Figure 52: Large Form Factor Baseboard Chassis CTF Dimensions (Rail Guide View)



SECTION B-B

Figure 53: Large Form Factor Baseboard Chassis CTF Dimensions (Rail Guide – Detail C)



DETAIL C

The right angle and straddle mount card guides are identical between the Small and Large form factor cards. 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>.

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 54: Small Card 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 $X \geq 5\text{mil}$ (0.127mm)
- 2D data matrix
- Data matrix shall use ECC200 error correction
- Minimum cell size $X \geq 10\text{mil}$ (0.254mm)
- 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/IER 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 (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.

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 9. The

constructed label is shown in Figure 55. 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 9: 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 55: 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 10. The constructed label is shown in Figure 56. 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 10: 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	P:8	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 56: 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 11 and Figure 57.

Table 11: 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 #0

Figure 57: 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 P_x+1. This is shown in Table 12 and Figure 58.

Table 12: 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 58: MAC Address Label Example 4 – Single Port with Quad Host, Single Managed Controller



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 13: NIC Implementation Examples and 3D CAD

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

Small Size cards fit in the Primary Connector. Primary Connector compliant cards are 76mm x 115mm 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 1.57mm. The gold finger dimensions for the Primary Connector compliant cards are shown below.

Large Size Cards support up to a x32 PCIe implementation and may use both the Primary and Secondary (4C) Connectors. Large Size Cards may implement a reduced PCIe lane count and optionally implement only the Primary Connector 4C+, or 2C OCP bay.

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 59: Small Size Primary Connector Gold Finger Dimensions – x16 – Top Side (“B” Pins)




Figure 60: Large Size Card Gold Finger Dimensions – x32 – Top Side (“B” Pins)





Figure 61: Large Size Card 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 recommended OCP NIC 3.0 card gold finger staging is shown in Table 14 for a two stage, first-mate, last-break functionality. 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 14 are used for first mate/second mate reference only. Full pin definitions are described in Sections 3.3 and 3.4.

Table 14: Contact Mating Positions for the Primary and Secondary Connectors

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		

OCN B14	RBT_CRS_DV	OCN 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	BIFO#	A7	SMCLK
B8	BIF1#	A8	SMDAT
B9	BIF2#	A9	SMRST#
B10	PERST0#	A10	PRNSTA#
B11	+3.3V_EDGE	A11	PERST1#
B12	AUX_PWR_EN	A12	PRNSTB2#
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	PRNSTB0#	A42	PRNSTB1#
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	UART_RXRFU1_N/C		A68	RFU-23, N/C	
B69	UART_TXRFU2_N/C		A69	RFU-43, N/C	
B70	PRSNTB3#		A70	PWRBRK#	

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 15: Right Angle Connector Options

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

Figure 62: 168-pin Base Board Primary Connector – Right Angle






Figure 63: 140-pin Base Board Secondary Connector – Right Angle



3.2.2 Right Angle Offset

The OCP NIC 3.0 right angle connectors have a 4.05mm offset from the baseboard (pending SI simulation results). This is shown in Figure 64.

Figure 64: 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 16: Straddle Mount Connector Options

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

Secondary Connector – 4C	140 pins	Straddle Mount for 0.093"	Coplanar (0mm)
--------------------------	----------	---------------------------	----------------

Figure 65: 168-pin Base Board Primary Connector – Straddle Mount





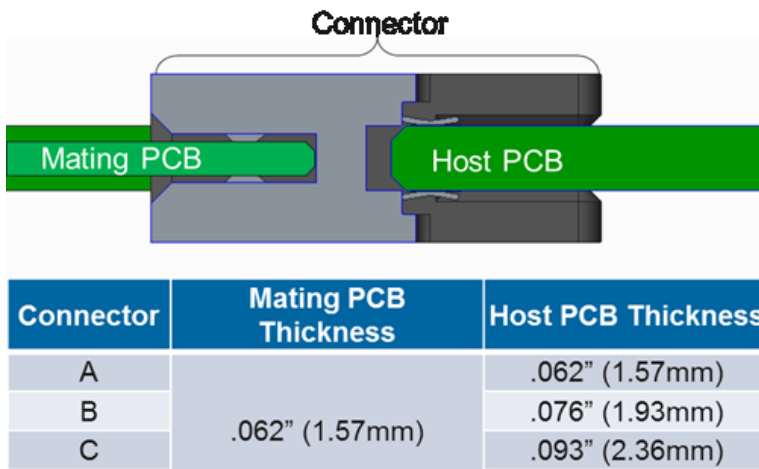
Figure 66: 140-pin Base Board Secondary Connector – Straddle Mount



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 67. 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 67: 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 68. Additionally, the connectors are also capable of having a 0.3mm offset from the centerline of the host board as shown in Figure 69.

Figure 68: 0mm Offset (Coplanar) for 0.062" Thick Baseboards





Figure 69: 0.3mm Offset for 0.076" Thick Baseboards



3.2.5 Large Card Connector Locations

In order to support the large form factor, systems must locate the Primary and Secondary Connectors per the mechanical drawing shown in Figure 70 and Figure 71.

Figure 70: Primary and Secondary Connector Locations for Large Card Support with Right Angle Connectors





Figure 71: Primary and Secondary Connector Locations for Large Card 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 17 and Table 18. 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 pins common to the Primary and Secondary Connectors are shown in Section 3.4. The OCP Bay pins on the Primary Connector only are explicitly called out with the prefix "OCP_" in pin location column.

Cards or systems that do not require the use of a PCIe x16 connection may optionally implement a subset 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.

Table 17: Primary Connector Pin Definition (x16) (4C+)

Side B		Side A	
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	PRSNTB2#	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

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)

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	UART_RX RFU1, N/C	RFU23, N/C	A68
B69	UART_TX RFU2, N/C	RFU43, N/C	A69
B70	PRSNB3#	PWRBRK#	A70

Table 18: 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	PERST0#	PRSENTA#	A10
B11	+3.3V_EDGE	PERST1#	A11
B12	AUX_PWR_EN	PRSENTB2#	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	PRSENTB0#	PRSENTB1#	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

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

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	UART-RXRFU1, N/C	RFU2RFU3, N/C	A68
B69	UART-TXRFU2, N/C	RFU3RFU4, N/C	A69
B70	PRNTB3#	PWRBRK#	A70

3.4 Signal Descriptions

The pins shown in this section are for both the Primary and Secondary Connectors. Pins that exist only for the Primary Connector OCP Bay are explicitly called out in the pin location column with the prefix “OCP_xxx”. 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 AC/DC specifications are defined in the PCIe CEM Specification, Rev 4.0. Example connection diagrams for are shown in Figure 85 and Figure 86.

Table 19: Pin Descriptions – PCIe

Signal Name	Pin #	Baseboard Direction	Signal Description
REFCLKn0 REFCLKp0	B14 B15	Output	PCIe compliant differential reference clock #0, #1, #2 and #3. 100MHz reference clocks are used for the OCP NIC 3.0 card PCIe core logic. REFCLK0 is always available to all OCP NIC 3.0 cards. The card should not assume REFCLK1, REFCLK2 or
REFCLKn1 REFCLKp1	A14 A15	Output	
REFCLKn2 REFCLKp2	OCP_B11 OCP_B12	Output	

REFCLKn3 REFCLKp3	OCP_A11 OCP_A12	Output	<p>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 connector for supported designs. Separate REFCLK0 and REFCLK1 instances are available for the Primary and Secondary connectors. REFCLK2 and REFCLK3 are only available on the Primary connector in the OCP Bay.</p> <ul style="list-style-type: none"> REFCLK0 is required for all designs. REFCLK1, REFCLK2 and REFCLK3 are required for designs that support 2 <i>xn</i>, 4 <i>xn</i>, 8 <i>xn</i> bifurcation implementations. <p>Baseboards that implement REFCLK1, REFCLK2 and REFCLK3, 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 REFCLK1, REFCLK2 or REFCLK3 are not implemented.</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> <p>Refer to Section 2.1 in the PCIe CEM Specification, Rev 4.0 for electrical details.</p>
PETn0 PETp0	B17 B18	Output	<p>Transmitter differential pairs [0:15]. 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.</p>
PETn1 PETp1	B20 B21	Output	
PETn2 PETp2	B23 B24	Output	
PETn3 PETp3	B26 B27	Output	
PETn4	B30	Output	

PETp4	B31		<p>For baseboards, the PET[0:15] signals are required at the connector.</p> <p>For 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>Refer to Section 6.1 in the PCIe CEM Specification, Rev 4.0 for details.</p>
PETn5	B33	Output	
PETp5	B34		
PETn6	B36	Output	
PETp6	B37		
PETn7	B39	Output	
PETp7	B40		
PETn8	B44	Output	
PETp8	B45		
PETn9	B47	Output	
PETp9	B48		
PETn10	B50	Output	
PETp10	B51		
PETn11	B53	Output	
PETp11	B54		
PETn12	B56	Output	
PETp12	B57		
PETn13	B59	Output	
PETp13	B60		
PETn14	B62	Output	
PETp14	B63		
PETn15	B65	Output	
PETp15	B66		
PERn0	A17	Input	<p>Receiver differential pairs [0:15]. These pins are connected from the OCP NIC 3.0 card transmitter differential pairs to the receiver differential pairs on the baseboard.</p> <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.</p> <p>For baseboards, the PER[0:15] signals are required at the connector.</p> <p>For 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>Refer to Section 6.1 in the PCIe CEM Specification, Rev 4.0 for details.</p>
PERp0	A18		
PERn1	A20	Input	
PERp1	A21		
PERn2	A23	Input	
PERp2	A24		
PERn3	A26	Input	
PERp3	A27		
PERn4	A30	Input	
PERp4	A31		
PERn5	A33	Input	
PERp5	A34		
PERn6	A36	Input	
PERp6	A37		
PERn7	A39	Input	
PERp7	A40		
PERn8	A44	Input	
PERp8	A45		
PERn9	A47	Input	
PERp9	A48		
PERn10	A50	Input	
PERp10	A51		
PERn11	A53	Input	
PERp11	A54		

PERn12 PERp12	A56 A57	Input	
PERn13 PERp13	A59 A60	Input	
PERn14 PERp14	A62 A63	Input	
PERn15 PERp15	A65 A66	Input	
PERST0# PERST1# PERST2# PERST3#	B10 A11 OCP_A1 OCP_A2	Output	

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 10kOhm 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 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>
PWRBRK#	A70	Output, OD	<p>Power break. Active low, open drain.</p> <p>This signal shall be pulled up to +3.3V_EDGE on the OCP NIC 3.0 card with a minimum of 95kOhm. The pull up on the baseboard shall be a stiffer resistance in-order to meet the timing specs 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 PWRBRK# pin shall be implemented and available on the Primary Connector.</p> <p>For OCP NIC 3.0 cards, the PWRBRK# pin usage is optional. If used, the PWRBRK# should be connected to the network silicon to enable reduced power state. If not used, the PWRBRK# signal shall be left as a no connect.</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 3.12. Example connection diagrams are shown in Figure 72 and Figure 73.

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 latched before PWR_EN assertion to ensure the correct values are detected by the system. Changing the pin states after this timing window is not allowed. Refer to the AC timing diagram in Section 3.12 for details.

Table 20: 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 1kOhm 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>

<p>BIF0# BIF1# BIF2#</p>	<p>B7 B8 B9</p>	<p>Output</p>	<p>Bifurcation [0:2]# pins allow the baseboard to force configure the OCP NIC 3.0 card bifurcation.</p> <p>For baseboards, these pins shall be outputs driven from the baseboard I/O hub and allow the system to force configure 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 are described in Section 3.5 if no dynamic bifurcation configuration is required.</p> <p>The BIF[0:2]# pins shall be low until AUX_PWR_EN is asserted to prevent leakage paths into an unpowered card.</p> <p>For baseboards that allow dynamic bifurcation, the BIF[0:2] pins are driven low prior to AUX_PWR_EN. Refer to Figure 72 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 73 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.</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 72: PCIe Present and Bifurcation Control Pins (Baseboard Controlled BIF[0:2]#)





Figure 73: 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 and I²C bus specifications. An example connection diagram is shown in Figure 74.

