



OPEN
Compute Project

OCP NIC 3.0 Design Specification

Version 0.70

Author: OCP Server Workgroup, OCP NIC subgroup

Table of Contents

1 Overview	9
1.1 License	9
1.2 Acknowledgements	10
1.3 Background	11
1.4 Overview	13
1.4.1 Mechanical Form factor overview	13
1.4.2 Electrical overview	15
1.5 References	17
1.5.1 Trademarks	18
2 Card Form Factor	19
2.1 Form Factor Options	19
2.2 I/O bracket	22
2.2.1 Small Form Factor OCP NIC 3.0 Card I/O Bracket	22
2.2.2 Small Form Factor OCP NIC 3.0 Card with Thumbscrew Critical-to-Function (CTF) Dimensions	27
2.2.3 Small Form Factor OCP NIC 3.0 Card with Ejector Latch Critical-to-Function (CTF) Dimensions	29
2.2.4 Small Form Factor OCP NIC 3.0 Baseboard Critical-to-Function (CTF) Dimensions	31
2.2.5 Large Form Factor OCP NIC 3.0 Card I/O Bracket	33
2.2.6 Large Form Factor OCP NIC 3.0 Card Critical-to-Function (CTF) Dimensions	37
2.2.7 Large Form Factor OCP NIC 3.0 Baseboard Critical-to-Function (CTF) Dimensions	38
2.3 Line Side I/O Implementations	40
2.4 Port Numbering and LED Implementations	40
2.4.1 OCP NIC 3.0 Port Naming and Port Numbering	40
2.4.2 OCP NIC 3.0 Card LED Configuration	40
2.4.3 OCP NIC 3.0 Card LED Ordering	41
2.4.4 Baseboard LEDs Configuration over the Scan Chain	42
2.5 Mechanical Keep Out Zones	43
2.5.1 Small Card Form Factor Keep Out Zones	43
2.5.2 Large Card Form Factor Keep Out Zones	45
2.5.3 Baseboard Keep Out Zones	47
2.6 Insulation Requirements	48
2.6.1 Small Card Insulator	48
2.6.2 Large Card Insulator	49
2.7 Labeling Requirements	50
2.7.1 General Guidelines	50
2.7.2 Small Card Label Areas	50
2.7.3 Large Card Label Areas	51
2.7.4 NIC Vendor Factory Label Example	51
2.7.5 NIC Vendor Serial Number Label Example	51
2.7.6 Baseboard MAC and Serial Number Label Example	52
2.7.7 Regulatory Label Example	52
2.7.8 System Vendor Part Number Label Example	53
2.7.9 NIC Vendor Part Number Label Example	53
2.8 NIC Implementation Examples	54
2.9 Non-NIC Use Cases	54
3 Card Edge and Baseboard Connector Interface	55
3.1 Gold Finger Requirements	55
3.1.1 Gold Finger Mating Sequence	56
3.2 Baseboard Connector Requirements	58
3.2.1 Right Angle Connector	58
3.2.2 Right Angle Offset	59
3.2.3 Straddle Mount Connector	59

3.2.4	Straddle Mount Offset and PCB Thickness Options	61
3.2.5	Large Card Connector Locations.....	62
3.3	Pin definition	62
3.4	Signal Descriptions – Common	66
3.4.1	PCIe Interface Pins.....	66
3.4.2	PCIe Present and Bifurcation Control Pins	69
3.4.3	SMBus Interface Pins.....	71
3.4.4	Power Supply Pins	73
3.4.5	Miscellaneous Pins	75
3.5	Signal Descriptions – OCP Bay (Primary Connector)	76
3.5.1	PCIe Interface Pins – OCP Bay (Primary Connector).....	76
3.5.2	NC-SI Over RBT Interface Pins – OCP Bay (Primary Connector)	78
3.5.3	Scan Chain Pins – OCP Bay (Primary Connector).....	83
3.5.4	Primary Connector Miscellaneous Pins – OCP Bay (Primary Connector)	89
3.6	PCIe Bifurcation Mechanism	91
3.6.1	PCIe OCP NIC 3.0 Card to Baseboard Bifurcation Configuration (PRSNTA#, PRSNTB[3:0]#)	91
3.6.2	PCIe Baseboard to OCP NIC 3.0 Card Bifurcation Configuration (BIF[2:0]#)	91
3.6.3	PCIe Bifurcation Decoder	92
3.6.4	Bifurcation Detection Flow.....	94
3.6.5	PCIe Bifurcation Examples.....	95
3.7	PCIe Clocking Topology.....	100
3.8	PCIe Bifurcation Results and REFCLK Mapping.....	101
3.9	Power Capacity and Power Delivery.....	111
3.9.1	NIC Power Off.....	111
3.9.2	ID Mode.....	111
3.9.3	Aux Power Mode (S5).....	111
3.9.4	Main Power Mode (S0).....	112
3.10	Power Supply Rail Requirements and Slot Power Envelopes	112
3.11	Hot Swap Considerations for +12V_EDGE and +3.3V_EDGE Rails.....	112
3.12	Power Sequence Timing Requirements.....	113
4	Management and Pre-OS Requirements	115
4.1	Sideband Management Interface and Transport	115
4.2	NC-SI Traffic.....	116
4.3	Management Controller (MC) MAC Address Provisioning	116
4.4	Temperature Reporting.....	118
4.5	Power Consumption Reporting	119
4.6	Pluggable Transceiver Module Status and Temperature Reporting	119
4.7	Management and Pre-OS Firmware Inventory and Update.....	120
4.7.1	Secure Firmware.....	120
4.7.2	Firmware Inventory.....	120
4.7.3	Firmware Inventory and Update in Multi-Host Environments.....	121
4.8	NC-SI Package Addressing and Hardware Arbitration Requirements	121
4.8.1	NC-SI over RBT Package Addressing	121
4.8.2	Arbitration Ring Connections	121
4.9	SMBus 2.0 Addressing Requirements.....	121
4.9.1	SMBus Address Map.....	122
4.10	FRU EEPROM	122
4.10.1	FRU EEPROM Address, Size and Availability.....	122
4.10.2	FRU EEPROM Content Requirements.....	122
5	Routing Guidelines and Signal Integrity Considerations	125
5.1	NC-SI Over RBT	125
5.2	SMBus 2.0.....	125
5.3	PCIe.....	125

6	Thermal and Environmental	126
6.1	Airflow Direction	126
6.1.1	Hot Aisle Cooling	126
6.1.2	Cold Aisle Cooling	127
6.2	Design Guidelines	128
6.2.1	ASIC Cooling – Hot Aisle	128
6.2.2	ASIC Cooling – Cold Aisle	130
6.3	Thermal Simulation (CFD) Modeling	132
6.3.1	CFD Geometry – Small Card	132
6.3.2	Transceiver Simulation Modeling	133
6.4	Thermal Test Fixture – Small Card	134
6.5	Sensor Requirements	134
6.6	Card Cooling Tiers	135
6.6.1	Hot Aisle Cooling Tiers	135
6.6.2	Cold Aisle Cooling Tiers	135
6.7	Shock & Vibration	136
6.8	Gold Finger Plating Requirements	136
6.8.1	Host Side Gold Finger Plating Requirements	136
6.8.2	Line Side Gold Finger Plating Requirements	136
7	Regulatory	137
7.1	Required Compliance	137
7.1.1	Required Environmental Compliance	137
7.1.2	Required EMC Compliance	137
7.1.3	Required Product Safety Compliance	138
7.2	Recommended Compliance	138
7.2.1	Recommended Environmental Compliance	138
7.2.2	Recommended EMC Compliance	138
8	Revision History	140

List of Figures

Figure 1: Representative Small OCP NIC 3.0 Card with Dual QSFP Ports	11
Figure 2: Representative Large OCP NIC 3.0 Card with Dual QSFP Ports and on-board DRAM	12
Figure 3: Small and Large Card Form-Factors (not to scale)	13
Figure 4: Primary Connector (4C+) and Secondary Connector (4C) (Large) OCP NIC 3.0 Cards	19
Figure 5: Primary Connector (4C+) Only (Large) OCP NIC 3.0 Cards	20
Figure 6: Primary Connector (4C+) with 4C and 2C (Small) OCP NIC 3.0 Cards.....	20
Figure 7: Small Card Standard I/O Bracket with Thumbscrew and Pulltab (3D View)	22
Figure 8: Small Card Standard I/O Bracket with Thumbscrew and Pulltab (2D View)	22
Figure 9: Small Card Generic I/O Bracket with a Latching Lever (3D View)	23
Figure 10: Small Card Generic I/O Bracket with a Latching Lever (2D View)	23
Figure 11: Small Card Ejector Lever (2D View).....	24
Figure 12: Ejector Lock.....	24
Figure 13: Side EMI Finger	25
Figure 14: Small Card Assembly (Thumbscrew + Pull Tab Version)	26
Figure 15: Small Card Assembly (Ejector Lever Version).....	27
Figure 16: Small Form Factor OCP NIC 3.0 Card with Thumbscrew CTF Dimensions (Top View).....	28
Figure 17: Small Form Factor OCP NIC 3.0 Card with Thumbscrew CTF Dimensions (Front View)	28
Figure 18: Small Form Factor OCP NIC 3.0 Card with Thumbscrew CTF Dimensions (Side View)	29
Figure 19: Small Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Top View)	30
Figure 20: Small Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Front View)	30
Figure 21: Small Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Side View)	30
Figure 22: Small Form Factor Baseboard Chassis CTF Dimensions (Rear View).....	31
Figure 23: Small Form Factor Baseboard Chassis to Card Thumb Screw CTF Dimensions (Side View)	31
Figure 24: Small Form Factor Baseboard Chassis to Ejector lever Card CTF Dimensions (Side View).....	32
Figure 25: Small Form Factor Baseboard Chassis CTF Dimensions (Rear Rail Guide View).....	32
Figure 26: Small Form Factor Baseboard Chassis CTF Dimensions (Rail Guide Detail) – Detail C	32
Figure 27: Large Card I/O Bracket with Ejector (3D View)	33
Figure 28: Large I/O Bracket (2D View).....	34
Figure 29: Large Card Standard I/O Bracket – Left Ejector Lever.....	34
Figure 30: Large Card Standard I/O Bracket – Right Ejector Lever.....	35
Figure 31: Large Card Assembly (Dual QSFP)	36
Figure 32: Large Card Assembly (Quad SFP)	36
Figure 33: Large Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Top View)	37
Figure 34: Large Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Front View)	38
Figure 35: Large Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Side View)	38
Figure 36: Large Form Factor Baseboard Chassis CTF Dimensions (Rear View).....	38
Figure 37: Large Form Factor Baseboard Chassis CTF Dimensions (Side View)	39
Figure 38: Large Form Factor Baseboard Chassis CTF Dimensions (Rail Guide View)	39
Figure 39: Large Form Factor Baseboard Chassis CTF Dimensions (Rail Guide – Detail C).....	39
Figure 40: Port and LED Ordering – Example Small Card Link/Activity and Speed LED Placement.....	42
Figure 41: Small Form Factor Keep Out Zone – Top View.....	43
Figure 42: Small Form Factor Keep Out Zone – Top View – Detail A	43
Figure 43: Small Form Factor Keep Out Zone – Bottom View.....	44
Figure 44: Small Form Factor Keep Out Zone – Side View	44
Figure 45: Small Form Factor Keep Out Zone – Side View – Detail D.....	45
Figure 46: Large Form Factor Keep Out Zone – Top View.....	45
Figure 47: Large Form Factor Keep Out Zone – Top View – Detail A	46
Figure 48: Large Form Factor Keep Out Zone – Bottom View.....	46
Figure 49: Large Form Factor Keep Out Zone – Side View	47
Figure 50: Large Form Factor Keep Out Zone – Side View – Detail D	47
Figure 51: Small Card Bottom Side Insulator (3D View).....	48
Figure 52: Small Card Bottom Side Insulator (Top and Side View).....	48
Figure 53: Large Card Bottom Side Insulator (3D View).....	49
Figure 54: Large Card Bottom Side Insulator (Top and Side View)	49
Figure 55: Small Card Label Areas	50
Figure 56: Large Card Label Placement Example	51
Figure 57: NIC Vendor Factory Label Example	51

