



**OPEN**  
Compute Project

## OCP NIC 3.0 Design Specification

Version 0.01

Author: OCP Server Workgroup, OCP NIC subgroup



## Table of Contents

<b>1</b>	<b>Overview</b> .....	<b><u>77</u></b>
1.1	License .....	<u>77</u>
1.2	Background.....	<u>88</u>
1.3	Acknowledgements.....	<u>99</u>
1.4	Overview .....	<u>99</u>
1.4.1	Mechanical Form factor overview.....	<u>99</u>
1.4.2	Electrical overview.....	<u>1212</u>
1.5	References.....	<u>1212</u>
<b>2</b>	<b>Card Form Factor</b> .....	<b><u>1414</u></b>
2.1	Overview .....	<u>1414</u>
2.2	Form Factor Options .....	<u>1414</u>
2.3	I/O bracket.....	<u>1616</u>
2.4	Line Side I/O Implementations .....	<u>1616</u>
2.5	LED Implementation .....	<u>1717</u>
2.6	Mechanical Keepout Zones .....	<u>1818</u>
2.6.1	Baseboard Keep Out Zone.....	<u>1818</u>
2.6.2	Add-in Card Keep Out Zone.....	<u>1818</u>
2.7	Labeling Requirements.....	<u>1919</u>
2.8	Insulation Requirements.....	<u>1919</u>
2.9	NIC Implementation Examples.....	<u>1919</u>
2.10	Non-NIC Use Cases .....	<u>1919</u>
2.10.1	PCIe Retimer card .....	<u>1919</u>
2.10.2	Accelerator card.....	<u>1919</u>
2.10.3	Storage HBA / RAID card .....	<u>1919</u>
<b>3</b>	<b>Card Edge and Baseboard Connector Interface</b> .....	<b><u>1919</u></b>
3.1	Gold Finger Requirement.....	<u>1919</u>
3.2	Baseboard Connector Requirement .....	<u>2121</u>
3.3	Pin definition.....	<u>2222</u>
3.4	Signal Descriptions – Common .....	<u>2928</u>

3.4.1	PCIe Interface Pins .....	<del>292</del> <del>8</del>
3.4.2	PCIe Present and Bifurcation Control Pins.....	<del>323</del> <del>1</del>
3.4.3	SMBus Interface Pins .....	<del>343</del> <del>3</del>
3.4.4	Power Supply Pins .....	<del>353</del> <del>4</del>
3.4.5	Miscellaneous Pins.....	<del>363</del> <del>5</del>
3.5	Signal Descriptions – OCP Bay (Primary Connector) .....	<del>363</del> <del>5</del>
3.5.1	PCIe Interface Pins – OCP Bay (Primary Connector) .....	<del>363</del> <del>5</del>
3.5.2	NC-SI Over RBT Interface Pins – OCP Bay (Primary Connector).....	<del>383</del> <del>7</del>
3.5.3	Scan Chain Pins – OCP Bay (Primary Connector).....	<del>403</del> <del>9</del>
3.5.4	Primary Connector Miscellaneous Pins – OCP Bay (Primary Connector).....	<del>474</del> <del>6</del>
3.6	PCIe Bifurcation Mechanism .....	<del>484</del> <del>7</del>
3.6.1	PCIe Add-in Card to Baseboard Bifurcation Configuration (PRSNTA#, PRSNTB[3:0]#).....	<del>494</del> <del>8</del>
3.6.2	PCIe Baseboard to Add-in Card Bifurcation Configuration (BIF[2:0]#) .....	<del>494</del> <del>8</del>
3.6.3	PCIe Bifurcation Decoder.....	<del>504</del> <del>9</del>
3.6.4	Bifurcation Detection Flow.....	<del>525</del> <del>1</del>
3.6.5	PCIe Bifurcation Examples .....	<del>535</del> <del>2</del>
3.7	PCIe Clocking Topology.....	<del>575</del> <del>6</del>
3.8	PCIe Bifurcation and REFCLK Mapping.....	<del>595</del> <del>8</del>
3.9	Power Capacity and Power Delivery.....	<del>686</del> <del>7</del>
3.9.1	AC Power Off .....	<del>686</del> <del>7</del>
3.9.2	Management (FRU Only Mode).....	<del>686</del> <del>7</del>
3.9.3	Aux Power Mode (S5) .....	<del>696</del> <del>8</del>
3.9.4	Main Power Mode (S0).....	<del>696</del> <del>8</del>
3.10	Power Sequence Timing Requirements .....	<del>696</del> <del>8</del>
<b>4</b>	<b>Management.....</b>	<del>706</del> <del>9</del>
4.1	SMBus Interface .....	<del>706</del> <del>9</del>
4.2	NC-SI Sideband Interface .....	<del>706</del> <del>9</del>
4.2.1	NC-SI addressing and Arb#.....	<del>706</del> <del>9</del>
4.3	MAC Address Requirement.....	<del>706</del> <del>9</del>
4.4	FRU EEPROM.....	<del>706</del> <del>9</del>
4.4.1	Addressing(TBD).....	<del>706</del> <del>9</del>
4.5	FW Requirement (TBD).....	<del>706</del> <del>9</del>



4.6	Thermal Reporting Interface .....	<del>7069</del>
<b>5</b>	<b>Data Network Requirement.....</b>	<del>7170</del>
5.1	Network Booting (collect view from OEMs and hyperscale) .....	<del>7170</del>
<b>6</b>	<b>Routing Guidelines and Signal Integrity Considerations.....</b>	<del>7271</del>
6.1	NC-SI Over RBT.....	<del>7271</del>
<b>7</b>	<b>Thermal and Environmental .....</b>	<del>7271</del>
7.1	Environmental Requirements.....	<del>7271</del>
7.1.1	Thermal Simulation Boundary Example .....	<del>7271</del>
7.2	Shock & Vibration .....	<del>7271</del>
7.3	Regulation .....	<del>7271</del>
<b>8</b>	<b>Revision History .....</b>	<del>7372</del>

## List of Figures

Figure 1: Representative Small OCP 3.0 NIC Card with Quad SFP Ports .....	<del>88</del>
Figure 2: Representative Large OCP 3.0 NIC Card with Dual QSFP Ports and on-board DRAM.....	<del>99</del>
Figure 3: Small and Large Card Form-Factors (not to scale).....	<del>101</del>
Figure 4: Example Small Card Form Factor .....	<del>141</del>
Figure 5: Example Large Card Form Factor.....	<del>151</del>
Figure 6: Primary Connector (4C + OCP Bay) with 4C and 2C Add in Cards.....	<del>161</del>
Figure 7: Small Size Primary Connector Gold Finger Mating Card Dimensions – x16 – Top Side.....	<del>202</del>
Figure 8: Small Size Primary Connector Gold Finger Mating Card Dimensions – x16 – Bottom Side .....	<del>202</del>
Figure 9: Small Size Secondary Connector Gold Finger Mating Card Dimensions – x16 – Top Side.....	<del>212</del>
Figure 10: Small Size Secondary Connector Gold Finger Mating Card Dimensions – x16 – Bottom Side .....	<del>212</del>
Figure 11: Large Size Card Gold Finger Mating Card Dimensions – x32 – Top Side .....	<del>212</del>
Figure 12: Large Size Card Gold Finger Mating Card Dimensions – x32 – Bottom Side.....	<del>212</del>
Figure 13: 168-pin Base Board Primary Connector – Right Angle.....	<del>222</del>
Figure 14: 140-pin Base Board Secondary Connector – Right Angle .....	<del>222</del>
Figure 15: 168-pin Base Board Primary Connector – Straddle Mount.....	<del>222</del>
Figure 16: 140-pin Base Board Secondary Connector – Straddle Mount.....	<del>222</del>
Figure 17: Primary and Secondary Connector Locations for Large Card Support.....	<del>222</del>
Figure 18: PCIe Present and Bifurcation Control Pins.....	<del>333</del>
Figure 19: NC-SI Over RBT Connection Example .....	<del>393</del>
Figure 20: Scan Bus Connection Example .....	<del>464</del>
Figure 21: PCIe Bifurcation Pin Connections Support.....	<del>484</del>
Figure 22: Single Host (1 x16) and 1 x16 Add-in Card (Single Controller) .....	<del>535</del>
Figure 23: Single Host (2 x8) and 2 x8 Add-in Card (Dual Controllers).....	<del>545</del>
Figure 24: Four Hosts (4 x4) and 4 x4 Add-in Card (Single Controller).....	<del>555</del>
Figure 25: Four Hosts (4 x4) and 4 x4 Add-in Card (Four Controllers).....	<del>565</del>
Figure 26: Single Host with no Bifurcation (1 x16) and 2 x8 Add-in Card (Two Controllers).....	<del>575</del>
Figure 27: PCIe Interface Connections for 1 x16 and 2 x8 Add-in Cards.....	<del>585</del>
Figure 28: PCIe Interface Connections for a 4 x4 Add-in Card.....	<del>595</del>
Figure 29: Baseboard Power Sequencing.....	<del>686</del>
Figure 30: Baseboard Power Sequencing.....	<del>696</del>



## List of Tables

Table 1: OCP 3.0 Form Factor Dimensions.....	<u>1010</u>
Table 2: Baseboard to OCP NIC Form factor Compatibility Chart.....	<u>1111</u>
Table 3: OCP 3.0 Card Definitions.....	<u>1616</u>
Table 4: OCP 3.0 Line Side I/O Implementations.....	<u>1716</u>
Table 5: Default LED Configuration.....	<u>1717</u>
Table 6: Primary Connector Card Pin Definition (x16) (4C + OCP Bay).....	<u>2323</u>
Table 7: Secondary Connector Card Pin Definition (x16) (4C).....	<u>2726</u>
Table 8: Card Pin Descriptions – PCIe.....	<u>2928</u>
Table 9: Card Pin Descriptions – PCIe Present and Bifurcation Control Pins.....	<u>3231</u>
Table 10: Card Pin Descriptions – SMBus.....	<u>3433</u>
Table 11: Card Pin Descriptions – Power.....	<u>3534</u>
Table 12: Card Pin Descriptions – Miscellaneous.....	<u>3635</u>
Table 13: Card Pin Descriptions – PCIe.....	<u>3635</u>
Table 14: Card Pin Descriptions – NC-SI Over RBT.....	<u>3837</u>
Table 15: Card Pin Descriptions – Scan Bus.....	<u>4039</u>
Table 16: Card Pin Descriptions – Scan Bus DATA_OUT Bit Definition.....	<u>4241</u>
Table 17: Card Pin Descriptions – Scan Bus DATA_IN Bit Definition.....	<u>4342</u>
Table 18: Card Pin Descriptions – Miscellaneous.....	<u>4746</u>
Table 19: PCIe Bifurcation Decoder for x16 and x8 Card Widths.....	<u>5150</u>
Table 20: PCIe Clock Associations.....	<u>5757</u>
Table 21: Bifurcation for Single Host, Single Upstream Socket and Single Upstream Link (BIF[2:0]#=0b000).....	<u>6059</u>
Table 22: Bifurcation for Single Host, Single Upstream Socket and Single/Dual Upstream Links (BIF[2:0]#=0b000) .....	<u>6160</u>
Table 23: Bifurcation for Single Host, Single Upstream Socket and Single/Dual/Quad Upstream Links (BIF[2:0]#=0b000).....	<u>6261</u>
Table 24: Bifurcation for Single Host, Dual Upstream Sockets and Dual Upstream Links (BIF[2:0]#=0b001).....	<u>6362</u>
Table 25: Bifurcation for Single Host, Four Upstream Sockets and Dual Upstream Links (BIF[2:0]#=0b010).....	<u>6463</u>
Table 26: Bifurcation for Dual Host, Dual Upstream Sockets and Dual Upstream Links (BIF[2:0]#=0b101).....	<u>6564</u>
Table 27: Bifurcation for Quad Host, Quad Upstream Sockets and Quad Upstream Links (BIF[2:0]#=0b110).....	<u>6665</u>
Table 28: Bifurcation for Quad Host, Quad Upstream Sockets and Quad Upstream Links (BIF[2:0]#=0b110).....	<u>6766</u>
Table 29: Power Envelopes.....	<u>6867</u>
Table 30: Power Sequencing Parameters.....	<u>6968</u>

## 1 Overview

### 1.1 License

As of **April 7, 2011**, the following persons or entities have made this Specification available under the Open Web Foundation Final Specification Agreement (OWFa 1.0), which is available at <http://www.openwebfoundation.org/legal/the-owf-1-0-agreements/owfa-1-0>:  
Facebook, Inc.