Table 21: 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. The SMRST# signal shall be left as a no connect if it is not used on the OCP NIC 3.0 card.</p>

Figure 74: 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. An example connection diagram is shown in Figure 75 and Figure 76.

Note: The RBT pins must provide the ability to be isolated on the baseboard side when AUX_PWR_EN is not asserted. This prevents a leakage path through unpowered silicon. The RBT REF_CLK must also be disabled until AUX_PWR_EN is asserted. Example buffering implementations are shown in Figure 75 and Figure 76. The isolator shall be controlled on the baseboard with a signal called RBT_ISOLATE#.

Table 22: Pin Descriptions – NC-SI Over RBT

Signal Name	Pin #	Baseboard Direction	Signal Description
RBT_REF_CLK	OCP_A14	Output	Reference clock input. Synchronous clock reference for receive, transmit and control interface. The clock shall have a typical frequency of 50MHz. For baseboards, this pin shall be connected between the baseboard NC-SI over RBT PHY and the Primary Connector OCP bay. This signal requires a 100kOhm pull down resistor on the baseboard. If the baseboard does not support NC-SI over RBT, then this signal shall be terminated to ground through a 100kOhm pull down resistor. The RBT_REF_CLK shall

			<p>not be driven until the card has transitioned into AUX Power Mode.</p> <p>For OCP NIC 3.0 cards, this pin shall be connected between the gold finger to 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 100kOhm pull down resistor on the baseboard. If the baseboard does not support NC-SI over RBT, then this signal shall be terminated to ground through a 100kOhm pull down resistor.</p> <p>For OCP NIC 3.0 cards, this pin shall be connected between the gold finger to 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 100kOhm pull down resistor to GND on the baseboard. If the baseboard does not support NC-SI over RBT, then this signal shall be terminated to GND through a 100kOhm pull down.</p> <p>For OCP NIC 3.0 cards, this pin shall be connected between the gold finger 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 100kOhm pull down resistor to ground on the baseboard. If the baseboard does not support NC-SI over RBT, then this signal shall be terminated to ground through a 100kOhm pull down.</p>

			For OCP NIC 3.0 cards, this pin shall be connected between the gold finger to the endpoint silicon. This pin shall be left as a no connect if NC-SI over RBT is not supported.
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 100kOhm pull down resistor to GND on the baseboard. If the baseboard does not support NC-SI over RBT, then this signal shall be terminated to GND through a 100kOhm pull down.</p> <p>For OCP NIC 3.0 cards, this pin shall be connected between the gold finger to 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.</p>
RBT_ARB_OUT	OCP_A5	Output	<p>NC-SI hardware arbitration output. This pin shall only be used if the endpoint silicon supports hardware arbitration. This pin shall be connected to the RBT_ARB_IN signal of an adjacent device in the hardware arbitration ring.</p> <p>The baseboard shall implement a multiplexing implementation that directs the RBT_ARB_OUT to the RBT_ARB_IN pin of the next NC-SI over RBT capable device in the ring, or back to the RBT_ARB_IN pin of the source device if there is a single device on the ring.</p> <p>For baseboards, this pin shall be connected between the baseboard OCP connector(s) to complete the hardware arbitration ring. If the baseboard does not support NC-SI over RBT, this signal shall be directly connected to the RBT_ARB_IN pin to allow a complete hardware arbitration ring on the OCP NIC 3.0 card.</p> <p>For OCP NIC 3.0 cards, this pin shall be connected from the gold finger to the RBT_ARB_IN pin on the endpoint silicon. This pin shall be directly connected to the card edge RBT_ARB_IN pin if NC-SI is not supported. 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. This pin shall only be used if the endpoint silicon supports hardware arbitration. This pin shall be connected to the RBT_ARB_OUT signal of an adjacent device in the hardware arbitration ring.</p> <p>The baseboard shall implement a multiplexing implementation that directs the RBT_ARB_IN to the RBT_ARB_OUT pin of the next NC-SI over RBT capable device in the ring, or back to the RBT_ARB_OUT pin of the source device if there is a single device on the ring.</p> <p>For baseboards, this pin shall be connected between the baseboard OCP connector(s) to complete the hardware arbitration ring. If the baseboard does not support NC-SI over RBT, this signal shall be directly connected to the RBT_ARB_OUT pin to allow a complete hardware arbitration ring on the OCP NIC 3.0 card.</p> <p>For OCP NIC 3.0 cards, this pin shall be connected between the gold finger to the RBT_ARB_OUT pin on the endpoint silicon. This pin shall be directly connected to the card edge RBT_ARB_OUT pin if NC-SI is not supported. This allows the hardware arbitration signals to pass through in a multi-Primary Connector baseboard.</p>															
SLOT_ID0 SLOT_ID1	OCP_B7 OCP_A6	Output	<p>NC-SI / FRU EEPROM Address 0/1.</p> <p>The SLOT_ID[1:0] pins shall be used to set the RBT Package ID field. This pin is also used in setting the FRU EEPROM address.</p> <p>For baseboards, the SLOT_ID[1:0] pins shall be physically tied to GND or to +3.3V_EDGE. The SLOT[1:0] values are based on the following mapping:</p> <table border="1" data-bbox="592 1451 1068 1619"> <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>	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, SLOT_ID0 shall be connected to the endpoint device GPIO associated with the Package ID[0] field. SLOT_ID1 shall be associated with the Package ID[1] field. Refer to Section 4.8.1 and the device datasheet for details.</p> <p>For OCP NIC 3.0 cards with multiple endpoint devices, the Package ID[2] field may 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 with a FET switch (or a similar implementation) to prevent a leakage path when the OCP NIC 3.0 card is in ID mode.</p> <p>For FRU EEPROM addressing, the SLOT_ID0 pin shall be directly connected to the EEPROM A1 address field; SLOT_ID1 shall be connected to the EEPROM A2 address field.</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 75: NC-SI Over RBT Connection Example – Single Primary Connector

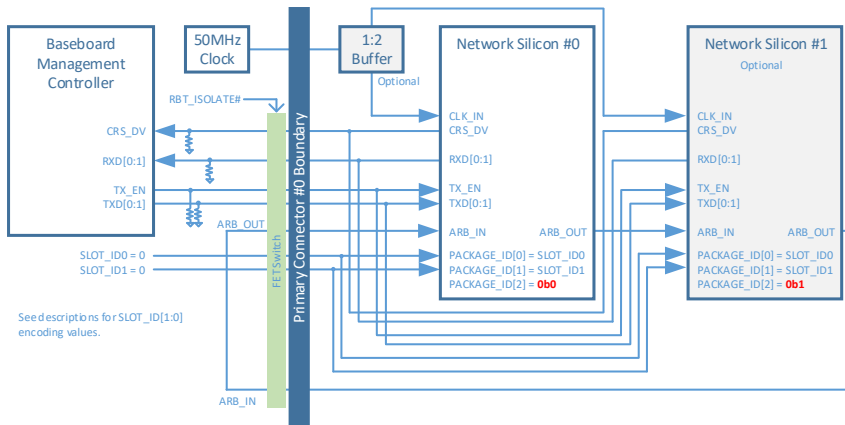
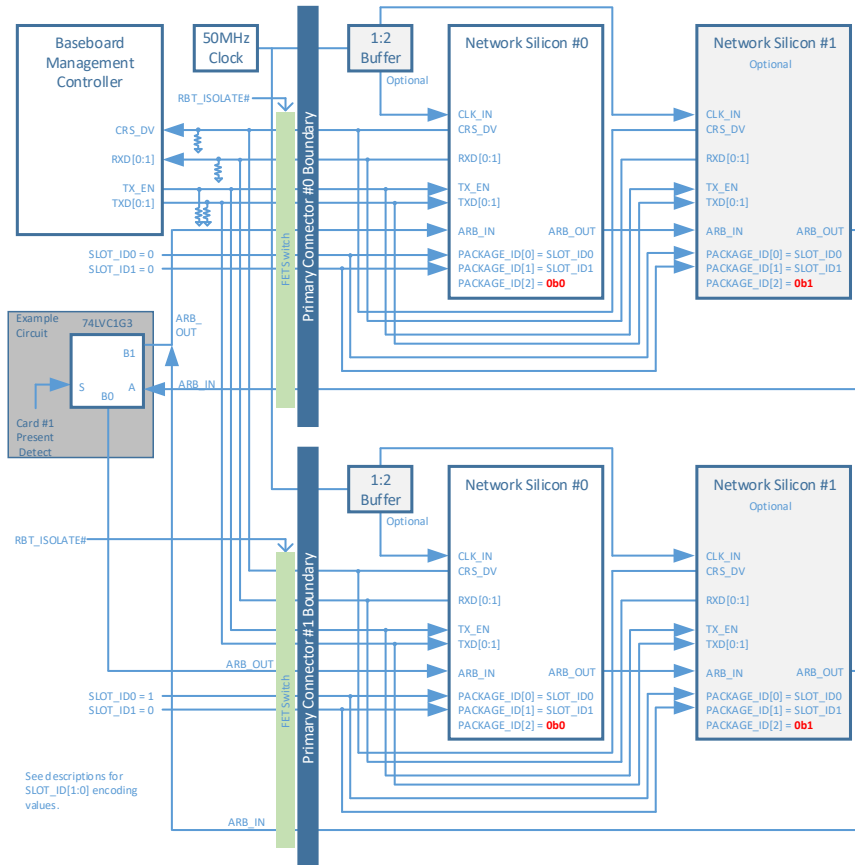


Figure 76: NC-SI Over RBT Connection Example – Dual Primary Connectors



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 76.

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 logical implementation of the hardware arbitration ring shall maintain the arbitration ring integrity when there exists one or more cards that are plugged in, but are powered off (e.g 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 77. An example connection diagram is shown in Figure 78.


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

Table 23: 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 78, below. The CLK pin shall be pulled up to +3.3V_EDGE through a 1kOhm resistor.</p>
DATA_OUT	OCP_B5	Output	<p>Scan clock data output from the baseboard to the OCP NIC 3.0 card. This bit stream is used to shift in NIC configuration data.</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 1kOhm resistor if the scan chain is not used.</p>

			For NIC implementations, the DATA_OUT pin shall be pulled up to +3.3V_EDGE on the OCP NIC 3.0 card through a 1kOhm resistor.
DATA_IN	OCP_B4	Input	<p>Scan clock data input to the baseboard. This bit stream is used to shift out NIC status bits.</p> <p>For baseboard implementations, the DATA_IN pin shall be pulled up to +3.3V_EDGE through a 10kOhm 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 Registers 0 & 1, as defined in the text and Figure 78.</p>
LD#	OCP_B3	Output	<p>Scan clock 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 1kOhm 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 as defined in the text and Figure 78. The LD# pin shall be pulled up to +3.3V_EDGE through a 1kOhm resistor.</p>

Figure 77: 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 24 and Table 25 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 24: Pin Descriptions – Scan Chain DATA_OUT Bit Definition

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

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 1kOhm 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.