Figure 58: NIC Vendor Serial Number Label Example	52
Figure 59: Baseboard MAC and Serial Number Label Example	52
Figure 60: OCP NIC 3.0 Card Regulatory Label Example	52
Figure 61: System Vendor Part Number Label Example	53
Figure 62: OCP NIC 3.0 Card Vendor Part Number Label	53
Figure 63: Small Size Primary Connector Gold Finger Dimensions – x16 – Top Side (“B” Pins)	55
Figure 64: Large Size Card Gold Finger Dimensions – x32 – Top Side (“B” Pins)	56
Figure 65: Large Size Card Gold Finger Dimensions – x32 – Bottom Side (“A” Pins)	56
Figure 66: 168-pin Base Board Primary Connector – Right Angle	58
Figure 67: 140-pin Base Board Secondary Connector – Right Angle	59
Figure 68: OCP NIC 3.0 Card and Host Offset for Right Angle Connectors	59
Figure 69: 168-pin Base Board Primary Connector – Straddle Mount	60
Figure 70: 140-pin Base Board Secondary Connector – Straddle Mount	60
Figure 71: OCP NIC 3.0 Card and Baseboard PCB Thickness Options for Straddle Mount Connectors	61
Figure 72: 0mm Offset (Coplanar) for 0.062” Thick Baseboards	61
Figure 73: 0.3mm Offset for 0.076” Thick Baseboards	62
Figure 74: Primary and Secondary Connector Locations for Large Card Support with Right Angle Connectors	62
Figure 75: Primary and Secondary Connector Locations for Large Card Support with Straddle Mount Connectors	62
Figure 76: PCIe Present and Bifurcation Control Pins	71
Figure 77: Example SMBus Connections	72
Figure 78: Example Power Supply Topology	74
Figure 79: NC-SI Over RBT Connection Example – Single Primary Connector	81
Figure 80: NC-SI Over RBT Connection Example – Dual Primary Connector	82
Figure 81: Example Scan Chain Timing Diagram	84
Figure 82: Scan Bus Connection Example	88
Figure 83: Single Host (1 x16) and 1 x16 OCP NIC 3.0 Card (Single Controller)	95
Figure 84: Single Host (2 x8) and 2 x8 OCP NIC 3.0 Card (Dual Controllers)	96
Figure 85: Quad Hosts (4 x4) and 4 x4 OCP NIC 3.0 Card (Single Controller)	97
Figure 86: Quad Hosts (4 x4) and 4 x4 OCP NIC 3.0 Card (Quad Controllers)	98
Figure 87: Single Host with no Bifurcation (1 x16) and 2 x8 OCP NIC 3.0 Card (Dual Controllers)	99
Figure 88: PCIe Interface Connections for 1 x16 and 2 x8 OCP NIC 3.0 Cards	100
Figure 89: PCIe Interface Connections for a 4 x4 OCP NIC 3.0 Card	101
Figure 90: Baseboard Power States	111
Figure 91: Power-Up Sequencing	113
Figure 92: Power-Down Sequencing	114
Figure 93: Airflow Direction for Hot Aisle Cooling	126
Figure 94: Airflow Direction for Cold Aisle Cooling	127
Figure 95: ASIC Supportable Power for Hot Aisle Cooling – Small Card Form Factor	128
Figure 96: OCP NIC 3.0 Reference Geometry CAD & CFD	129
Figure 97: Server System Airflow Capability – Hot Aisle Cooling	130
Figure 98: ASIC Supportable Power for Cold Aisle Cooling – Small Card Form Factor	131
Figure 99: Server System Airflow Capability – Cold Aisle Cooling	131
Figure 100: ASIC Supportable Power Comparison – Small Card Form Factor	132
Figure 101: Small Card Thermal Test Fixture Preliminary Design	134
Figure 102: Small Card Thermal Test Fixture Preliminary Design – Transparent View	134

List of Tables

Table 1: Acknowledgements – By Company	10
Table 2: OCP 3.0 Form Factor Dimensions	14
Table 3: Baseboard to OCP NIC Form factor Compatibility Chart	14
Table 4: OCP NIC 3.0 Card Definitions.....	21
Table 5: Mechanical BOM for the Small Card Assembly	25
Table 6: CTF Default Tolerances.....	27
Table 7: CTF Default Tolerances.....	29
Table 8: Mechanical BOM for the Large Card Assembly	35
Table 9: CTF Default Tolerances.....	37
Table 10: OCP NIC 3.0 Line Side I/O Implementations.....	40
Table 11: OCP NIC 3.0 Card LED Configuration with Two Physical LEDs per Port	41
Table 12: NIC Implementation Examples and 3D CAD	54
Table 13: Example Non-NIC Use Cases	54
Table 14: Contact Mating Positions for the Primary and Secondary Connectors	56
Table 15: Right Angle Connector Options	58
Table 16: Straddle Mount Connector Options	59
Table 17: Primary Connector Pin Definition (x16) (4C+)	63
Table 18: Secondary Connector Pin Definition (x16) (4C)	65
Table 19: Pin Descriptions – PCIe 1	66
Table 20: Pin Descriptions – PCIe Present and Bifurcation Control Pins	69
Table 21: Pin Descriptions – SMBus	71
Table 22: Pin Descriptions – Power	73
Table 23: Pin Descriptions – Miscellaneous 1	75
Table 24: Pin Descriptions – PCIe 2	76
Table 25: Pin Descriptions – NC-SI Over RBT	78
Table 26: Pin Descriptions – Scan Chain	83
Table 27: Pin Descriptions – Scan Chain DATA_OUT Bit Definition.....	85
Table 28: Pin Descriptions – Scan Bus DATA_IN Bit Definition	85
Table 29: Pin Descriptions – Miscellaneous 2	89
Table 30: PCIe Bifurcation Decoder for x16 and x8 Card Widths	93
Table 31: PCIe Clock Associations	100
Table 32: Bifurcation for Single Host, Single Socket and Single Upstream Link (BIF[2:0]#=0b000)	102
Table 33: Bifurcation for Single Host, Single Socket and Single/Dual Upstream Links (BIF[2:0]#=0b000)	103
Table 34: Bifurcation for Single Host, Single Socket and Single/Dual/Quad Upstream Links (BIF[2:0]#=0b000).....	104
Table 35: Bifurcation for Single Host, Dual Sockets and Dual Upstream Links (BIF[2:0]#=0b001).....	105
Table 36: Bifurcation for Single Host, Quad Sockets and Quad Upstream Links (BIF[2:0]#=0b010)	106
Table 37: Bifurcation for Single Host, Quad Sockets and Quad Upstream Links – First 8 PCIe Lanes (BIF[2:0]#=0b011)	107
Table 38: Bifurcation for Dual Host, Dual Sockets and Dual Upstream Links (BIF[2:0]#=0b101)	108
Table 39: Bifurcation for Quad Host, Quad Sockets and Quad Upstream Links (BIF[2:0]#=0b110).....	109
Table 40: Bifurcation for Quad Host, Quad Sockets and Quad Upstream Links – First 8 lanes (BIF[2:0]#=0b111).....	110
Table 41: Power States	111
Table 42: Baseboard Power Supply Rail Requirements – Slot Power Envelopes	112
Table 43: Power Sequencing Parameters	114
Table 44: OCP NIC 3.0 Management Implementation Definitions.....	115
Table 45: Sideband Management Interface and Transport Requirements.....	115
Table 46: NC-SI Traffic Requirements	116
Table 47: MC MAC Address Provisioning Requirements	116
Table 48: Temperature Reporting Requirements	118
Table 49: Power Consumption Reporting Requirements.....	119
Table 50: Pluggable Module Status Reporting Requirements	119
Table 51: Management and Pre-OS Firmware Inventory and Update Requirements	120
Table 52: SMBus Address Map	122
Table 53: FRU EEPROM Record – OEM Record 0xC0, Offset 0x00	123
Table 54: Hot Aisle Air Temperature Boundary Conditions	127
Table 55: Hot Aisle Airflow Boundary Conditions	127
Table 56: Cold Aisle Air Temperature Boundary Conditions	127
Table 57: Cold Aisle Airflow Boundary Conditions.....	127

Table 58: Reference OCP NIC 3.0 Small Card Geometry 129

Table 59: Hot Aisle Card Cooling Tier Definitions (LFM) 135

Table 60: Cold Aisle Card Cooling Tier Definitions (LFM) 136

Table 61: FCC Class A Radiated and Conducted Emissions Requirements Based on Geographical Location 137

Table 62: Safety Requirements 138

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 and 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 3.5 for a complete list of pin and function descriptions for the OCP Bay portion of the Primary Connector.

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 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. *DSP0248, 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.
- IPMI Platform Management FRU Information Storage Definition, v1.2, February 28th, 2013.
- 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.

1.5.1 Trademarks

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

2 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 4 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 4 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 4: 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 (x24)		2C	4C+	OCP Bay
Large (x32)	4C		4C+	OCP Bay

2.2 I/O bracket

The following section defines the standard I/O bracket and standard chassis opening required for both the Small and Large form-factor cards.

2.2.1 Small Form Factor OCP NIC 3.0 Card I/O Bracket

Figure 7 and Figure 8 shows the standard Small Card form factor I/O bracket with a thumbscrew and pull tab assembly.

Figure 7: Small Card Standard I/O Bracket with Thumbscrew and Pulltab (3D View)




Figure 8: Small Card Standard I/O Bracket with Thumbscrew and Pulltab (2D View)




Figure 9 and Figure 10 shows the standard Small Card form factor I/O bracket with a latching lever assembly

Figure 9: Small Card Generic I/O Bracket with a Latching Lever (3D View)




Figure 10: Small Card Generic I/O Bracket with a Latching Lever (2D View)






Figure 11 shows the Small Card form factor ejector lever.

Figure 11: Small Card Ejector Lever (2D View)




The Small Card ejector uses a locking mechanism at the end of the handle to retain the lever position. This is shown in Figure 12.

