OCP NIC 3.0 Design Specification

Version 0.50

Author: OCP Server Workgroup, OCP NIC subgroup

Table of Contents

List of Figures

List of Tables

1 Overview

1.1 License

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. Compared to the OCP Mezz Card 2.0 Design Specification, the updated specification provides a broader solution space for NIC and system vendors to support the following use case scenarios:

- NICs with a higher TDP
- Support up to 80W of power delivery to a single connector (Small) card; and 150W to a dual connector (Large) card
- Support up to PCIe Gen5 on the system and add-in card
- Support for up to 32 lanes of PCIe per add-in card
- Support for single host, multi-root complex 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 and accelerators
- 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 NIC 3.0 Card with Quad SFP Ports

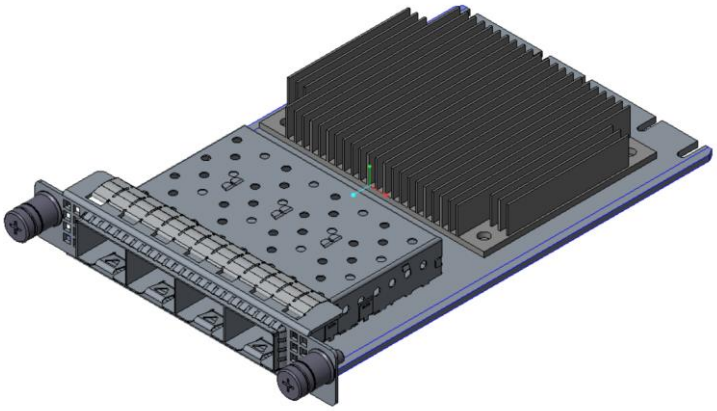

Figure 2: Representative Large OCP NIC 3.0 Card with Dual QSFP Ports and on-board DRAM

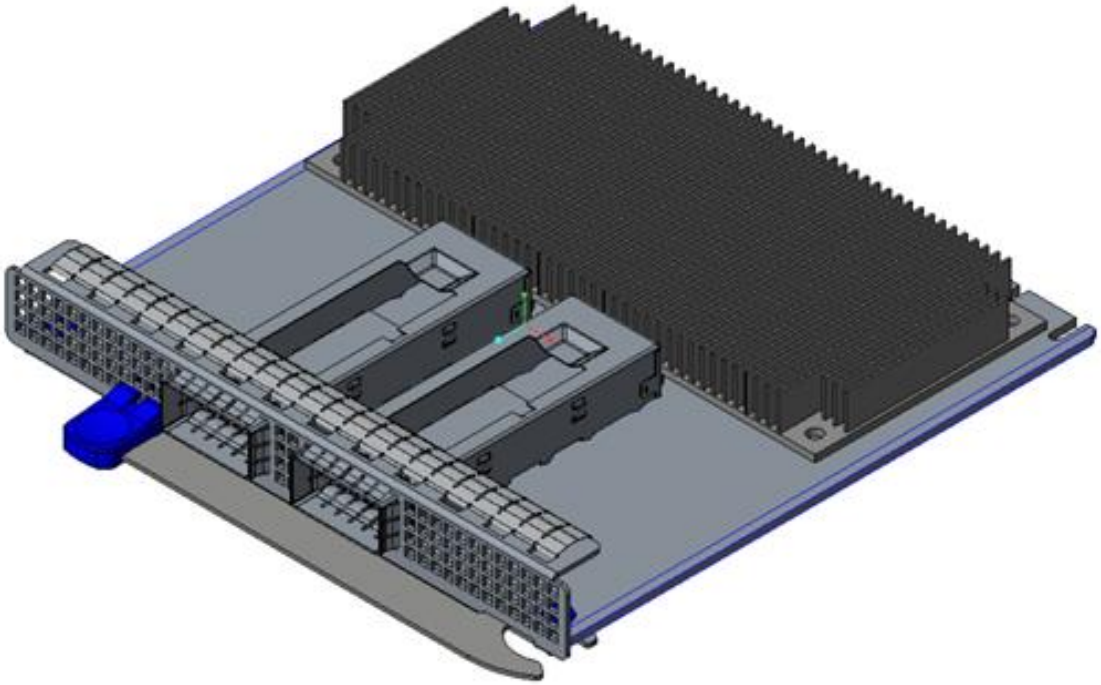

In order to achieve the features outlined in this specification, OCP 3.0 compliant cards are not backwards compatible to OCP Mezz 2.0 cards.

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 website:

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

1.3 Acknowledgements

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 TCS	Intel Corporation	
Broadcom	Lenovo	
Dell	Mellanox	
Facebook	Netronome	
Hewlett Packard Enterprise	TE	

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 add-in 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 one or two connectors (Primary Connector only and both the Primary and Secondary Connectors) on the baseboard.

Both the Primary and Secondary connectors are defined in and compliant to SFF-TA-1002. On the add-in 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. The gold finger design follows SFF-TA-1002 as well.

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

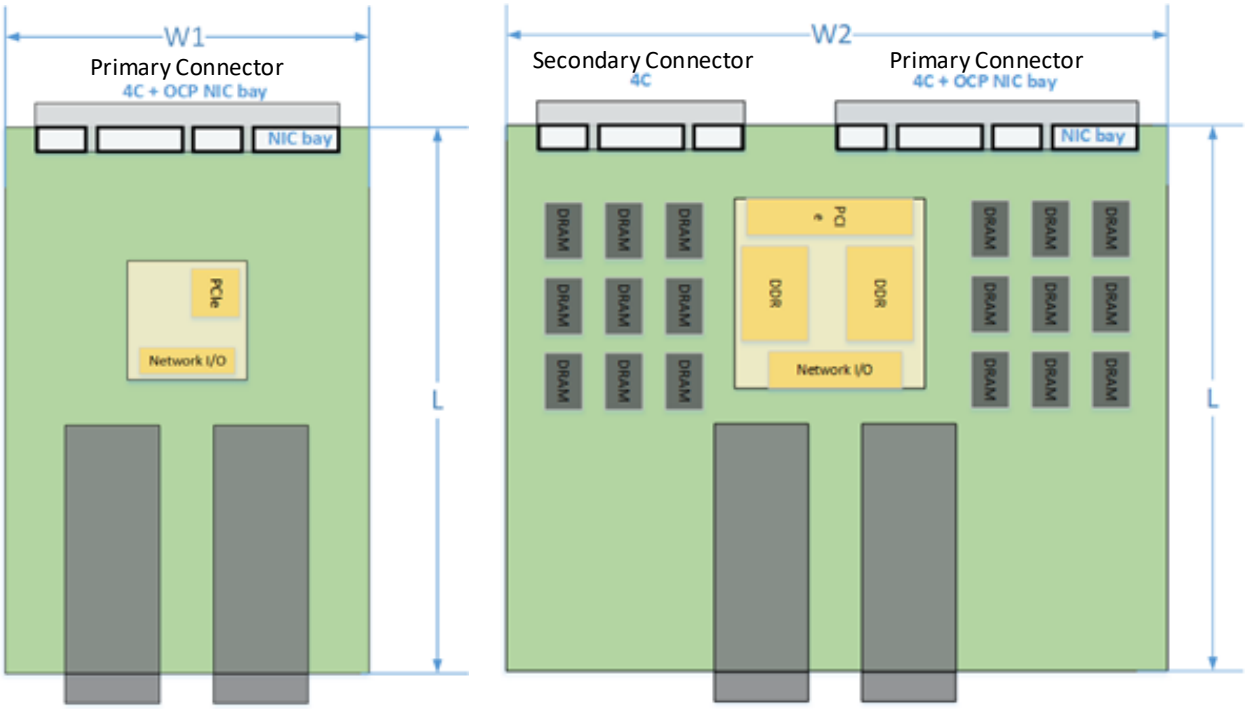

The two form factor dimensions are shown in Table 2.

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 + OCP sideband 168 pins	N/A	Low profile and NIC with a similar profile as an OCP NIC 2.0 add-in card; up to 16 PCIe lanes.
Large	W2 = 139 mm	L = 115 mm	4C + OCP sideband 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 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 card, this specification defines the component and routing keep out areas. Refer to Section 2.6 for details.

Both the straddle mount and right angle implementations shall accept the same OCP add-in 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 side and add-in card sizes are a supported combination as shown in Table 3.

This specification defines the form factor at the add-in 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 add-in card. 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 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):

- DSP0222 1.1 compliant Network Controller Sideband Interface (NC-SI) RMII Based Transport (RBT) Physical Interface
- Power management and status reporting
    - Power disable
    - State change control
- SMBus 2.0
- 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 5 support
- 2x 100 MHz differential reference clocks
- Control signals
    - 2x PCIe Resets
    - Link Bifurcation Control
    - Card power disable/enable
- Power

- 12V /12V AUX
- 3.3V AUX

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 5 support
- 2x 100 MHz differential reference clocks
- Control signals
  - 2x PCIe Resets
  - Link Bifurcation Control
  - Card power disable/enable
- Power
  - 12V /12V AUX
- 3.3V AUX

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, Rev 1.1.0, September 23rd, 2015.
- DMTF Standard. *DSP0222, Network Controller Sideband Interface (NC-SI) Specification.* Distributed Management Task Force, Rev 1.2.0, Work-in-progress.
- 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.
- 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 4.0 (draft).
- 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 1.0, December 12th, 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 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.

Figure 4: Example Small Card Form Factor

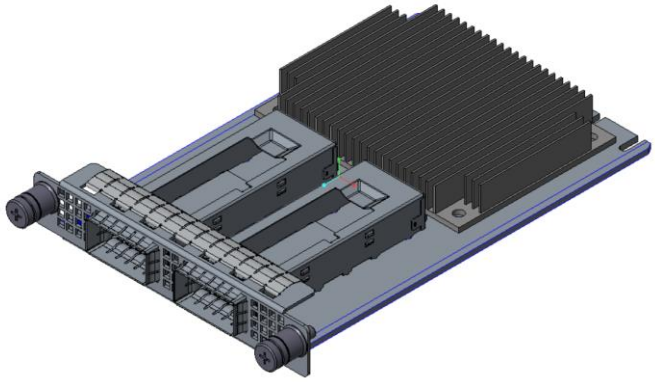

The large card uses the Primary 4C + OCP bay connector to provide the same functionality as the small card along with an additional Secondary 4C connector to provide up to a x32 PCIe interface. The large card may utilize both the Primary and Secondary connectors, or just the Primary connector for lower PCIe lane count applications. Table 4 summarizes the large card permutations. The large size 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.

Figure 5: Example Large Card Form Factor

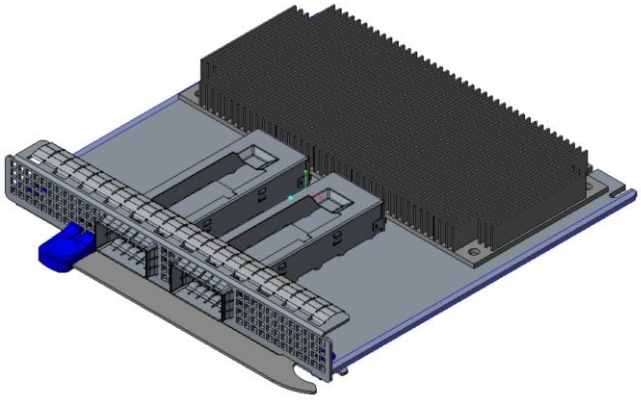

For both form-factors, an add-in card may optionally implement a subset of pins to support up to a x8 PCIe connection. This is implemented using a 2C card edge per SFF-TA-1002. The Primary Connector may support a 2C sized add-in card along with the 28 pin OCP bay. The following diagram from the SFF-TA-1002 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 (Small) Add-in Cards

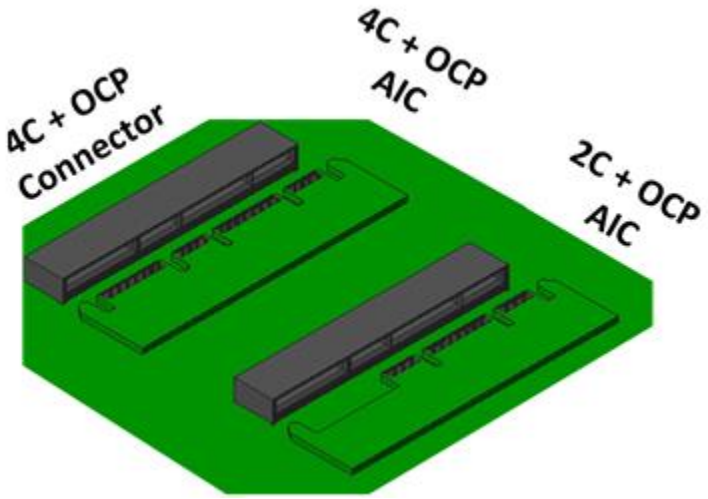

Figure 7: Primary Connector (4C + OCP Bay) and Secondary Connector (4C) (Large) Add-in Cards

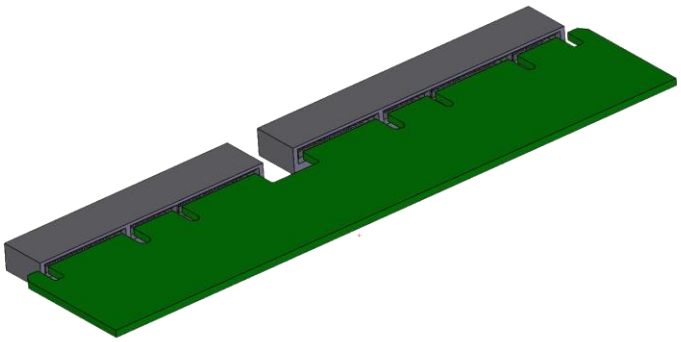

Table 4 summarizes the supported card form factors. Small form factors cards support the Primary Connector and up to 16 PCIe lanes. Large form factor cards support implementations with both the Primary and Secondary Connectors and up to 32 PCIe lanes, or a Primary Connector only implementation with up 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		OCP Bay
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

2.3 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.