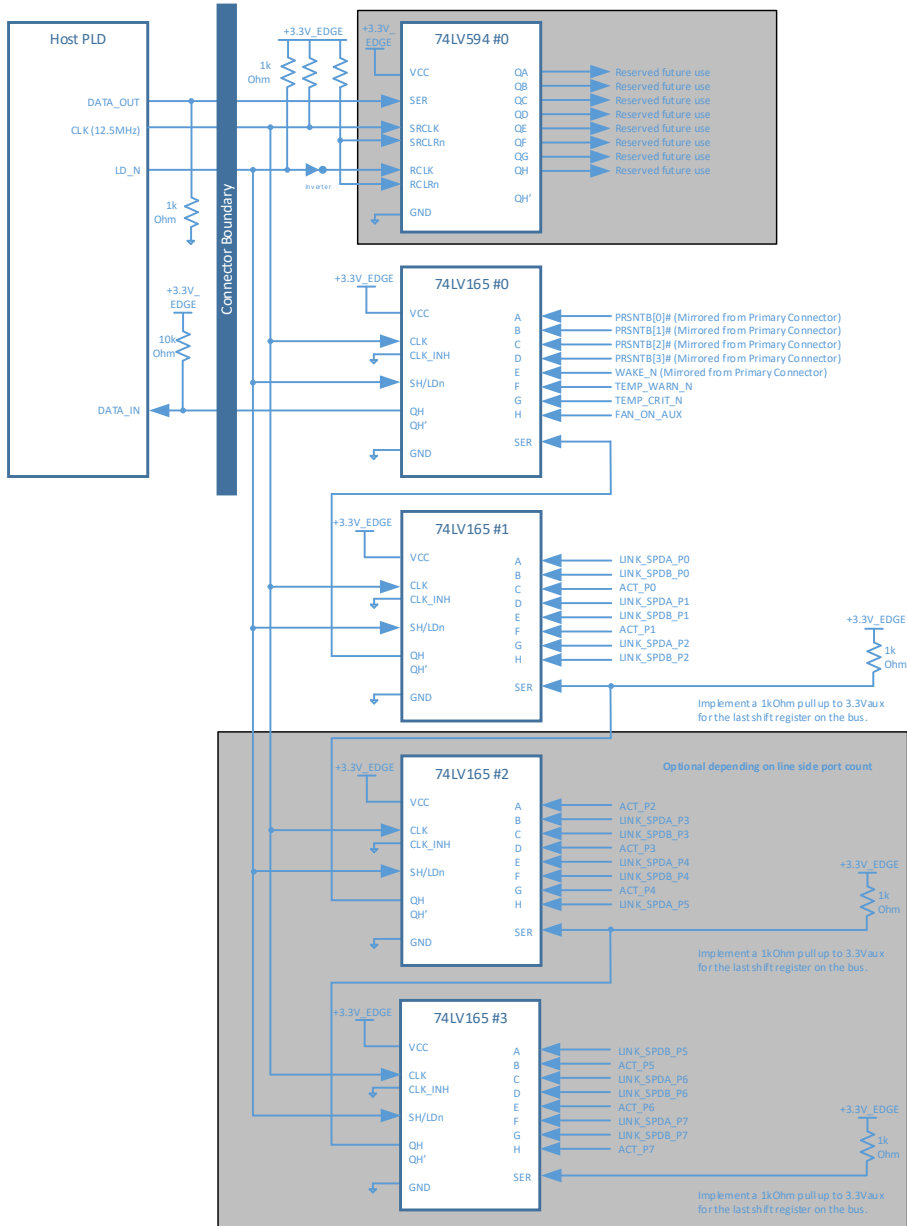
Table 25: Pin Descriptions – Scan Chain DATA_IN Bit Definition

Byte.bit	DATA_IN Field Name	Default Value	Description
0.0	PRSNTB[0]#	0bX	PRSNTB[3:0]# bits shall reflect the same state as the signals on the Primary Connector.
0.1	PRSNTB[1]#	0bX	
0.2	PRSNTB[2]#	0bX	
0.3	PRSNTB[3]#	0bX	
0.4	WAKE_N	0bX	PCIe WAKE_N signal shall reflect the same state as the signal on the Primary Connector.
0.5	TEMP_WARN_N	0b1	Temperature monitoring pin from the on-card thermal solution. This pin shall be asserted low when temperature sensor exceeds the temperature warning threshold.
0.6	TEMP_CRIT_N	0b1	Temperature monitoring pin from the on-card thermal solution. This pin shall be asserted low

			when temperature sensor exceeds the temperature critical threshold.
0.7	FAN_ON_AUX	0b0	When high, FAN_ON_AUX shall request the system fan to be enabled for extra cooling in the S5 state. 0b0 – The system fan is not requested/off in S5. 0b1 – The system fan is requested/on in S5.
1.0	LINK_SPDA_P0	0b1	Port 0 link and speed A indication (max speed). Active low. 0b0 – Link LED is illuminated on the host platform. 0b1 – Link LED is not illuminated on the host platform. Steady = link is detected on the port and is at the maximum speed. Off = the physical link is down, not at the maximum speed or is disabled. Note: The link and speed A LED may also be blinked for use as port identification.
1.1	LINK_SPDB_P0	0b1	Port 0 link and speed B indication (not max speed). Active low. 0b0 – Link LED is illuminated on the host platform. 0b1 – Link LED is not illuminated on the host platform. Steady = link is detected on the port and is not at the max speed. Off = the physical link is down, or is disabled. Note: The link and speed B LED may also be blinked for use as port identification.
1.2	ACT_P0	0b1	Port 0 activity indication. Active low. 0b0 – ACT LED is illuminated on the host platform. 0b1 – ACT LED is not illuminated on the host platform. Steady = no activity is detected on the port. Blinking = activity is detected on the port. The blink rate should blink low for 50-500ms during activity periods. Off = the physical link is down or disabled.
1.3	LINK_SPDA_P1	0b1	Port 1 link and speed A indication. Active low.

1.4	LINK_SPDB_P1	0b1	Port 1 link and speed B indication. 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. Active low.
1.7	LINK_SPDB_P2	0b1	Port 2 link and speed B indication. 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. Active low.
2.2	LINK_SPDB_P3	0b1	Port 3 link and speed B indication. 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. Active low.
2.5	LINK_SPDB_P4	0b1	Port 4 link and speed B indication. 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. Active low.
3.0	LINK_SPDB_P5	0b1	Port 5 link and speed B indication. 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. Active low.
3.3	LINK_SPDB_P6	0b1	Port 6 link and speed B indication. 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. Active low.
3.6	LINK_SPDB_P7	0b1	Port 7 link and speed B indication. Active low.
3.7	ACT_P7	0b1	Port 7 activity indication. Active low.

Figure 78: Scan Chain Connection Example



3.4.6—USB 2.0 Pins


This section provides the pin assignments for the USB 2.0 interface signals. The USB 2.0 pins may be used for Smart NIC applications with end point silicon that requires a USB connection to the baseboard. Implementations may allow for a USB JTAG application to program and update the Smart NIC FPGA. An example connection diagram is shown in Figure 80.

Commented [NT1]: Re-evaluate pin usage. Further discussion is required. Remove from the Ov80 release. Return to RFU pins.

Table 26: Pin Descriptions—USB 2.0

Signal Name	Pin #	Baseboard Direction	Signal Description
USB_DATp USB_DATn	A68 A69	Bi-directional	<p>USB 2.0 Differential Pair</p> <p>All baseboard implementations shall provide a USB connection to the OCP NIC 3.0 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 buffered on the NIC to prevent a leakage path into unpowered silicon. The USB_ISOLATE# signal may be deasserted when the card enters the Aux Power (S5) state.</p>

Figure 80: USB 2.0 Connection Example



3.4.7—UART Pins

This section provides the pin assignments for the UART interface signals. The UART pins may be used for Smart NIC applications with end point silicon that requires a serial console to the baseboard. An example connection diagram is shown in Figure 81.


Commented [NT2]: Re-evaluate pin usage. Further discussion is required. Remove from the Ov80 release. Return to RFU pins.

Table 27: Pin Descriptions—Miscellaneous UART

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

UART_TX	B69	Output	<p>UART Transmit. +3.3V_{AUX} signaling levels.</p> <p>All baseboard implementations 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.</p>
UART_RX	B68	Input	<p>UART Receive. +3.3V_{AUX} signaling levels.</p> <p>All baseboard implementations shall provide a UART receive connection from the OCP NIC 3.0 connector. The UART_RX pin shall be pulled up to +3.3V_{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_TX pin to the UART_RX 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 buffered on the NIC to prevent a leakage path into unpowered silicon. The UART_ISOLATE# signal may be deasserted when the card enters the Aux Power (S5) state.</p>

Figure 80: UART Connection Example



3.4.83.4.6 Power Supply Pins

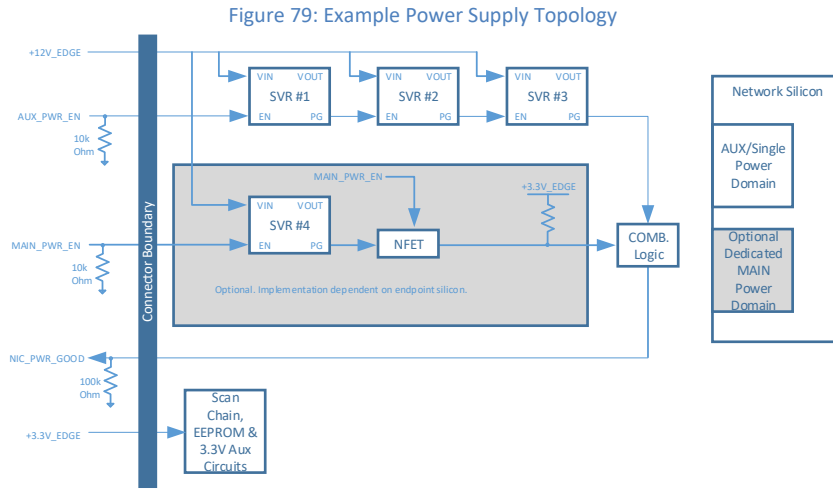
This section provides the pin assignments for the power supply interface signals. The AC/DC specifications are defined in the PCIe CEM Specification, Rev 4.0 and amended in Section 3.10. An example connection diagram is shown in Figure 79.

Table 26: 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 5 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>+12V main or +12V aux power; total of 6 pins per connector. The +12V_EDGE pins shall be rated to 1.1A per pin with a maximum derated power delivery of 80W.</p> <p>The +12V_EDGE power pins shall be within the rail tolerances as defined in Section 3.10 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>+3.3V main or +3.3V aux power; total of 1 pin per connector. The +3.3V_EDGE pin shall be rated to 1.1A for a maximum derated power delivery of 3.63W.</p> <p>The +3.3V_EDGE power pin shall be within the rail tolerances as defined in Section 3.10 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 +12V_EDGE and +3.3V_EDGE power is from the baseboard aux power rails.</p> <p>This signal shall be pulled down to GND through a 10kOhm 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>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 +12_EDGE and +3.3V_EDGE power is from the baseboard main power rails. 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. This pin must be implemented on baseboard systems, but may optionally be used 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 baseboard implementations, this signal shall be pulled down to GND through a 10kOhm 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="574 533 1057 726"> <thead> <tr> <th>AUX_PWR_EN</th> <th>MAIN_PWR_EN</th> <th>NIC_PWR_GOOD Nominal Steady State Value</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>0</td> <td>1</td> <td>Invalid</td> </tr> <tr> <td>1</td> <td>1</td> <td>1</td> </tr> </tbody> </table> <p>Refer to the power up and power down sequencing diagrams (Figure 89 and Figure 90) 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 79 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 tolerances or are in a fault condition after the power ramp times (T_{APL} and T_{MPL}) have expired.</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 100kOhm 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>	AUX_PWR_EN	MAIN_PWR_EN	NIC_PWR_GOOD Nominal Steady State Value	0	0	0	1	0	1	0	1	Invalid	1	1	1
AUX_PWR_EN	MAIN_PWR_EN	NIC_PWR_GOOD Nominal Steady State Value															
0	0	0															
1	0	1															
0	1	Invalid															
1	1	1															



3.4.93.4.7 Miscellaneous Pins

This section provides the pin assignments for the miscellaneous interface signals.

Table 27: Pin Descriptions – Miscellaneous

Signal Name	Pin #	Baseboard Direction	Signal Description
RFU1, N/C	B68	Input /	Reserved future use pins. These pins shall be left as no connect. These pins may also be used as a differential pair.
RFU2, N/C	B69	Output	
RFU3, N/C	A68	Input /	Reserved future use pins. These pins shall be left as no connect. These pins may also be used as a differential pair.
RFU4, N/C	A69	Output	

Commented [NT3]: The RFU pin usage is up for debate. These were defined as serial and USB in prior draft versions, but additional discussion is required on the usage. And if implemented, they are optional pins.

Need to define mechanism to control interface buffering (on the NIC) and FRU EEPROM entry to identify interface usage.

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.
- 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.

A high level bifurcation connection diagram is shown in Figure 72.

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. 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 system. 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 PRSNTB[3:0]# bits is shown in Table 28 for x16 and x8 PCIe cards.

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.

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 28 and indicate to the 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 28 shows the resulting number of PCIe links and its width for known combinations of baseboards and OCP NIC 3.0 cards.