Figure 12: Ejector Lock



The side EMI finger is defined in Figure 13. The top and bottom EMI fingers are commercial off the shelf components and are listed in the mechanical BOM in Table 5.

Figure 13: Side EMI Finger



1. DIMENSION ARE IN MILIMETER
2. MATERIAL: 0.05mm BeCu, C17200 DARDENED, BRIGHT TIN PLATING
3. TOLERANCE UNLESS OTHERWISE SPECIFIED: ± 0.25mm, ±1.0°

In addition to the sheet metal, Table 5 lists the additional hardware components used for the Small Card assembly.

Table 5: Mechanical BOM for the Small Card Assembly

Item #	Item description	Supplier Name	Supplier Part Number
1	Bracket		See Figure 8 and Figure 10. NIC_OCPv3_SFF_Bracket_1tab_20180124.pdf NIC_OCPv3_SFF_Bracket_latch_20180124.pdf
2	Top and bottom EMI fingers		TF187VE32F11
3	Rivet		1-AC-2421-03_2.4x2.1
4	Side EMI Finger		See Figure 13 and drawing NIC_OCPv3_sideEMI_20180124.pdf
5	Thumbscrew		J-4C-99-343-KEEE_rev04
6	Pull Tab	TBD	TBD
7	Screw for securing pull tab (M2 x 5mm)		ICTB0D200509B-ZD01
8	Ejector Compression Washer	TBD	TBD

9	Ejector Handle		See Figure 11 and drawing NIC_OCPv3_EjectorHandle_20180124.pdf
10	Ejector lock		See Figure 12 and drawing NIC_OCPv3_EjectorLock_20180124.pdf
11	Ejector Bushing	TBD	TBD
12	Screw (used for attaching bracket to NIC)		FCMMQ200503N
13	Screw (used for attaching bracket and ejector to NIC)		ICMMAJ200403N3
14	SMT Nut (on NIC)		82-950-22-010-01-RL

Dimensionally identical parts may be substituted in the assembly. Substituted parts shall meet or exceed the tolerances specified by the supplier part numbers in the BOM table above.

Note: The “Pull Tab” shown in the 3D drawings and in Table 5 are tentative. Alternate designs are under evaluation and therefore the BOM may change in the next revision of the specification.

Figure 14 shows the thumbscrew + pull tab assembly and Figure 15 shows the card assembly with the ejector.

Figure 14: Small Card Assembly (Thumbscrew + Pull Tab Version)





Figure 15: Small Card Assembly (Ejector Lever Version)



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

2.2.2 Small Form Factor OCP NIC 3.0 Card with Thumbscrew Critical-to-Function (CTF) Dimensions

The following dimensions are considered critical-to-function (CTF) for each small form factor OCP NIC 3.0 card with thumbscrew. The CTF default tolerances are shown in Table 6.

Table 6: CTF Default Tolerances

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

Figure 16: Small Form Factor OCP NIC 3.0 Card with Thumbscrew CTF Dimensions (Top View)




Figure 17: Small Form Factor OCP NIC 3.0 Card with Thumbscrew CTF Dimensions (Front View)





Figure 18: Small Form Factor OCP NIC 3.0 Card with Thumbscrew CTF Dimensions (Side View)



2.2.3 Small Form Factor OCP NIC 3.0 Card with Ejector Latch Critical-to-Function (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 Table 7.

Table 7: CTF Default Tolerances

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

Figure 19: Small Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Top View)




Figure 20: Small Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Front View)





Figure 21: Small Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Side View)



2.2.4 Small Form Factor OCP NIC 3.0 Baseboard Critical-to-Function (CTF) Dimensions

The following dimensions are considered critical-to-function (CTF) for each small form factor baseboard chassis.

Figure 22: Small Form Factor Baseboard Chassis CTF Dimensions (Rear View)




Figure 23: Small Form Factor Baseboard Chassis to Card Thumb Screw CTF Dimensions (Side View)




Figure 24: Small Form Factor Baseboard Chassis to Ejector lever Card CTF Dimensions (Side View)




Figure 25: Small Form Factor Baseboard Chassis CTF Dimensions (Rear Rail Guide View)





Figure 26: Small Form Factor Baseboard Chassis CTF Dimensions (Rail Guide Detail) – Detail C



Card guides are identical between the Small and Large form factor cards. The card guide 3D CAD packages may be downloaded from the OCP NIC 3.0 Wiki site:

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

2.2.5 Large Form Factor OCP NIC 3.0 Card I/O Bracket

Figure 27 and Figure 28 shows the standard Large Card form factor I/O bracket with ejector levers.

Figure 27: Large Card I/O Bracket with Ejector (3D View)





Figure 30: Large Card Standard I/O Bracket – Right Ejector Lever



In addition to the sheet metal, Table 8 lists the additional hardware components used for the Small Card bracket assembly.

Table 8: Mechanical BOM for the Large Card Assembly

Item #	Item description	Supplier Name	Supplier Part Number
1	Bracket		See Figure 28 & Drawing: NIC_OCPv3_LFF_Bracket_latch_20180124.pdf
2	Top and bottom EMI fingers		TF187VE32F11
3	Rivet		1-AC-2421-03_2.4x2.1
4	Side EMI Finger		See Figure 13 from Small Card & Drawing NIC_OCPv3_sideEMI_20180124.pdf
5	Ejector Compression Washer	TBD	TBD
6	Ejector Lever – Left		See Figure 29 & Drawing NIC_OCPv3_EjectorLever_Left_20180124.pdf
7	Ejector Lever – Right		See Figure 30 & Drawing NIC_OCPv3_EjectorLever_Right_20180124.pdf
8	Ejector Lock		See Figure 12 from Small Card & Drawing NIC_OCPv3_EjectorLock_20180124.pdf
9	Ejector Bushing	TBD	TBD
10	Screw (for attaching bracket & ejector to NIC)		ICMMAJ200403N3

11	SMT Nut (on NIC)		82-950-22-010-01-RL
----	------------------	--	---------------------

Dimensionally identical parts may be substituted in the assembly. Substituted parts shall meet or exceed the tolerances specified by the supplier part numbers in the BOM table above.

Figure 31 shows a dual QSFP assembly and Figure 32 shows a quad SFP assembly on the OCP NIC 3.0 card.

Figure 31: Large Card Assembly (Dual QSFP)





Figure 32: Large Card Assembly (Quad SFP)



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

2.2.6 Large Form Factor OCP NIC 3.0 Card Critical-to-Function (CTF) Dimensions

The following dimensions are considered critical-to-function (CTF) for each large form factor OCP NIC 3.0 card.

Table 9: CTF Default Tolerances

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

Figure 33: Large Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Top View)




Figure 34: Large Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Front View)





Figure 35: Large Form Factor OCP NIC 3.0 Card with Ejector CTF Dimensions (Side View)



2.2.7 Large Form Factor OCP NIC 3.0 Baseboard Critical-to-Function (CTF) Dimensions

The following dimensions are considered critical-to-function (CTF) for each large form factor baseboard chassis.

Figure 36: Large Form Factor Baseboard Chassis CTF Dimensions (Rear View)




Figure 37: Large Form Factor Baseboard Chassis CTF Dimensions (Side View)






Figure 38: Large Form Factor Baseboard Chassis CTF Dimensions (Rail Guide View)



SECTION B-B

Figure 39: Large Form Factor Baseboard Chassis CTF Dimensions (Rail Guide – Detail C)



DETAIL C

Card guides are identical between the Small and Large form factor cards. The card guide 3D CAD packages may be downloaded from the OCP NIC 3.0 Wiki site:
<http://www.opencompute.org/wiki/Server/Mezz>.

2.3 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 10: 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.

2.4 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 3.5.3) 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.

2.4.1 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 40 as an example implementation.

2.4.2 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. The recommended local (on-card) LED implementation uses two physical LEDs (a bi-colored Speed A/Speed B Link LED and a discrete Activity LED). Table 11 describes the OCP NIC 3.0 card LED implementations.

Table 11: 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	This LED shall be used to indicate link. 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. 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. 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. The illuminated Link LED indicator may be blinked and used for port identification through vendor specific link diagnostic software. 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. For serviceability, green LEDs shall emit light at a wavelength between 513nm and 537nm while yellow LEDs shall emit light at a wavelength between 580nm and 589nm.
	Off	
Activity	Green	Active low.
	Off	When the link is up and there is no activity, this LED shall be lit and solid. When the link is up and there is link activity, then this LED should blink at the interval of 50-500ms during link activity. 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. For serviceability, green LEDs shall emit light at a wavelength between 513nm and 537nm.


2.4.3 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.

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 for each port and the activity LED shall be located on the right side for each port. The port ordering shall increase from left to right.

The 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. For 4xSFP and 2xQSFP configurations, no LEDs are expected on the OCP NIC 3.0 card.

Figure 40: Port and LED Ordering – Example Small Card Link/Activity and Speed LED Placement



Note: The example port and LED ordering diagrams shown in Figure 40 are viewed with the card in the horizontal position and the primary side is facing up.

2.4.4 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 discrete on-board (faceplate) LED indicators. In this case, the line side link and activity LED indicators are implemented on the baseboard system via the Scan Chain. The Scan Chain bit stream is defined in Section 3.5.3.

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 yellow 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).


Figure 43: Small Form Factor Keep Out Zone – Bottom View



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


Figure 44: 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 45: Small Form Factor Keep Out Zone – Side View – Detail D



2.5.2 Large Card Form Factor Keep Out Zones

Figure 46: Large Form Factor Keep Out Zone – Top View




Figure 47: Large Form Factor Keep Out Zone – Top View – Detail A




Figure 48: Large Form Factor Keep Out Zone – Bottom View






Figure 49: Large Form Factor Keep Out Zone – Side View



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

Figure 50: Large Form Factor Keep Out Zone – Side View – Detail D



2.5.3 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>

2.6 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.

2.6.1 Small Card Insulator

Figure 51: Small Card Bottom Side Insulator (3D View)





Figure 52: Small Card Bottom Side Insulator (Top and Side View)



2.6.2 Large Card Insulator

Figure 53: Large Card Bottom Side Insulator (3D View)





Figure 54: 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: $\pm 0.30\text{mm}$, $\pm 1.0^\circ$



2.7 Labeling Requirements

OCP NIC 3.0 cards shall implement all (or a subset of) label items listed below as deemed necessary by each customer. All labels shall be placed on the secondary side of the insulator and within their designated areas or zones.

There are four label areas for the OCP NIC 3.0 cards:

- Serial Number Labels
- Part Number Labels
- MAC Labels
- Regulatory Labels

Note: regulatory marks may printed on the insulator instead of affixed via a label.

Additional labels can 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) within each label zone.

2.7.1 General Guidelines


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 label may include:

- Serial number
- Date Code
- Manufacturing Site Code

The label size and typeface may vary based on each customer's label content and requirements. The following sections show representative label examples for each label area.

2.7.2 Small Card Label Areas


Small form factor OCP NIC 3.0 card labels shall be placed in the indicated areas below.



2.7.3 Large Card Label Areas

Large form factor OCP NIC 3.0 card labels shall be placed in the indicated areas below.