2.3.1 Small Form Factor Add-in Card I/O Bracket

Figure 8 defines the standard Small Card form factor I/O bracket.

Figure 8: Small Card Standard I/O Bracket

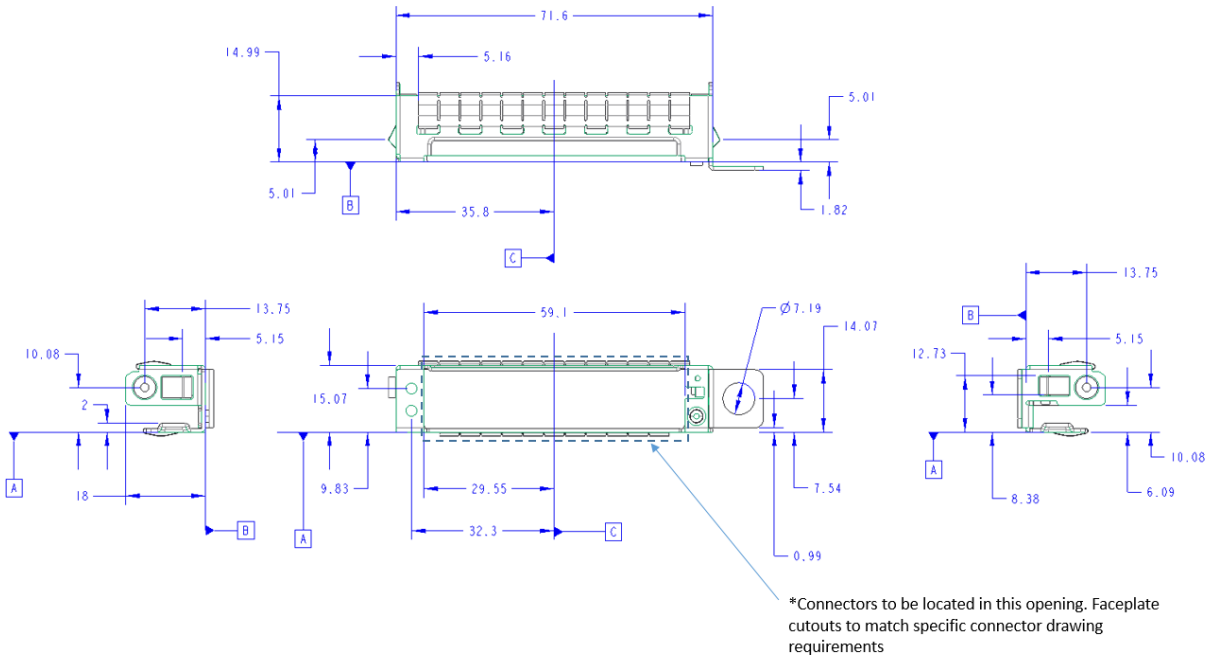

Note: The add-in card supplier shall add port identification on bracket that meet their manufacturing and customer requirements.

For RJ-45 implementations, a customized bracket must be created. Figure 9 shows an implementation example.

Figure 9: Small Card Customized bracket for RJ-45 Connector

Drawing to be inserted

Figure 10 shows the standalone bracket assembly and Figure 11 shows the bracket assembly on the add-in card.

Figure 10: Small Card 3D Bracket Assembly (Standalone)

TBD

Figure 11: Small Card 3D Bracket Assembly (Installed on Add-in Card)

TBD

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

Table 5: Mechanical BOM for the Small Card Bracket

Item description	**Supplier Part Number**
Top and bottom EMI fingers	TF187VE32F11
Screw / Rivet (part of bracket assy)?	TBD
Side EMI Finger	TBD
Thumb screw	TBD
Pull Tab	TBD
Latch	TBD
Screw (attaching Bracket & NIC)	TBD
SMT Nut (on NIC)	TBD

2.3.2 Small Form Factor Add-in Card Critical-to-Function (CTF) Dimensions

The following dimensions are considered critical-to-function (CTF) for each small form factor add-in card.

Figure 12: Small Form Factor Add-in Card Critical-to-Function (CTF) Dimensions (Top View)

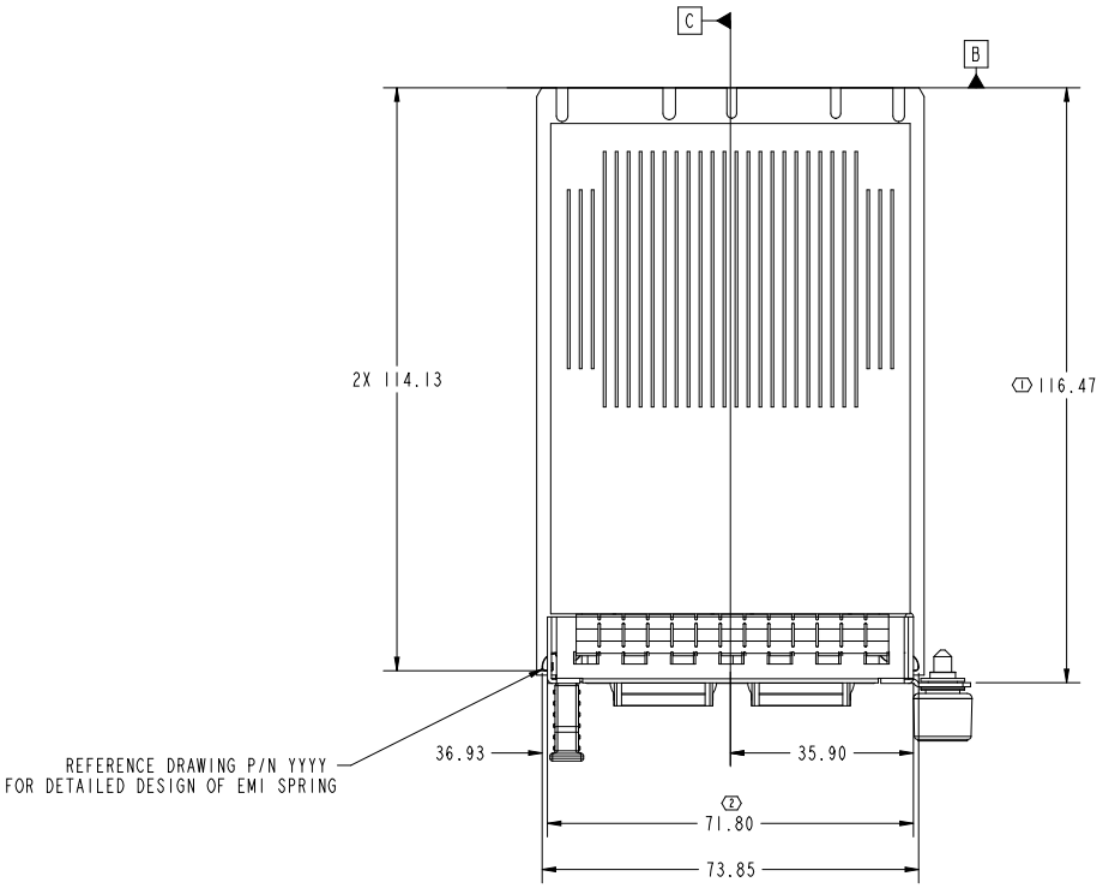

Figure 13: Small Form Factor Add-in Card Critical-to-Function (CTF) Dimensions (Front View)

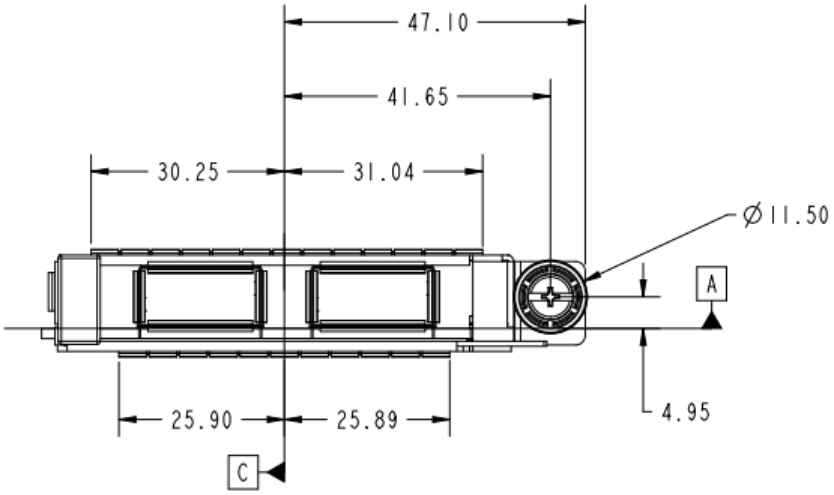

Figure 14: Small Form Factor Add-in Card Critical-to-Function (CTF) Dimensions (Side View – Left)

Figure 15: Small Form Factor Add-in Card Critical-to-Function (CTF) Dimensions (Side View – Right)

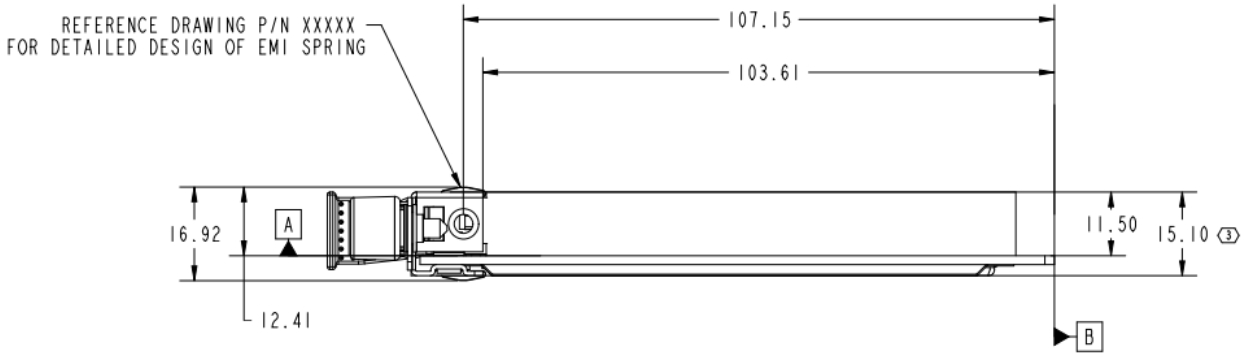

2.3.3 Small Form Factor Baseboard Critical-to-Function (CTF) Dimensions

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

Figure 16: Small Form Factor Baseboard Chassis Critical-to-Function (CTF) Dimensions (Rear View)

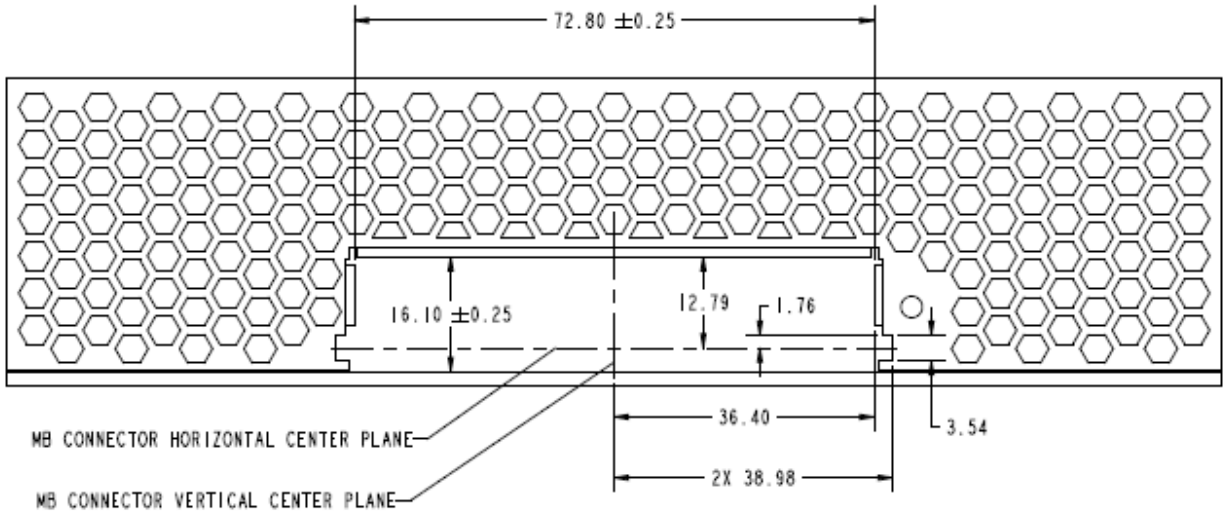

Figure 17: Small Form Factor Baseboard Chassis Critical-to-Function (CTF) Dimensions (Side View)