You can review the signed copies of the Open Web Foundation Agreement Version 1.0 for this Specification at <http://opencompute.org/licensing/>, which may also include additional parties to those listed above.

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 entire risk as to implementing or otherwise using the Specification is assumed by the Specification implementer and user. 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 Background

The OCP NIC 3.0 specification is a follow-on to the OCP 2.0 form-factor for PCIe add-in cards. 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 a Large Card supports up to 32 PCIe lanes on the card edge. Compared to the OCP NIC 2.0 release, the updated specification provides a broader solution space for NIC and system vendors to support the following use cases scenarios:

- NICs with a higher TDP
- Support up to 79W of power delivery to a single connector (Small) card; and 158W to a dual connector (Large) card
- PCIe Gen4 and Gen5 on the system and card
- Support for upto 32 lanes of PCIe per card
- Support for single host and multi-host environments
- Support a greater board area for more complex add-in card designs
- Support for Smart NIC implementations with on-board DRAM
- Simplification of FRU installation and removal while reducing overall down time

A representative Small Card OCP 3.0 NIC mezzanine card is shown in Figure 1 and a representative Large Card is shown in Figure 2

Figure 1: Representative Small OCP 3.0 NIC Card with Quad SFP Ports







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



OCP 3.0 compliant cards are not backwards compatible to the 2.0 cards in order to achieve the features outlined in this specification.

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 website:

[http://www.opencompute.org/wiki/Server/Mezz#Specifications\\_and\\_Designs](http://www.opencompute.org/wiki/Server/Mezz#Specifications_and_Designs)

### 1.3 Acknowledgements

Placeholder

### 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 specification compliant systems and 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 baseboard. The Large form factor card has two connectors (Primary Connector and Secondary Connector) on the baseboard. Both the Primary and Secondary connectors are defined in and compliant to SFF-TA-1002. [Editor's note: plan to submit change back to SFF-TA-1002]. On the NIC side, the card edge is implemented with gold fingers. The gold finger design follows SFF-TA-1002 as well.

Figure 3: Small and Large Card Form-Factors (not to scale)



The two form factor dimensions are shown in Table 1.

Table 1: OCP 3.0 Form Factor Dimensions

Form Factor	Width	Depth	Primary Connector	Secondary Connector	Typical Use Case
Small	W1=76	115mm	4C+ OCP sideband	N/A	Low profile and general NIC in similar profile as OCP

					NIC 2.0; up to x16 PCIe lanes
Large	W2=139 [to be confirmed]	115mm	4C+ OCP sideband 168 pins	4C	Largest PCB width to support feature rich NICs, and up to x32 PCIe lanes

The OCP 3.0 NIC design allows downward compatibility between the two sizes. Table 2 shows the compatibility between the baseboard and NIC combinations. A Large size baseboard slot may accept a small or large sized NIC. A small size baseboard slot may only accept a small sized NIC.

Table 2: Baseboard to OCP NIC Form factor Compatibility Chart

Baseboard Slot Size	NIC Size / Supported PCIe Width	
	Small	Large
Small	Up to x16	Not Supported
Large	Up to x16	Up to x32

There are two baseboard connector 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 base board to use this area for additional baseboard routing. The straddle mount and right angle connectors are shown in Section [3.24.2](#).

For both the baseboard and OCP card, this specification defines the component and routing keep out areas. Refer to Section [2.63.6](#) for details.

Both the straddle mount and right angle implementations shall use the same OCP NIC 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 size and NIC size are a supported combination as shown in Table 2.



This specification defines form factor at NIC module level, including the front panel, latches and card guide features [TBD; pending on the Mechanical work across stakeholders].

More details about the card form-factor is shown in Section [23](#).

#### 1.4.2 Electrical overview

This specification defines electrical interface between baseboard/system and card/NIC module.

The electrical interface is implemented with a right angle or straddle mount connector on baseboard and gold finger on the add-in 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 x16 lanes of PCIe. The Secondary connector in conjunction with the Primary allows Large form-factor implementations and may support up to 32 lanes of PCIe.

##### 1.4.2.1 Primary Connector

##### 1.4.2.2 Secondary connector

## 1.5 References

DMTF Standard. *DSP0222, Network Controller Sideband Interface (NC-SI) Specification*. Distributed Management Task Force, Inc, Rev 1.0.1, January 24<sup>th</sup>, 2013.

Open Compute Project. *OCP NIC Subgroup*. Online.

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

PCIe Base Specification. *PCI Express Base Specification, Revision 4.0 (draft)*.

PCIe CEM Specification. *PCI Express Card Electromechanical Specification, Revision 4.0 (draft)*.

SMBus specification

**Commented [TN1]:** References need to be correctly sited per MLA standards.

Q: How do we handle references to unpublished (draft) specifications?

SNIA. *SFF-TA-1002, Specification for Protocol Agnostic Multi-Lane High Speed Connector.*

SNIA SFF TWG Technology Affiliate, Rev 0.0.9.1, September 9<sup>th</sup>, 2017.



## 2 Card Form Factor

### 2.1 Overview


### 2.2 Form Factor Options

OCP3.0 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 plus a 28-pin bay for OCP 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 add-in cards.

The small card uses the 4C connector for up to a x16 PCIe interface. The small cards implement the 28 pin OCP bay for management functions and support for up to a four PCIe hosts. 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 79W of delivered power to the card edge.


Figure 4: Example Small Card Form Factor



**Commented [TN2]:** Are we going to uprev the SFF-TA spec to include these 28 pins?

The large card uses provides the same functionality as the small card, but with support up to a x32 PCIe interface. The large card utilizes both the Primary and Secondary connectors. The large size card supports higher power envelopes and provides additional board area for more complex designs. The small card form factor supports up to 158W of delivered power to the card edge at 79W per connector.


Figure 5: Example Large Card Form Factor



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



Figure 6: Primary Connector (4C + OCP Bay) with 4C and 2C Add in Cards



**Commented [TN3]:** This needs to be updated to show the OCP bay along with the secondary connector location..

Table 3

Table 3 summarizes the supported card form factors. Small form factor cards support the Primary Connector and up to 16 PCIe lanes. Large form factor cards support both the Primary and Secondary Connectors and up to 32 PCIe lanes.

Table 3: OCP 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	OCP Bay
Small (x8)			2C	OCP Bay
Small (x16)			4C	OCP Bay
Large 1 (x24)		2C	4C	OCP Bay
Large 2 (x32)	4C		4C	OCP Bay

### 2.3 I/O bracket

TBD <need input from OCP mechanical groups>

### 2.4 Line Side I/O Implementations

At the time of this writing, the Small and Large form-factor cards may support the following standard line side I/O implementations:



Table 4: OCP 3.0 Line Side I/O Implementations

Form Factor	Max Topology Connector Count
Small	2x QSFP28
Small	4x SFP28
Small	4x RJ-45
Large	TBD
Large	TBD
Large	TBD

Additional combinations are permissible as I/O technologies and thermal capabilities evolve.

## 2.5 LED Implementation

A small form-factor OCP 3.0 NIC with a fully populated I/O bracket (2x QSFP28, 4xSFP28, or 4xRJ-45), there is insufficient space for on-board 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 [3.5.34.5.3](#).

For small form-factor low I/O count cards (such as 1x QSFP28, 2xSFP28, or 2xRJ-45), or a large form-factor OCP 3.0 NIC, where additional I/O bracket area is available, the card may optionally implement on-board link/activity indications in addition to the Scan Chain LED.

For both cases, the OCP3.0 specification recommends the following LED definitions:

Table 5: Default LED Configuration

LED Pin	LED Color	Description
Link	Green	Active low. Multifunction LED.  When lit and solid, this LED is used to indicate the link is up at the MAC level. Local and Remote Faults are clear and the link is ready for data transmission. When the LED is off, the physical link is down or disabled.

**Commented [TN4]:** The current Link/activity LED definition does not communicate the idea of "linked at max rate" vs "linked at less than max rate"

Do we optionally want to change the LED stream definition such that there is "Speed A (max)" and "Speed B (not max)" then multiplex the activity (via blink) on to the illuminated LED?



		<p>This LED indicator may also be used for port identification through vendor specific link diagnostic software.</p> <p>The link LED shall be located on the left hand side of each port.</p>
Activity	Green	<p>Active low.</p> <p>The Activity LED shall only be illuminated when the Link LED is illuminated.</p> <p>When lit and solid, this LED is used to indicate the port is "idle" and no data is being transmitted or received.</p> <p>When lit and blinking, this LED is used to indicate the port is "active" and data is either being transmitted or received.</p> <p>When the LED is off, no link is detected.</p> <p>The activity LED shall be located on the right hand side of each port.</p>

At the time of this writing, the Scan Chain definition allows for up to one 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 OCP3.0 compliant chassis (e.g. front LED indicators with rear I/O cards).

## 2.6 Mechanical Keepout Zones

### 2.6.1 Baseboard Keep Out Zone

TBD – Need keepout drawings and envelopes for small / large size baseboard including primary/secondary/rail keepouts/cutout for straddle mount/keepout for right angle.

### 2.6.2 Add-in Card Keep Out Zone

TBD – need keepout drawings and envelopes for small / large size NIC including primary/secondary/rail keepouts.

## 2.7 Labeling Requirements

TBD

## 2.8 Insulation Requirements

All cards must implement a secondary side insulator to prevent the bottom side card components from shorting out to the chassis. The recommended insulator thickness is 0.25mm and must reside within the following mechanical envelope for the Small and Large size cards:

TBD <need 2D drawings>

## 2.9 NIC Implementation Examples

TBD

### 2.10 Non-NIC Use Cases

“PCIe interface with extra management sideband”

#### 2.10.1 PCIe Retimer card

#### 2.10.2 Accelerator card

#### 2.10.3 Storage HBA / RAID card

## 3 Card Edge and Baseboard Connector Interface

### 3.1 Gold Finger Requirement

**Editor’s note:** Connector vendors to provide input and all detailed views from the mechanical drawing. First stab at it is below. Diagrams are copied from SFF-TA-1002.