Figure 56: Large Card Label Placement Example



2.7.4 NIC Vendor Factory Label Example

An example NIC vendor factory label is shown in Figure 57.


Figure 57: NIC Vendor Factory Label Example



2.7.5 NIC Vendor Serial Number Label Example

An example NIC vendor serial number label is shown in Figure 58.

Figure 58: NIC Vendor Serial Number Label Example



2.7.6 Baseboard MAC and Serial Number Label Example

An example baseboard MAC and serial number label is shown in Figure 59.

Figure 59: Baseboard MAC and Serial Number Label Example



2.7.7 Regulatory Label Example

An example regulatory label is shown in Figure 60. The regulatory markings information may be directly printed on to the insulator.

Figure 60: OCP NIC 3.0 Card Regulatory Label Example





Image of label is for reference only;
actual label will have different data.

2.7.8 System Vendor Part Number Label Example


An example system vendor part number label is shown in Figure 61.



1. Font Verdana or equivalent
2. Human readable text – 6 pt.
3. Barcode – code 93, height 6.0 mm
4. Print resolution – 300 dpi
5. 18.0mm x 10.0mm label size (0.7" x 0.375")
6. Label material shall be white litho paper, label stock should be 89 g/m² (equivalent to a basis weight of 24# bond or 60# offset paper and approximately 0.086mm thick)
7. The area occupied by the human readable text and barcode should be visually aligned to the center of the label. This alignment is for reference only, but must facilitate 100% scanning capability. See example below.

2.7.9 NIC Vendor Part Number Label Example

An example NIC vendor part number label is shown in Figure 62.



1. Font Verdana or equivalent
2. Human readable text – 6 pt.
3. Barcode – code 93, height 6.0 mm
4. Print resolution – 300 dpi
5. 18 mm x 10 mm label size (0.7" x 0.375")
6. Label material shall be white litho paper, label stock should be 89 g/m² (equivalent to a basis weight of 24# bond or 60# offset paper and approximately 0.086mm thick)
7. The area occupied by the human readable text and barcode should be visually aligned to the center of the label. This alignment is for reference only, but must facilitate 100% scanning capability. See example below.

2.8 NIC Implementation 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 12: 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

2.9 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 13.

Table 13: 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

3 Card Edge and Baseboard Connector Interface

3.1 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 63: Small Size Primary Connector Gold Finger Dimensions – x16 – Top Side (“B” Pins)




Figure 64: Large Size Card Gold Finger Dimensions – x32 – Top Side (“B” Pins)





Figure 65: 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	GND		
OCP B7	SLOT_ID		OCP A7	RBT_TX_EN		
OCP B8	RBT_RXD1		OCP A8	RBT_TXD1		
OCP B9	RBT_RXD0		OCP A9	RBT_TXD0		
OCP B10	GND		OCP A10	GND		
OCP B11	REFCLKn2		OCP A11	REFCLKn3		
OCP B12	REFCLKp2		OCP A12	REFCLKp3		

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

B67	GND		A67	GND	
B68	PWRBRK#		A68	RFU 2, N/C	
B69	RFU 1, N/C		A69	RFU 3, N/C	
B70	PRSNB3#		A70	RFU 4, N/C	

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 66: 168-pin Base Board Primary Connector – Right Angle






Figure 67: 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 68.

Figure 68: 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 69: 168-pin Base Board Primary Connector – Straddle Mount






Figure 70: 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 four baseboard PCB thicknesses they can accept. The available options are shown in Figure 71. 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 71: 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 72. Additionally, the connectors are also capable of having a 0.3mm offset from the centerline of the host board as shown in Figure 73.

Figure 72: 0mm Offset (Coplanar) for 0.062" Thick Baseboards





Figure 73: 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 74 and Figure 75.

Figure 74: Primary and Secondary Connector Locations for Large Card Support with Right Angle Connectors





Figure 75: 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 shown in Section 3.5.

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		Primary Connector (4C+, x16, 168-pin OCP NIC 3.0 card with OCP Bay)	Primary Connector (2C+, x8, 112-pin OCP NIC 3.0 card with OCP Bay)
OCP_B1	NIC_PWR_GOOD	PERST2#	OCP_A1		
OCP_B2	MAIN_PWR_EN	PERST3#	OCP_A2		
OCP_B3	LD#	WAKE#	OCP_A3		
OCP_B4	DATA_IN	RBT_ARB_IN	OCP_A4		
OCP_B5	DATA_OUT	RBT_ARB_OUT	OCP_A5		
OCP_B6	CLK	GND	OCP_A6		
OCP_B7	SLOT_ID	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	PERST0#	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		

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	PWRBRK#	RFU2, N/C	A68
B69	RFU1, N/C	RFU3, N/C	A69
B70	PRSNB3#	RFU4, N/C	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#	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
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	PRSNTB0#	PRSNTB1#	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

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

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	PWRBRK#	RFU2, N/C	A68
B69	RFU1, N/C	RFU3, N/C	A69
B70	PRSNB3#	RFU4, N/C	A70

3.4 Signal Descriptions – Common

The pins shown in this section are common to both the Primary and Secondary Connectors. All pin directions are from the perspective of the baseboard.

The OCP NIC 3.0 card shall implement protection methods to prevent leakage paths between the V_{AUX} and V_{MAIN} power domains in the event that a NIC is powered down in a powered up baseboard.

Note: Pins that are only used on Primary Connector 28-pin OCP bay are defined in Section 3.5.

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 89.

Table 19: Pin Descriptions – PCIe 1

Signal Name	Pin #	Baseboard Direction	Signal Description
REFCLKn0 REFCLKp0	B14 B15	Output	PCIe compliant differential reference clock #0, and #1. 100MHz reference clocks are used for the OCP NIC 3.0 card PCIe core logic.
REFCLKn1 REFCLKp1	A14 A15	Output	
			For baseboards, the REFCLK0 and REFCLK1 signals shall be available at the connector. Baseboards should disable REFCLK1 if it is not used by the OCP NIC 3.0 card.
			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>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.</p> <p>REFCLK0 is always available to all OCP NIC 3.0 cards. The card should not assume REFCLK1 is available until the bifurcation negotiation process is completed.</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> <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>
PETn1 PETp1	B20 B21	Output	
PETn2 PETp2	B23 B24	Output	
PETn3 PETp3	B26 B27	Output	
PETn4 PETp4	B30 B31	Output	
PETn5 PETp5	B33 B34	Output	
PETn6 PETp6	B36 B37	Output	
PETn7 PETp7	B39 B40	Output	
PETn8 PETp8	B44 B45	Output	
PETn9 PETp9	B47 B48	Output	
PETn10 PETp10	B50 B51	Output	
PETn11 PETp11	B53 B54	Output	
PETn12 PETp12	B56 B57	Output	
PETn13 PETp13	B59 B60	Output	
PETn14 PETp14	B62 B63	Output	
PETn15 PETp15	B65 B66	Output	
PERn0 PERp0	A17 A18	Input	<p>Receiver differential pairs [0:15]. These pins are connected from the OCP NIC 3.0 card transmitter</p>
PERn1	A20	Input	

PERp1	A21		<p>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>
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	A56	Input	
PERp12	A57		
PERn13	A59	Input	
PERp13	A60		
PERn14	A62	Input	
PERp14	A63		
PERn15	A65	Input	
PERp15	A66		
PERST0# PERST1#	B10 A11	Output	<p>PCIe Reset #0, #1. Active low.</p> <p>When PERSTn# is deasserted, the signal shall indicate the applied power is within tolerance and stable for the OCP NIC 3.0 card.</p> <p>PERST# shall be deasserted at least 100ms after the power rails are within the operating limits per the PCIe CEM Specification. The PCIe REFCLKs shall also become stable within this period of time.</p> <p>PERST shall be pulled high to +3.3V_EDGE on the baseboard.</p> <p>For OCP NIC 3.0, PERST deassertion shall also indicate the full card power envelope is available to the OCP NIC 3.0 card.</p>

			<p>For baseboards, the PERST[0:1]# signals are required at the connector.</p> <p>For OCP NIC 3.0 cards, the required PERST[0:1]# signals shall be connected to the endpoint silicon. Unused PERST[0:1]# signals shall be left as a no connect.</p> <p>Note: For cards that only support 1 x16, PERST0# is used. For cards that support 2 x8, PERST0# is used for the first eight PCIe lanes, and PERST1# is used for the second eight PCIe lanes.</p> <p>PERST0# is always available to all OCP NIC 3.0 cards. The card should not assume PERST1# is available until the bifurcation negotiation process is completed.</p> <p>Refer to Section 2.2 in the PCIe CEM Specification, Rev 4.0 for details.</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. An example connection diagram is shown in Figure 76.


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 3.6.</p> <p>Note: PRSNTB3# is located at the bottom of the 4C connector and is only applicable for OCP NIC 3.0 cards with a PCIe width of x16 (or greater). OCP NIC 3.0 cards that implement a 2C card edge do not use the PRSNTB3# pin for capabilities or present detection.</p>
BIF0# BIF1# BIF2#	B7 B8 B9	Output	<p>Bifurcation [0:2]# pins allow the baseboard to force configure the OCP NIC 3.0 card bifurcation.</p> <p>For baseboards, 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 tie the BIF[0:2]# signals to +3.3V_EDGE or to ground per the definitions are described in Section 3.6 if no dynamic bifurcation configuration is required.</p> <p>For OCP NIC 3.0 cards, these signals shall connect to the endpoint bifurcation pins if it is 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 76: PCIe Present and Bifurcation Control Pins



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 77.

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>

			For OCP NIC 3.0 cards, the SMDAT from the endpoint silicon shall be connected to the card edge gold fingers.
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 77: Example SMBus Connections



3.4.4 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 78.

Table 22: 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. 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>
+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>
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, the</p>

			<p>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>
PWRBRK#	B68	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>

Figure 78: Example Power Supply Topology



3.4.5 Miscellaneous Pins

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

Table 23: Pin Descriptions – Miscellaneous 1

Signal Name	Pin #	Baseboard Direction	Signal Description
RFU1, N/C	B69	Input / Output	Reserved future use pins. These pins shall be left as no connect.
RFU2, N/C	A68		
RFU3, N/C	A69		
RFU4, N/C	A70		

3.5 Signal Descriptions – OCP Bay (Primary Connector)

The following section describes the functions in the Primary Connector 28-pin OCP bay. This 28-pin bay is shown in Section 3.3 and have pin numbers designated as OCP_B[1:14], and OCP_A[1:14]. All pin directions on this OCP bay are from the perspective of the baseboard.

The OCP NIC 3.0 card shall implement protection methods to prevent leakage paths between the V_{aux} and V_{main} power domains in the event that a NIC is powered down in a powered up baseboard.

Note: The pins that are common to both the Primary and Secondary Connectors are defined in Section 3.4.

3.5.1 PCIe Interface Pins – OCP Bay (Primary Connector)

This section provides the pin assignments for the PCIe interface signals on the Primary Connector OCP bay. The AC/DC specifications are defined in the PCIe CEM Specification. An example connection diagram that shows REFCLK2, REFCLK3, PERST2# and PERST3# is shown in Section 3.7.