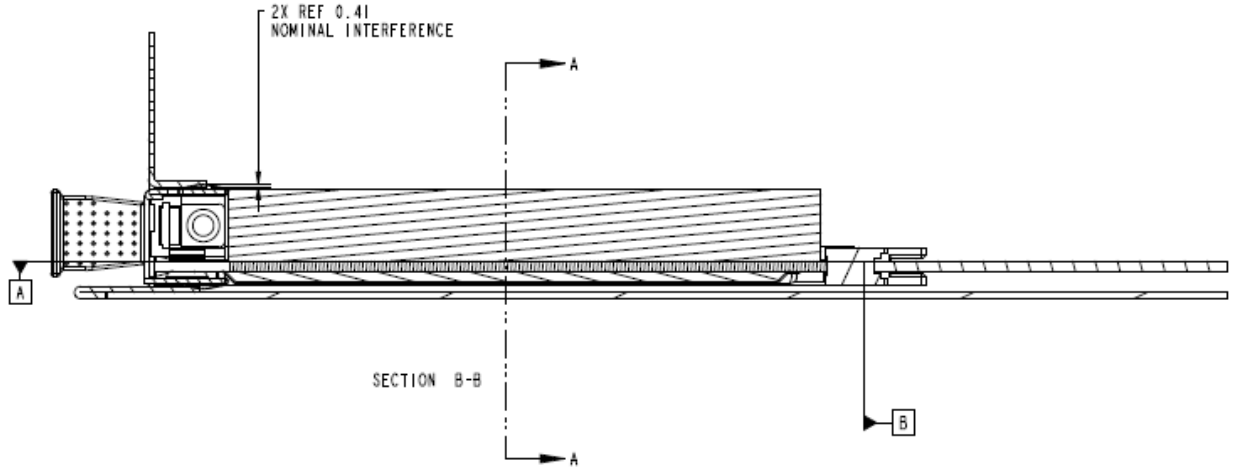

Figure 18: Small Form Factor Baseboard Chassis Critical-to-Function (CTF) Dimensions (Rear Rail Guide View)

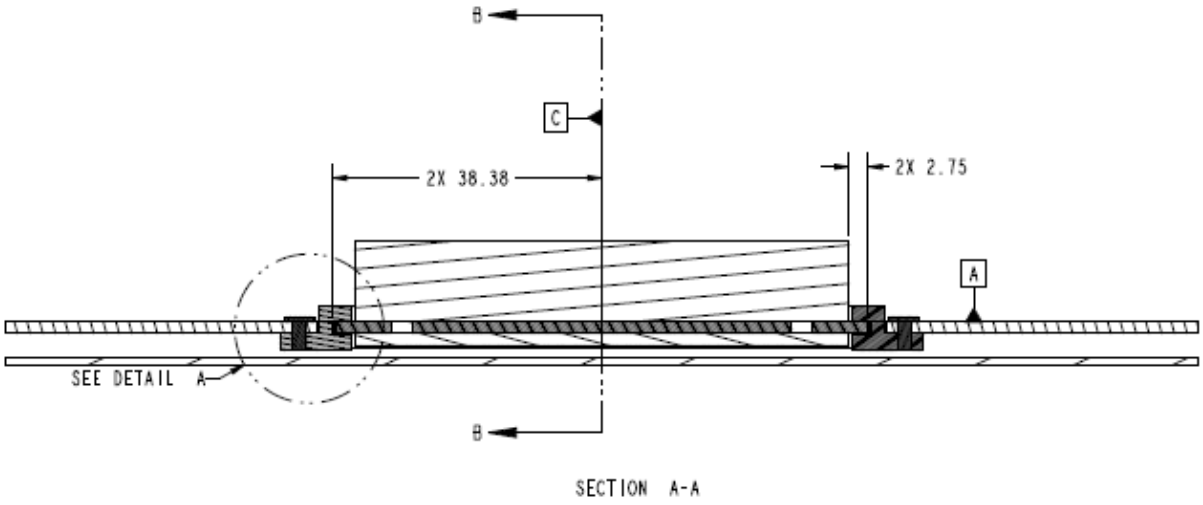

Figure 19: Small Form Factor Baseboard Chassis Critical-to-Function (CTF) Dimensions (Rail Guide Detail)

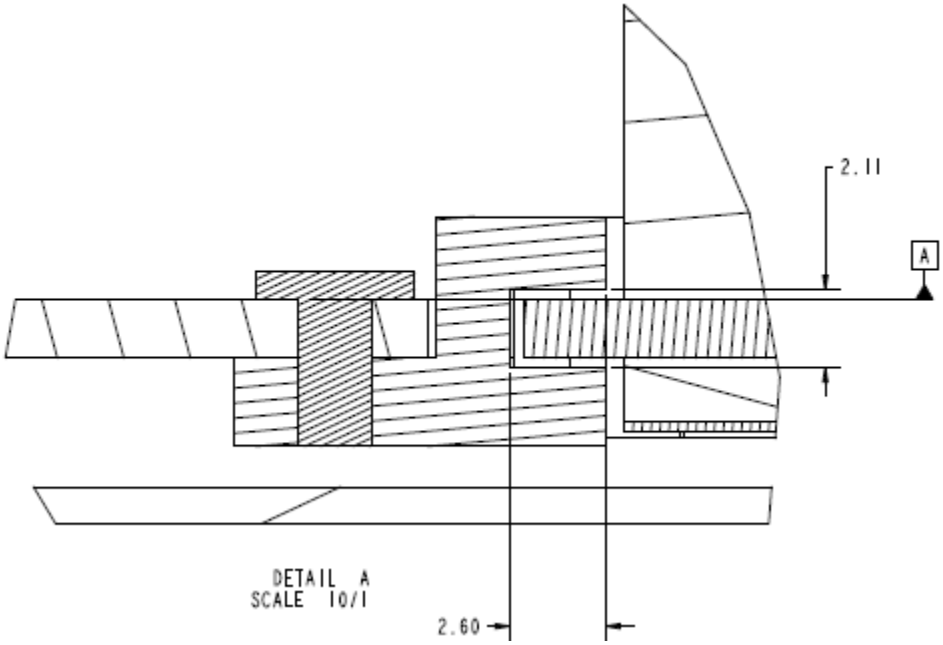

On the baseboard side, the following mechanical dimensions shall be met to support a small form factor add-in card:

Figure 20: Baseboard and Rail Assembly Drawing for Small Card

TBD; need 3D baseboard and rail assembly drawing.

2.3.4 Large Form Factor Add-in Card I/O Bracket

TBD <need input from OCP mechanical groups>. All drawings from the Small Form-Factor implementation need to be replicated for the Large form-factor.

2.3.5 Large Form Factor Add-in Card Critical-to-Function (CTF) Dimensions

2.3.6 Large Form Factor Baseboard Critical-to-Function (CTF) Dimensions

2.4 Line Side I/O Implementations

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

Table 6: OCP 3.0 Line Side I/O Implementations

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

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

2.5 LED Implementations

LEDs may be implemented on the card Scan Chain (as defined in Section 0) for remote link/activity indication on the baseboard or optionally be implemented on the OCP NIC 3.0 I/O bracket if there is sufficient space for local indication.. These two cases are described below. In both cases, the actual link rate may be directly queried through the management interface.

2.5.1 Baseboard LEDs Configuration Over the Scan Chain

A small form-factor OCP NIC 3.0 with a fully populated I/O bracket (2x QSFP28, 4x SFP28 or 4x RJ-45) has insufficient 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 0.

The baseboard LED implementation uses two discrete LEDs (Link/Activity and Speed indication). Table 7 describes the baseboard LED configuration for baseboard implementations.

This LED implementation is required for all add-in cards. The LED implementation is optional for baseboards.

Table 7: Baseboard LED Configurations with Two Physical LEDs per Port

LED Pin	LED Color	Description
Link / Activity	Green	Active low. Multifunction LED. This LED shall be used to indicate link and link activity. When the link is up and no link activity is present, 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. 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 baseboard Link/Activity LED location is not mandated in this specification and will be defined by the system vendor.
Speed	Green Off	Active low. Multifunction LED. The LED is Green when the port is linked at its maximum speed. The LED is off when the device is linked at a speed lower than the highest capable speed, or no link is present. The baseboard bicolor speed LED location is not mandated in this specification and will be defined by the system vendor.

At the time of this writing, the Scan Chain definition allows for up to one link/activity and one speed 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 3.0 compliant chassis (e.g. front LED indicators with rear I/O cards).

2.5.2 Add-in Card LED Configuration

For low I/O count small form-factor cards without built in light pipes (such as 1x QSFP28, 2x SFP28, or 2x RJ-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 place of the Scan Chain LED stream. The recommended local (on-card) LED implementation uses two physical LEDs (a discrete Link/Activity LED and a bi-colored Speed A/Speed B LED). Table 8 describes the add-in card LED implementations.

Table 8: Add-in Card LED Configuration with Two Physical LEDs per Port

LED Pin	LED Color	Description
	Green	Active low. Multifunction LED.

Link / Activity		This LED shall be used to indicate link and link activity. When the link is up and no link activity is present, 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. 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 Link/Activity LED shall be located on the left hand side or located on the top for each port when the add-in card is viewed in the horizontal plane.
Speed	Green Amber Off	Active low. Bicolor multifunction LED. The LED is Green when the port is linked at its maximum speed. The LED is Amber when the port is linked at it second highest speed. The LED is off when the device is linked at a speed lower than the second highest capable speed, or no link is present. The Amber Speed LED indicator may be used for port identification through vendor specific link diagnostic software. The bicolor speed LED shall be located on the right hand side or located on the bottom for each port when the add-in card is viewed in the horizontal plane.

2.5.3 Add-in Card LED Ordering

For all add-in card use cases, each port shall implement the green Link/Activity LED and a bicolor green/amber speed A/B LED. For all baseboards, each port shall implement the green Link/Activity LED and a green speed A LED.

For horizontal LED positions, the Link/Activity LED shall be located on the left side for each port and the speed LED shall be located on the right side for each port.

For vertical LED positions, the Link/Activity LED shall be located on top and the speed LED shall be located on the bottom.

The placement of the LEDs on the faceplate may be left up to the discretion of the add-in card and baseboard designers. The LED port association shall be clearly labeled on the add-in card and on the baseboard.