***Note:** 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.

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. An OCP NIC 3.0 card is present in the system if the resulting value is not 0b1111.
2. Firmware determines the OCP NIC 3.0 card PCIe lane width capabilities per Table 28 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.
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 >1s after NIC_PWR_GOOD assertion as defined in Figure 89. Refer to Section 3.12 for timing details.

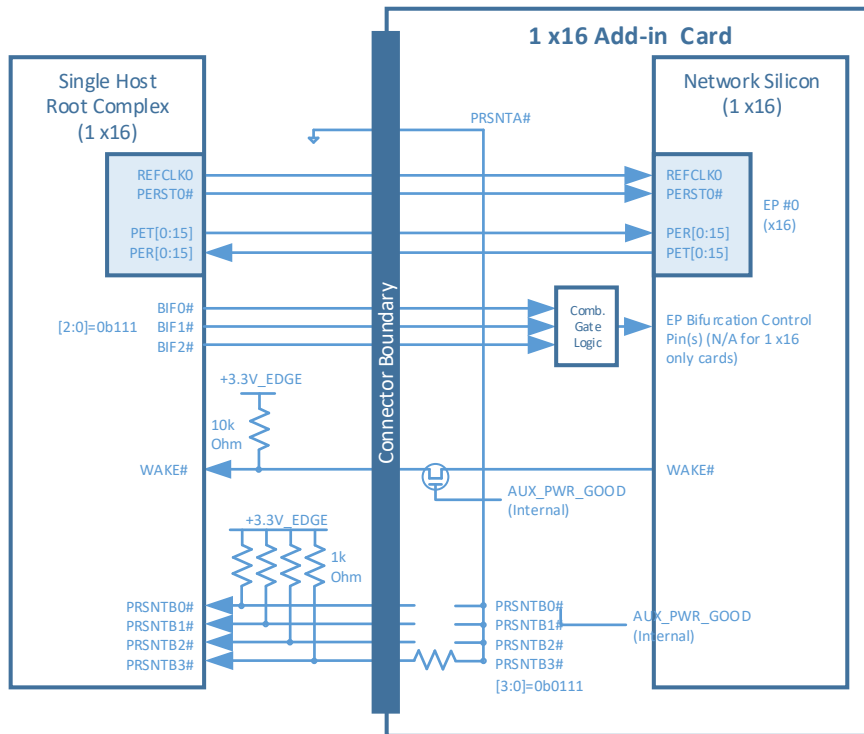
3.5.5 PCIe Bifurcation Examples

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

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

Figure 80 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 as there is no need to instruct the end-point network controller to a specific bifurcation. 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 80: 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 81 illustrates a single host baseboard that supports 2 x8 with a single controller OCP NIC 3.0 card that also supports 2 x8 with dual controllers. The PRSTNB[3:0]# state is 0b0110. The BIF[2:0]# state is 0b111 as there is no need to instruct the end-point network controllers to a specific bifurcation. 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.

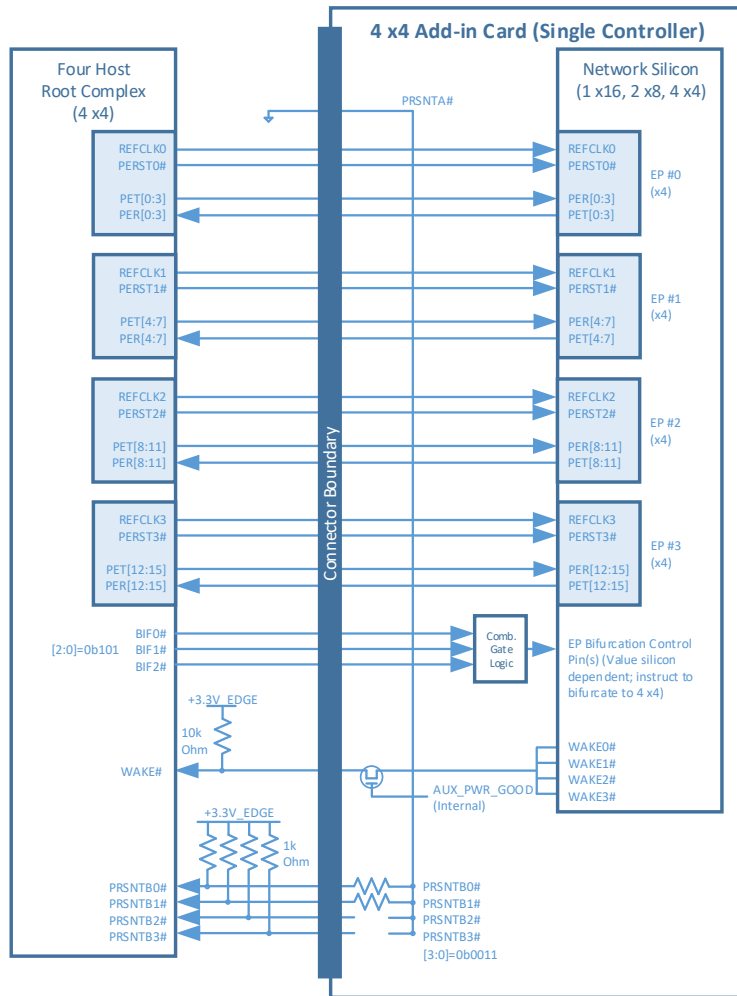
Figure 81: Single Host (2 x8) and 2 x8 OCP NIC 3.0 Card (Dual Controllers)



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

Figure 82 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 0b0011. The BIF[2:0]# state is 0b101 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.

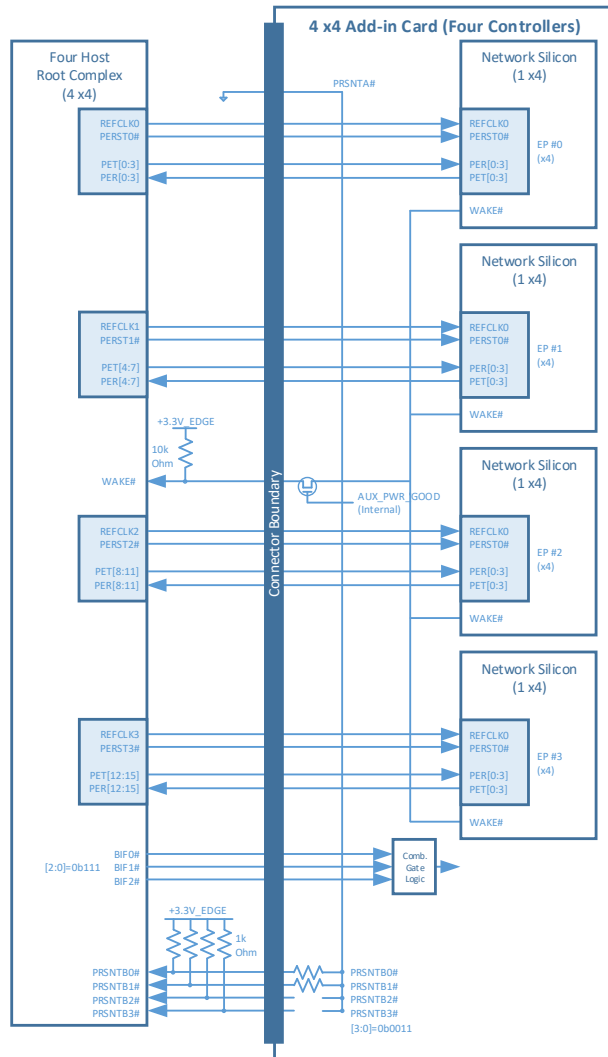
Figure 82: Quad Hosts (4 x4) and 4 x4 OCP NIC 3.0 Card (Single Controller)



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

Figure 83 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 0b111 as there is no need to instruct the end-point network controllers to a specific bifurcation. 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.

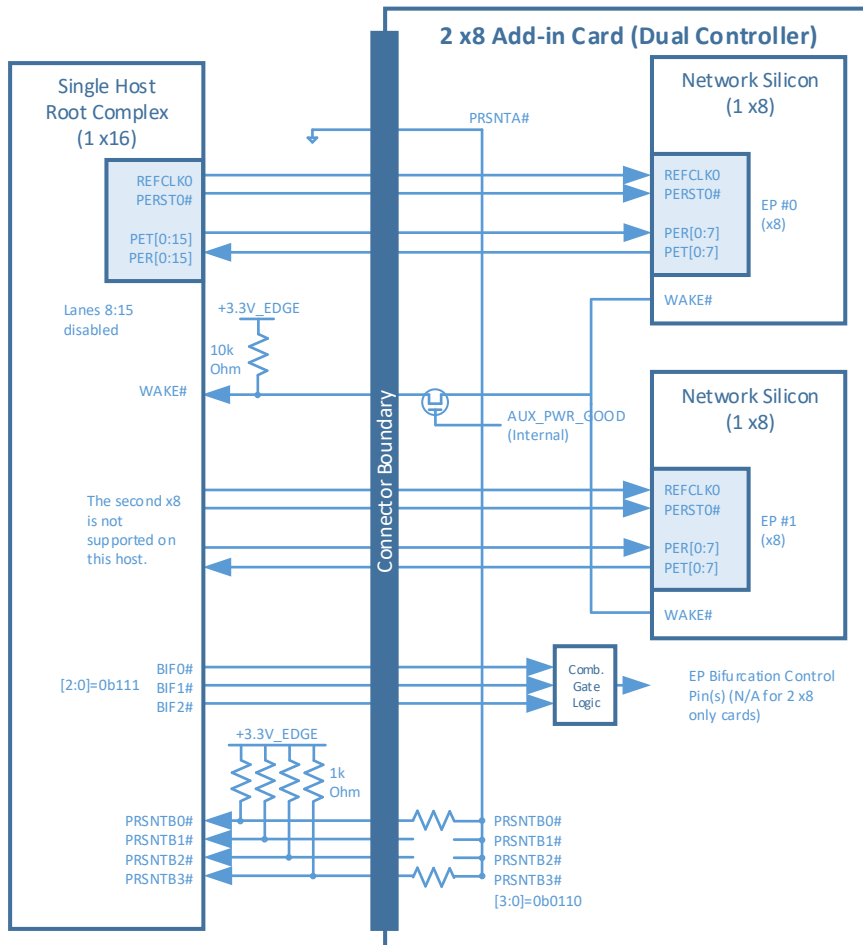
Figure 83: Quad Hosts (4 x4) and 4 x4 OCP NIC 3.0 Card (Quad Controllers)



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

Figure 84 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 0b111 as there is no need to instruct the end-point network controllers to a specific bifurcation. 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 84: Single Host with no Bifurcation (1 x16) and 2 x8 OCP NIC 3.0 Card (Dual Controllers)



3.6 PCIe Clock Topology

The OCP NIC 3.0 specification allows for up to four PCIe REFCLKs on the Primary Connector and up to two PCIe REFCLKs on the Secondary Connector. In general, the association of each REFCLK is based on the PCIe Link number on a per connector basis and is shown in Table 29. Cards that implement both the Primary and Secondary Connectors have a total of up to 6 REFCLKs.

Table 29: PCIe Clock Associations

REFCLK #	Description	Availability (Connector)
REFCLK0	REFCLK associated with Link 0.	Primary and Secondary Connectors.
REFCLK1	REFCLK associated with Link 1.	Primary and Secondary Connectors.
REFCLK2	REFCLK associated with Link 2.	Primary Connector only.
REFCLK3	REFCLK associated with Link 3.	Primary Connector only.

For each OCP NIC 3.0 card, the following REFCLK connection rules must be followed:

- For a 1 x16 capable OCP NIC 3.0 card, REFCLK0 shall be used for lanes [0:15].
- For a 2 x8 capable OCP NIC 3.0 card, REFCLK0 shall be used for lanes [0:7] and REFCLK1 shall be used for lanes [8:15].
- For a 4 x4 capable OCP NIC 3.0 card, REFCLK0 shall be used for lanes [0:3], REFCLK1 shall be used for lanes [4:7], REFCLK2 shall be used for lanes [8:11] and REFCLK3 shall be used for lanes [12:15]. Pins for REFCLK2 and REFCLK3 are located on the 28-pin OCP bay.