Table 24: Pin Descriptions – PCIe 2

Signal Name	Pin #	Baseboard Direction	Signal Description
REFCLKn2 REFCLKp2	OCP_B11 OCP_B12	Output	PCIe compliant differential reference clock #2, and #3. 100MHz reference clocks are used for the OCP NIC 3.0 card PCIe core logic.
REFCLKn3 REFCLKp3	OCP_A11 OCP_A12	Output	<p>For baseboards, the REFCLK2 and REFCLK3 signals are required at the Primary Connector. Baseboards may disable REFCLK2 and REFCLK3 if they are not used by the OCP NIC 3.0 card.</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: REFCLK2 and REFCLK3 are only used for cards that only support a four link PCIe bifurcation mode.</p> <p>The card should not assume REFCLK2 and REFCLK3 are available until the bifurcation negotiation process is completed.</p> <p>Refer to Section 2.1 in the PCIe CEM Specification, Rev 4.0 for details.</p>
PERST2# PERST3#	OCP_A1 OCP_A2	Output	<p>PCIe Reset #2, #3. Active low.</p> <p>When PERSTn# is deasserted, the signal shall indicate the applied power is within tolerance and stable for the OCP NIC 3.0 card.</p>

			<p>PERST# shall be deasserted at least 100ms after the power rails are within the operating limits per the PCIe CEM Specification. The PCIe REFCLKs shall also become stable within this period of time.</p> <p>PERST shall be pulled high to +3.3V_EDGE on the baseboard.</p> <p>For OCP NIC 3.0, PERST deassertion shall also indicate the full card power envelope is available to the OCP NIC 3.0 card.</p> <p>For baseboards, the PERST[2:3]# signals are required at the connector.</p> <p>For OCP NIC 3.0 cards, the required PERST[2:3]# signals shall be connected to the endpoint silicon. Unused PERST[2:3]# signals shall be left as a no connect.</p> <p>Note: PERST2# and PERST3# are only used for cards that support a four link PCIe bifurcation mode.</p> <p>The card should not assume PERST2# and PERST3# are available until the bifurcation negotiation process is completed.</p> <p>Refer to Section 2.2 in the PCIe CEM Specification, Rev 4.0 for details.</p>
WAKE#	OCP_A3	Input, OD	<p>WAKE#. Open drain. Active low.</p> <p>This signal shall be driven by the OCP NIC 3.0 card to notify the baseboard to restore PCIe link. For OCP NIC 3.0 cards that support multiple WAKE# signals, their respective WAKE# pins may be tied together as the signal is open-drain to form a wired-OR.</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 directly connected to the endpoint silicon WAKE# pin(s). This pin shall be left as a no connect if WAKE# is not supported by the silicon.</p>

			Refer to Section 2.3 in the PCIe CEM Specification, Rev 4.0 for details.
--	--	--	--

3.5.2 NC-SI Over RBT Interface Pins – OCP Bay (Primary Connector)

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 79.

Table 25: Pin Descriptions – NC-SI Over RBT

Signal Name	Pin #	Baseboard Direction	Signal Description
RBT_REF_CLK	OCP_A14	Output	<p>Reference clock input. Synchronous clock reference for receive, transmit and control interface. The clock shall have a typical frequency of 50MHz.</p> <p>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.</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-up resistor to +3.3V_EDGE on the baseboard. If the baseboard does not support NC-SI over RBT, then this signal shall be terminated to +3.3V_EDGE through a 100kOhm pull-up.</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> <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_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-up resistor to +3.3V_EDGE on the baseboard. If the baseboard does not support NC-SI over RBT, then this signal shall be terminated to +3.3V_EDGE through a 100kOhm pull-up.</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</p>

			arbitration signals to pass through in a multi-Primary Connector baseboard.
SLOT_ID	OCP_B7	Output	<p>NC-SI Address pin. This pin shall only be used if the end point silicon supports package identification.</p> <p>For baseboards, this pin shall be used to set the slot ID value. This pin shall be directly to GND for SlotID = 0. This pin shall be pulled up to +3.3V_EDGE for SlotID = 1.</p> <p>For OCP NIC 3.0 cards, this pin shall be connected to the endpoint device GPIO 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 SLOT_ID pin may be used to configure a different Package ID value so long as the resulting combination does not cause addressing interferences.</p> <p>For endpoint devices without NC-SI over RBT support, this pin shall be left as a no connect on the OCP NIC 3.0 card.</p>

Figure 79: NC-SI Over RBT Connection Example – Single Primary Connector





Figure 80: NC-SI Over RBT Connection Example – Dual Primary Connector



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 80.

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

3.5.3 Scan Chain Pins – OCP Bay (Primary Connector)

This section provides the pin assignments for the Scan Bus interface signals on the Primary Connector OCP Bay. The scan chain consists of two unidirectional busses, a clock and a 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, allows loading of the data on to the shift registers. An example timing diagram is shown in Figure 81. An example connection diagram is shown in Figure 82.


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

Table 26: 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 82, 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 tied directly to GND if the scan chain is not used.</p> <p>For NIC implementations, the DATA_OUT pin may be left floating if it is not used for OCP NIC 3.0 card configuration. The DATA_OUT pin shall be pulled up to +3.3V_EDGE through a 1kOhm resistor.</p>
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 82.</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 82. The LD# pin shall be pulled up to +3.3V_EDGE through a 1kOhm resistor.</p>

Figure 81: Example Scan Chain Timing Diagram



The scan chain provides side band 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 is mandatory 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 default implementation does not use the DATA_OUT stream and is not implemented on the NIC. However, all baseboard systems that implement the Scan Chain shall connect DATA_OUT between the platform and the Primary Connector for future-proofing NIC implementations and subsequent revisions of this specification.

Table 27: Pin Descriptions – Scan Chain DATA_OUT Bit Definition

Byte.bit	DATA_OUT Field Name	Default Value	Description
0.[0..7]	RSVD	0b000000	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 registers 0 & 1 shall be mandatory for scan chain implementations. DATA_IN shift registers 2 & 3 are optional depending on the card type and fields being reported to the host. DATA_IN shift register 2 may be used to indicate future definitions of the scan chain bit stream. DATA_IN shift registers 3 (in conjunction with shift register 2) are required for reporting link/activity indication on card implementations with 5-8 ports.

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.

A 1kOhm pull up resistor shall be implemented on the NIC to the SER input of the last shift register on the DATA_IN scan chain to maintain a default bit value of 0b1 for unused bits for implementations using less than four shift registers.


Table 28: Pin Descriptions – Scan Bus DATA_IN Bit Definition

Byte.bit	DATA_OUT 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.
1.0	LINK_P0	0b1	Port 0..3 link indication (max speed). Active low.

1.1	LINK_P1	0b1	<p>0b0 – Link LED is illuminated on the host platform. 0b1 – Link LED is not illuminated on the host platform.</p> <p>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.</p> <p>Note: The link LED may also be blinked for use as port identification.</p>
1.2	LINK_P2	0b1	
1.3	LINK_P3	0b1	
1.4	ACT_P0	0b1	<p>Port 0..3 activity indication. Active low.</p> <p>0b0 – ACT LED is illuminated on the host platform. 0b1 – ACT LED is not illuminated on the host platform.</p> <p>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.</p>
1.5	ACT_P1	0b1	
1.6	ACT_P2	0b1	
1.7	ACT_P3	0b1	
2.0	LINK_B_P0	0b1	<p>Port 0..3 link indication (not max speed). Active low.</p> <p>0b0 – Link LED is illuminated on the host platform. 0b1 – Link LED is not illuminated on the host platform.</p> <p>Steady = link is detected on the port and is not at the max speed. Off = the physical link is down, or is disabled.</p> <p>Note: The LINK_B LED may also be blinked for use as port identification.</p>
2.1	LINK_B_P1	0b1	
2.2	LINK_B_P2	0b1	
2.3	LINK_B_P3	0b1	
2.4	LINK_P4	0b1	<p>Port 4..7 link indication (max speed). Active low.</p> <p>0b0 – Link LED is illuminated on the host platform. 0b1 – Link LED is not illuminated on the host platform.</p> <p>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.</p>
2.5	LINK_P5	0b1	
2.6	LINK_P6	0b1	
2.7	LINK_P7	0b1	

			Note: The link LED may also be blinked for use as port identification.
3.0	ACT_P4	0b1	Port 4..7 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.
3.1	ACT_P5	0b1	
3.2	ACT_P6	0b1	
3.3	ACT_P7	0b1	
3.4	LINK_B_P4	0b1	Port 4..7 link 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_B LED may also be blinked for use as port identification.
3.5	LINK_B_P5	0b1	
3.6	LINK_B_P6	0b1	
3.7	LINK_B_P7	0b1	

Figure 82: Scan Bus Connection Example



3.5.4 Primary Connector Miscellaneous Pins – OCP Bay (Primary Connector)

This section provides the miscellaneous pin assignments for the pins on the Primary Connector OCP Bay. The AC/DC specifications are defined in the PCIe CEM Specification, Rev 4.0 and Section 3.12. An example NIC_PWR_GOOD connection diagram is shown in Figure 78.

Table 29: Pin Descriptions – Miscellaneous 2

Signal Name	Pin #	Baseboard Direction	Signal Description
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="802 373 1385 630"> <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 91 and Figure 92) 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 78 in Section 3.4.4 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																

GND	OCP_A6 OCP_A10 OCP_A13 OCP_B10 OCP_B13	GND	Ground return; a total of 5 ground pins are on the OCP bay area.
-----	--	-----	--

3.6 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 a 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 76.

3.6.1 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 PRSTNB[3:0]# bits is shown in Table 30 for x16 and x8 PCIe cards.

3.6.2 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 30 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.

3.6.3 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 30 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.