Figure 21: LED Ordering – Example Small Card Link/Activity and Speed LED Placement

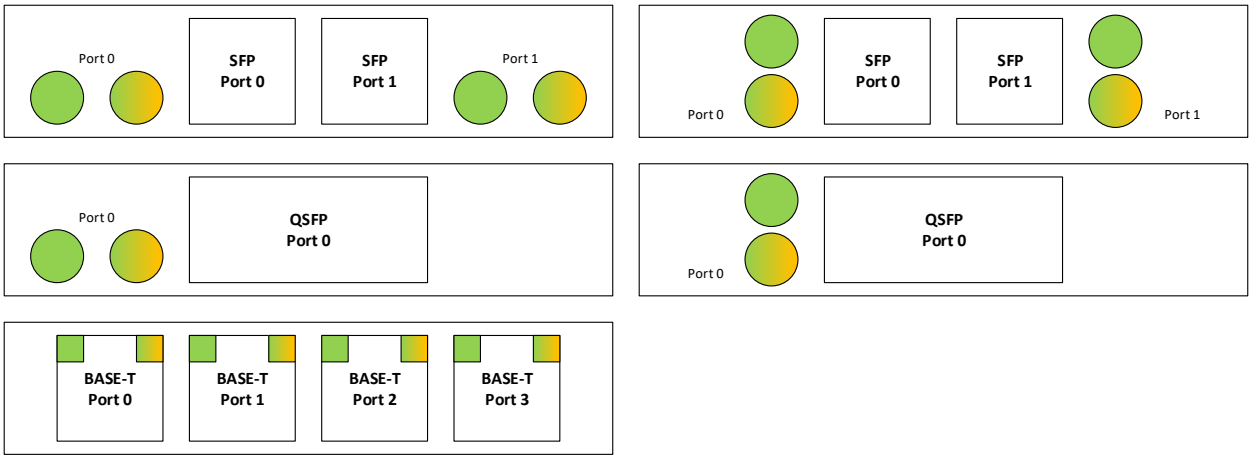

2.6 Mechanical Keepout Zones

2.6.1 Baseboard Keep Out Zones – Small Card Form Factor

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 Baseboard Keep Out Zones – Large Card Form Factor