Figure 85: PCIe Interface Connections for 1 x16 and 2 x8 OCP NIC 3.0 Cards





Figure 86: PCIe Interface Connections for a 4 x4 OCP NIC 3.0 Card



3.7 PCIe Bifurcation Results and REFCLK Mapping

For the cases where the baseboard and OCP NIC 3.0 card bifurcation are permissible, this section 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.

Table 33: Bifurcation for Single Host, Dual Sockets and Dual Upstream Links (BIF[2:0]#=0b001)

Single Host, Two Upstream Sockets, Two Upstream Links		1x8, 1x4, 1x2, 1x1 2x8, 2x4, 2x2, 2x1		Host	Upstream Devices	Upstream Links	BIF [2:0]#	Resulting Link	Key: Cells shown as Link Lane (e.g. Lk0/Ln0); HD = Host Disabled Lane														
Card Width	Card Short Name	Supported Modes	Bifurcation Encodings						Lk0	Lk1	Lk2	Lk3	Lk4	Lk5	Lk6	Lk7	Lk8	Lk9	Lk10	Lk11	Lk12	Lk13	Lk14
1x8	Option A	Card Not Present	0b1111	1 Host	2 Upstream Sockets	2 Links	0b001	1x8 (Socket 0 only)	Lk0, Lk1, Lk2, Lk3, Lk4, Lk5, Lk6, Lk7														
2x2	Option A	1x4, 1x2, 1x1	0b1110	1 Host	2 Upstream Sockets	2 Links	0b001	1x4 (Socket 0 only)	Lk0, Lk1, Lk2, Lk3														
2x2	Option A	1x2, 1x1	0b1110	1 Host	2 Upstream Sockets	2 Links	0b001	1x2 (Socket 0 only)	Lk0, Lk1														
2x2	Option A	1x1	0b1110	1 Host	2 Upstream Sockets	2 Links	0b001	1x1 (Socket 0 only)	Lk0														
2x2	Option B	1x8, 1x4, 1x2, 1x1	0b1101	1 Host	2 Upstream Sockets	2 Links	0b001	1x8 (Socket 0 only)	Lk0, Lk1, Lk2, Lk3, Lk4, Lk5, Lk6, Lk7														
2x2	Option B	2x4, 2x2, 2x1	0b1101	1 Host	2 Upstream Sockets	2 Links	0b001	2x4 (Socket 0 only)	Lk0, Lk1, Lk2, Lk3														
4x2	Option B	4x4, 4x2, 4x1	0b1101	1 Host	2 Upstream Sockets	2 Links	0b001	2x4 (Socket 0 only)	Lk0, Lk1, Lk2, Lk3, Lk4, Lk5, Lk6, Lk7														
2x2	Option D	1x8, 1x4, 2x4	0b1100	1 Host	2 Upstream Sockets	2 Links	0b001	2x4 (Socket 0 only)	Lk0, Lk1, Lk2, Lk3, Lk4, Lk5, Lk6, Lk7														
4x2	Option D	2x8, 2x4, 4x2 (First 8 lanes), 4x1	0b1100	1 Host	2 Upstream Sockets	2 Links	0b001	2x8 (Socket 0 only)	Lk0, Lk1, Lk2, Lk3, Lk4, Lk5, Lk6, Lk7														
2x2	Option D	2x4, 2x2, 2x1	0b1011	1 Host	2 Upstream Sockets	2 Links	0b001	1x4 (Socket 0 only)	Lk0, Lk1, Lk2, Lk3														
2x2	Option D	1x4, 1x2, 1x1	0b1011	1 Host	2 Upstream Sockets	2 Links	0b001	1x2 (Socket 0 only)	Lk0, Lk1														
2x2	Option D	2x2, 2x1	0b1011	1 Host	2 Upstream Sockets	2 Links	0b001	1x2 (Socket 0 only)	Lk0, Lk1														
4x2	Option A	1x8, 1x4, 1x2, 1x1	0b0111	1 Host	2 Upstream Sockets	2 Links	0b001	1x8 (Socket 0 only)	Lk0, Lk1, Lk2, Lk3, Lk4, Lk5, Lk6, Lk7														
4x2	Option A	2x8, 2x4, 2x2, 2x1	0b0110	1 Host	2 Upstream Sockets	2 Links	0b001	2x8 (Socket 0 only)	Lk0, Lk1, Lk2, Lk3, Lk4, Lk5, Lk6, Lk7														
4x2	Option B	1x8, 1x4, 1x2, 1x1	0b0101	1 Host	2 Upstream Sockets	2 Links	0b001	2x4 (Socket 0 only)	Lk0, Lk1, Lk2, Lk3, Lk4, Lk5, Lk6, Lk7														
4x2	Option B	2x8, 2x4, 2x2, 2x1	0b0101	1 Host	2 Upstream Sockets	2 Links	0b001	2x4 (Socket 0 only)	Lk0, Lk1, Lk2, Lk3, Lk4, Lk5, Lk6, Lk7														
4x2	Option C	4x4, 4x2, 4x1	0b0100	1 Host	2 Upstream Sockets	2 Links	0b001	2x4 (Socket 0 only)	Lk0, Lk1, Lk2, Lk3, Lk4, Lk5, Lk6, Lk7														
4x2	Option D	4x4, 4x2, 4x1	0b0011	1 Host	2 Upstream Sockets	2 Links	0b001	2x4 (BP 0 and 2 only)	Lk2, Lk3, Lk4, Lk5, Lk6, Lk7														
RSVD	RSVD	RSVD	0b0010	1 Host	2 Upstream Sockets	2 Links	0b001	--															
RSVD	RSVD	RSVD	0b0001	1 Host	2 Upstream Sockets	2 Links	0b001	--															
RSVD	RSVD	RSVD	0b0000	1 Host	2 Upstream Sockets	2 Links	0b001	--															

Table 34: Bifurcation for Single Host, Quad Sockets and Quad Upstream Links (BIF[2:0]#=0b010)

Single Host, Four Upstream Sockets, Four Upstream Links										Key: Cells shown as Link/Lane (e.g. Lk0/Ln0); HD = Host Disabled Lane															
Card Width	Card Name	Card Short Name	Supported Modes	Bifurcation	Host	Upstream Devices	Upstream Links	BIF [2:0]#	Resulting Link	Lk0	Lk1	Lk2	Lk3	Lk4	Lk5	Lk6	Lk7	Lk8	Lk9	Lk10	Lk11	Lk12	Lk13	Lk14	Lk15
1x8	Option A	1x8	1x8, 1x4, 1x2, 1x1	0b1111	1Host	4 Upstream Sockets	4 Links	0b010	1x4 (Socket0 only)	Lk0, Lk1, Lk2, Lk3															
2x4	Option B	1x4	1x4, 1x2, 1x1	0b1110	1Host	4 Upstream Sockets	4 Links	0b010	1x2 (Socket0 only)	Lk0, Lk1, Lk2, Lk3															
4x4	Option C	1x2	1x2, 1x1	0b1110	1Host	4 Upstream Sockets	4 Links	0b010	1x1 (Socket0 only)	Lk0, Lk1															
1x8	Option A	1x1	1x1	0b1110	1Host	4 Upstream Sockets	4 Links	0b010	1x1 (Socket0 only)	Lk0, Lk1															
2x4	Option B	1x8	1x8, 1x4, 1x2, 1x1	0b1101	1Host	4 Upstream Sockets	4 Links	0b010	2x4	Lk0, Lk1, Lk2, Lk3															
4x4	Option C	2x4	2x4, 2x2, 2x1, 4x4, 4x2, 4x1	0b1101	1Host	4 Upstream Sockets	4 Links	0b010	4x4	Lk0, Lk1, Lk2, Lk3															
1x8	Option A	2x4	2x4, 2x2, 2x1	0b1100	1Host	4 Upstream Sockets	4 Links	0b010	2x4	Lk0, Lk1, Lk2, Lk3															
2x4	Option B	4x4	4x4, 4x2, 4x1	0b1100	1Host	4 Upstream Sockets	4 Links	0b010	4x4	Lk0, Lk1, Lk2, Lk3															
4x4	Option C	4x2	4x2, 4x1	0b1100	1Host	4 Upstream Sockets	4 Links	0b010	4x2	Lk0, Lk1, Lk2, Lk3															
1x8	Option A	RSVD	RSVD	0b1010	1Host	4 Upstream Sockets	4 Links	0b010	RSVD																
2x4	Option B	RSVD	RSVD	0b1010	1Host	4 Upstream Sockets	4 Links	0b010	RSVD																
4x4	Option C	RSVD	RSVD	0b1010	1Host	4 Upstream Sockets	4 Links	0b010	RSVD																
1x8	Option A	RSVD	RSVD	0b0111	1Host	4 Upstream Sockets	4 Links	0b010	RSVD																
2x4	Option B	RSVD	RSVD	0b0111	1Host	4 Upstream Sockets	4 Links	0b010	RSVD																
4x4	Option C	RSVD	RSVD	0b0111	1Host	4 Upstream Sockets	4 Links	0b010	RSVD																
1x8	Option A	RSVD	RSVD	0b0110	1Host	4 Upstream Sockets	4 Links	0b010	RSVD																
2x4	Option B	RSVD	RSVD	0b0110	1Host	4 Upstream Sockets	4 Links	0b010	RSVD																
4x4	Option C	RSVD	RSVD	0b0110	1Host	4 Upstream Sockets	4 Links	0b010	RSVD																

3.8 Port Numbering and LED Implementations

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

Additionally, LEDs shall be implemented on the OCP NIC 3.0 I/O bracket when there is sufficient space for local indication. LEDs may also be implemented on the card Scan Chain (as defined in Section 3.4.5) for remote link/activity indication on the baseboard. The LED configuration is described for both cases in the sections below. In both cases, the actual link rate may be directly queried through the management interface.

3.8.1 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 87 as an example implementation.

3.8.2 OCP NIC 3.0 Card LED Configuration

For low I/O count small form-factor cards without built in light pipes (such as 1x QSFP, 2x SFP, or 2x RJ-45), or a large form-factor cards, where additional I/O bracket area is available, the card shall implement on-board link/activity indications in place of the Scan Chain LED stream.

For 4x SFP and 2x QSFP designs, 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.

Note: Depending on the end faceplate implementation (e.g. with an ejector latch), the secondary side LED implementation may be obstructed.

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 39 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 its surrounding chassis components.

Commented [TN4]: Mechanical proposals are currently in progress for secondary side SMT LEDs. Need placement recommendations/text if implemented.

Table 39: OCP NIC 3.0 Card LED Configuration with Two Physical LEDs per Port

LED Pin	LED Color	Description
Link	Green	Active low. Bicolor multifunction LED.
	Amber	<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>The LED is Green when the port is linked at its maximum speed. The LED is Amber when the port is not linked at the highest speed. The LED is off when no link is present.</p> <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 513nm and 537nm while amber LEDs shall emit light at a wavelength between 580nm and 589nm.</p> <p>For uniformity across OCP NIC 3.0 products, all link LEDs shall have its luminance across the total surface area measured in nits (cd/m²) to an average value between 150 nits to 450 nits.</p>
	Off	
Activity	Green	Active low.
	Off	<p>When the link is up and there is no activity, this LED shall be lit and solid.</p> <p>When the link is up and there is link activity, then this LED should blink at the interval of 50-500ms during link activity.</p> <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>

		For serviceability, green LEDs shall emit light at a wavelength between 513nm and 537nm. For uniformity across OCP NIC 3.0 products, all activity LEDs shall have its luminance across the total surface area measured in nits (cd/m ²) to an average value between 150 nits to 450 nits.
--	--	--

3.8.3 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 shall be clearly labeled on the OCP NIC 3.0 card. Similarly, the LED for link and the LED for Activity indication shall also be marked on the faceplate.

For 4xSFP and 2xQSFP 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 87: Port and LED Ordering – Example Small Card Link/Activity and Speed LED Placement



Note: The example port and LED ordering diagrams shown in Figure 87 are viewed with the card in the horizontal position and the primary side is facing up.

3.8.4 Baseboard LEDs Configuration over the Scan Chain

A small form-factor OCP NIC 3.0 with a fully populated I/O bracket (2x QSFP, 4x SFP or 4x RJ-45) does not have sufficient space for primary-side discrete on-board (faceplate) LED indicators. Section 3.8.2 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 513nm and 537nm while amber LEDs shall emit light at a wavelength between 580nm and 589nm.