Table 30: PCIe Bifurcation Decoder for x16 and x8 Card Widths

Minimum Required Card Edge	Host CPU Sockets	Single Host										Dual Host	Quad Host	Quad Host		
		1 Host 1 Upstream Socket	1 Host 1 Upstream Socket	1 Host 1 Upstream Socket	1 Host 2 Upstream Sockets	1 Host 4 Upstream Sockets	1 Host 4 Sockets (1 Socket per Host) First 8 PCIe lanes	RSVD	2 Hosts 2 Upstream Sockets (1 Socket per Host)	4 Hosts 4 Upstream Sockets (1 Socket per Host)	4 Hosts 4 Sockets (1 Socket per Host) First 8 PCIe lanes					
Network Card - Supported PCIe Configurations																
	Total PCIe Links	1 Link (No Bifurcation)	1 or 2 Links	1, 2, or 4 Links	2 Links	4 Links	4 Links	4 x2 links	RSVD	2 Links	4 Links	4 Links	4 x2 links	RSVD	4 x2 links	
	System Support	1x16, 1x8, 1x4, 1x2, 1x1	2x8, 2x4, 2x2, 2x1	4x4, 4x2, 4x1	4x2, 4x1	RSVD	2x8, 2x4, 2x2, 2x1	4x4, 4x2, 4x1	4x4, 4x2, 4x1	4x2, 4x1	RSVD	4x2, 4x1				
Minimum Required Card Edge	System Encoding BIF(2,0)#	0b0000	0b0000	0b0000	0b0001	0b0010	0b0010	0b0011	0b1000	0b1001	0b1010	0b1011	0b1011	0b1000	0b1001	0b1011
	Card Short Name	Not Present	Not Present	Not Present	Not Present	Not Present	Not Present	Not Present	Not Present	Not Present	Not Present	Not Present	Not Present	Not Present	Not Present	Not Present
	Supported Bifurcation Modes	Card Not Present	Card Not Present	Card Not Present	Card Not Present	Card Not Present	Card Not Present	Card Not Present	Card Not Present	Card Not Present	Card Not Present	Card Not Present	Card Not Present	Card Not Present	Card Not Present	Card Not Present
1x16	Option A	1x8	1x8	1x8	1x8	1x8	1x8	1x8	1x8	1x8	1x8	1x8	1x8	1x8	1x8	1x8
2x8	Option B	1x4	1x4	1x4	1x4	1x4	1x4	1x4	1x4	1x4	1x4	1x4	1x4	1x4	1x4	1x4
4x4	Option C	1x2	1x2	1x2	1x2	1x2	1x2	1x2	1x2	1x2	1x2	1x2	1x2	1x2	1x2	1x2
8x2	Option D	1x1	1x1	1x1	1x1	1x1	1x1	1x1	1x1	1x1	1x1	1x1	1x1	1x1	1x1	1x1
1x16	Option B	1x8	1x8	1x8	1x8	1x8	1x8	1x8	1x8	1x8	1x8	1x8	1x8	1x8	1x8	1x8
2x8	Option B	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8
4x4	Option D	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4
1x16	Option D	4x2 (First 8 lanes), 4x1	4x2 (First 8 lanes), 4x1	4x2 (First 8 lanes), 4x1	4x2 (First 8 lanes), 4x1	4x2 (First 8 lanes), 4x1	4x2 (First 8 lanes), 4x1	4x2 (First 8 lanes), 4x1	4x2 (First 8 lanes), 4x1	4x2 (First 8 lanes), 4x1	4x2 (First 8 lanes), 4x1	4x2 (First 8 lanes), 4x1	4x2 (First 8 lanes), 4x1			
2x8	Option D	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8
4x4	Option D	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4
RSVD	Option D	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD
2x4	Option A	2x4	2x4	2x4	2x4	2x4	2x4	2x4	2x4	2x4	2x4	2x4	2x4	2x4	2x4	2x4
4x2	Option A	4x2	4x2	4x2	4x2	4x2	4x2	4x2	4x2	4x2	4x2	4x2	4x2	4x2	4x2	4x2
1x16	Option A	1x16	1x16	1x16	1x16	1x16	1x16	1x16	1x16	1x16	1x16	1x16	1x16	1x16	1x16	1x16
2x8	Option A	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8
1x16	Option B	1x16	1x16	1x16	1x16	1x16	1x16	1x16	1x16	1x16	1x16	1x16	1x16	1x16	1x16	1x16
2x8	Option B	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8	2x8
4x4	Option C	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4
4x4	Option C	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4	4x4
RSVD	Option C	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD	RSVD
0b0000	Option C	0b0000	0b0000	0b0000	0b0000	0b0000	0b0000	0b0000	0b0000	0b0000	0b0000	0b0000	0b0000	0b0000	0b0000	0b0000

3.6.4 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 30 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 91. Refer to Section 3.12 for timing details.


3.6.5 PCIe Bifurcation Examples

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

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

Figure 83 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 83: Single Host (1 x16) and 1 x16 OCP NIC 3.0 Card (Single Controller)



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

Figure 84 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 84: Single Host (2 x8) and 2 x8 OCP NIC 3.0 Card (Dual Controllers)



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

Figure 85 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 85: Quad Hosts (4 x4) and 4 x4 OCP NIC 3.0 Card (Single Controller)



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

Figure 86 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 86: Quad Hosts (4 x4) and 4 x4 OCP NIC 3.0 Card (Quad Controllers)



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

Figure 87 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 87: Single Host with no Bifurcation (1 x16) and 2 x8 OCP NIC 3.0 Card (Dual Controllers)



3.7 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 31. Cards that implement both the Primary and Secondary Connectors have a total of up to 6 REFCLKs.

Table 31: 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 described in Section 3.5.1 and are located on the 28-pin OCP bay.

Figure 88: PCIe Interface Connections for 1 x16 and 2 x8 OCP NIC 3.0 Cards





Figure 89: PCIe Interface Connections for a 4 x4 OCP NIC 3.0 Card



3.8 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 3.6.3.

Table 32: Bifurcation for Single Host, Single Socket and Single Upstream Link (BIF[2:0]#=0b000)

Single Host, Single Upstream Socket, One Upstream Link, no bifurcation																								
Min Card Width	Card Short Name	Supported Bifurcation Modes	Add-in Card Encoding PRSNTB[3:0]#	Host	Upstream Devices	Upstream Links	BIF[2:0]#	Resulting Link	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	Not Present	Card Not Present	0b1111	1 Host	1 Upstream Socket	1 Link	0b0000	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
2C	1x8 Option A	1x8, 1x4, 1x2, 1x1	0b1110	1 Host	1 Upstream Socket	1 Link	0b0000	1x8	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
2C	1x4	1x4, 1x2, 1x1	0b1110	1 Host	1 Upstream Socket	1 Link	0b0000	1x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
2C	1x2	1x2, 1x1	0b1110	1 Host	1 Upstream Socket	1 Link	0b0000	1x2	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
2C	1x1	1x1	0b1110	1 Host	1 Upstream Socket	1 Link	0b0000	1x1	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
2C	1x8 Option B	1x8, 1x4, 1x2, 1x1	0b1101	1 Host	1 Upstream Socket	1 Link	0b0000	1x8	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
2C	2x8 Option B	2x8, 2x4, 2x2, 2x1	0b1101	1 Host	1 Upstream Socket	1 Link	0b0000	1x8*	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
4C	2x8 Option B	4x8, 4x4, 4x2, 4x1	0b1100	1 Host	1 Upstream Socket	1 Link	0b0000	1x8	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
2C	1x8 Option D	2x8, 1x8, 1x4	0b1100	1 Host	1 Upstream Socket	1 Link	0b0000	1x15	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7	Link 0, Lane 8	Link 0, Lane 9	Link 0, Lane 10	Link 0, Lane 11	Link 0, Lane 12	Link 0, Lane 13	Link 0, Lane 14	Link 0, Lane 15
4C	1x16 Option D	4x8, 2x8, 2x4, 2x2, 2x1	0b1100	1 Host	1 Upstream Socket	1 Link	0b0000	1x15	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7	Link 0, Lane 8	Link 0, Lane 9	Link 0, Lane 10	Link 0, Lane 11	Link 0, Lane 12	Link 0, Lane 13	Link 0, Lane 14	Link 0, Lane 15
RSVD	RSVD	RSVD	0b1011	1 Host	1 Upstream Socket	1 Link	0b0000	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	2x4	2x4, 2x2, 2x1	0b1010	1 Host	1 Upstream Socket	1 Link	0b0000	1x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	4x2	4x2 (First 8 lanes), 4x1	0b1001	1 Host	1 Upstream Socket	1 Link	0b0000	1x2	Link 0, Lane 0	Link 0, Lane 1														
2C	4x2	1x2, 1x1	0b1001	1 Host	1 Upstream Socket	1 Link	0b0000	1x2	Link 0, Lane 0	Link 0, Lane 1														
RSVD	RSVD	RSVD for future 8 encoding	0b1000	1 Host	1 Upstream Socket	1 Link	0b0000	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7	Link 0, Lane 8	Link 0, Lane 9	Link 0, Lane 10	Link 0, Lane 11	Link 0, Lane 12	Link 0, Lane 13	Link 0, Lane 14	Link 0, Lane 15
4C	1x16 Option A	1x16, 1x8, 1x4, 1x2, 1x1	0b0111	1 Host	1 Upstream Socket	1 Link	0b0000	1x16	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7	Link 0, Lane 8	Link 0, Lane 9	Link 0, Lane 10	Link 0, Lane 11	Link 0, Lane 12	Link 0, Lane 13	Link 0, Lane 14	Link 0, Lane 15
4C	2x8 Option A	2x8, 2x4, 2x2, 2x1	0b0110	1 Host	1 Upstream Socket	1 Link	0b0000	1x8*	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7	Link 0, Lane 8	Link 0, Lane 9	Link 0, Lane 10	Link 0, Lane 11	Link 0, Lane 12	Link 0, Lane 13	Link 0, Lane 14	Link 0, Lane 15
4C	1x16 Option B	1x16, 1x8, 1x4, 1x2, 1x1	0b0101	1 Host	1 Upstream Socket	1 Link	0b0000	1x16	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7	Link 0, Lane 8	Link 0, Lane 9	Link 0, Lane 10	Link 0, Lane 11	Link 0, Lane 12	Link 0, Lane 13	Link 0, Lane 14	Link 0, Lane 15
4C	1x16 Option C	2x8, 2x4, 2x2, 2x1	0b0100	1 Host	1 Upstream Socket	1 Link	0b0000	1x16	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7	Link 0, Lane 8	Link 0, Lane 9	Link 0, Lane 10	Link 0, Lane 11	Link 0, Lane 12	Link 0, Lane 13	Link 0, Lane 14	Link 0, Lane 15
4C	1x16 Option C	4x8, 4x4, 4x2, 4x1	0b0011	1 Host	1 Upstream Socket	1 Link	0b0000	1x4*	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7	Link 0, Lane 8	Link 0, Lane 9	Link 0, Lane 10	Link 0, Lane 11	Link 0, Lane 12	Link 0, Lane 13	Link 0, Lane 14	Link 0, Lane 15
4C	4x4	4x4	0b0010	1 Host	1 Upstream Socket	1 Link	0b0000	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7	Link 0, Lane 8	Link 0, Lane 9	Link 0, Lane 10	Link 0, Lane 11	Link 0, Lane 12	Link 0, Lane 13	Link 0, Lane 14	Link 0, Lane 15
RSVD	RSVD	RSVD	0b0001	1 Host	1 Upstream Socket	1 Link	0b0000	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7	Link 0, Lane 8	Link 0, Lane 9	Link 0, Lane 10	Link 0, Lane 11	Link 0, Lane 12	Link 0, Lane 13	Link 0, Lane 14	Link 0, Lane 15
RSVD	RSVD	RSVD	0b0000	1 Host	1 Upstream Socket	1 Link	0b0000	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7	Link 0, Lane 8	Link 0, Lane 9	Link 0, Lane 10	Link 0, Lane 11	Link 0, Lane 12	Link 0, Lane 13	Link 0, Lane 14	Link 0, Lane 15

Table 33: Bifurcation for Single Host, Single Socket and Single/Dual Upstream Links (BIF[2:0]#=#0b000)