TBD. – need input from mechanical engineering

2.6.3 Small Card Form Factor Keep Out Zones

Figure 22: Small Form Factor Keep Out Zone – Top View

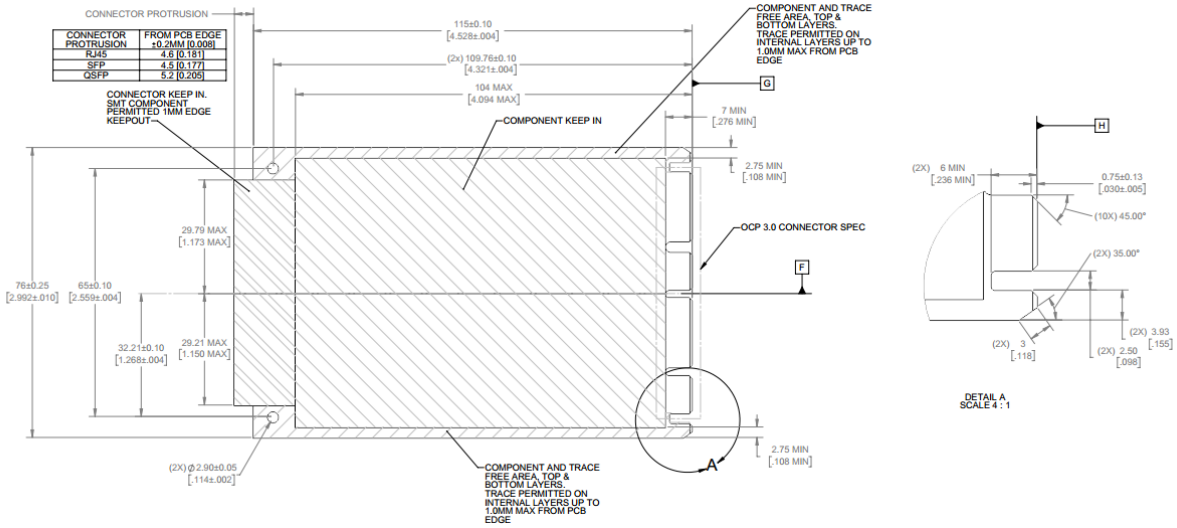

Figure 23: Small Form Factor Keep Out Zone – Bottom View

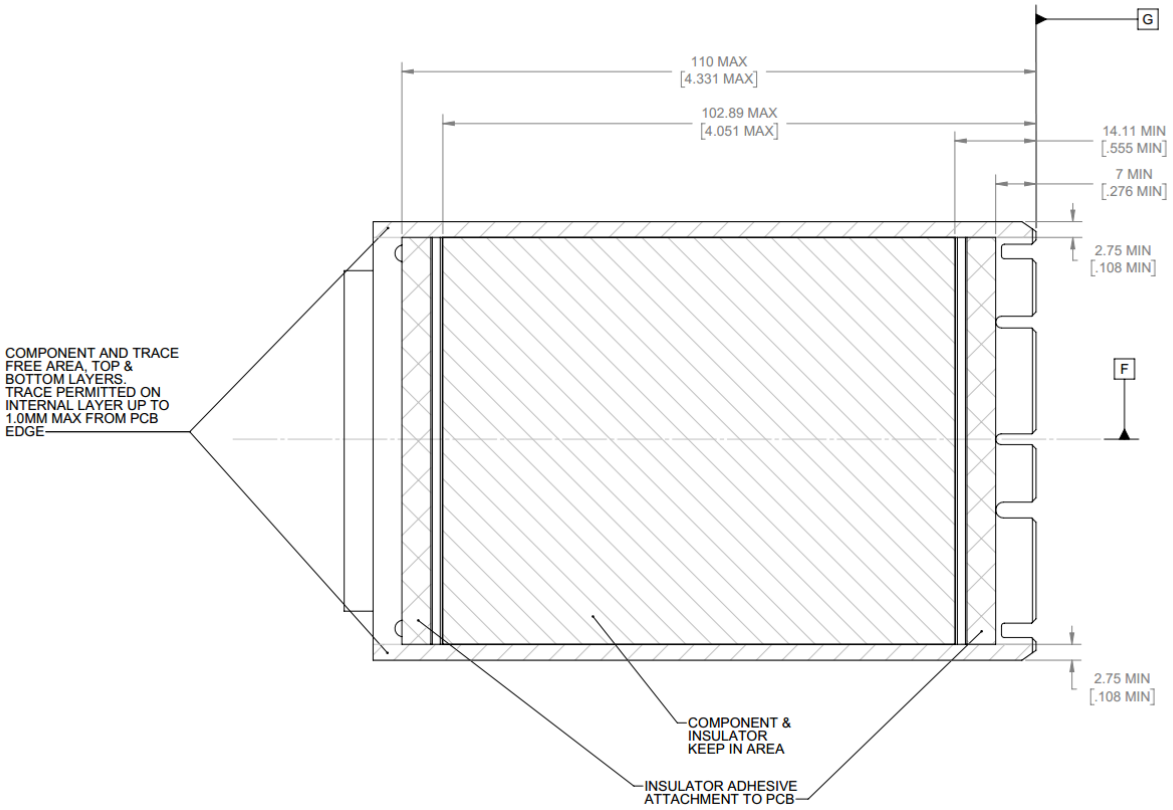

Figure 24: Small Form Factor Keep Out Zone – Side View

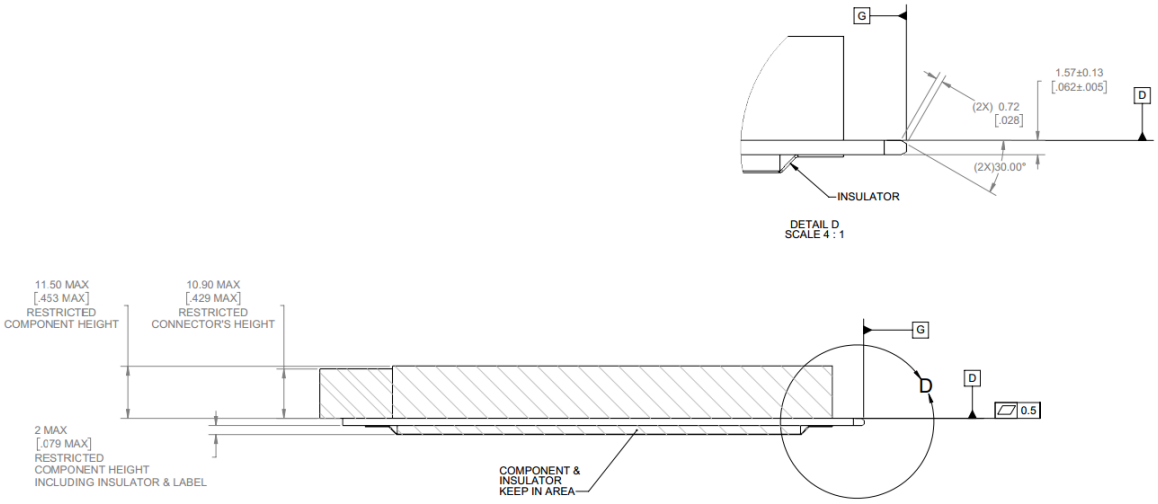

2.6.4 Large Card Form Factor Keep Out Zones

Figure 25: Large Form Factor Keep Out Zone – Top View

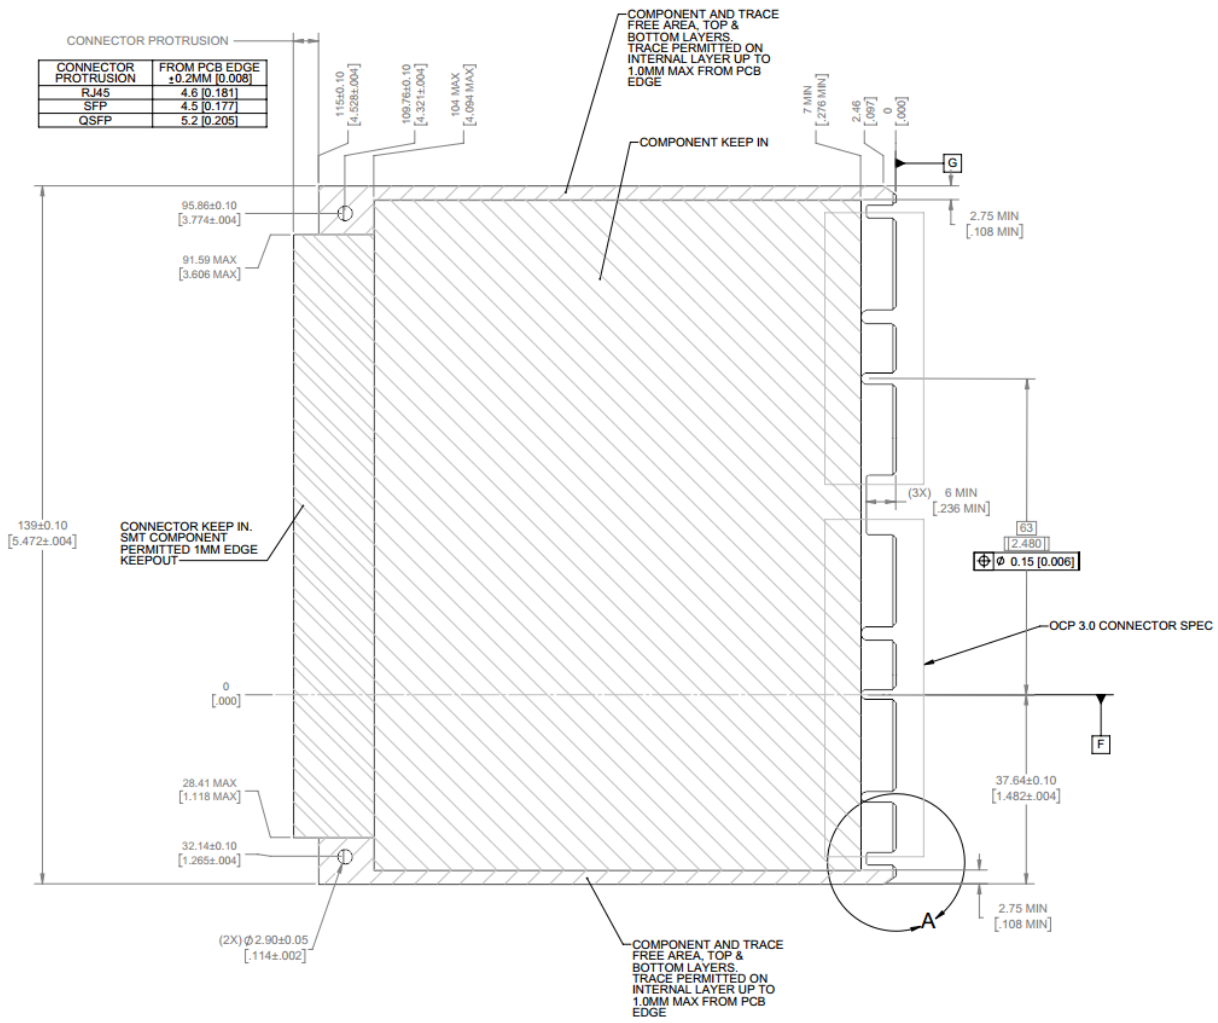

Figure 26: Large Form Factor Keep Out Zone – Bottom View

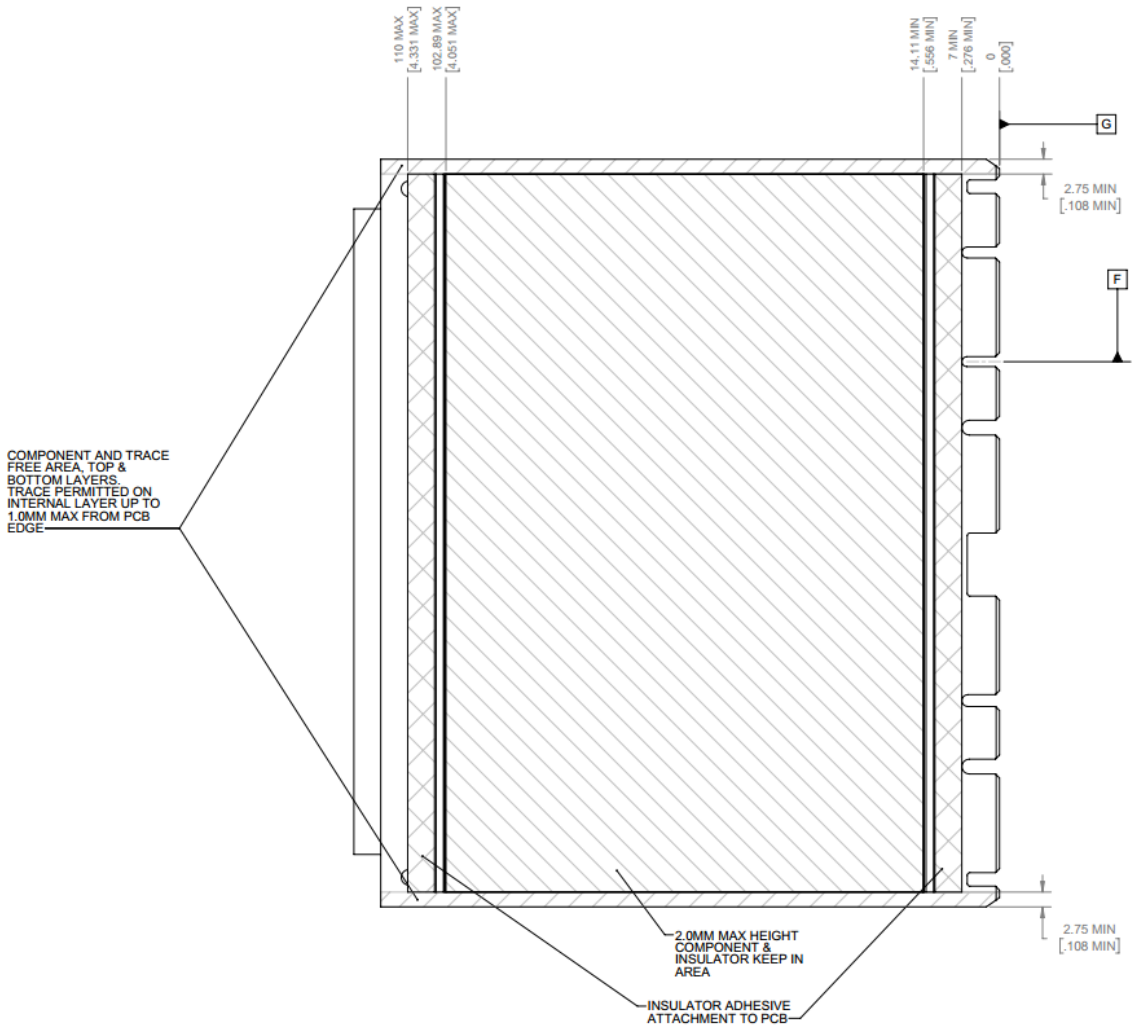

Figure 27: Large Form Factor Keep Out Zone – Side View

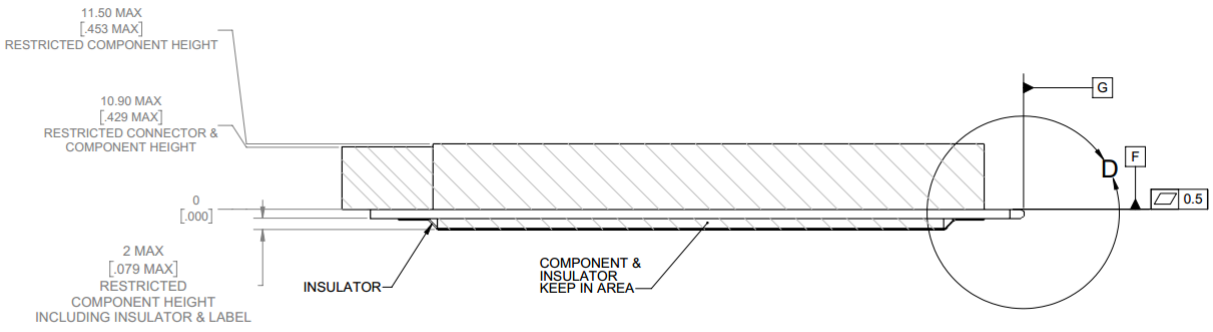

2.7 Labeling Requirements

TBD

Editor's note [TN 20171214]: Consider the following label attributes:

- Label attributes are human (e.g. ASCII) and machine readable (e.g. barcode)
- Add-in card MAC address shall be visible (used MAC address range, or base value)
- Board serial number

2.8 Insulation Requirements

All 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.

Figure 28: Small Card Bottom Side Insulator (Top and 3/4 View)

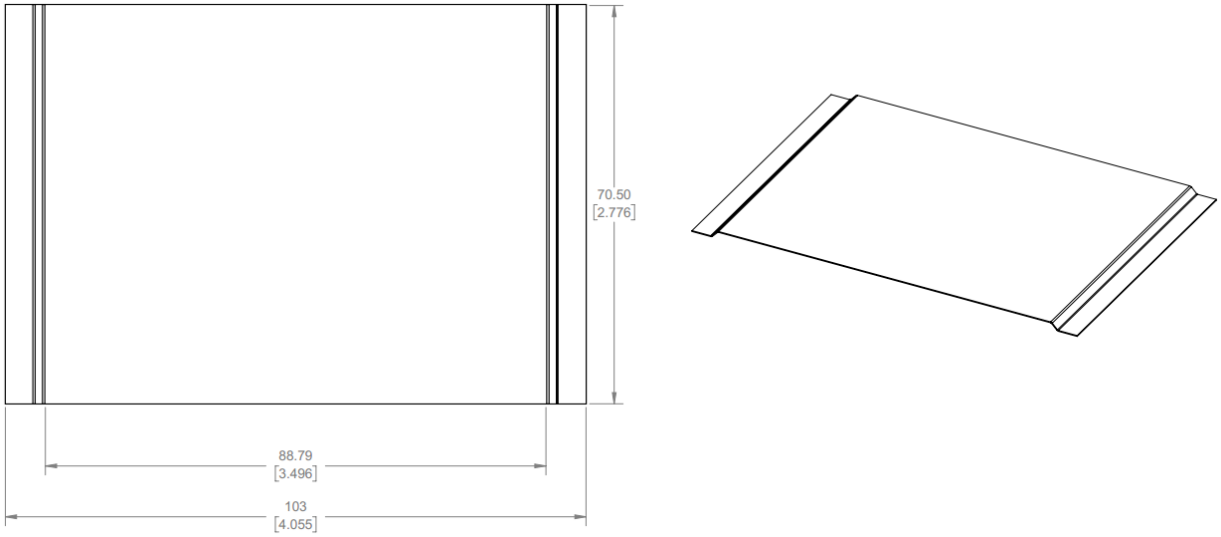

Figure 29: Small Card Bottom Side Insulator (Side View)

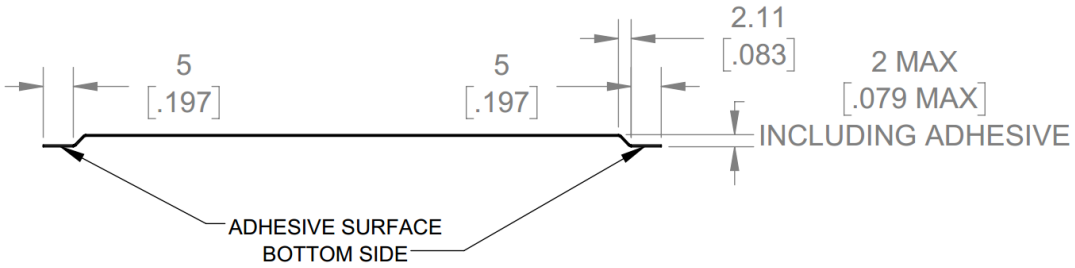

Figure 30: Large Card Bottom Side Insulator (Top and 3/4 View)

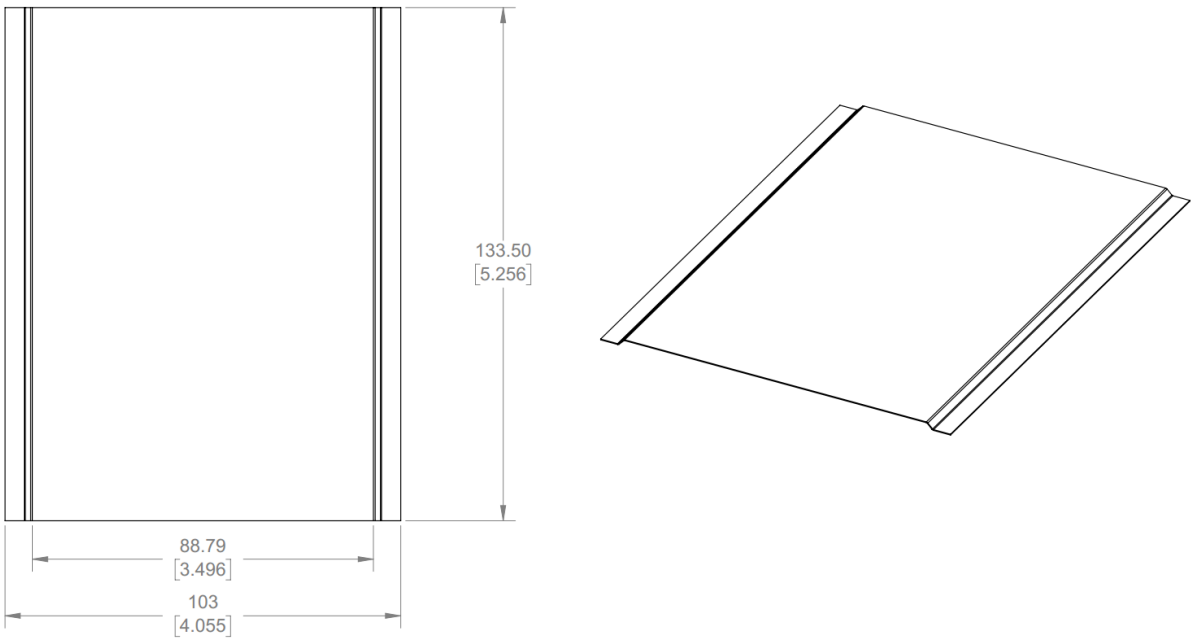

Figure 31: Large Card Bottom Side Insulator (Side View)

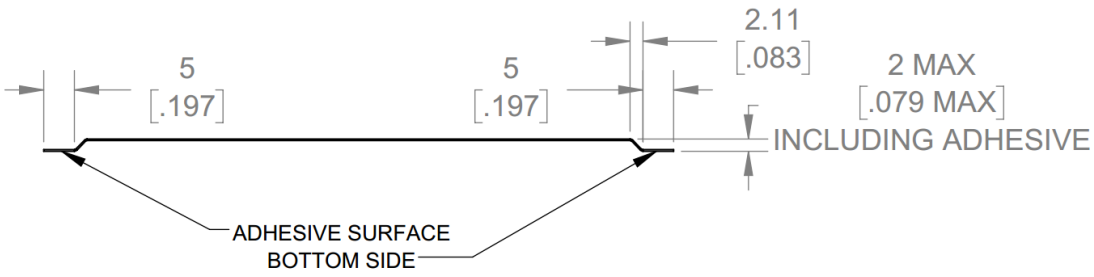

2.9 NIC Implementation Examples

TBD

2.10 Non-NIC Use Cases

"PCIe interface with extra management sideband"

2.10.1 PCIe Retimer card

TBD

2.10.2 Accelerator card

TBD

2.10.3 Storage HBA / RAID card

TBD

3 Card Edge and Baseboard Connector Interface

3.1 Gold Finger Requirements

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 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 connectors. Large Size Cards may implement a reduced PCIe lane count and optionally implement only the Primary Connector 4C, or 2C plus OCP bay.