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).

3.9 Power Capacity and Power Delivery

There are four permissible power states: NIC Power Off, ID Mode, Aux Power Mode (S5), and Main Power Mode (S0). The transition of these states is shown in Figure 88. The max available power envelopes for each of these states are defined in Table 40.

Figure 88: Baseboard Power States

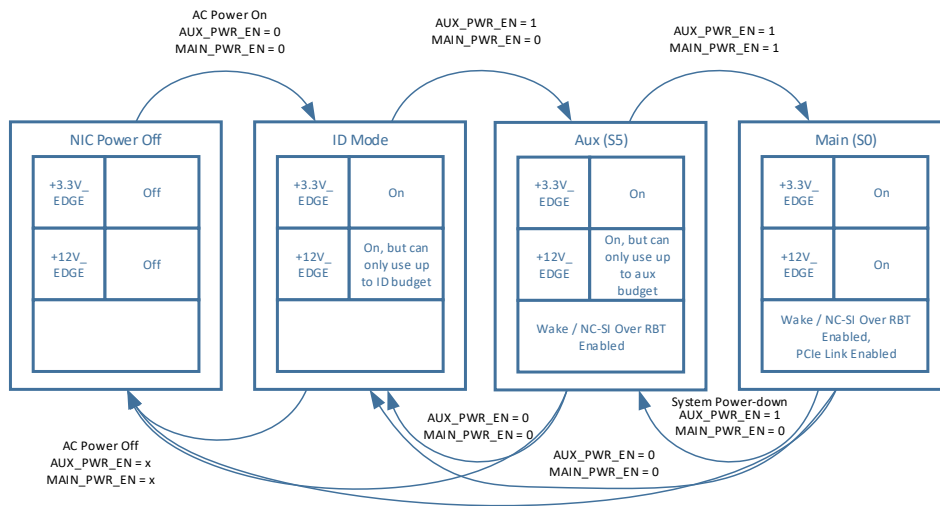


Table 40: Power States

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
Aux Power Mode (S5)	High	Low	Low	X	X	X	X		X	X
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.

3.9.1 NIC Power Off

In NIC power off mode, all power delivery has been turned off or disconnected from the baseboard. Transition to this state can be from any other state.

3.9.2 ID Mode

In the ID Mode, only +3.3V_EDGE is available for powering up management only functions. Only FRU and scan chain accesses are allowed in this mode. Only the card PRSNTB[0:3]# bits are valid on the 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 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 3.10. An OCP NIC 3.0 card shall transition to this mode when AUX_PWR_EN=0 and MAIN_PWR_EN=0.

3.9.3 Aux Power Mode (S5)

In 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 41. An OCP NIC 3.0 card shall transition to this mode when AUX_PWR_EN=1 and MAIN_PWR_EN=0.

3.9.4 Main Power Mode (S0)

In 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 80W may be delivered on +12V_EDGE for a Small Card and up to 150W for a Large Card. Additionally, up to 3.63W is delivered on each +3.3V_EDGE pin. An OCP NIC 3.0 card shall transition to this mode when AUX_PWR_EN=1 and MAIN_PWR_EN=1.

3.10 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 cards, the requirements are as follows:

Table 41: Baseboard Power Supply Rail Requirements – Slot Power Envelopes

Power Rail	15W Slot Small Card Hot Aisle	25W Slot Small Card Hot Aisle	35W Slot Small Card Hot Aisle	80W Slot Small Card Cold Aisle	150W Large Card Cold Aisle
+3.3V_EDGE					

Voltage Tolerance	±9% (max)	±9% (max)	±9% (max)	±9% (max)	±9% (max)
Supply Current					
ID Mode	100mA (max)	100mA (max)	100mA (max)	100mA (max)	100mA (max)
Aux Mode	1.1A (max)	1.1A (max)	1.1A (max)	1.1A (max)	2.2A (max)
Main Mode	1.1A (max)	1.1A (max)	1.1A (max)	1.1A (max)	2.2A (max)
Capacitive Load	150µF (max)	150µF (max)	150µF (max)	150µF (max)	300µF (max)
+12V_EDGE					
Voltage Tolerance	+8%/-12% (max)	+8/-12% (max)	+8/12% (max)	+8/-12% (max)	+8/-12% (max)
Supply Current					
ID Mode	50mA (max)	50mA (max)	50mA (max)	50mA (max)	50mA (max)
Aux Mode	0.7A (max)	1.1A (max)	1.5A (max)	3.3A (max)	6.3A (max)
Main Mode	1.25A (max)	2.1A (max)	2.9A (max)	6.6A (max)	12.5A (max)
Capacitive Load	500µF (max)	500µF (max)	500µF (max)	500µF (max)	1000µF (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 0.1A/µs per the PCIe CEM specification.

Note 3: Each OCP NIC 3.0 card shall limit the bulk capacitance to the max values published (500µF for a Small Form-Factor card, 1000µF for a Large Form-Factor 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 4.10.2 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.

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

For baseboards that support system hot (powered on) OCP NIC 3.0 card insertions and extractions, the system implementer shall consider the use of hotswap controllers on both the +12V_EDGE and +3.3V_EDGE pins to prevent damage to the baseboard or the OCP NIC 3.0 card. Hotswap controllers help with in-rush current limiting while also providing overcurrent protection, undervoltage and overvoltage protection capabilities.

Commented [TN5]: Hot swap/Hot plug considerations are still being discussed within the OCP NIC 3.0 working group. Please follow the OCP Wiki site for updates.

The hotswap 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 hotswap 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.

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

3.12 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, 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 3.4.6.1.1 for the NIC_PWR_GOOD definition. Refer to DMTF DSP0222 for details on the NC-SI clock startup requirements.

Figure 89: Power-Up Sequencing

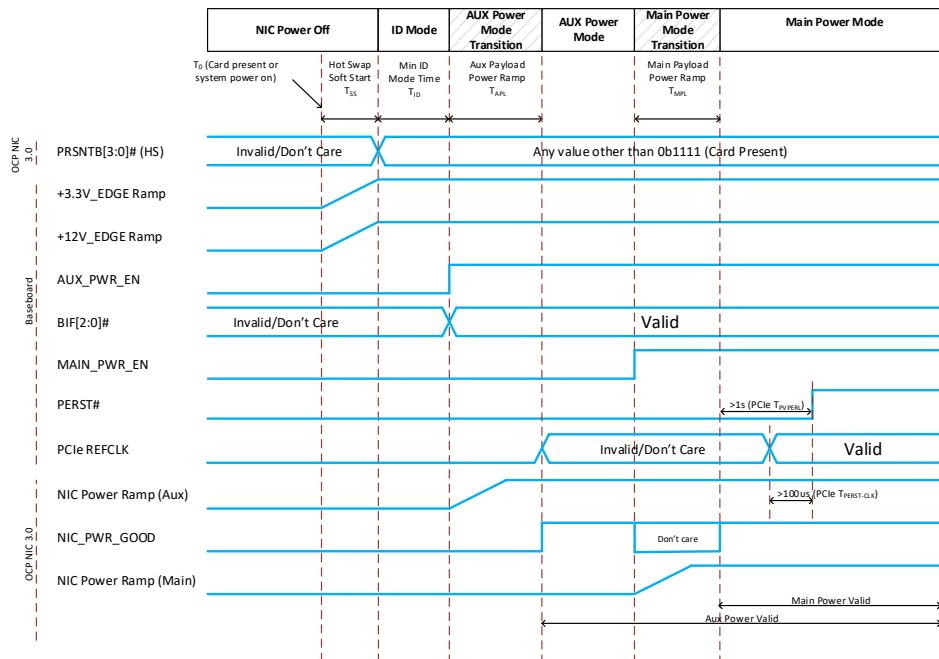
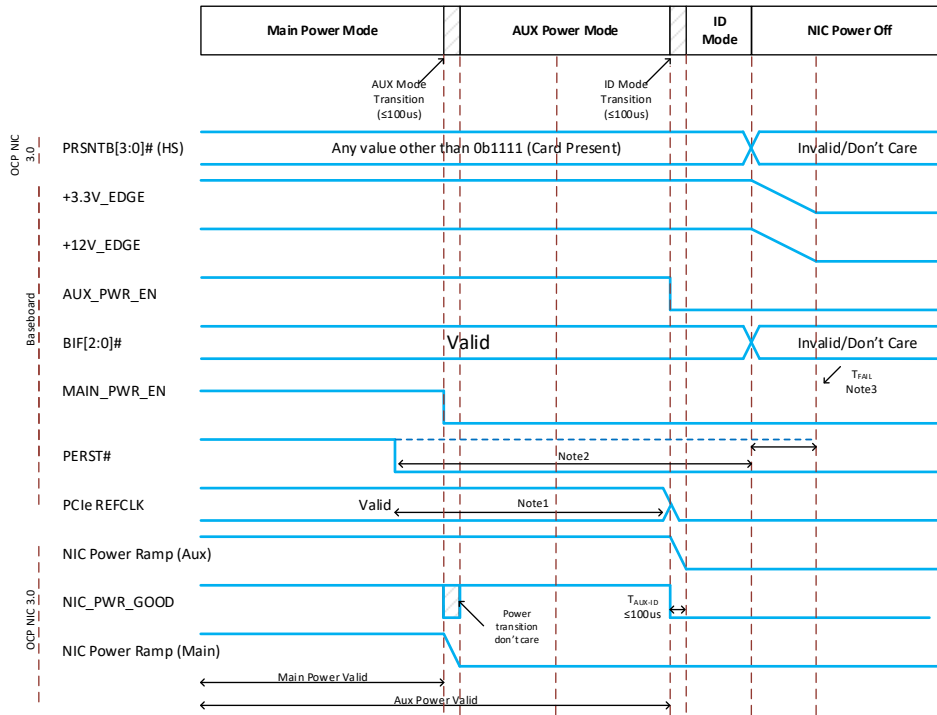


Figure 90: Power-Down Sequencing



- Note1: REFCLK go inactive after PERST# goes active. (PCIe CEM Section 2.2.3)
 Note2: PERST# goes active before the power on the connector is removed. (PCIe CEM Section 2.2.3)
 Note3: In the case of a surprise power down, PERST# goes active T_{FALL} after power is no longer stable.

Table 42: 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.
T_{ID}	20	ms	Minimum guaranteed time per spec to spend in ID mode.
T_{APL}	25	ms	Maximum time between AUX_PWR_EN assertion to NIC_PWR_GOOD assertion.
T_{MPL}	25	ms	Maximum time between MAIN_PWR_EN assertion to 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.
$T_{PERST-CLK}$	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.

3.13 Digital I/O Specifications

All digital I/O pins on the connector boundary are +3.3V signaling levels. Table 43 following tables provide the absolute max levels. Refer to the appropriate specifications for the RBT, PCIe and SMBus DC/AC specifications.

Table 43: 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 44: 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 45: 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 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 DSP0236 Management Component Transport Protocol (MCTP) Base Specification and the associated binding specifications. The PMCI Platform Layer Data Model (PLDM) will be the primary payload (or “MCTP Message”) to convey information from the OCP 3.0 NIC to the management controller. The NC-SI over MCTP Message Type may also be used monitoring and pass-through communication.

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 NIC. Table 46 summarizes the sideband management interface and transport requirements.

Table 46: Sideband Management Interface and Transport Requirements

Requirement	RBT+MCTP Type	RBT Type	MCTP Type
NC-SI 1.1 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 1.3 (DSP0236 1.3 compliant) over MCTP/SMBus Binding (DSP0237 1.1 compliant)	Required	N/A	Required
PCIe VDM compliant physical interface	Optional	Optional	Optional

Management Component Transport Protocol (MCTP) Base 1.3 (DSP0236 1.3 compliant) over MCTP/PCIe VDM Binding (DSP0238 1.0 compliant)	Optional	Optional	Optional
--	----------	----------	----------

4.2 NC-SI Traffic

DMTF DSP0222 defines two types of NC-SI traffic: Pass-Through and Control. Table 47 summarizes the NC-SI traffic requirements.