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

Small Size cards may fit in the Primary Connector or the Secondary Connector. Primary Connector compliant cards are 76mm x 115mm and may implement the full 168-pins.



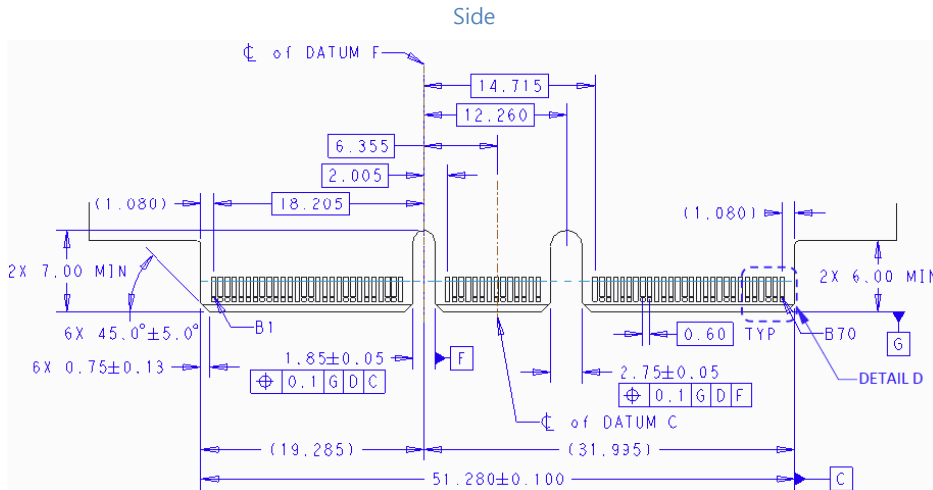
Secondary Connector compliant cards are XXXmm x 115mm and may implement the 140-pin gold finger. Both the Primary and Secondary Connector cards may optionally implement a subset of gold finger pins if there is a reduced PCIe width requirement (such as 1 x8 and below). In this case, the card edge gold finger may implement a 2C design. The overall board thickness is 1.60mm. The gold finger dimensions for the Primary Connector and Secondary Connector compliant cards are shown below.

Large Size Cards support up to a x32 PCIe implementation and uses both the Primary and Secondary connectors.

For additional details, refer to the card and connector mechanical drawings located in XXX.

Note: The "B" pins on the connector are associated with the top side of the add-in card. The "A" pins on the connector are associated with the bottom side of the add-in card.

Figure 7: Small Size Primary Connector Gold Finger Mating Card Dimensions – x16 – Top Side



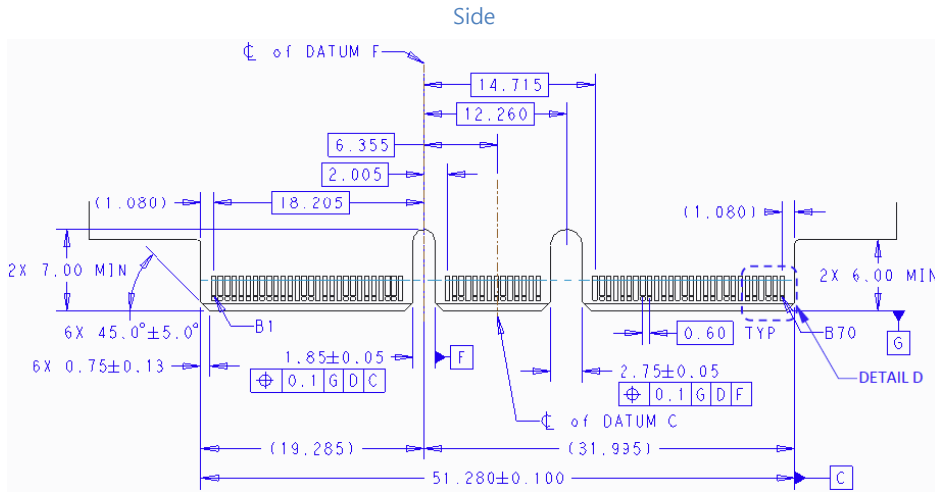
**Commented [TN5]:** Drawing was pulled from SFF-TA-1002 FIGURE 5-18.

Figure 8: Small Size Primary Connector Gold Finger Mating Card Dimensions – x16 – Bottom Side

**Commented [TN6]:** Drawing was pulled from SFF-TA-1002 FIGURE 5-18.

TBD

Figure 9: Small Size Secondary Connector Gold Finger Mating Card Dimensions – x16 – Top



Commented [TN7]: Drawing was pulled from SFF-TA-1002 FIGURE 5-18.

Figure 10: Small Size Secondary Connector Gold Finger Mating Card Dimensions – x16 – Bottom Side

Bottom Side

TBD

Figure 11: Large Size Card Gold Finger Mating Card Dimensions – x32 – Top Side

TBD

Figure 12: Large Size Card Gold Finger Mating Card Dimensions – x32 – Bottom Side

TBD

Commented [TN8]: Drawing was pulled from SFF-TA-1002 FIGURE 5-18.

Commented [TN9]: Drawing was pulled from SFF-TA-1002 FIGURE 5-18.

Commented [TN10]: Drawing was pulled from SFF-TA-1002 FIGURE 5-18.

### 3.2 Baseboard Connector Requirement

**Editor's note:** Connector vendors to provide input.


The OCP3.0 connector is compliant to the "4C connector" as defined in the SFF-TA-1002 specification for a right angle or straddle mount form-factor. The 4C connector is 140-pins in width and includes support for up to 32 differential pairs to support a x16 PCIe

Commented [TN11]: Note: Perhaps the SFF-TA-1002 spec needs to be updated with the 168 pin and straddle-mount definitions.




connection. The connector also provides 6 pins of 12V for payload power. This implementation is common between both the Primary and Secondary Connector. In addition, the Primary Connector has a 28-pin OCP Bay to the right of pin 1. These pins are used for management and support for up to a 4 x4 multi-host configuration on the Primary Connector. The Primary and Secondary Connector drawings are shown in Figure 13 and Figure 14, below.

Figure 13: 168-pin Base Board Primary Connector – Right Angle



Commented [TN12]: Get updated drawing from TE with the straddle mount option.

Figure 14: 140-pin Base Board Secondary Connector – Right Angle



Commented [TN13]: Get updated drawing from TE with the straddle mount option.

Figure 15: 168-pin Base Board Primary Connector – Straddle Mount

TBD

Commented [TN14]: Get updated drawing from TE with the straddle mount option.

Figure 16: 140-pin Base Board Secondary Connector – Straddle Mount

TBD

Commented [TN15]: Get updated drawing from TE with the straddle mount option.

In order to support the large form factor, systems must locate the Primary and Secondary Connectors per the mechanical drawing shown in Figure 17.

Figure 17: Primary and Secondary Connector Locations for Large Card Support

TBD

### 3.3 Pin definition

**Editor's note:** The pin map aligns with OCP 3.0 Pinout Proposal 20171121a\_TN.xlsx.

The pin definitions of an OCP NIC 3.0 card with up to a x32 PCIe interface are shown in Table 6 and Table 7. 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 add-in cards. Both connectors share common functionality with power, SMBus, x16 PCIe Gen4 and bifurcation control. The Primary Connector 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 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 Connector are shown in Section [3.44.4](#). The OCP Bay for the Primary Connector only are shown in Section [3.54.5](#).

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 [3.14.1](#) and [3.24.2](#) 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 + OCP Bay) and Secondary (4C) connectors.

Table 6: Primary Connector Card Pin Definition (x16) (4C + OCP Bay)

Side B		Side A		Primary Connector (x16, 168-pin add-in card with OCP Bay)	Primary Connector (x8, 112-pin add-in card with OCP bay)
OCP_B1	NIC_PWR_GOOD	WAKE_N	OCP_A1		
OCP_B2	PWRBRK#	PERST2#	OCP_A2		
OCP_B3	LD#	PERST3#	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_ID0	RBT_TX_EN	OCP_A7		
OCP_B8	RBT_RXD1	RBT_TXD1	OCP_A8		
OCP_B9	RBT_RXD0	RBT_TXD0	OCP_A9		
OCP_B10	GND	GND	OCP_A10		



OCP_B11	REFCLKn2	REFCLKn3	OCP_A11
OCP_B12	REFCLKp2	REFCLKp3	OCP_A12
OCP_B13	GND	GND	OCP_A13
OCP_B14	RBT_CRS_DV	RBT_CLK_IN	OCP_A14
<b>Mechanical Key</b>			
B1	+12V/+12V_AUX	GND	A1
B2	+12V/+12V_AUX	GND	A2
B3	+12V/+12V_AUX	GND	A3
B4	+12V/+12V_AUX	GND	A4
B5	+12V/+12V_AUX	GND	A5
B6	+12V/+12V_AUX	GND	A6
B7	BIF0#	SMCLK	A7
B8	BIF1#	SMDAT	A8
B9	BIF2#	SMRST#	A9
B10	PERST0#	PRSNTA#	A10
B11	+3.3V/+3.3V_AUX	PERST1#	A11
B12	PWRDIS	PRSNB2#	A12
B13	GND	GND	A13
B14	REFCLKn0	REFCLKn1	A14
B15	REFCLKp0	REFCLKp1	A15
B16	GND	GND	A16
B17	PETn0	PERn0	A17
B18	PETp0	PERp0	A18
B19	GND	GND	A19
B20	PETn1	PERn1	A20
B21	PETp1	PERp1	A21
B22	GND	GND	A22
B23	PETn2	PERn2	A23
B24	PETp2	PERp2	A24
B25	GND	GND	A25
B26	PETn3	PERn3	A26
B27	PETp3	PERp3	A27
B28	GND	GND	A28
<b>Mechanical Key</b>			
B29	GND	GND	A29
B30	PETn4	PERn4	A30
B31	PETp4	PERp4	A31
B32	GND	GND	A32
B33	PETn5	PERn5	A33

**Commented [TN16]:** Jia brought up a good point about PRSNTA#

Unlike PCIe, the PRSNTA / PRSNB indication doesn't necessarily help us with an x-axis alignment. The Present signals may still connect.

Perhaps we could connect PRSNB pins to GND instead of PRSNTA? This would free up a pin for use. (As a bonus, pin A10 is a bidirectional pin (based on function) per EDSFF).