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 32: Small Size Primary Connector Gold Finger Dimensions – x16 – Top Side ("B" Pins)

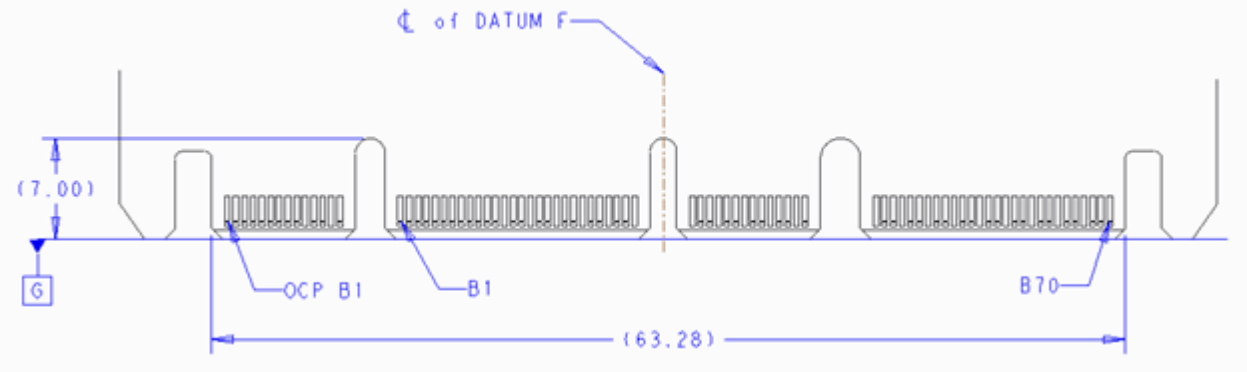

Figure 33: Small Size Primary Connector Gold Finger Dimensions – x16 – Bottom Side ("A" Pins)

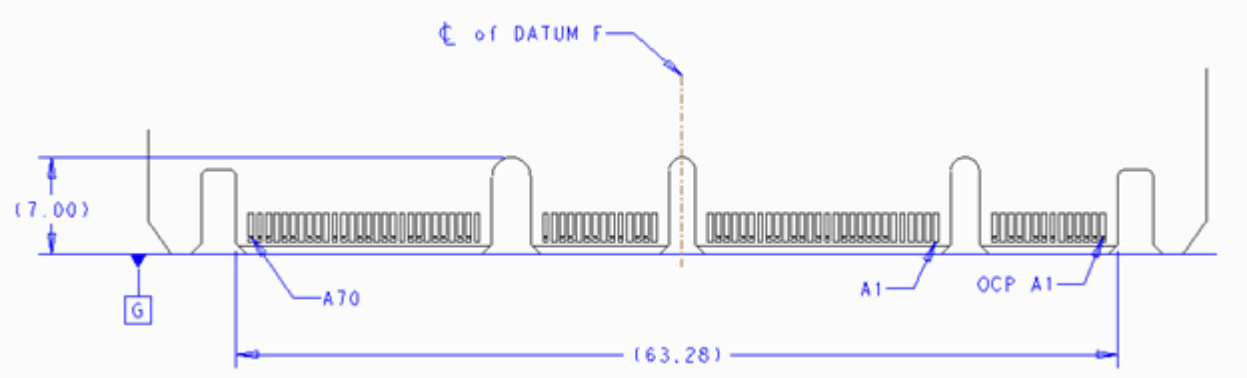

Figure 34: Large Size Card Gold Finger Dimensions – x32 – Top Side ("B" Pins)

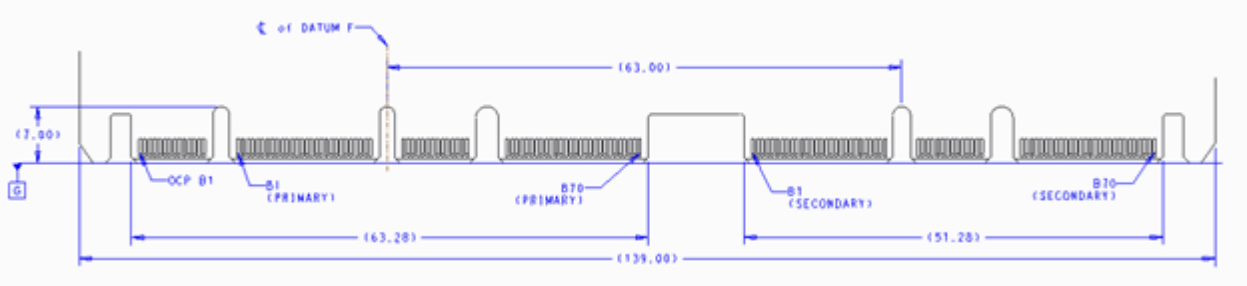

Figure 35: Large Size Card Gold Finger Dimensions – x32 – Bottom Side ("A" Pins)

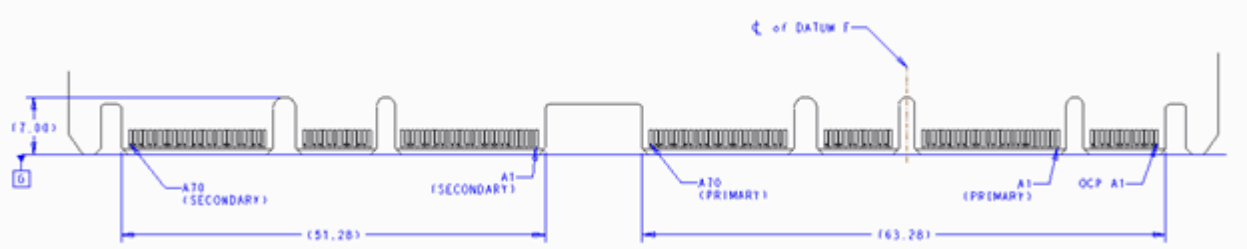

3.1.1 Gold Finger Mating Sequence

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 add-in card gold finger staging is shown in Table 9 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 add-in 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.

Table 9: Contact Mating Positions for the Primary and Secondary Connectors

Side B	Side A

	AIC Plug (Free)		Receptacle (Fixed)		AIC Plug (Free)		Receptacle (Fixed)
	2nd Mate	1st Mate			2nd Mate	1st Mate	
OCP B1				OCP A1			
OCP B2				OCP A2			
OCP B3				OCP A3			
OCP B4				OCP A4			
OCP B5				OCP A5			
OCP B6				OCP A6			
OCP B7				OCP A7			
OCP B8				OCP A8			
OCP B9				OCP A9			
OCP B10				OCP A10			
OCP B11				OCP A11			
OCP B12				OCP A12			
OCP B13				OCP A13			
OCP B14				OCP A14			
Mechanical Key							
B1				A1			
B2				A2			
B3				A3			
B4				A4			
B5				A5			
B6				A6			
B7				A7			
B8				A8			
B9				A9			
B10				A10			
B11				A11			
B12				A12			
B13				A13			
B14				A14			
B15				A15			
B16				A16			
B17				A17			
B18				A18			
B19				A19			
B20				A20			
B21				A21			
B22				A22			
B23				A23			
B24				A24			
B25				A25			
B26				A26			
B27				A27			
B28				A28			
Mechanical Key							
B29				A29			
B30				A30			
B31				A31			
B32				A32			
B33				A33			
B34				A34			
B35				A35			
B36				A36			
B37				A37			
B38				A38			
B39				A39			
B40				A40			
B41				A41			
B42				A42			
Mechanical Key							
B43				A43			
B44				A44			
B45				A45			
B46				A46			
B47				A47			
B48				A48			
B49				A49			
B50				A50			
B51				A51			
B52				A52			
B53				A53			
B54				A54			

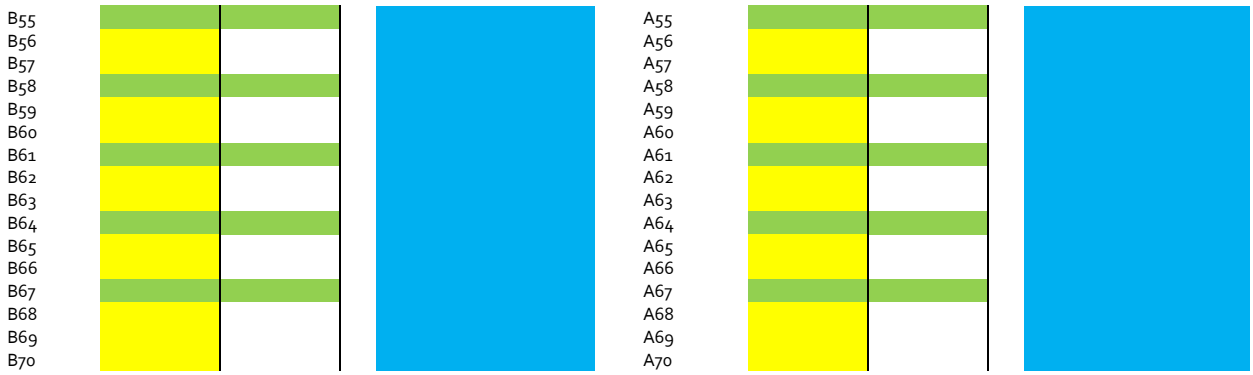

3.2 Baseboard Connector Requirements

The OCP NIC 3.0 connectors are 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 connection. The connector also provides 6 pins of 12V for payload power. This implementation is common between both the Primary and Secondary Connectors. 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 x2 and 4 x4 multi-host configuration on the Primary Connector. The Primary and Secondary Connector drawings are shown in Figure 36, Figure 37, Figure 38 and Figure 39 below.

All diagram units are in mm unless otherwise noted.

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

Table 10: Right Angle Connector Options

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

Figure 36: 168-pin Base Board Primary Connector – Right Angle

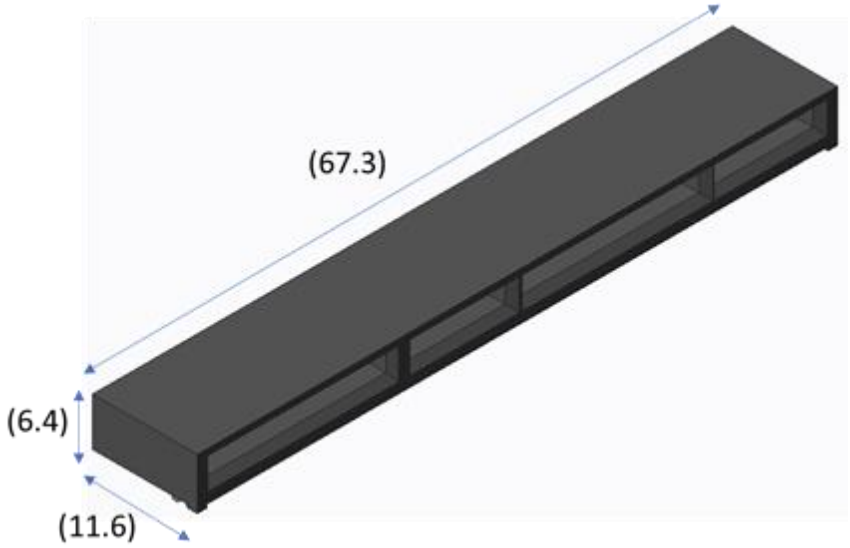

Figure 37: 140-pin Base Board Secondary Connector – Right Angle

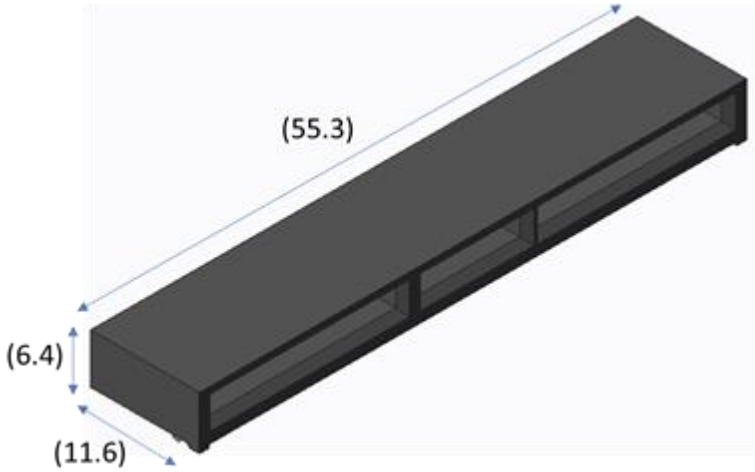

The following offset and height options are available for the straddle mount Primary and Secondary Connectors.