Single Host, Single Upstream Socket, One or Two Upstream Links		Supported Bifurcation Modes		Add-in Card		1 x15, 1 x8, 1 x4, 1 x2, 1 x1 2 x8, 2 x4, 2 x2, 2 x1																			
Min Card Width	Card Name	Card Not Present	Host	Encoding	Host	Upstream Devices	Upstream Links	BIF[2:0]#	Resubing Link	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	Not Present	Card Not Present	1 Host	0b1111	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
2C	1 x8 Option A	1 x8, 1 x4, 1 x2, 1 x1	1 Host	0b1110	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	1 x8	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
2C	1 x4	1 x4, 1 x2, 1 x1	1 Host	0b1110	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	1 x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
2C	1 x2	1 x2, 1 x1	1 Host	0b1110	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	1 x2	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
2C	1 x1	1 x1	1 Host	0b1110	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	1 x1	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
2C	1 x8 Option B	1 x8, 1 x4, 1 x2, 1 x1 2 x8, 2 x4, 2 x2, 2 x1	1 Host	0b1101	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	1 x8	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
4C	1 x8 Option B	2 x8, 2 x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	1 Host	0b1101	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	2 x8	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
2C	1 x8 Option D	1 x8, 1 x4 4 x2 (First 8 lanes), 4 x1	1 Host	0b1100	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	1 x8	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
4C	1 x16 Option D	1 x16, 1 x8, 1 x4 2 x8, 2 x4, 4 x4, 4 x2 (First 8 lanes), 4 x1	1 Host	0b1100	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	1 x16	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
RSVD	RSVD	RSVD	1 Host	0b1011	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	2 x4	2 x4, 2 x2, 2 x1 1 x8, 1 x4, 1 x2, 1 x1	1 Host	0b1010	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	1 x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	4 x2	4 x2 (First 8 lanes), 4 x1 1 x4, 1 x2, 1 x1	1 Host	0b1001	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	1 x2	Link 0, Lane 0	Link 0, Lane 1														
RSVD	RSVD	RSVD for future x8 encoding	1 Host	0b1000	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7								
4C	1 x16 Option A	1 x16, 1 x8, 1 x4, 1 x2, 1 x1	1 Host	0b0111	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	1 x16	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7	Link 0, Lane 8	Link 0, Lane 9	Link 0, Lane 10	Link 0, Lane 11	Link 0, Lane 12	Link 0, Lane 13	Link 0, Lane 14	Link 0, Lane 15
4C	2 x8 Option A	2 x8, 2 x4, 2 x2, 2 x1	1 Host	0b0110	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	2 x8	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7	Link 0, Lane 8	Link 0, Lane 9	Link 0, Lane 10	Link 0, Lane 11	Link 0, Lane 12	Link 0, Lane 13	Link 0, Lane 14	Link 0, Lane 15
4C	1 x16 Option B	1 x16, 1 x8, 1 x4, 1 x2, 1 x1 2 x8, 2 x4, 2 x2, 2 x1	1 Host	0b0101	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	1 x16	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7	Link 0, Lane 8	Link 0, Lane 9	Link 0, Lane 10	Link 0, Lane 11	Link 0, Lane 12	Link 0, Lane 13	Link 0, Lane 14	Link 0, Lane 15
4C	1 x16 Option C	1 x16, 1 x8, 1 x4 2 x8, 2 x4, 2 x2, 2 x1	1 Host	0b0100	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	1 x16	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7	Link 0, Lane 8	Link 0, Lane 9	Link 0, Lane 10	Link 0, Lane 11	Link 0, Lane 12	Link 0, Lane 13	Link 0, Lane 14	Link 0, Lane 15
4C	4 x4	4 x4, 4 x2, 4 x1	1 Host	0b0011	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	2 x4*	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7	Link 0, Lane 8	Link 0, Lane 9	Link 0, Lane 10	Link 0, Lane 11	Link 0, Lane 12	Link 0, Lane 13	Link 0, Lane 14	Link 0, Lane 15
RSVD	RSVD	RSVD	1 Host	0b0010	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
RSVD	RSVD	RSVD	1 Host	0b0001	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
RSVD	RSVD	RSVD	1 Host	0b0000	1 Host	1 Upstream Socket	1 or 2 Links	0b0000	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 0, Lane 4	Link 0, Lane 5	Link 0, Lane 6	Link 0, Lane 7	Link 0, Lane 8	Link 0, Lane 9	Link 0, Lane 10	Link 0, Lane 11	Link 0, Lane 12	Link 0, Lane 13	Link 0, Lane 14	Link 0, Lane 15

Table 34: Bifurcation for Single Host, Single Socket and Single/Dual/Quad Upstream Links
(BIF[2:0]#=0b000)

Single Host, Single Upstream Socket, One, Two or Four Upstream Links		Supported Bifurcation Modes		Add-in-Card		Residual Link																
Min Card Width	Card Short Name	Host	Upstream Devices	Upstream Links	BIF[2:0]#	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15	
2C	1x8, 1x4, 1x2, 1x1	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
2C	1x8, 1x4, 1x2, 1x1	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
2C	1x4, 1x2, 1x1	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
2C	1x4	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
2C	1x4	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
2C	1x8 Option B	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
4C	2x8 Option B	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
2C	1x8 Option D	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
4C	1x16 Option D	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
RSVD	RSVD	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
2C	2x4	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
2C	4x2	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
RSVD	RSVD	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
4C	1x16 Option A	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
4C	2x8 Option A	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
4C	1x16 Option B	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
4C	1x16 Option C	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
4C	4x4	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
RSVD	RSVD	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
RSVD	RSVD	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0
RSVD	RSVD	1 Host	1 Upstream Socket	1, 2, or 4 Links	0b000	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0	Link 0

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 90. The max available power envelopes for each of these states are defined in Table 41.

Figure 90: Baseboard Power States




Table 41: 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				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

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. FRU and scan chain accesses are only allowed in this 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 42. 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 42: 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	375mA (max)	375mA (max)	375mA (max)	375mA (max)	375mA (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% (max)	±8% (max)	±8% (max)	±8% (max)	±8% (max)
Supply Current					
ID Mode	100mA (max)	100mA (max)	100mA (max)	100mA (max)	100mA (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)	1000µF (max)	1000µF (max)	2000µF (max)

Note: While cards may draw up to the published power ratings, the baseboard vendor shall evaluate its cooling capacity for each slot power envelope.

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.

The hotswap controller may gate the +12V_EDGE and +3.3V_EDGE based on the PRSNTB[3:0]# value. Per Section 3.6.3, 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 3.5.4 for the NIC_PWR_GOOD definition. Refer to DMTF DSP0222 for details on the NC-SI clock startup requirements.

Figure 91: Power-Up Sequencing





Figure 92: 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_{FAIL} after power is no longer stable.

Table 43: 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.

4 Management and Pre-OS Requirements

OCN NIC 3.0 card management is an important aspect to overall system management. This section specifies a common set of management requirements for OCN 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 OCN NIC 3.0 implementation shall support at least one type of implementation for card management. For a given type of implementation, an OCN NIC 3.0 card shall support type specific requirements described in Sections 4.1 through 4.7.

Table 44: OCN 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 OCN 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

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

Table 45: Sideband Management Interface and Transport Requirements

Requirement	RBT+MCTP Type	RBT Type	MCTP Type
NC-SI 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 46 summarizes the NC-SI traffic requirements.

Table 46: NC-SI Traffic Requirements

Requirement	RBT+MCTP Type	RBT Type	MCTP Type
NC-SI Control over RBT (DMTF DSP0222 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 for Out-Of-Band (OOB) management traffic. The number of MC MAC addresses provisioned is implementation dependent. These MAC addresses are not exposed to the host(s) as available MAC addresses. The MC is not required to use these provisioned MAC addresses. Table 47 summarizes the MC MAC address provisioning requirements.

Table 47: MC MAC Address Provisioning Requirements

Requirement	RBT+MCTP Type	RBT Type	MCTP Type
<p>One or more MAC Addresses shall be provisioned for the MC.</p> <p>The OCP NIC 3.0 platform may use the NIC vendor allocated MAC addresses for the BMC. Each management channel requires a dedicated MAC address. Some platforms may employ multiple BMCs (or virtual BMCs) each with a dedicated MAC address. The NIC may also support multiple partitions on a physical port.</p> <p>The recommended MAC address allocation scheme is stated below.</p>	Required	Required	Optional

<p>Assumptions:</p> <ol style="list-style-type: none"> 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). 2. The maximum number of partitions on each port is the same. <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) <p>Note: This capability is planned to be included in revision 1.2 of the NC-SI specification.</p> <ul style="list-style-type: none"> • NC-SI Control/MCTP (DMTF DSP0261 1.2 compliant) 	Required	Required	Optional

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 48 summarizes temperature reporting requirements. These requirements improve the system thermal management and allow the baseboard management device to access key component temperatures on an OCP NIC 3.0 card. When the temperature reporting function is implemented, it is recommended that the temperature reporting accuracy is within $\pm 3^{\circ}\text{C}$.

Table 48: Temperature Reporting Requirements

Requirement	RBT+MCTP Type	RBT Type	MCTP Type
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. 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 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 recommended that the self-shutdown threshold value is higher than the maximum	Required	Required	Required

<p>junction temperature of the ASIC implementing the NIC function and this value is between the critical and fatal temperature thresholds.</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. An AC power cycle of the system may be required to bring the NIC back to an operational state. In order to recover the NIC from the self-shutdown state, the OCP NIC 3.0 card should go through the NIC power off state as described in Section 3.9.1.</p>			
---	--	--	--

4.5 Power Consumption Reporting

An OCP NIC 3.0 implementation may be able to report the power consumed by one or more component implementing NIC functions. It is important for the system management that the information about the power consumption can be retrieved over sideband interfaces. Table 49 summarizes power consumption reporting requirements.

Table 49: Power Consumption Reporting Requirements

Requirement	RBT+MCTP Type	RBT Type	MCTP Type
Component Estimated Power Consumption Reporting	Required	Required	Required
Component Runtime Power Consumption Reporting	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 50 summarizes pluggable module status reporting requirements.

Table 50: Pluggable Module Status Reporting Requirements

Requirement	RBT+MCTP Type	RBT Type	MCTP Type
Pluggable Transceiver modules Presence Status and Temperature Reporting	Required	Required	Required

PLDM for Platform Monitoring and Control (DSP0248 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 51 summarizes the firmware inventory and update requirements.

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

Requirement	RBT+MCTP Type	RBT Type	MCTP Type
Network boot in UEFI driver (supporting both IPv4 and IPv6 addressing for network boot)	Required	Required	Required
UEFI secure boot for UEFI drivers	Required	Required	Required
UEFI secure firmware update	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 pin on the Primary Connector for this identification. The Slot_ID value may be directly connected to GND (Slot ID = 0), or pulled up to +3.3V_EDGE (Slot ID = 1).