Table 47: NC-SI Traffic Requirements

Requirement	RBT+MCTP Type	RBT Type	MCTP Type
NC-SI Control over RBT (DMTF DSP0222 1.1 or later compliant)	Required	Required	N/A
NC-SI Control over MCTP (DMTF DSP0261 1.2 compliant)	Required	N/A	Required
NC-SI Pass-Through over RBT (DMTF DSP0222 1.1 compliant)	Required	Required	N/A
NC-SI Pass-Through over MCTP (DMTF DSP0261 1.2 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 48 summarizes the MC MAC address provisioning requirements.

Table 48: 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.</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 • <code>num_hosts</code> – Number of hosts supported by the NIC • <code>first_addr</code> – The MAC address of the first port on the first host for the first partition on that 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>Formulae:</p> <ul style="list-style-type: none"> • $\text{host_addr}[i] = \text{first_addr} + i * \text{num_ports} * (\text{max_parts} + 1)$ • 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. • $\text{bmc_addr}[i] = \text{host_addr}[i] + \text{num_ports} * \text{max_parts}$ • 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>Support at least one of the following mechanism for provisioned MC MAC Address retrieval:</p> <ul style="list-style-type: none"> • NC-SI Control/RBT (DMTF DSP0222 1.1 or later compliant) • NC-SI Control/MCTP (DMTF DSP0261 1.2 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</p>	Required	Required	Optional

implementation, that is compliant with DMTF DSP0222 that defines standard NC-SI commands for MC MAC address retrieval, shall support those NC-SI commands.			
--	--	--	--

4.4 Temperature Reporting

An OCP NIC 3.0 implementation can have several silicon components including one or more ASICs implementing NIC functions and one or more transceiver modules providing physical network media connectivity. For the system management, it is important that temperatures of these components can be retrieved over sideband interfaces.

The temperature reporting interface shall be accessible in Aux Power Mode (S5), and Main Power Mode (S0). Table 49 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 49: Temperature Reporting Requirements

Requirement	RBT+MCTP Type	RBT Type	MCTP Type
Component Temperature Reporting for a component with TDP $\geq 8\text{W}$	Required	Required	Required
Component Temperature Reporting for a component with TDP $< 8\text{W}$	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 1.1 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: For definitions of the warning, critical, and fatal thresholds, refer to DSP0248 1.1.	Required	Required	Required
When the temperature reporting function is implemented using PLDM numeric sensors, the temperature tolerance shall be reported.	Required	Required	Required
Support for NIC self-shutdown.	Optional	Optional	Optional
The purpose of this feature is to “self-protect” the NIC from permanent damage due to high operating temperature experienced by the NIC. The NIC can accomplish this by reducing the power consumed by the device. The NIC shall monitor its temperature and shut-down itself as soon as the threshold value is reached. The value of the self-shutdown threshold is implementation specific. It is			

Commented [HS6]: Add a table for warning, critical, and fatal temps in terms of the maximum operating temperature.

For example:
Upper warning = Omax;
Upper critical = 1.1 Omax;
Upper fatal > 1.1 Omax.

The OCP Mezz sub-group could not agree on relationship between upper warning, upper critical, and upper fatal and the maximum operating temperature.

The setting of upper warning, upper critical, and upper fatal thresholds are implementation dependent and should be compliant with the severity levels defined in DMTF DSP0248 1.1.

Commented [MC7]: 8W seems high. Are there known components that need this? A target of 3W would be more acceptable. Is it planned that an ambient sensor near the component is required?

Commented [NT8R7]: Carl Massey will send e-mail out to the distribution list.

Commented [TN9]: Intel proposes removal of the NIC self-shutdown requirement or changing it.

As written, the NIC will asynchronously shutdown without host intervention. This may cause the system to freeze/blue screen as the PCIe endpoint is removed unexpectedly.

I suggest removing this requirement in favor of having the BMC implementation read sensors and disable functions if we cross the upper temperature thresholds.

Commented [TN10R9]: 20180425 – open.
Working group notes:
-FB more biased to remove requirement (original proposal).
-Call participants okay with this as optional.
-Is this requirement supposed to be the card gates power to itself (e.g. ASIC is no longer powered); or is this requirement supposed to be the ASIC goes into a low power state?
-Could be the ASIC GPIO connected to PWRBRK# → ASIC PWRBRK# input.

Commented [TN11R9]: Hemel@Broadcom: to follow up with proposed text update.

<p>recommended that the self-shutdown threshold value is higher than the maximum junction temperature of the ASIC implementing the NIC function and this value is between the critical and fatal temperature thresholds. The self-shutdown feature is a final effort in preventing permanent card damage at the expense of potential data loss.</p> <p>If this feature is implemented, care shall be taken to ensure that the board power down state is latched and that the board does not autonomously resume normal operation.</p> <p>Note: It is assumed that a system management function will prevent a component from reaching its fatal threshold temperature.</p> <p>The OCP NIC 3.0 card does not need to know the reason for the self-shutdown threshold crossing (e.g. fan failure). After entering the self-shutdown state, the OCP NIC 3.0 card is not required to 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 from the self-shutdown state, the OCP NIC 3.0 card shall go through the NIC ID Mode state as described in Section 3.9.1.</p>			
---	--	--	--

Commented [MC12]: The self-shutdown should be 5C or more higher than than Tjmax.

Commented [MC13]: Please replace "this value" with the actual name to eliminate interpretation issues.

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 50 summarizes power consumption reporting requirements.

Table 50: 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. This field reports the board max power consumption value without transceivers plugged into the line side receptacles.	Required	Required	Required
Pluggable Transceiver Module Power Reporting. The pluggable transceivers plugged into the line side receptacles shall be inventoried (via an EEPROM query) and the total module power consumption is 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 1.1 compliant) for component 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 51 summarizes pluggable module status reporting requirements.

Table 51: 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 1.1 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 52 summarizes the firmware inventory and update requirements.

Table 52: 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 1.0 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 (draft) 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 53. 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 53: 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

Commented [HS14]: Current firmware inventory definition is vague. Need to define what it means in each environment including UEFI, OOB via PLDM, and NC-SI ctrl. Need to define what is the minimum set for firmware inventory.

There is no change in text needed. Firmware inventory information is implementation dependent.

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 may be connected to each other. The arbitration ring shall support operation with one card, or multiple cards installed. Figure 76 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 connected on this bus. Additionally, 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 SMBus Address Resolution Protocol (ARP) to allow each device to be dynamically assigned an addresses for MCTP communication. This method automatically resolves address conflicts and eliminate the need for manual configuration of addresses. The address type of dynamic addresses can be either a dynamic and persistent address device or a dynamic and volatile address device. Refer to SMBus 2.0 specification and Section 6.11 of DSP0237 1.1 for details on SMBus address assignment.

A system implementation may choose to only use fixed addresses for an OCP NIC 3.0 card on the system. The assignment of these fixed addresses is system dependent and outside the scope of this specification. When fixed addresses are assigned to OCP NIC 3.0 card, then the OCP NIC 3.0 card shall be a fixed and discoverable SMBus device. Refer to SMBus 2.0 specification for more details.

All predefined SMBus addresses for OCP NIC 3.0 are shown in Table 54. 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 54: 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 Address, Size and Availability

The FRU EEPROM provided for the baseboard to determine the card type and 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 EEPROM is addressable at the addresses indicated in Table 54. The write/read pair is presented in 8-bit format. The size of EEPROM shall be at least 4Kbits for the base EEPROM map. OCP NIC 3.0 card suppliers may use a larger size EEPROM if needed to store vendor specific information. The FRU EEPROM shall use double byte addressing. The FRU EEPROM shall be write protected for production cards by pulling the EEPROM WP pin high to +3.3V_EDGE. The FRU shall be writable for manufacturing test and during card development by pulling the EEPROM WP pin low to ground.

The FRU EEPROM is readable in all three power states (ID mode, AUX(S5) mode, and MAIN(S0) mode).

4.10.2 FRU EEPROM Content Requirements

The FRU EEPROM shall follow the data format specified in the IPMI Platform Management FRU Information Storage Definition v1.0 Document Revision 1.3. Both the Product Info and Board Info records shall be populated in the FRU EEPROM. Where applicable, fields common to the Product Info and Board 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. For an OCP NIC 3.0 card, the FRU EEPROM OEM record content based on the format defined in [Table 55](#) shall be populated.

Note: [Table 55](#) 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 55: FRU EEPROM Record – OEM Record 0xC0, Offset 0x00

Offset	Length	Description
0	3	Manufacturer ID. For OCP NIC 3.0 compliant cards, the value of this field shall be set to the OCP IANA assigned number. This value is 0x7FA600, LS byte first. (42623 in decimal)
3	1	OCP NIC 3.0 FRU OEM Record Version. For OCP NIC 3.0 cards compliant to this specification, the value of this field shall be set to 0x01.
4	1	Card Max power (in Watts) in MAIN (S0) mode. 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>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. 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). 0x00 – 0xFE – Card power rounded up to the nearest Watt for fractional values. 0xFF – Unknown</p>
6	1	<p>Hot Aisle Card Cooling Tier. The encoded value reports the OCP NIC 3.0 Card Hot Card Cooling Tier as defined in Section 6.6.1. 0x00 – RSVD 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 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. The encoded value reports the OCP NIC 3.0 Card Cold Aisle Cooling Tier as defined in Section 6.6.2. 0x00 – RSVD 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. 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. Refer to Section 6 for more information about the thermal and environmental requirements.</p> <p>Byte 9 is the LS byte, byte 10 is the MS byte.</p> <p>0x0000-0xFFFE – 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. Refer to Section 6 for more information about the thermal and environmental requirements.</p> <p>Byte 11 is the LS byte, byte 12 is the MS byte.</p> <p>0x0000-0xFFFE – LFM required for cooling card in Cold Aisle Operation. 0xFFFF – Unknown.</p>
13	1	<p>UART Configuration 1.</p> <p>This byte denotes the UART configuration 1. A value 0x00 means no serial connection is required.</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.</p>

		<p>This byte denotes the UART configuration 2. A value 0x00 means no serial connection is required.</p> <p>Bits [1:0] denotes the number of stop bits. 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. 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.</p> <p>This byte denotes a USB 2.0 connection is present on the card edge. 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>
16:30	15	<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 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. OCP NIC 3.0 cards may implement up to six physical controllers (N=6) for a Large Form Factor card.</p>
32:47	16	<p>Controller 1 UDID (if applicable).</p> <p>This field reports the Controller 1 Universal Device Identifier (UDID) and is used to aid in the dynamic slave address assignment over the SMBus Address Resolution Protocol.</p> <p>This field shall list the MS 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>
48:63	16	Controller 2 UDID (if applicable).
64:79	16	Controller 3 UDID (if applicable).
80:95	16	Controller 4 UDID (if applicable).
96:111	16	Controller 5 UDID (if applicable).
112:127	16	Controller 6 UDID (if applicable).

Commented [NT15]: Need to define a byte for optional interface (USB, UART) location on the scan chain DATA_OUT chain to enable buffers

4.10.3 FRU Template

A FRU template is provided as a baseline implementation example. This FRU template contains the IPMI Platform Management FRU Information Storage Definition v1.2 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

For the purposes of this specification, the min and max electrical trace length of the NC-SI signals shall be between 2 inches and 4 inches. The traces shall be implemented as 50 Ohm impedance controlled nets. This requirement applies to both the small and large form factor OCP NIC 3.0 cards.

5.1.1 Channel Budget Requirements

5.1.1.1 Budget impact requirements using isolation buffers

5.1.1.2 Add-in Card Channel Budget

5.1.1.3 Baseboard Channel Budget

Total capacitive load
Etc.

5.1.1.4 SFF-TA-1002 Connector Channel Budget

5.1.1.5 Timing Budget

5.1.1.6 Impedance

5.2 SMBus 2.0

This section is a placeholder for SMBus 2.0 related routing guidelines and SI considerations. The OCP NIC 3.0 subgroup intends to define the bus operational speed range, capacitive loading, 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

This section is a placeholder for the PCIe routing guidelines and SI considerations.