Table 11: Straddle Mount Connector Options

Name	Pins	Style and Baseboard Thickness	Offset (mm)
Primary Connector – 4C + OCP	168 pins	Straddle Mount for 0.062"	Coplanar (0mm)
Primary Connector – 4C + OCP	168 pins	Straddle Mount for 0.076"	-0.3mm
Primary Connector – 4C + OCP	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 38: 168-pin Base Board Primary Connector – Straddle Mount

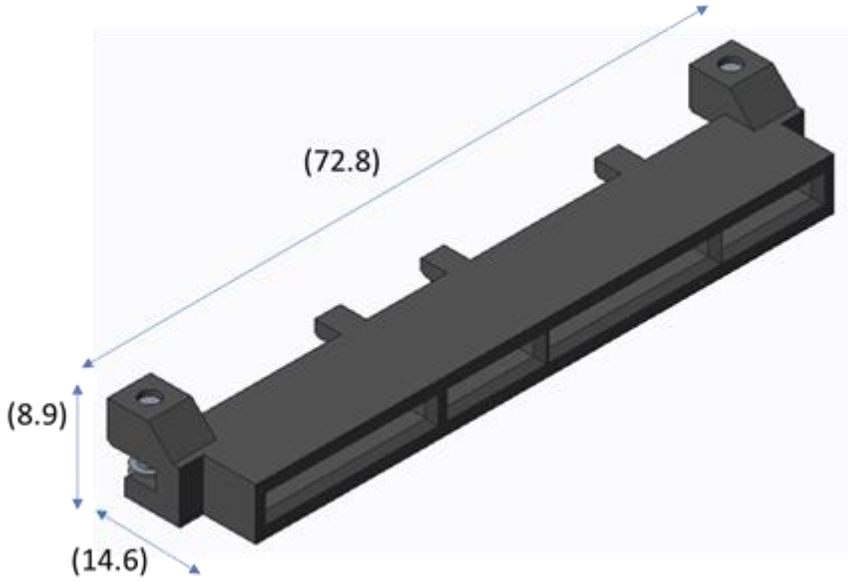

Figure 39: 140-pin Base Board Secondary Connector – Straddle Mount

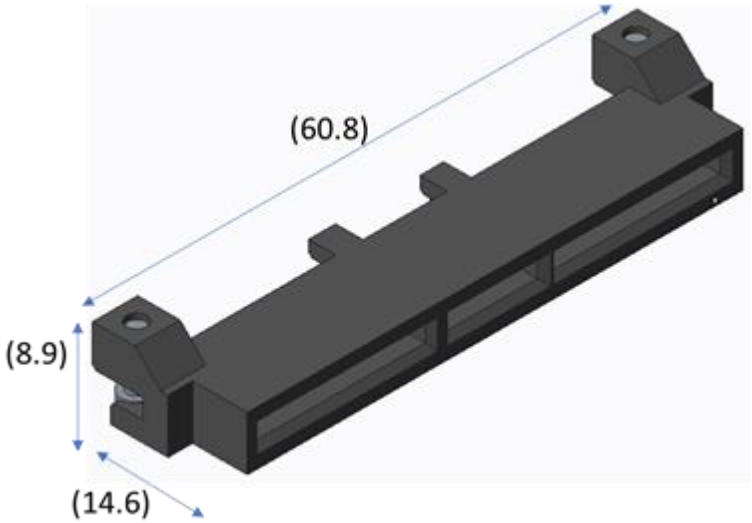

In order to the support the large form factor, systems must locate the Primary and Secondary Connectors per the mechanical drawing shown in Figure 40 and Figure 41.

Figure 40: Primary and Secondary Connector Locations for Large Card Support For Right Angle Connectors

Figure 41: Primary and Secondary Connector Locations for Large Card Support For Straddle Mount Connectors

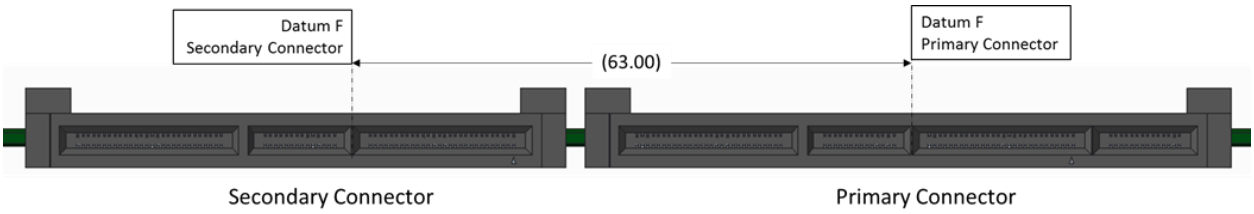

3.2.1 Straddle Mount Offset and PCB Thickness Options

The OCP NIC 3.0 straddle mount connectors have four PCB thicknesses they can accept. The available options are shown in Figure 42. The thicknesses are 0.062", 0.076", 0.093", and 0.105". These PCBs must be controlled to a thickness of ±8%. 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" host board thickness.

Figure 42: Add-in Card and Host PCB Thickness Options for Straddle Mount Connectors

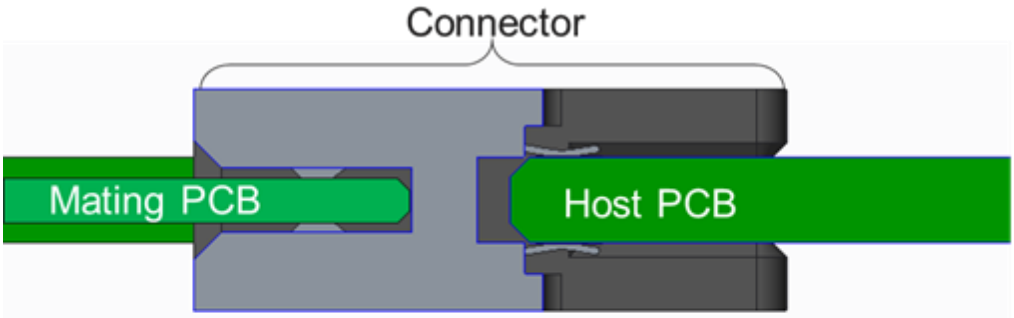

Connector	Mating PCB Thickness	Host PCB Thickness
A	.062" (1.57mm)	.062" (1.57mm)
B		.076" (1.93mm)
C		.093" (2.36mm)
D		.105" (2.67mm)

The connectors are capable of being used coplanar as shown in Figure 43. Additionally, the connectors are also capable of having a 0.3mm offset from the centerline of the host board as shown in Figure 44.

Figure 43: 0mm Offset (Coplanar) for 0.062" Thick Baseboards

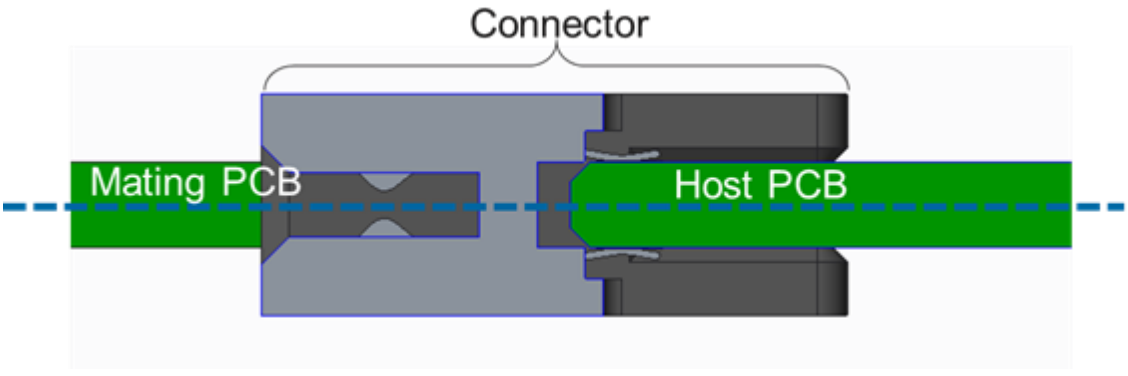

Figure 44: 0.3mm Offset for 0.076" Thick Baseboards

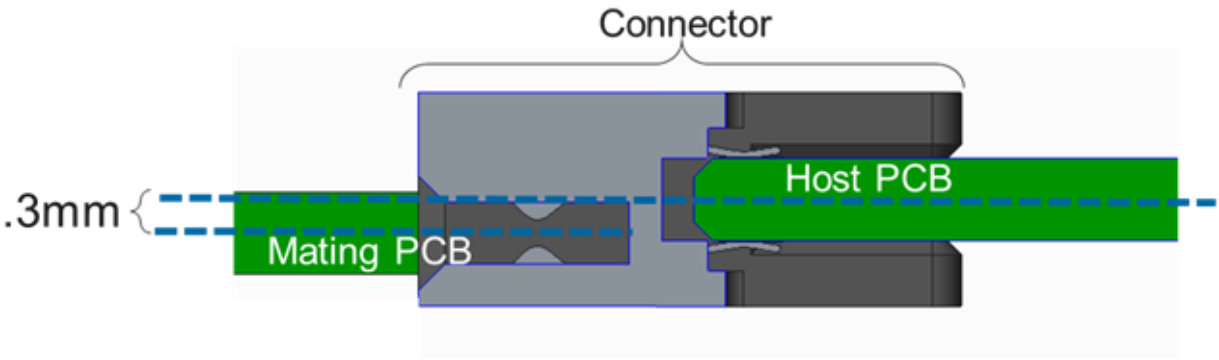

3.2.2 Right Angle Offset

The OCP NIC 3.0 right angle connectors have a 4.0mm offset from the baseboard (pending SI simulation results). This is shown in Figure 45.

Figure 45: Add-in Card and Host Offset for Right Angle Connectors

TBD

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 12 and Table 13. 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 add-in cards. Both connectors share common functionality with power, SMBus 2.0, x16 PCIe 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 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 + OCP Bay) and Secondary (4C) connectors.

Table 12: Primary Connector Pin Definition (x16) (4C + OCP Bay)

Side B	Side A

OCP_B1	NIC_PWR_GOOD	PERST2#	OCP_A1	Primary Connector (x16, 168-pin add-in card with OCP Bay)	Primary Connector (x8, 112-pin add-in card with OCP bay)
OCP_B2	PWRBRK#	PERST3#	OCP_A2		
OCP_B3	LD#	WAKE_N#	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/+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		
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		
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 13: Secondary Connector Pin Definition (x16) (4C)

Side B		Side A			
B1	+12V/+12V_AUX	GND	A1	Secondary Connector (x16, 140-pin add-in card)	Secondary Connector (x8, 84-pin add-in card)
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		
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		

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	

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 add-in card shall implement protection methods to prevent leakage paths between the Vaux and Vmain 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 58.

Table 14: 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 add-in card PCIe core logic. For baseboards, the REFCLK0 and REFCLK1 signals shall be available at the connector. For add-in cards, the required REFCLKs shall be connected per the endpoint datasheet.
REFCLKn1 REFCLKp1	A14 A15	Output	

			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. Refer to Section 2.1 in the PCIe CEM Specification, Rev 4.0 for electrical details.
PETn0 PETp0	B17 B18	Output	Transmitter differential pairs [0:15]. These pins are connected from the baseboard transmitter differential pairs to the receiver differential pairs on the add-in card. The PCIe transmit pins shall be AC coupled on the baseboard with capacitors. The capacitors shall be placed next to the baseboard transmitters. The AC coupling capacitor value shall be between 176nF (min) and 265nF (max). For baseboards, the PET[0:15] signals are required at the connector. For add-in 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. Refer to Section 6.1 in the PCIe CEM Specification, Rev 4.0 for details.
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	Receiver differential pairs [0:15]. These pins are connected from the add-in card transmitter differential pairs to the receiver differential pairs on the baseboard. The PCIe receive pins shall be AC coupled on the add-in card with capacitors. The capacitors shall be placed next to the add-in card transmitters. The AC
PERn1 PERp1	A20 A21	Input	
PERn2 PERp2	A23 A24	Input	
PERn3 PERp3	A26 A27	Input	

PERn4 PERp4	A30 A31	Input	coupling capacitor value shall be between 176nF (min) and 265nF (max). For baseboards, the PER[0:15] signals are required at the connector. For add-in 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. Refer to Section 6.1 in the PCIe CEM Specification, Rev 4.0 for details.
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. When PERSTn# is deasserted, the signal shall indicate the applied power is within tolerance and stable for the add-in card. 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. PERST shall be pulled high to 3.3Vaux on the baseboard. For OCP NIC 3.0, PERST deassertion shall also indicate the full card power envelope is available to the add-in card. For baseboards, the PERST[0:1]# signals are required at the connector.

			For add-in cards, the required PERST[0:1]# signals shall be connected to the endpoint silicon. **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. Refer to Section 2.2 in the PCIe CEM Specification, Rev 4.0 for details.

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 46.

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 PWRDIS deassertion 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 XXX for details.

Table 15: Pin Descriptions – PCIe Present and Bifurcation Control Pins

Signal Name	Pin #	Baseboard Direction	Signal Description
PRSNTA#	A12	Output	Present A is used for add-in card presence and PCIe capabilities detection. For baseboards, this pin shall be directly connected to GND. For add-in cards, this pin shall be directly connected to the PRSNTB[3:0]# pins.
PRSNTB0# PRSNTB1# PRSNTB2# PRSNTB3#	B42 A42 A10 B70	Input	Present B [0:3]# are used for add-in card presence and PCIe capabilities detection. For baseboards, these pins shall be connected to the I/O hub and pulled up to +3.3Vaux using 1kOhm resistors. For add-in cards, these pins shall be strapped to PRSNTA# per the encoding definitions described in Section 3.6. Note: 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.
BIF0# BIF1# BIF2#	A7 A8 A9	Output	Bifurcation [0:2]# pins allow the baseboard to force configure the add-in card bifurcation. For baseboards, these pins shall be outputs driven from the baseboard I/O hub and allow the system to force configure the add-in card bifurcation. The baseboard may optionally tie the BIF[0:2]# signals to 3.3Vaux or to ground per the definitions are described in Section 3.6 if no dynamic bifurcation configuration is required. For add-in cards, these signals shall connect to the endpoint bifurcation pins if it is supported. Note: the required combinatorial logic output for endpoint bifurcation is dependent on the specific silicon and is not defined in this specification.