Package ID[2:0] is a 3-bit field and is encoded in the NC-SI Channel ID as bits [7:5]. Package ID[2] defaults to 0b0 in the NC-SI specification, but is optionally configurable if the target silicon supports configuring this bit. Package ID[1] is directly connected to the SLOT_ID pin. Package ID[0] is set to 0b0 for Network Controller ASIC #0. For an OCP NIC 3.0 card with two discrete silicon instances, Package ID[0] shall be set to 0b1 for Network Controller ASIC #1. Refer to the specific endpoint device datasheet for details on the Package ID configuration options.

Up to four silicon devices are supported on the bus if only Package ID[1:0] is configurable (e.g. Package ID[2] is statically set to 0b0). Up to eight silicon devices are supported on the NC-SI bus if Package ID[2:0] are all configurable.

Refer to the DMTF DSP0222 standard for more information on package addressing, Slot ID 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 a one card, or both cards installed. Figure 80 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 and on-board temperature sensors are 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

OCN NIC 3.0 cards shall support SMBus ARP (be ARP-capable) to allow the cards to be dynamically assigned addresses for MCTP communications to avoid address conflicts and eliminate the need for manual configuration of addresses. The address type of dynamic addresses can be either dynamic and persistent address device or 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 OCN 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 OCN NIC 3.0 card, then the OCN 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 OCN NIC 3.0 are shown in Table 52. Baseboard and OCN 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 52: SMBus Address Map

Address (8-bit)	Device	Notes
0xA0 / 0xA1 – SLOT0 0xA2 / 0xA3 – SLOT1	EEPROM	On-board FRU EEPROM. Mandatory. Powered from Aux power domain. The EEPROM ADDR0 pin shall be connected to the SLOT_ID pin on the OCN NIC 3.0 card gold finger to allow up to two OCN NIC 3.0 cards to exist on the same I ² C bus.

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 OCN NIC 3.0 card regardless of the PCIe width or number of physical card edge connectors it occupies. The FRU EEPROM shall be connected to the Primary Connector SMBus.

The EEPROM is addressable at the addresses indicated in Table 52. The write/read pair is presented in 8-bit format. The size of EEPROM shall be at least 4Kbits for the base EEPROM map. OCN NIC 3.0 card suppliers may use a larger size EEPROM if needed to store vendor specific information.

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 v1.2. 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 OCN NIC 3.0. For an OCN NIC 3.0 card, the FRU EEPROM OEM record content based on the format defined in Table 53 shall be populated.

Table 53: 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. 0x00 – 0xFE – Card power rounded up to the nearest Watt for fractional values. 0xFF – Unknown
5	1	Card Max power (in Watts) in AUX(S5) mode. 0x00 – 0xFE – Card power rounded up to the nearest Watt for fractional values. 0xFF – Unknown
6	1	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
7	1	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
8	1	Card active/passive cooling. This bit defines if the card has passive cooling (there is no fan on the card) or active cooling (a fan is located on the card). 0x00 – Passive Cooling 0x01 – Active Cooling 0x02 – 0xFE – Reserved 0xFF – Unknown
9	2	Hot aisle standby airflow requirement. 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. Byte 9 is the LS byte, byte 10 is the MS byte. 0x0000-0xFFFE – LFM required for cooling card in Hot Aisle Operation. 0xFFFF – Unknown.
11	2	Cold aisle standby airflow requirement. 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. Byte 11 is the LS byte, byte 12 is the MS byte. 0x0000-0xFFFE – LFM required for cooling card in Cold Aisle Operation. 0xFFFF – Unknown.
13	1	Temperature Target Max – ASIC 0. 0x00 – 0xFE – The T_{max} value of ASIC 0. The temperature value is in degrees Celsius. 0xFF – Unknown
14	1	Temperature Target Max – ASIC 1. 0x00 – 0xFE – The T_{max} value of ASIC 0. The temperature value is in degrees Celsius. 0xFF – Unknown
15:30	16	Reserved for future use. Set each byte to 0xFF for this version of the specification.
31	1	Number of physical controllers (N). 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. If $N \geq 1$, then the controller UDID records below shall be included for each controller N.
32+16*(N-1):16*N+31	16	Controller N UDID. MS Byte First (to align the FRU order to the reported UDID order on the SMBus). This field is populated for values of $N \geq 1$ for each controller N.

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.

5.2 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.

6 Thermal and Environmental

6.1 Airflow Direction


The OCP NIC 3.0 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 no airflow will exist on the bottom side of the card. The local approach air temperature and speed 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 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 93. 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 93: Airflow Direction for Hot Aisle Cooling



The boundary conditions for Hot Aisle cooling are shown below in Table 54 and Table 55. 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. 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 55 represent the airflow velocities typical in mainstream servers. Higher airflow velocities are available within the Hot Aisle cooling tiers listed in Table 59 but card designers must be sure to understand the system level implications of such high card LFM requirements.

Table 54: 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 55: 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 94. 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 94: Airflow Direction for Cold Aisle Cooling



The boundary conditions for Cold Aisle cooling are shown below in Table 56 and Table 57. The temperature values listed in Table 56 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 56: 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 57: Cold Aisle Airflow Boundary Conditions

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

6.2 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 documentation in Section 6.4.

6.2.1 ASIC Cooling – Hot Aisle

The ASIC or controller chip is typically the highest power consumer on the card. Thus, as OCP NIC 3.0 cards are developed it is important to understand the ASIC cooling capability. Figure 95 below provides an estimate of the maximum ASIC power that can be supported as a function of the local inlet velocity for the small card form factor. Each curve in Figure 95 represents a different local inlet air temperature from 45°C to 65°C.

The curves shown in Figure 95 were obtained using CFD analysis of a reference OCP NIC 3.0 small form factor card. The reference card has a 20mm x 20mm ASIC with two QSFP connectors. Figure 96 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 58. The OCP NIC 3.0 simulation was conducted within a virtual version of the test fixture defined in Section 6.4.

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 95.

Figure 95: ASIC Supportable Power for Hot Aisle Cooling – Small Card Form Factor




Figure 96: OCP NIC 3.0 Reference Geometry CAD & CFD





Table 58: Reference OCP NIC 3.0 Small Card Geometry

OCP NIC 3.0 Form Factor	Small Card
Heatsink Width	65mm
Heatsink Length	54mm
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

It is important to point out that the curves shown in Figure 95 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 97 below shows the 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 97: Server System Airflow Capability – Hot Aisle Cooling



6.2.2 ASIC Cooling – Cold Aisle


Compared to the Hot Aisle cooling 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 Cold Aisle was conducted utilizing the same geometry and boundary conditions described in Figure 96 and Table 58 with airflow moving from I/O connector to ASIC (opposite to the Hot Aisle analysis). Figure 98 below shows the results of this analysis for the Cold Aisle cooling configuration. Each curve in Figure 98 represents a different system inlet air temperature from 25°C to 45°C.

Figure 98: ASIC Supportable Power for Cold Aisle Cooling – Small Card Form Factor




Similar to Figure 97 for Hot Aisle cooling, Figure 99 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 99: Server System Airflow Capability – Cold Aisle Cooling



A comparison of Hot Aisle (55°C) and Cold Aisle (35°C) ASIC cooling capability curves is shown below in Figure 100. 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 100: ASIC Supportable Power Comparison – Small Card Form Factor



6.3 Thermal Simulation (CFD) Modeling

Thermal simulation of OCP NIC 3.0 cards using CFD is recommended. The information that follows includes details of the geometry that should be used for CFD modeling of the OCP NIC 3.0 Small form factor. The geometry described below was developed to ensure consistency across card vendors when analyzing the card cooling and thermal solution. The geometry to be used for CFD analysis is based on the OCP NIC 3.0 thermal test fixture detailed in Section 6.4.

6.3.1 CFD Geometry – Small Card

The geometry to be used for CFD analysis is defined by the following parameters:

- Sheet metal enclosure
- Internal width: 128mm
- Internal height: 40.6mm
- Internal length: 256.7mm
- Fixture Faceplate Open Area Ratio: 25% (as shown in Figure 101)
- Internal height between top side of board and fixture cover: 34.94mm
- OCP Card is centered on the width of the host PCB.
- Inlet temperature boundary condition: desired approach temperature, e.g. 55°C
- Airflow boundary condition: Desired volume flow in the range of 1 to 20 CFM

- OCP NIC 3.0 local velocity monitor:
- Hot Aisle Cooling – monitor plane 25mm upstream from ASIC heatsink
- Cold Aisle Cooling – monitor planes upstream and downstream of ASIC heatsink depending on I/O connector proximity to ASIC heatsink.

CAD step files for the Hot Aisle CFD and Cold Aisle CFD geometry, and CFD thermal models are available for download on the OCP NIC 3.0 Wiki: <http://www.opencompute.org/wiki/Server/Mezz>.

6.3.2 Transceiver Simulation Modeling

The OCP NIC 3.0 subgroup plans to provide transceiver (both optical and active copper) thermal models to aid in simulating card operational conditions in the Hot Aisle and Cold Aisle.

This section is a placeholder and will be updated in a future revision of this specification.

6.4 Thermal Test Fixture – Small Card

Full definition of the thermal test fixture will be included in a future specification release. Images of preliminary design are shown in Figure 101 and Figure 102.

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>.

Figure 101: Small Card Thermal Test Fixture Preliminary Design





Figure 102: Small Card Thermal Test Fixture Preliminary Design – Transparent View



6.5 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 59. 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 59: 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 60. 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 60: 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 Shock & Vibration

This specification does not cover the shock and vibration testing requirements for an OCP NIC 3.0 add in card or its associated baseboard systems. OCP NIC 3.0 components are deployed in various environments. It is up to each OCP NIC 3.0 card and baseboard vendor to decide how the shock and vibration tests shall be done.

6.8 Gold Finger Plating Requirements

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

6.8.1 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.

6.8.2 Line Side Gold Finger Plating Requirements

This section is a placeholder and will be updated in a future revision of the specification.

For the line side connector contact plating, the minimum requirement is to plate 50 microinches of gold over 50 microinches of nickel.

7 Regulatory

7.1 Required Compliance

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

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

7.1.1 Required Environmental Compliance

- China RoHS Directive
- **EU RoHS 2 Directive (2011/65/EU)** aims to reduce the environmental impact of electronic and electrical equipment (EEE) by restricting the use of certain hazardous materials. The substances banned under RoHS are lead, mercury, cadmium, hexavalent chromium, polybrominated biphenyls, polybrominated diphenyl ether, and four phthalates.
- **EU REACH Regulation (EC) No 1907/2006** addresses the production and use of chemical substances and their potential impact on human health and the environment.
- **EU Waste Electrical and Electronic Equipment (“WEEE”) Directive (2012/19/EU)** - mandates the treatment, recovery and recycling of EEE.
- **The Persistent Organic Pollutants Regulation (EC) No. 850/2004** bans production, placing on the market and use of certain persistent organic pollutants.
- **The California Safe Drinking Water and Toxic Enforcement Act of 1986 (“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 61 for details.

Table 61: 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 62.

Table 62: 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 63.

Table 63: Immunity (ESD) Requirements

Targeted Category	Applicable Specifications
Immunity (ESD)	EN 55024 2010, and IEC 61000-4-2 2008 for ESD. Required ± 4 kV contact charge and ± 8 kV air discharge

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.