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
<b>Mechanical Key</b>			
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	RFU, N/C	RFU, N/C	A68
B69	RFU, N/C	RFU, N/C	A69
B70	PRSNTB3#	RFU, N/C	A70




Table 7: Secondary Connector Card Pin Definition (x16) (4C)

Side B		Side A	
B1	+12V/+12V_AUX	GND	A1
B2	+12V/+12V_AUX	GND	A2
B3	+12V/+12V_AUX	GND	A3
B4	+12V/+12V_AUX	GND	A4
B5	+12V/+12V_AUX	GND	A5
B6	+12V/+12V_AUX	GND	A6
B7	BIF0#	SMCLK	A7
B8	BIF1#	SMDAT	A8
B9	BIF2#	SMRST#	A9
B10	PERST0#	PRSNTA#	A10
B11	+3.3V/+3.3V_AUX	PERST1#	A11
B12	PWRDIS	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

Secondary Connector (x16, 140-pin add-in card)

Secondary Connector (x8, 84-pin add-in card)



B37	PETp6	PERp6	A37
B38	GND	GND	A38
B39	PETn7	PERn7	A39
B40	PETp7	PERp7	A40
B41	GND	GND	A41
B42	PRSNB0#	PRSNB1#	A42
<b>Mechanical Key</b>			
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	RFU, N/C	RFU, N/C	A68
B69	RFU, N/C	RFU, N/C	A69
B70	PRSNB3#	RFU, 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.

**Note:** Pins that are only used on Primary Connector 28-pin OCP bay are defined in Section [3.54-5](#).

#### 3.4.1 PCIe Interface Pins

This section provides the pin assignments for the PCIe interface signals. The AC/DC specifications are defined in Section XXX. Example connection diagrams for are shown in Figure 28.

Table 8: Card Pin Descriptions – PCIe

Signal Name	Pin #	Baseboard Direction	Signal Description
REFCLKn0 REFCLKp0	B14 B15	Output	PCIe compliant differential reference clock #0, and #1. 100MHz HCSL reference clocks are used for the add-in card PCIe core logic.  <b>Note:</b> 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.  Refer to Section 2.1 in the PCIe CEM Specification, Rev 4.0 for electrical details.
REFCLKn1 REFCLKp1	A14 A15	Output	
PETn0 PETp0	B17 B18	Output	Transmitter differential pair [0:15]. These pins are connected from the baseboard transmitter differential pairs to the receiver differential pairs on the add-in card.
PETn1 PETp1	B20 B21	Output	
PETn2 PETp2	B23 B24	Output	



PETn3 PETp3	B26 B27	Output	<p>The PCIe Transmit pins are AC coupled on the baseboard with capacitors and are placed next to the baseboard transmitters. The AC coupling capacitor must be between 176nF (min) and 265nF (max).</p> <p>Refer to Section 6.1 in the PCIe CEM Specification, Rev 4.0 for details.</p>	
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 pair [0:15]. These pins are connected from the add-in card transmitter differential pairs to the receiver differential pairs on the baseboard.</p>
PERn1 PERp1	A20 A21	Input		
PERn2 PERp2	A23 A24	Input		

PERn3 PERp3	A26 A27	Input	The PCIe Receive pins are AC coupled on the add-in card with capacitors and are placed next to the add-in card transmitters. The AC coupling capacitor must be between 176nF (min) and 265nF (max).  Refer to Section 6.1 in the PCIe CEM Specification, Rev 4.0 for details.	
PERn4 PERp4	A30 A31	Input		
PERn5 PERp5	A33 A34	Input		
PERn6 PERp6	A36 A37	Input		
PERn7 PERp7	A39 A40	Input		
PERn8 PERp8	A44 A45	Input		
PERn9 PERp9	A47 A48	Input		
PERn10 PERp10	A50 A51	Input		
PERn11 PERp11	A53 A54	Input		
PERn12 PERp12	A56 A57	Input		
PERn13 PERp13	A59 A60	Input		
PERn14 PERp14	A62 A63	Input		
PERn15 PERp15	A65 A66	Input		
PERST0# PERST1#	B10 A11	Output		PCIe Reset #0, #1. Active low.  Indicates when the applied power is within tolerance and stable for the add-in card. PERST# goes high after 100ms per the PCI CEM Specification when the power rails are within



			<p>operating limits. The PCIe REFCLKs also become stable within this period of time. PERST is pulled high on the baseboard.</p> <p><b>Note:</b> 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>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 XXX. An example connection diagram is shown in Figure 18.

The PRSNTA#/PRSNTB[0:3]#/BIF[0:2]# pins must be latched within **TBD ms** of the system AC power on to ensure the correct values are latched by the system. Changing the pin states after this timing window is not allowed. Refer to the AC timing diagram in Section XXX for details.

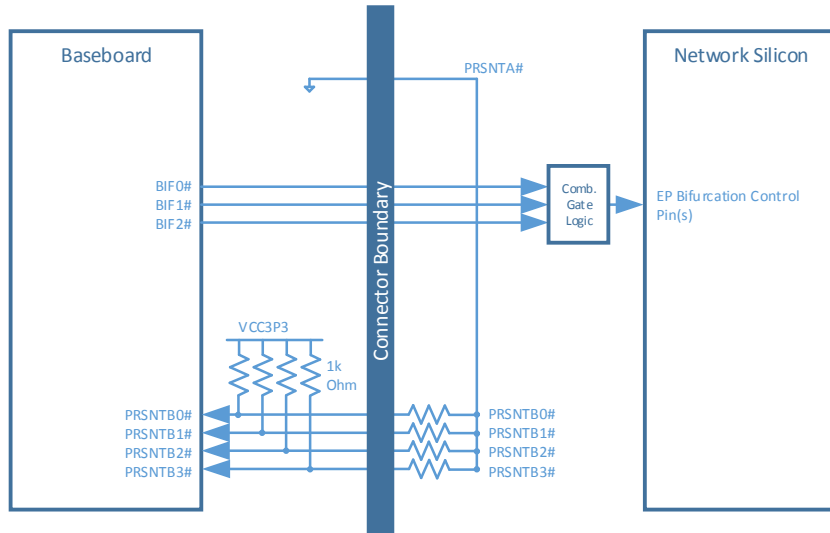
Table 9: Card Pin Descriptions – PCIe Present and Bifurcation Control Pins

Signal Name	Pin #	Baseboard Direction	Signal Description
PRSNTA#	A12	Output	Present A is used for card presence and add-in card PCIe capabilities detection. This pin is connected to GND on the baseboard. This pin is connected to the Present B pins on the add-in card.



PRSNTB0# PRSNTB1# PRSNTB2# PRSNTB3#	B42 A42 A10 B70	Input	<p>Present B [0:3]# are used for card presence detection and PCIe capabilities detection.</p> <p>For baseboards, these pins are connected to the I/O hub and are pulled up to +3.3V using 1kOhm resistors.</p> <p>For add-in cards, these pins are strapped to PRSNTA#. The encoding definitions are described in Section <a href="#">3.64.6</a>.</p> <p>PRSNTB3# is located at the bottom of the 4C connector and is only applicable for add-in cards with a PCIe width of x16 (or greater). Add-in cards that implement a 2C card edge do not use the PRSNTB3# pin for capabilities or present detection.</p>
BIF0# BIF1# BIF2#	A7 A8 A9	Output	<p>Bifurcation [0:2]# are outputs driven from the baseboard I/O hub and allows the system to force configure the add-in card bifurcation.</p> <p>The BIF[0:2]# encoding definitions are described in Section <a href="#">3.64.6</a>.</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 18: 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 Section XXX. An example connection diagram is shown in Figure XXX.

Table 10: Card Pin Descriptions – SMBus

Signal Name	Pin #	Baseboard Direction	Signal Description
SMCLK	A7	Output	SMBus clock. Open drain, pulled up to +3.3V on the baseboard.
SMDAT	A8	Input / Output	SMBus Data. Open drain, pulled up to +3.3V on the baseboard.
SMRST#	A9	Output	SMBus reset. Open drain, pulled up to +3.3V on the baseboard. Used to reset optional downstream SMBus devices (such as temperature sensors).

### 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 Section XXX. An example connection diagram is shown in Figure XXX.

Table 11: Card 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.
+12V/+12V_AUX	B1, B2, B3, B4, B5, B6	Power	+12V main or 12V Aux power; total of 6 pins per connector. The 12V pins are rated to 1.1A per pin with a maximum derated power delivery of 79.2W.  The +12V power pins must be within the rail tolerances when the PWRDIS pin is driven low by the baseboard.
+3.3V/3.3V_AUX	B11	Power	+3.3V main or +3.3V Aux power; total of 1 pin per connector. The 3.3V pin is rated to 1.1A for a maximum derated power delivery of 3.63W.  The +3.3V power pin must be within the rail tolerances when the PWRDIS pin is driven low by the baseboard.
PWRDIS	B12	Output	Power disable. Active high.  This signal is driven by the baseboard.  When high, this signal notifies the add-in card to turn off all systems connected to +12V power.



			When low, this signal notifies the add-in card to enable the on-card power supplies.
--	--	--	--

### 3.4.5 Miscellaneous Pins

This section provides the pin assignments for the miscellaneous interface signals. The AC/DC specifications are defined in Section XXX.

Table 12: Card Pin Descriptions – Miscellaneous

Signal Name	Pin #	Baseboard Direction	Signal Description
RFU, N/C	B68, B69, A68, A69, A70	Input / Output	Reserved future use pin. Leave these pins as no connect.

## 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 [3.34.3](#) 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.

**Note:** The pins that are common to both the Primary and Secondary Connectors are defined in Section [3.44.4](#).

### 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 Section XXX. An example connection diagram that shows REFCLK2, REFCLK3, PERST2# and PERST3# is shown in Section [3.74.7](#).

Table 13: Card Pin Descriptions – PCIe

Signal Name	Pin #	Baseboard Direction	Signal Description
REFCLKn2 REFCLKp2	OCP_B11 OCP_B12	Output	<p>PCIe compliant differential reference clock #2, and #3. 100MHz HCSL reference clocks are used for the add-in card PCIe core logic.</p> <p><b>Note:</b> REFCLK2 and REFCLK3 are not used for cards that only support a 1 x16 or 2 x8 connection.</p> <p>Refer to Section 2.1 in the PCIe CEM Specification, Rev 4.0 for details.</p>
REFCLKn3 REFCLKp3	OCP_A11 OCP_A12	Output	
PERST2# PERST3#	OCP_A2 OCP_A3	Output	<p>PCIe Reset #2, #3. Active low.</p> <p>Indicates when the applied power is within tolerance and stable for the add-in card. PERST# goes high after 100ms per the PCI CEM Specification when the power rails are within operating limits. The PCIe REFCLKs also become stable within this period of time. PERST is pulled high on the baseboard.</p> <p><b>Note:</b> PERST2# and PERST3# are not used for cards that only support a 1 x16 or 2 x8 connection.</p> <p>Refer to Section 2.2 in the PCIe CEM Specification, Rev 4.0 for details.</p>
WAKE#	OCP_A1	Input	<p>WAKE#. Active low. This signal is pulled up to +3.3V on the baseboard with a 10kOhm resistor.</p>



			<p>This signal is driven by the add-in card to notify the baseboard restore the PCIe link. For add-in 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>Refer to Section 2.3 in the PCIe CEM Specification, Rev 4.0 for details.</p>
--	--	--	---

### 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 are defined in Section XXX. An example connection diagram is shown in Figure 19.