OCP NIC 3.0 card suppliers shall follow the PCIe routing specifications. At this time, the OCP NIC 3.0 subgroup is working to identify and agree to the channel budget for an OCP NIC 3.0 card and leave sufficient margin for the baseboard. Refer to the PCIe CEM and PCIe Base specifications for end-to-end channel signal integrity considerations.

5.3.1 Background

5.3.2 Channel Requirements

5.3.2.1 PCIe Gen3 Channel Budget and Crosstalk Requirements

Reference channel budgets for PCIe Gen3.

5.3.2.2 PCIe Gen4 Channel Budget and Crosstalk Requirements

Reference channel budgets for PCIe Gen4 – See Section 4.7 of the PCIe CEM 4.0 spec.

Commented [TN16]: The OCP NIC 3.0 SI Workgroup is currently contributing to this section. The contents of this section are a work in progress and is expected to be complete for version 0v90.

Commented [TN17]: - Refer to the SMBus specification for details / speed / voltage range.

-Max capacitance and location of pull ups.

- Refer to SMBus specification as appropriate. Differences/implementation specific items for OCP NIC 3.0 are called out here.

Commented [TN18]: Refer to CEM for gen3/4/5.

5.3.2.3 PCIe Gen5 Channel Budget and Crosstalk Requirements

The OCP NIC 3.0 specification uses SFF-TA-1002 compliant 4C and 4C+ connectors. The SFF-TA-1002 working group expects these connectors to work with PCIe Gen5 rates. This section shall be used as a placeholder for Gen5 cards.

5.3.2.4 REFCLK requirements

For the four REFCLKs – each REFCLK shall be treated per the PCIe CEM.

Commented [TN19]: The lane under test shall be coupled to the REFCLK associated with that lane. (e.g. 2x8 --> use the appropriate REFCLK for each x8.)

5.3.2.5 Add-in Card Channel Budget

This section defines the OCP NIC 3.0 card channel budget from the gold finger edge to the end point silicon.

Commented [TN20]: Align per CEM.

5.3.2.6 Baseboard Channel Budget

This section defines the baseboard channel budget from the root complex silicon to the pads of the OCP 4C and 4C+ connector. This definition does not include the channel budget of the SFF-TA-1002 connector (which is defined in the following section).

Commented [TN21]: Align per CEM.

5.3.2.7 SFF-TA-1002 Connector Channel Budget

Reference the SFF-TA-1002 spec.

5.3.2.8 Insertion Loss – Normative

5.3.2.9 Return Loss – Normative

5.3.2.10 Differential Skew – Normative

For PCIe transmit and receive differential pairs, the target differential skew is 5mils for the OCP NIC 3.0 card and 10 mil for the baseboard. This is the same requirement values set forth in the PCIe CEM specification to minimize the common-mode signal leading to a reduction in potential EMI impact on the system.

For the PCIe REFCLKs, the target differential skew is 10mils.

5.3.2.11 Lane-to-Lane skew

Reference PCIe CEM 4.0 section 4.7.5

5.3.2.12 Differential Impedance

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%.

5.3.3 Test Fixtures

5.3.3.1 Load Board

5.3.3.2 Baseboard

5.3.4 Test Methodology

- 5.3.4.1 DUT Control and Test Automation Recommendations
- 5.3.4.2 Transmitter Testing
- 5.3.4.3 Receiver Testing
- 5.3.4.4 PLL Test

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 91. 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 91: Airflow Direction for Hot Aisle Cooling (SFF and LFF)



The boundary conditions for Hot Aisle cooling are shown below in Table 56 and Table 57. 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 57 represent the airflow velocities typical in mainstream

servers. Higher airflow velocities are available within the Hot Aisle cooling tiers listed in Table 62 but card designers must be sure to understand the system level implications of such high card LFM requirements.

Table 56: 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 57: 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 RJ-45) side of the card. This airflow direction is referred to as Cold Aisle cooling and is illustrated below in Figure 92. 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 92: Airflow Direction for Cold Aisle Cooling (SFF and LFF)



The boundary conditions for Cold Aisle cooling are shown below in Table 58 and Table 59. The temperature values listed in Table 58 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 (i.e. A4).

Table 58: 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 59: 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 93 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 93 represents a different local inlet air temperature from 45°C to 65°C.

Figure 93: ASIC Supportable Power for Hot Aisle Cooling – Small Card Form Factor



The curves shown in Figure 93 were obtained using CFD analysis of a reference OCP NIC 3.0 SFF card. The reference card has a 20mm x 20mm ASIC with two QSFP connectors. Figure 94 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 60. The OCP NIC 3.0 simulation was conducted within a virtual version of the test fixture defined in Section 6.4.

Figure 94: OCP NIC 3.0 SFF Reference Design and CFD Geometry




Table 60: Reference OCP NIC 3.0 SFF Card Geometry

OCP NIC 3.0 Form Factor	SFF Card
Heatsink Width	65mm
Heatsink Length	45mm
Heatsink Height	9.24mm
Heatsink Base Thickness	1.5mm
Fin Count/Thickness	28/0.5mm
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 @ 3.5W 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 93.

It is important to point out that the curves shown in Figure 93 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.

Card designers must also consider the airflow capability of the server systems that the cards are targeted for use within. Figure 95 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 959597: Server System Airflow Capability – SFF Card Hot Aisle Cooling



6.2.2 LFF Card ASIC Cooling – Hot Aisle

Figure 96 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 96 represents a different local inlet air temperature from 45°C to 65°C.

Figure 96: ASIC Supportable Power for Hot Aisle Cooling – LFF Card



The curves shown in Figure 96 were obtained using CFD analysis of the reference OCP NIC 3.0 LFF card. The reference card has a 45mm x 45mm ASIC with two QSFP connectors. Additional card geometry parameters and boundary conditions used in the reference CFD analysis are summarized in Table 61. Figure 97 shows a comparison of the 3D CAD and CFD model geometry for the reference OCP NIC 3.0 card.

Figure 97: OCP NIC 3.0 LFF Reference Design and CFD Geometry





Table 61: Reference OCP NIC 3.0 LFF Card Geometry

OCP NIC 3.0 Form Factor	LFF Card
Heatsink Width	75mm
Heatsink Length	85mm
Heatsink Height	9.3mm
Heatsink Base Thickness	1.5mm
Fin Count/Thickness	33/0.5mm
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 @ 3.5W 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.

Figure 98 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 98: 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 94 and Table 60 with airflow moving from I/O connector to ASIC (opposite to the Hot Aisle analysis). Figure 99 below shows the results of this analysis for the Cold Aisle cooling configuration. Each curve in Figure 99 represents a different system inlet air temperature from 25°C to 45°C.

Figure 99: ASIC Supportable Power for Cold Aisle Cooling – SFF Card




Similar to Figure 95 for Hot Aisle cooling, Figure 100 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. 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 100: 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 101. The comparison shows the Hot Aisle ASIC cooling capability at 12W at 150LFM while the cold Aisle cooling capability shows support for 19W at 150LFM. 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 101: 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 97 and Table 61 with airflow moving from I/O connector to ASIC (opposite to the Hot Aisle analysis). Figure 102 below shows the results of this analysis for the Cold Aisle cooling configuration. Each curve in Figure 102 represents a different system inlet air temperature from 25°C to 45°C.

Figure 102: ASIC Supportable Power for Cold Aisle Cooling – LFF Card




Similar to Figure 100 for LFF Hot Aisle cooling, Figure 103 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. 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 103: 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 104. The comparison shows the Hot Aisle ASIC cooling capability at 19W at 150LFM while the cold aisle cooling capability shows support for 42W at 150LFM. 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 104: 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. One prototype is 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
 - Power connections for 3.3V, GND, GND, 12V (SFF)
 - Power connections for 3.3V, GND, GND, GND, 12V, 12V (LFF)
- RJ45 connector for NC-SI pass-through
- USB Type-X connector for microprocessor connectivity
- 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 board.
 - Quantity 4x for LFF – see Figure 110
 - Candlestick style sensors available at: <https://www.gats.com/Products/Instruments/Temperature-and-Velocity-Measurement/Sensors/Candlestick-Sensor>

- Candlestick sensors must be procured separately, not integrated with fixture PCB
- Blockage above OCP3 card to mimic system geometry and prevent airflow bypass
 - Low profile PCIe card for SFF fixture
- Block sheet metal obstruction built into the top cover for the LFF fixture

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 105 and Figure 106. The SFF fixture PCB is shown in Figure 107. Note the three candlestick sensor locations directly next to the OCP NIC 3.0 connectors.

Figure 105: Small Card Thermal Test Fixture Preliminary Design




Figure 106: SFF Card Thermal Test Fixture Preliminary Design – Cover Removed





Figure 107: 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 108 and Figure 109. The LFF fixture PCB is shown in Figure 110. Note the three candlestick sensor locations directly next to the OCP NIC 3.0 connectors.

Figure 108: LFF Card Thermal Test Fixture Design




Figure 109: LFF Card Thermal Test Fixture Design – Cover Removed






Figure 110: 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 111.

Figure 111: Thermal Test Fixture Airflow Direction



6.4.4 Thermal Test Fixture Candlestick Sensors

As noted in 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 112 and Figure 113 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 velocity obtained from the sensor locations vs. the velocity based on the duct cross-sectional area.

Figure 112: SFF Fixture, Hot Aisle Flow - Candlestick Air Velocity vs. Volume Flow





Figure 113: 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 4.10.2 defines a number of registers that may be read by the associated baseboard system. Two of these registers provide the Hot Aisle and Cold Aisle Card Cooling Tiers that may be used for fan speed control. The Card Cooling Tiers relate the card local inlet temperature to the required local inlet velocity which allows the system to set fan speeds according to the cooling requirements of the card.

The Card Cooling Tier registers are particularly useful for systems that do not implement temperature sensor monitoring. The register may also be used as a backup for cards that do implement temperature sensor monitoring.

6.6.1 Hot Aisle Cooling Tiers

Card Cooling Tiers for Hot Aisle Cooling are defined in Table 62. 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 62: Hot Aisle 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												
10												
15												
20												
25												
30												
35												
40												
45												
50												
55	50	100	150	200	250	300	350	400	450	500	750	1000
60												
65												

Work in Progress

6.6.2 Cold Aisle Cooling Tiers

Card Cooling Tiers for Cold Aisle Cooling are defined in Table 63. 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 63: Cold Aisle 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												
10												
15												
20												
25												
30												
35	50	100	150	200	250	300	350	400	450	500	750	1000
40												
45												
50												
55												
60												
65												

Work in Progress

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 fixtured in the standard test fixture as described in Section 6.7.1.

6.7.1 Shock & Vibe Test Fixture

TBD. Working group to provide description and mechanical details and figures.

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.
- 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 1.88G_{RMS} for a duration of 15 minutes per side for all six surfaces per Table 64.

Commented [TN22]: This section is a work in progress. Contact the OCP NIC 3.0 Work Group for updates.

Commented [TN23R22]: What about the other environmental testing requirements?

Table 64: Random Vibration Testing 1.88G_{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 71G ±5% with a 2ms duration. All six sides shall be tested.
- Non-operational square wave shock test at 32G ±5% at a rate of 270 inches/sec. All six sides shall be tested.
- 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 115) 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 115: 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 exceeding the 50% limit 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

Commented [TN24]: Line side plating / durability requirements are actively being discussed in the OCP NIC 3.0 Workgroup.

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 (“Prop 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 65 for details.

Table 65: 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 55024: 2010+A1:2015 Immunity requirements for European Union (EU)
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 66.

Table 66: Safety Requirements

Targeted Category	Applicable Specifications
Safety	UL 60950-1/CSA C22.2 No. 60950-1-07, 2nd 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. 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 67.

Table 67: Immunity (ESD) Requirements

Targeted Category	Applicable Specifications
Immunity (ESD)	EN 55024 2010, and IEC 61000-4-2 2008 for ESD. Required ±4kV contact charge and ±8kV air discharge
NEBS Level III (optional)	Optionally test devices to NEBS level 3 – Required ±8kV contact charge and ±16kV 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

An OCP NIC 3.0 card is recommended to meet below compliance requirements.

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

- 10dB margin to FCC sub-part 15 b class A emission requirements as specified in Section 7.1.2.

8 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	<ul style="list-style-type: none"> - 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	05/31/2018