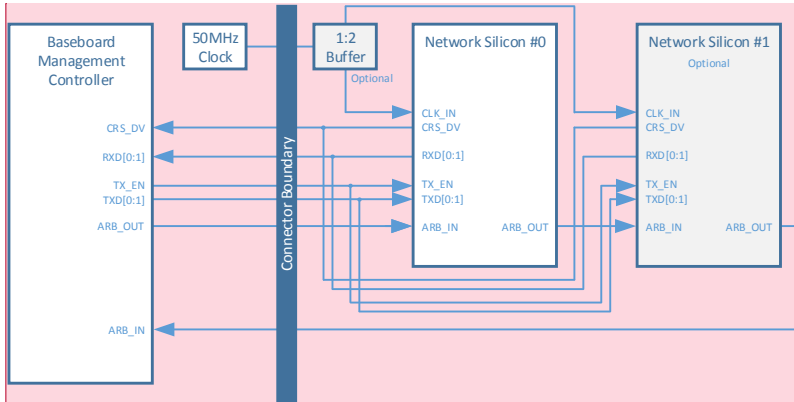
Refer to the NC-SI Specification for implementation and timing details. For the purposes of this specification, the min and max length of the NC-SI signals shall be between 2 inches and 4 inches. The traces shall be implemented as impedance controlled 50 Ohm nets.

Table 14: Card Pin Descriptions – NC-SI Over RBT

Signal Name	Pin #	Baseboard Direction	Signal Description
RBT_CLK_IN	OCP_A14	Output	Reference clock input. Synchronous clock reference for receive, transmit and control interface. The clock has a nominal frequency of 50MHz ±100ppm.
RBT_CRS_DV	OCP_B14	Input	Carrier sense/receive data valid. Signal is used to indicate to the baseboard that the carrier sense/receive data is valid.
RBT_RXD0 RBT_RXD1	OCP_B9 OCP_B8	Input	Receive data. Data signals from the network controller to the BMC.
RBT_TX_EN	OCP_A7	Output	Transmit enable.

RBT_TXD0 RBT_TXD1	OCP_A9 OCP_A8	Output	Transmit data. Data signals from the BMC to the network controller.
RBT_ARB_OUT	OCP_A5	Output	NC-SI hardware arbitration output. Used only if the end point silicon supports hardware arbitration. Connects to the ARB_IN signal of an adjacent device.  The ARB_IN pin is also routed to the card edge to allow multiple devices and OCP slots on the baseboard to share the NC-SI ring.
RBT_ARB_IN	OCP_A4	Input	NC-SI hardware arbitration input. Used only if the end point silicon supports hardware arbitration. Connects to the ARB_OUT signal of an adjacent device.  The ARB_OUT pin is also routed to the card edge to allow multiple devices and OCP slots on the baseboard to share the NC-SI ring.
SLOT_ID0	OCP_B7	Output	NC-SI Address pin. Used only if the end point silicon supports package identification. N/C on NIC if not supported.  Tie to GND for Slot ID = 0 Tie to +3.3Vaux for Slot ID = 1

Figure 19: NC-SI Over RBT Connection Example



**Commented [TN17]:** Update this diagram. The ARB\_IN and ARB\_OUT on the host side should either: 1) connect to each other to complete the ring, or 2) connect to an a second Primary Connector slot.

Also, show an implementation that has two Primary Connectors with RBT connected.

### 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 AC/DC specifications are defined in Section XXX. An example connection diagram is shown in Figure 20.

Table 15: Card Pin Descriptions – Scan Bus

Signal Name	Pin #	Baseboard Direction	Signal Description
CLK	OCP_B6	Output	Scan clock. The CLK is an output pin from the baseboard to the add-in card. The CLK may run up to 12.5MHz.  For baseboard implementations, tie the CLK pin directly to GND if the scan chain is not used.  For NIC implementations, the CLK pin must be connected to Shift Registers 0 & 1 as defined in the text and Figure 20, below.
DATA_OUT	OCP_B5	Output	Scan clock data output from the baseboard to the add-in card. This bit stream is used to shift in NIC configuration data.



			<p>For baseboard implementations, tie the DATA_OUT pin 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 on the add-in card.</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 3.3Vaux through a 10kOhm resistor to prevent the signal from floating if a card is not installed.</p> <p>For NIC implementations, the DATA_IN scan chain is required. Shift Registers 0 &amp; 1, as defined in the text and Figure 20 below are required.</p>
LD#	OCP_B3	Output	<p>Scan clock shift register load. Used to latch configuration data on the add-in card.</p> <p>For baseboard implementations, the LD# pin shall be pulled up to 3.3Vaux through a 10kOhm resistor if the scan chain is not used</p> <p>For NIC implementations, the LD# pin must be connected to Shift Registers 0 &amp; 1 as defined in the text and Figure 20, below.</p>

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 provides side band status indication between the add-in card and the baseboard.

The DATA\_OUT bus is an output from the host. The DATA\_OUT bus provides initial configuration options to the add-in 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.

Table 16: Card Pin Descriptions – Scan Bus 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. The DATA\_IN bus provides NIC status indication to the host. 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 are mandatory for all cards. DATA\_IN shift registers 2 & 3 are optional depending on the card type. DATA\_IN shift register 2 may be used for future revisions of the scan chain. DATA\_IN shift registers 2 & 3 are required for card implementations with 5-8 ports.

The host should read the DATA\_IN bus multiple (TBD) times to qualify the incoming data stream.

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 17: Card 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]# value mirrored from 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 mirrored from the Primary Connector.
0.5	TEMP_WARN	0b0	Temperature monitoring pin from on-card thermal solution. Asserted high when temperature sensor exceeds the warning threshold.
0.6	TEMP_CRIT	0b0	Temperature monitoring pin from on-card thermal solution. Asserted high when temperature sensor exceeds the critical threshold.
0.7	FAN_ON_AUX	0b0	When high, FAN_ON_AUX requests the system fan to be enabled for extra cooling when the card is in the S5 state.
1.0	LINK0	0b1	Port 0..3 link indication. 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. Off = no link is detected on the port.
1.1	LINK1	0b1	
1.2	LINK2	0b1	
1.3	LINK3	0b1	

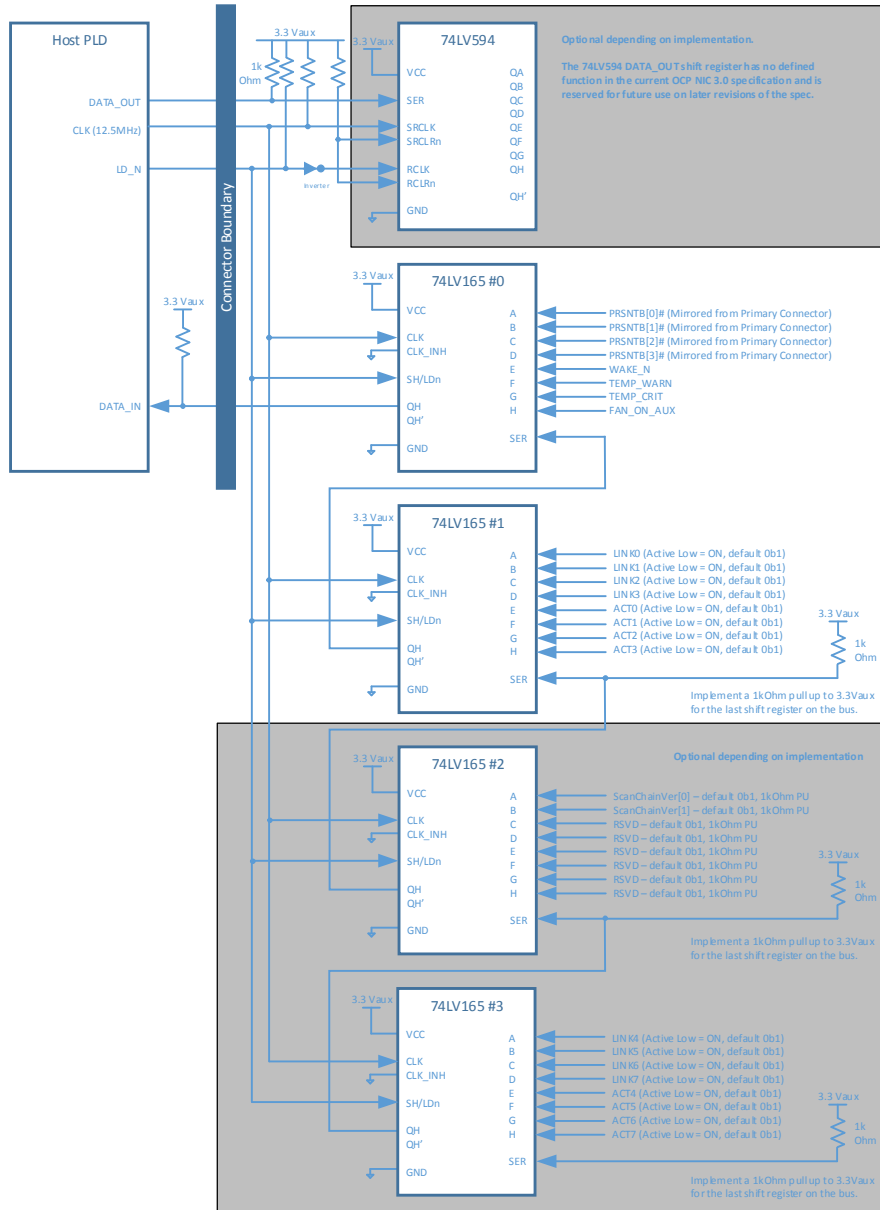


1.4	ACT0	0b1	<p>Port 0..3 activity indication. Active low.</p> <p>0b0 – Link LED is illuminated on the host platform. 0b1 – Link LED is not illuminated on the host platform.</p> <p>Steady = no activity is detected on the port Blink = activity is detected on the port. Off = no link, see also LINK[3:0] LED bits.</p> <p>The LED blink duty cycle is dependent on the add-in card implementation <b>TBD</b>.</p>
1.5	ACT1	0b1	
1.6	ACT2	0b1	
1.7	ACT3	0b1	
2.0	ScanChainVer[0]	0b1	<p>ScanChainVer[1:0] is used to indicate the scan chain bit definitions. The encoding is as follows:</p> <p>0b11 – Scan chain bit definitions version 1 corresponding to OCP 3.0 spec version 1.0.</p> <p>All other encodings are reserved.</p>
2.1	ScanChainVer[1]	0b1	
2.2	RSVD	0b1	<p>Byte 2 bits [2:7] are reserved. These bits shall default to the value of 0b1.</p>
2.3	RSVD	0b1	
2.4	RSVD	0b1	
2.5	RSVD	0b1	
2.6	RSVD	0b1	
2.7	RSVD	0b1	
3.0	LINK4	0b1	<p>Port 4..7 link indication. Active low.</p> <p>0b0 – Link LED is illuminated on the host platform. 0b1 – Link LED is not illuminated on the host platform.</p>
3.1	LINK5	0b1	
3.2	LINK6	0b1	
3.3	LINK7	0b1	

			Steady = link is detected on the port. Off = no link is detected on the port.
3.4	ACT4	0b1	Port 4..7 activity indication. Active low.  0b0 – Link LED is illuminated on the host platform. 0b1 – Link LED is not illuminated on the host platform.  Steady = no activity is detected on the port Blink = activity is detected on the port. Off = no link, see also LINK[3:0] LED bits.
3.5	ACT5	0b1	
3.6	ACT6	0b1	
3.7	ACT7	0b1	



Figure 20: 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 Section XXX. An example connection diagram is shown in Figure XXX.

Table 18: Card Pin Descriptions – Miscellaneous

Signal Name	Pin #	Baseboard Direction	Signal Description
PWRBRK#	OCP_B2	Output	<p>Power break. Active low, open drain.</p> <p>This signal is pulled up to +3.3V on the add-in card with a minimum of 95kOhm and the baseboard with a stiffer resistance in-order to meet the timing specs as shown in CEM.</p> <p>This signal is driven low by the baseboard and is used to notify that an Emergency Power Reduction State is requested.</p>
NIC_PWR_GOOD	OCP_B1	Input	<p>NIC power good. Active high. This signal is driven by the add-in card.</p> <p>When high, this signal indicates that all of the add-in card power rails are operating within nominal tolerances.</p> <p>When low the add-in card power supplies are not yet ready or are in a fault condition.</p> <p>Add-in cards may implemented a cascaded power good output or use a discrete power good monitor on the card. This signal is pulled down to ground with a 100kOhm resistor on the baseboard to prevent a false</p>



			power good indication if no add-in card is present.
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


OCP3.0 baseboards and add-in 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 cases are supported. The bifurcation mechanism also supports add-in 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 connects to the PRSNTB# pins as a hard coded value on the add-in card. The encoding of the PRSNTB[3:0]# pins allows the baseboard to determine the PCIe Links available on the add-in card.
- BIF[3:0]#. The BIF# pin states are controlled by the baseboard and allows 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 high level bifurcation connections are shown in Figure 21.

Figure 21: PCIe Bifurcation Pin Connections Support





### 3.6.1 PCIe Add-in Card to Baseboard Bifurcation Configuration (PRSNTA#, PRSNTB[3:0]#)

The add-in card to baseboard configuration mechanism consists of four dual use pins (PRSNTB[3:0]#) on the add-in 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 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 add-in 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 19 for x16 and x8 PCIe cards.

### 3.6.2 PCIe Baseboard to Add-in Card Bifurcation Configuration (BIF[2:0]#)

Three signals (BIF[2:0]#) are driven by the baseboard to notify requested bifurcation on the add-in card silicon. This allows the baseboard to set the lane configuration on the add-in 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 19 and indicate to the add-in 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 state combination of each of the PRSNTB[3:0]# and BIF[2:0]# pins deterministically sets the PCIe lane width for a given combination of baseboard and add-in cards. Table 19 shows the resulting number of PCIe links and its width for known combinations of baseboards and add-in 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 add-in 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 19: PCIe Bifurcation Decoder for x16 and x8 Card Widths

Minimum Required Card Capabilities	Card Short	Card Short	Host	1 Host	1 Host	Single Host	1 Host	1 Host	1 Host	RSVD	RSVD	Downstream	Downstream
PCIe	Not Present	Not Present	Host CPU Sockets	1 Host	1 Host	1 Host	1 Host	1 Host	1 Host	RSVD	RSVD	4 Hosts	4 or 8 Hosts
PCIe	Not Present	Not Present	Total PCIe Links	1 Host	1 Host	1 Host	1 Host	1 Host	1 Host	RSVD	RSVD	4 Upstream Sockets (1 Socket per Host)	4 or 8 Upstream Sockets (1 Socket per Host)
PCIe	Not Present	Not Present	System Support	1 Link	1 or 2 Links	1, 2, or 4 Links	2 Links	4 Links	4 Links	RSVD	RSVD	4 Links	4 or 8 Links
PCIe	Not Present	Not Present	System Encodings	0x000	0x000	0x000	0x001	0x010	0x011	0x001	0x000	0x010	0x011
PCIe	Not Present	Not Present	System Encodings	0x000	0x000	0x000	0x001	0x010	0x011	0x001	0x000	0x010	0x011
PCIe	Not Present	Not Present	System Encodings	0x000	0x000	0x000	0x001	0x010	0x011	0x001	0x000	0x010	0x011
2C	1x8	1x8, 1x4, 1x2, 1x1	0x110	1x8	1x8	1x8	1x8 (Socket 0 only)	1x4 (Socket 0 only)	1x4 (Socket 0 only)	1x8	1x8 (Host 0 only)	1x4 (Host 0 only)	1x2 (Host 0 only)
2C	1x4	1x4, 1x2, 1x1	0x110	1x4	1x4	1x4	1x4 (Socket 0 only)	1x2 (Socket 0 only)	1x2 (Socket 0 only)	1x4	1x4 (Host 0 only)	1x2 (Host 0 only)	1x2 (Host 0 only)
2C	1x2	1x2, 1x1	0x110	1x2	1x2	1x2	1x2 (Socket 0 only)	1x1 (Socket 0 only)	1x1 (Socket 0 only)	1x2	1x2 (Host 0 only)	1x1 (Host 0 only)	1x1 (Host 0 only)
2C	1x1	1x1	0x110	1x1	1x1	1x1	1x1 (Socket 0 only)	1x1 (Socket 0 only)	1x1 (Socket 0 only)	1x1	1x1 (Host 0 only)	1x1 (Host 0 only)	1x1 (Host 0 only)
2C	1x8 Option B	1x8, 1x4, 1x2, 1x1	0x101	1x8	1x8	1x8	1x8 (Socket 0 only)	2x4 (Socket 0 only)	2x4 (Socket 0 only)	1x8	1x8 (Host 0 only)	2x2 (Host 0.5, 1 only)	2x2 (Host 0.5, 1 only)
2C	2x8 Option B	4x4, 4x2, 4x1	0x101	2x8	2x8	2x8	2x8 (Socket 0 only)	2x4 (Socket 0 only)	2x4 (Socket 0 only)	2x8	2x8 (Host 0 only)	4x4 (Host 0.5, 1 only)	4x4 (Host 0.5, 1 only)
2C	1x8 Option D	4x2 (First 8 lanes), 4x1	0x100	1x8	2x4	2x4	1x8 (Socket 0 only)	2x4 (Socket 0 only)	2x4 (Socket 0 only)	1x8	1x8 (Host 0 only)	2x4 (Host 0.5, 1 only)	2x4 (Host 0.5, 1 only)
2C	1x8 Option D	4x2 (First 8 lanes), 4x1	0x100	1x8	1x8	1x8	1x8 (Socket 0 only)	2x4 (Socket 0 only)	2x4 (Socket 0 only)	1x8	1x8 (Host 0 only)	2x4 (Host 0.5, 1 only)	2x4 (Host 0.5, 1 only)
2C	2x4	2x4, 2x2, 2x1	0x101	1x4	1x4	1x4	1x4 (Socket 0 only)	2x4 (Socket 0 only)	2x4 (Socket 0 only)	1x4	1x4 (Host 0 only)	2x4 (Host 0.5, 1 only)	2x4 (Host 0.5, 1 only)
2C	1x8	1x8, 1x4, 1x2, 1x1	0x011	1x8	1x8	1x8	1x8 (Socket 0 only)	1x4 (Socket 0 only)	1x4 (Socket 0 only)	1x8	1x8 (Host 0 only)	1x4 (Host 0 only)	1x2 (Host 0 only)
2C	2x8 Option A	1x8, 1x4, 1x2, 1x1	0x011	1x8	1x8	1x8	1x8 (Socket 0 only)	2x4 (Socket 0.2 only)	2x4 (Socket 0.2 only)	1x8	1x8 (Host 0.5, 1 only)	2x4 (Host 0.5, 1 only)	2x4 (Host 0.5, 1 only)
2C	1x8 Option B	1x8, 1x4, 1x2, 1x1	0x011	1x8	1x8	1x8	1x8 (Socket 0.2 only)	2x4 (Socket 0.2 only)	2x4 (Socket 0.2 only)	1x8	1x8 (Host 0.5, 1 only)	2x4 (Host 0.5, 1 only)	2x4 (Host 0.5, 1 only)
2C	1x8 Option C	4x4, 4x2, 4x1	0x011	1x8	1x8	1x8	1x8 (Socket 0.2 only)	2x4 (Socket 0.2 only)	2x4 (Socket 0.2 only)	1x8	1x8 (Host 0.5, 1 only)	2x4 (Host 0.5, 1 only)	2x4 (Host 0.5, 1 only)
2C	4x4	4x4	0x011	1x4*	2x4*	2x4*	1x4 (BP 0 and 2 only)	2x4 (BP 0 and 2 only)	2x4 (BP 0 and 2 only)	1x4*	1x4* (BP 0 and 2 only)	2x4* (BP 0 and 2 only)	2x4* (BP 0 and 2 only)
2C	RSVD	RSVD	0x000	-	-	-	-	-	-	-	-	-	-
2C	RSVD	RSVD	0x000	-	-	-	-	-	-	-	-	-	-
2C	RSVD	RSVD	0x000	-	-	-	-	-	-	-	-	-	-



#### 3.6.4 Bifurcation Detection Flow

**[Need input and clarification from system vendors]**

The following detection flow shall be used to determine the resulting link count and lane width based on the baseboard and add-in card configurations.

1. The baseboard reads the state of the PRSNTB[3:0]# pins. If the resulting value is not 0b1111, a card is present.
2. Firmware determines the add-in card PCIe lane width capabilities per Table 19 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 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 would be asserted as appropriate. Doing so assumes the add-in card supports the requested link override.
5. PERST# is deasserted after the 100ms window as defined by the PCIe specification.  
Refer to **Section/Figure XXX** for timing details.

### 3.6.5 PCIe Bifurcation Examples

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

**Figure 22**

Figure 22 illustrates a single host baseboard that supports x16 with a single controller add-in card that also supports x16 (Type 6). 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 22: Single Host (1 x16) and 1 x16 Add-in Card (Single Controller)

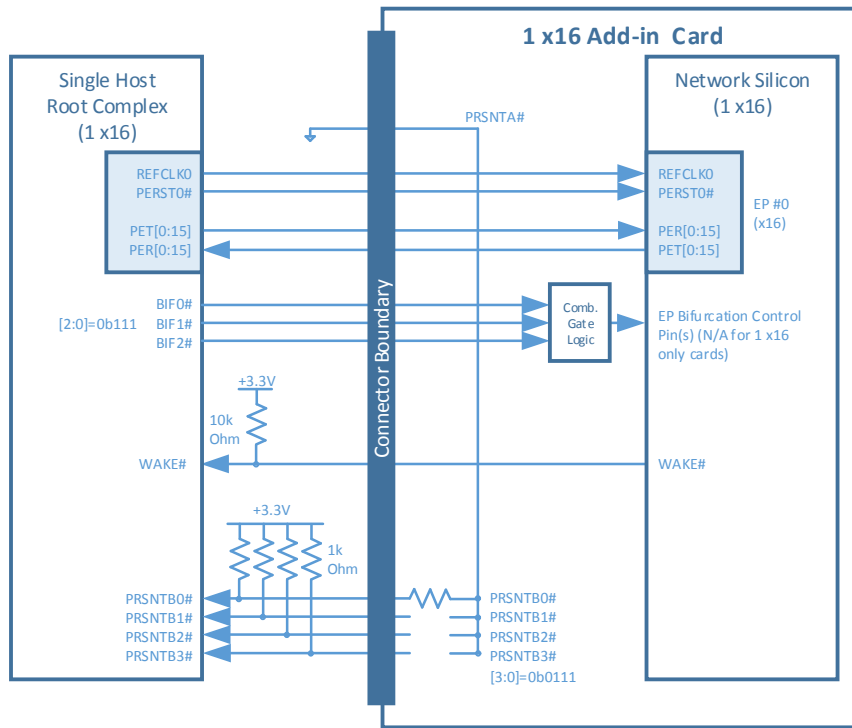




Figure 23 illustrates a single host baseboard that supports 2 x8 with a single controller add-in card that also supports 2 x8 (Type 2). The PRSTNB[3:0]# state is 0b0110. The BIF[2:0]# state is 0b111 as there is no need to instruct the end-point network controllers to a specific bifurcation. The PRSNTB 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 23: Single Host (2 x8) and 2 x8 Add-in Card (Dual Controllers)

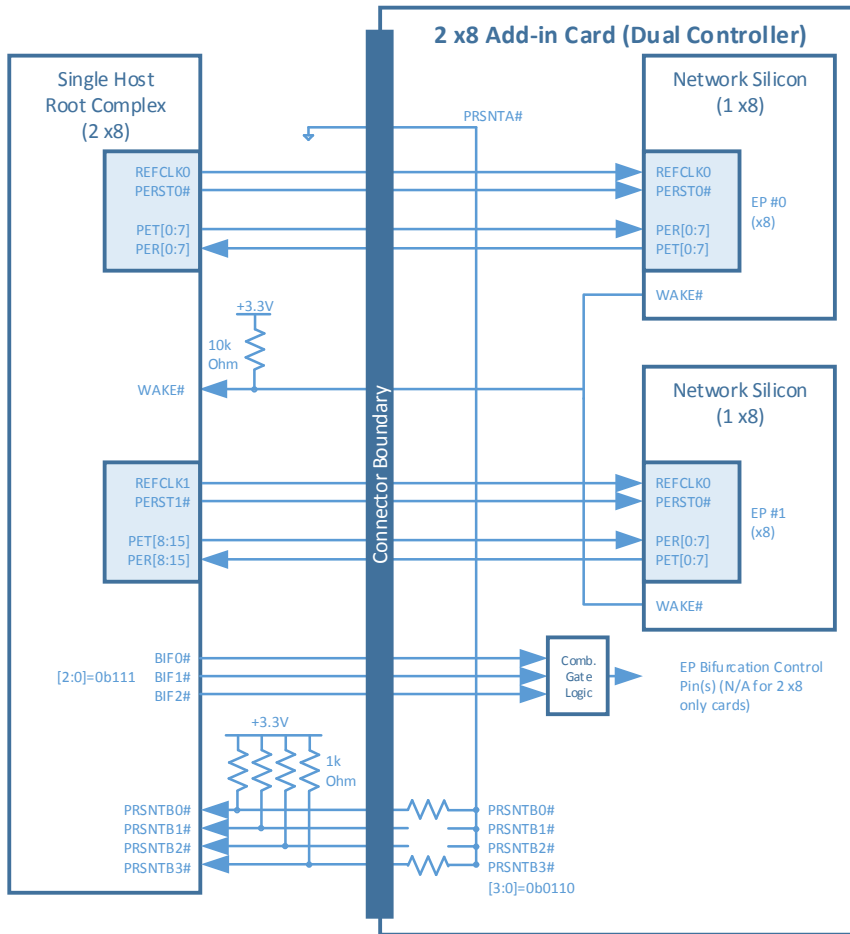


Figure 24 illustrates a four host baseboard that supports 4 x4 with a single controller add-in card that supports 1 x16, 2 x8 and 4 x4 (Type 4). The PRSTNB[3:0]# state is 0b0011. The BIF[2:0]# state is 0b101 as the end point network controller is forced to bifurcate to 4 x4. The PRSNTB encoding notifies the baseboard that this card is only capable of 1 x16, 2 x8 and 4 x4. The four host baseboard determines that it is also capable of supporting 4 x4. The resulting link width is 4 x4.

Figure 24: Four Hosts (4 x4) and 4 x4 Add-in Card (Single Controller)

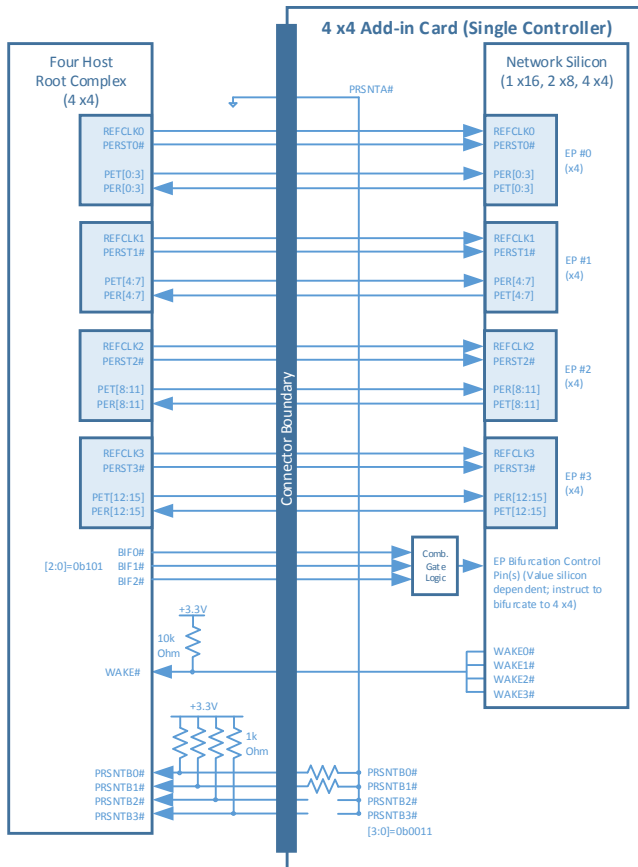




Figure 25 illustrates a four host baseboard that supports 4 x4 with a four controller add-in card that supports 4 x4 (Type 3). The PRSTNB[3:0]# state is 0b0011. The BIF[2:0]# state is 0b111 as there is no need to instruct the end-point network controllers to a specific bifurcation. The PRSNTB encoding notifies the baseboard that this card is only capable of 4 x4. The four host baseboard determines that it is also capable of supporting 4 x4. The resulting link width is 4 x4.

Figure 25: Four Hosts (4 x4) and 4 x4 Add-in Card (Four Controllers)

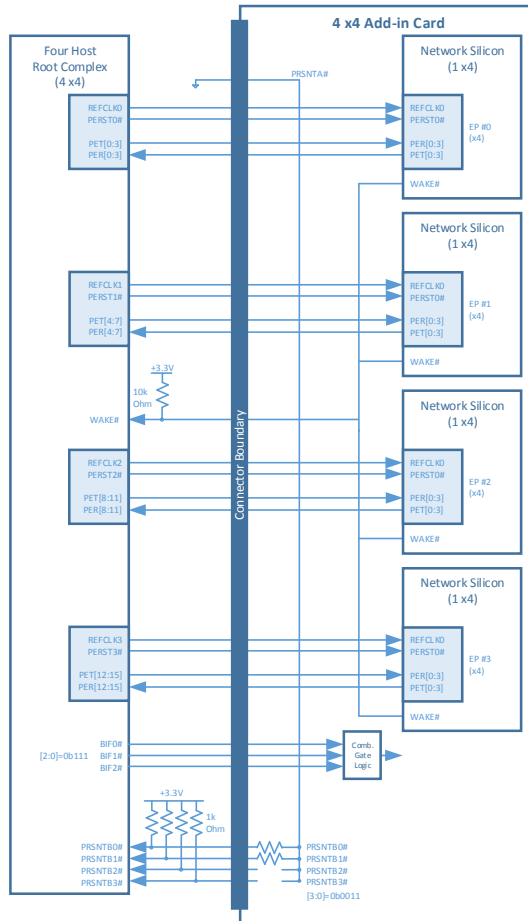
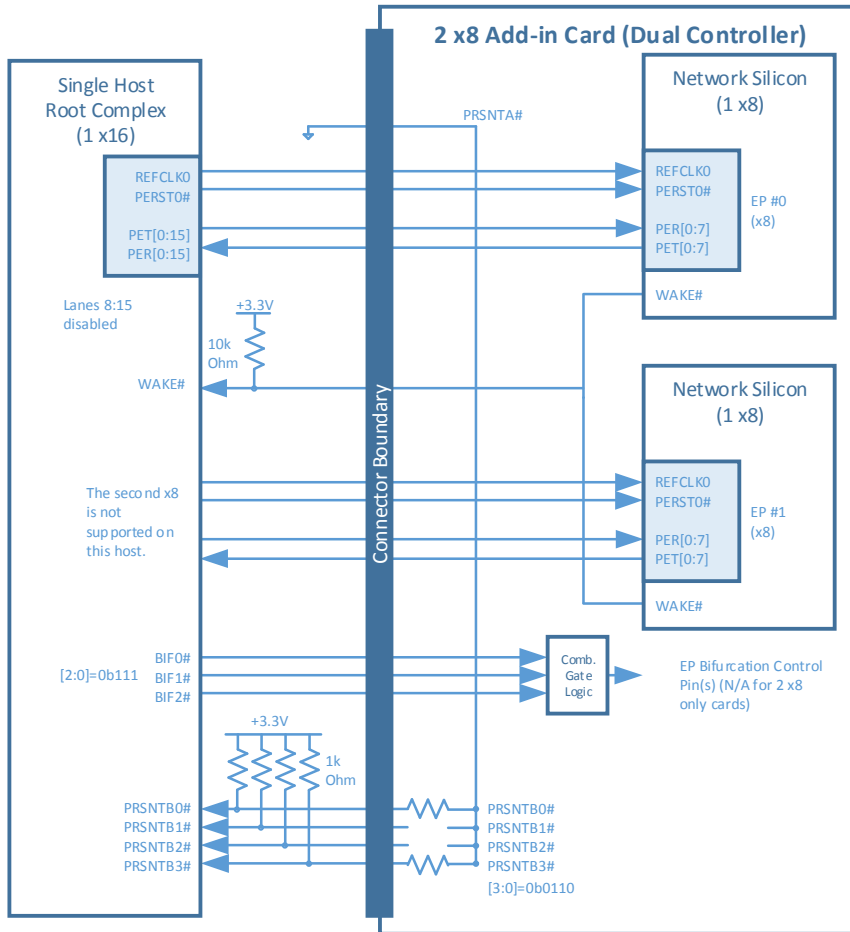




Figure 26: Single Host with no Bifurcation (1 x16) and 2 x8 Add-in Card (Two Controllers)



### 3.7 PCIe Clocking 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 20.

Table 20: 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.

As noted in the Pin Definition (Section [3.34.3](#)), cards that only implement the Primary Connector have up to four PCIe REFCLKS (0-3). Cards that implement both the Primary and Secondary connectors have a total of up to 6 REFCLKs.

For each add-in card, the following REFCLK connection rules must be followed:

- For a 1 x16 capable add-in card, REFCLK0 is used for lanes [0:15].
- For a 2 x8 capable add-in card, REFCLK0 is used for lanes [0:7] and REFCLK1 is used for lanes [8:15].
- For a 4 x4 capable add-in card, REFCLK0 is used for lanes [0:3], REFCLK1 is used for lanes [4:7], REFCLK2 is used for lanes [8:11] and REFCLK3 is used for lanes [12:15].

Pins for REFCLK2 and REFCLK3 are described in Section [3.5.14.5.1](#) and are located on the 28-pin OCP bay.

Figure 27: PCIe Interface Connections for 1 x16 and 2 x8 Add-in Cards

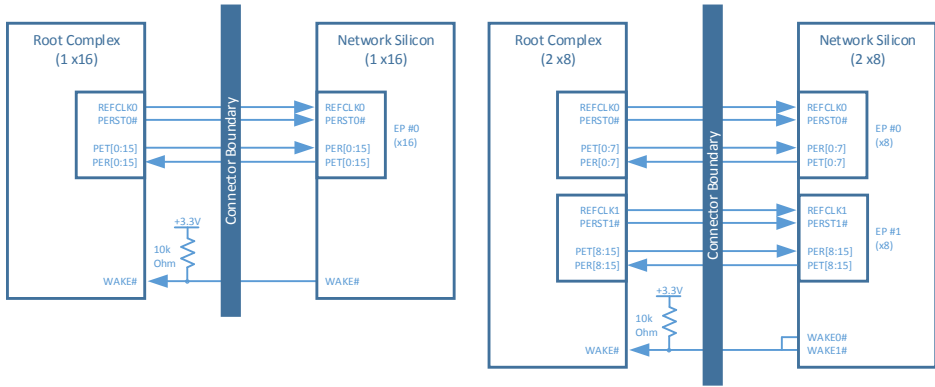
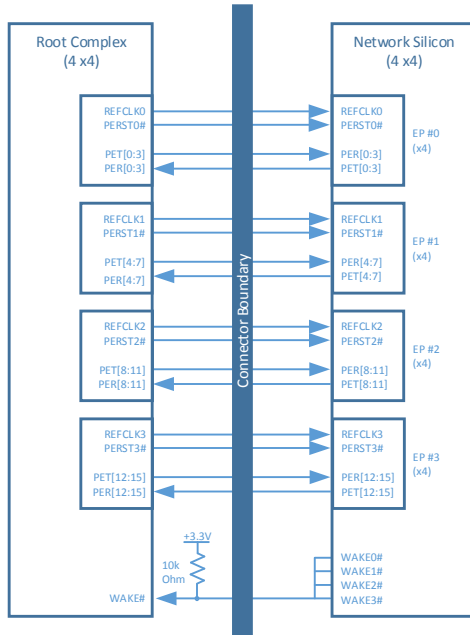


Figure 28: PCIe Interface Connections for a 4 x4 Add-in Card



### 3.8 PCIe Bifurcation Results and REFCLK Mapping

For the cases where bifurcation is 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 [3.6.34.6.3](#).








Table 23: Bifurcation for Single Host, Single Upstream Socket and Single/Dual/Quad Upstream Links (BIF[2:0]#=#0b000)

Host	Add-in-Card P/N	Supported Bifurcations		Host	Upstream Devices	Upstream Links	BIF[2:0] #	Receiving Link																
		Card/Port	Port					Link 0	Link 1	Link 2	Link 3	Link 4	Link 5	Link 6	Link 7	Link 8	Link 9	Link 10	Link 11	Link 12	Link 13	Link 14	Link 15	
2C	1x8	2 x 8	1x8	1Host	1Upstream Socket	1, 2, or 4 Lanes	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	1x8	1Host	1Upstream Socket	1, 2, or 4 Lanes	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	1x4	1Host	1Upstream Socket	1, 2, or 4 Lanes	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	1x2	1x1	1x2	1Host	1Upstream Socket	1, 2, or 4 Lanes	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	2 x 8, 2 x 4, 2 x 2, 2 x 1	1x8	1Host	1Upstream Socket	1, 2, or 4 Lanes	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	2 x 8	1x8, 1x4	2 x 4, 2 x 2, 2 x 1	1Host	1Upstream Socket	1, 2, or 4 Lanes	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	2 x 8, 2 x 4, 2 x 2, 2 x 1	1x8	1Host	1Upstream Socket	1, 2, or 4 Lanes	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	1x8	2 x 8, 2 x 4, 2 x 2, 2 x 1	1x8, 1x8, 1x4, 1x2, 1x1	1Host	1Upstream Socket	1, 2, or 4 Lanes	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	2 x 4	2 x 4, 2 x 2, 2 x 1	2 x 4	1Host	1Upstream Socket	1, 2, or 4 Lanes	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	2 x 4	RESVD	RESVD	1Host	1Upstream Socket	1, 2, or 4 Lanes	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
RESVD	RESVD	RESVD	RESVD	1Host	1Upstream Socket	1, 2, or 4 Lanes	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	1x8	2 x 8, 2 x 4, 2 x 2, 2 x 1	1x8, 1x8, 1x4, 1x2, 1x1	1Host	1Upstream Socket	1, 2, or 4 Lanes	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	1x8	2 x 8, 2 x 4, 2 x 2, 2 x 1	1x8, 1x8, 1x4, 1x2, 1x1	1Host	1Upstream Socket	1, 2, or 4 Lanes	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	1x8	2 x 8, 2 x 4, 2 x 2, 2 x 1	1x8, 1x8, 1x4, 1x2, 1x1	1Host	1Upstream Socket	1, 2, or 4 Lanes	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	4 x 4	4 x 4, 4 x 2, 4 x 1	4 x 4	1Host	1Upstream Socket	1, 2, or 4 Lanes	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	RESVD	RESVD	RESVD	1Host	1Upstream Socket	1, 2, or 4 Lanes	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
RESVD	RESVD	RESVD	RESVD	1Host	1Upstream Socket	1, 2, or 4 Lanes	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

















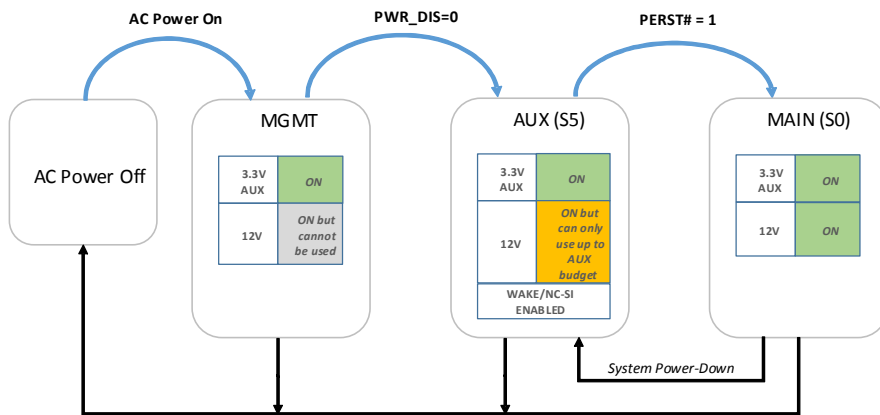
### 3.9 Power Capacity and Power Delivery

There are four permissible power states: AC Power Off, Management (FRU Only Mode), Aux Power Mode (S5), and Main Power Mode (S0). The transition of these states is shown in Figure 29. The main/aux power domains are switched on the baseboard and uses the power pins defined in Table 11. For each of these states, the max power envelopes are defined as follows:

Table 29: Power Envelopes

Power State	Max Power	Notes
AC Power Off	0W	AC power removed; board off
Management (FRU only mode)	TBD	Used only for board identification purposes.
Aux Power Mode (S5)	35W	
Main Power Mode (S0)	79.2W	Add-in card may use up to the 79.2W limit per connector.

Figure 29: Baseboard Power Sequencing



#### 3.9.1 AC Power Off

In AC power off mode, all power delivery has been turned off or disconnected from the baseboard.

#### 3.9.2 Management (FRU Only Mode)

**Commented [NT18]:** Per Jia:

What is the expectation of NIC in this mode? Is the NIC expected to respond to NC-SI command?

Some NIC may need core powered to run management FW, and the core power rail could be from VR under P12V\_AUX

In the Management (FRU Only Mode), only +3.3V Aux is available for powering up management only functions. FRU accesses are only allowed in this mode.

### 3.9.3 Aux Power Mode (S5)

In Aux Power Mode provides both +3.3V Aux as well as +12V Aux is available. +12V Aux may be used to deliver power to the add-in card, but only up to the Aux budget of 35W.

### 3.9.4 Main Power Mode (S0)

In Main Power Mode provides both +3.3V and +12V (Main) across the OCP connector. The add-in card operates in full capacity. Up to 79.2W may be delivered on +12V, and 3.63W on the +3.3V pins.

## 3.10 Power Sequence Timing Requirements

The following figure shows the power sequence of 3.3V/3.3V\_AUX, 12VMain/Aux relative to PWRDIS, PERSTn\* and NIC\_PWR\_GOOD.




Table 30: Power Sequencing Parameters



Parameter Name	Value	Units	Description

## 4 Management

### 4.1 SMBus Interface

The SMBus provides manageability of the add-in card.

### 4.2 NC-SI Sideband Interface

#### 4.2.1 NC-SI addressing and Arb#

### 4.3 MAC Address Requirement

### 4.4 FRU EEPROM

#### 4.4.1 Addressing(TBD)

### 4.5 FW Requirement (TBD)

### 4.6 Thermal Reporting Interface

## 5 Data Network Requirement

### 5.1 Network Booting (collect view from OEMs and hyperscale)

Mezzanine NIC shall support network booting in uEFI system environment. Mezzanine NIC shall support both IPv4 and IPv6 network booting.



## 6 Routing Guidelines and Signal Integrity Considerations

### 6.1 NC-SI Over RBT

Min Length: 2"

Max length: 4"

Impedance: 50 Ohm single ended

## 7 Thermal and Environmental

### 7.1 Environmental Requirements

The specific environment requirement is removed to allow the adoption of OCP Mezzanine NIC in systems with very different thermal requirement and boundary condition.

This Mezzanine card shall meet the same environmental requirements specified in the OCP systems that the Mezzanine card is in. The OCP system that uses OCP Mezzanine card shall define air flow direction, inlet air temperature, air flow (or speed) to the local area where Mezzanine card is at, and simulation boundary.

#### 7.1.1 Thermal Simulation Boundary Example

**Placeholder for Thermal Simulation Method. Using Facebook Intel® Motherboard V3.0 as example. Not covered by this update.**

### 7.2 Shock & Vibration

This Mezzanine card shall meet the same shock & vibration requirements specified in updated Facebook OCP Intel® Motherboard V2.0 and V3.0 Design Specification.

### 7.3 Regulation

This Mezzanine card shall meet CE, CB, FCC Class A, WEEE, ROHS requirements.



## 8 Revision History

Author	Description	Revision	Date
Thomas Ng	Initial draft	0.1	12/xx/2017