Figure 46: PCIe Present and Bifurcation Control Pins

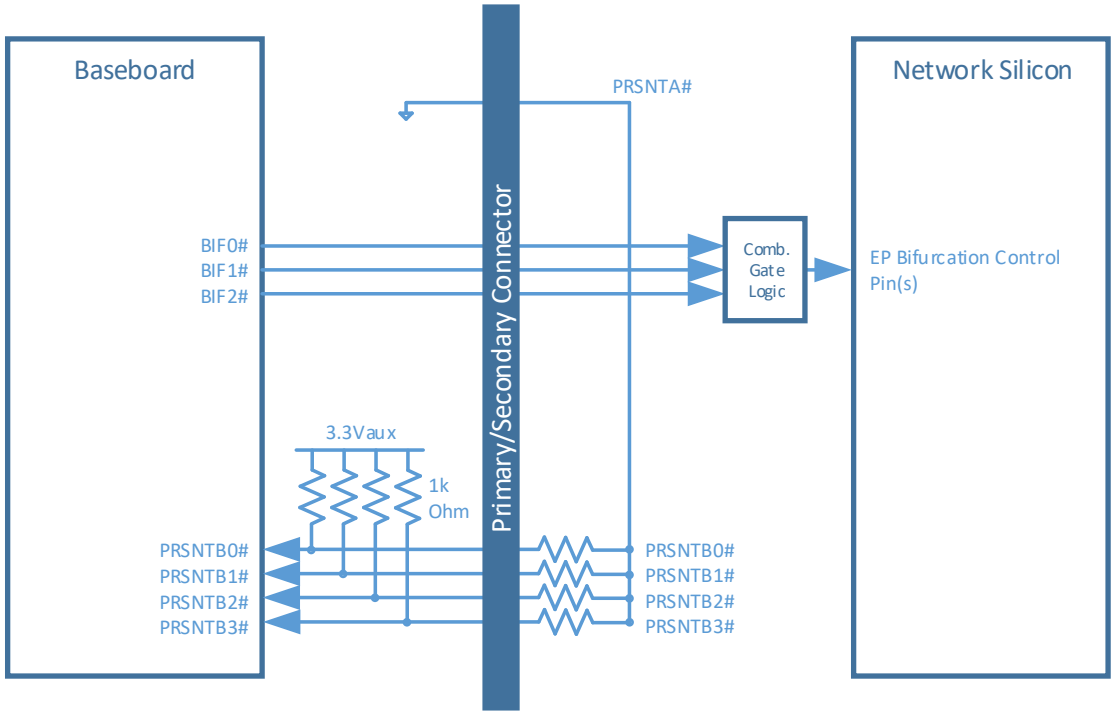

3.4.3 SMBus Interface Pins

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 XXX.

Table 16: Pin Descriptions – SMBus

Signal Name	Pin #	Baseboard Direction	Signal Description
SMCLK	A7	Output, OD	SMBus clock. Open drain, pulled up to 3.3Vaux on the baseboard. For baseboards, the SMCLK from the platform SMBus master shall be connected to the connector. For add-in cards, the SMCLK from the endpoint silicon shall be connected to the card edge gold fingers.
SMDAT	A8	Input / Output, OD	SMBus Data. Open drain, pulled up to 3.3Vaux on the baseboard. For baseboards, the SMDAT from the platform SMBus master shall be connected to the connector. For add-in cards, the SMDAT from the endpoint silicon shall be connected to the card edge gold fingers.
SMRST#	A9	Output, OD	SMBus reset. Open drain. For baseboards, this pin shall be pulled up to 3.3Vaux. 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. For add-in cards, SMRST# is optional and is dependent on the add-in card implementation.

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 47.

Table 17: 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 shall be rated to 1.1A per pin with a maximum derated power delivery of 80W.

			The +12V power pins shall be within the rail tolerances as defined in Section 3.10 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 shall be rated to 1.1A for a maximum derated power delivery of 3.63W. The 3.3Vaux/main power pin shall be within the rail tolerances as defined in Section 3.10 when the PWRDIS pin is driven low by the baseboard.
PWRDIS	B12	Output, O/D	Power disable. Active high. Open-drain This signal shall be pulled up to 3.3V through a 10kOhm resistor on the baseboard. When high, all add-in card supplies shall be disabled. When low, add-in card supplies shall be enabled.

Figure 47: Example Power Supply Topology

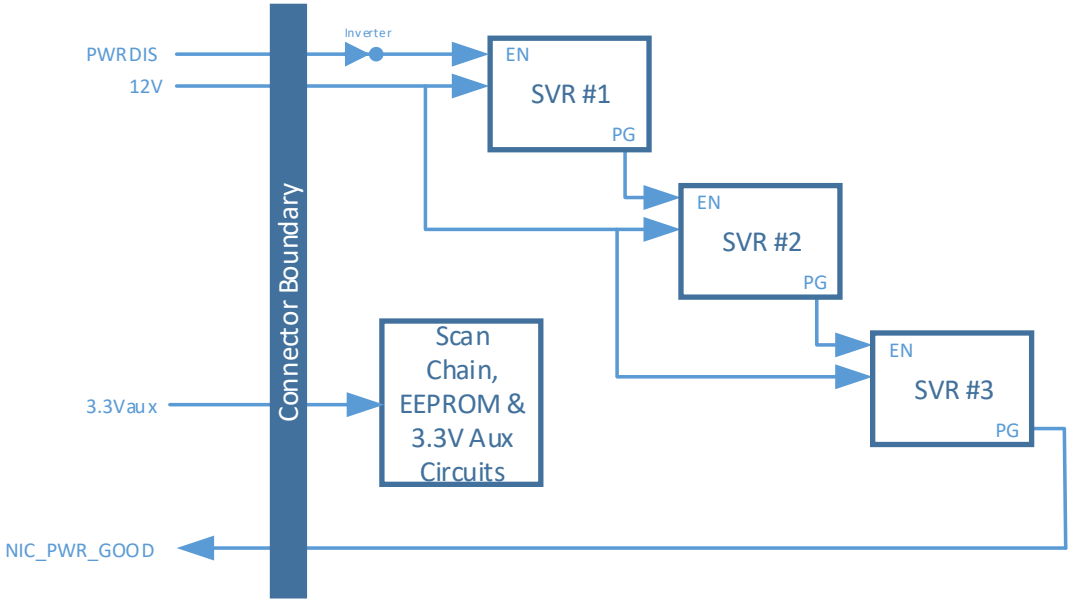

3.4.5 Miscellaneous Pins

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

Table 18: Pin Descriptions – Miscellaneous 1

Signal Name	Pin #	Baseboard Direction	Signal Description

RFU, N/C	B68, B69, A68, A69, A70	Input / Output	Reserved future use pins. These pins shall be left 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 0 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 add-in card shall implement protection methods to prevent leakage paths between the Vaux and Vmain 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 Section XXX. An example connection diagram that shows REFCLK2, REFCLK3, PERST2# and PERST3# is shown in Section 0.

Table 19: 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 add-in card PCIe core logic. For baseboards, the REFCLK2 and REFCLK3 signals are required at the Primary connector. For add-in cards, the required REFCLKs shall be connected per the endpoint datasheet. **Note:** REFCLK2 and REFCLK3 are not used for cards that only support a 1 x16, 1 x8 or 2 x8 connection. Refer to Section 2.1 in the PCIe CEM Specification, Rev 4.0 for details.
REFCLKn3 REFCLKp3	OCP_A11 OCP_A12	Output	
PERST2# PERST3#	OCP_A1 OCP_A2	Output	PCIe Reset #2, #3. Active low. When PERSTn# is deasserted, the signal shall indicate the applied power is within tolerance and stable for the add-in card. 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. PERST shall be pulled high to 3.3Vaux on the baseboard. For OCP NIC 3.0, PERST deassertion shall also indicate the full card power envelope is available to the add-in card. For baseboards, the PERST[0:1]# signals are required at the connector. For add-in cards, the required PERST[0:1]# signals shall be connected to the endpoint silicon. **Note:** PERST2# and PERST3# are not used for cards that only support a 1 x16 or 2 x8 connection. Refer to Section 2.2 in the PCIe CEM Specification, Rev 4.0 for details.
WAKE#	OCP_A3	Input, OD	WAKE#. Open drain. Active low. This signal shall be driven by the add-in card to notify the baseboard to restore 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. For baseboards, this signal shall be pulled up to +3.3V on the baseboard with a 10kOhm resistor. This signals shall be connected to the system WAKE# signal. For add-in cards, this signal shall be directly connected to the endpoint silicon WAKE# pin(s). 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 are defined in the NC-SI specification. An example connection diagram is shown in Figure 48.

Table 20: 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 shall have a nominal frequency of 50MHz ±100ppm. 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. For add-in 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 is not supported.
RBT_CRS_DV	OCP_B14	Input	Carrier sense/receive data valid. This signal is used to indicate to the baseboard that the carrier sense/receive data is valid. 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. For add-in 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 is not supported.
RBT_RXD0 RBT_RXD1	OCP_B9 OCP_B8	Input	Receive data. Data signals from the network controller to the BMC. 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 3.3Vaux on the baseboard. If the

			baseboard does not support NC-SI over RBT, then this signal shall be terminated to 3.3Vaux through a 100kOhm pull-up. For add-in 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 is not supported.
RBT_TX_EN	OCP_A7	Output	Transmit enable. 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. For add-in 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 is not supported.
RBT_TXD0 RBT_TXD1	OCP_A9 OCP_A8	Output	Transmit data. Data signals from the BMC to the network controller. 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 3.3Vaux on the baseboard. If the baseboard does not support NC-SI over RBT, then this signal shall be terminated to 3.3Vaux through a 100kOhm pull-up. For add-in 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 is not supported.
RBT_ARB_OUT	OCP_A5	Output	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. The baseboard shall implement a multiplexing implementation that directs the RBT_ARB_OUT to the RBT_ARB_IN pin of the next NC-SI 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. 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. For add-in cards, this pin shall be connected from the gold finger to the RBT_ARB_IN pin on the endpoint silicon. This pin shall be left as a no connect if NC-SI is not supported.
RBT_ARB_IN	OCP_A4	Input	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. The baseboard shall implement a multiplexing implementation that directs the RBT_ARB_IN to the RBT_ARB_OUT pin of the next NC-SI 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. 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. For add-in cards, this pin shall be connected between the gold finger to the RBT_ARB_OUT pin on the endpoint silicon. Leave this pin as a no connect if NC-SI is not supported.
SLOT_ID	OCP_B7	Output	NC-SI Address pin. This pin shall only be used if the end point silicon supports package identification. For baseboards, this pin shall be used to identify the slot ID value. This pin shall be directly to GND for SlotID = 0. This pin shall be pulled up to 3.3Vaux for SlotID = 1. For add-in cards, this pin shall be connected to the endpoint device GPIO associated with the Package ID[1] field. Refer to Section 4.9.1 and the device datasheet for details.

			For add-in 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. For endpoint devices without NC-SI support, this pin shall be left as a no connect on the add-in card.

Figure 48: NC-SI Over RBT Connection Example – Single Primary Connector

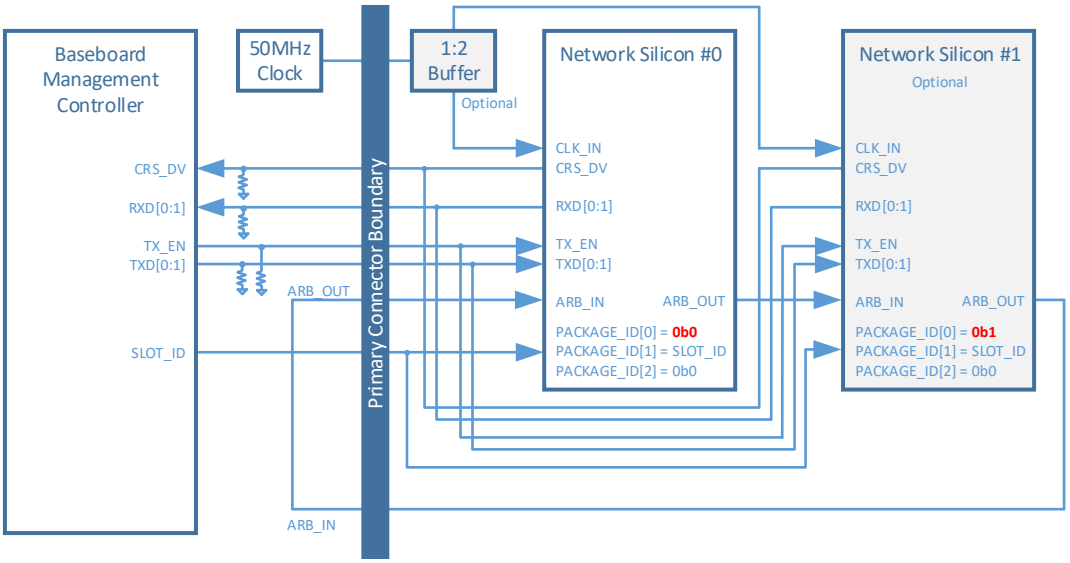

Figure 49: NC-SI Over RBT Connection Example – Dual Primary Connector

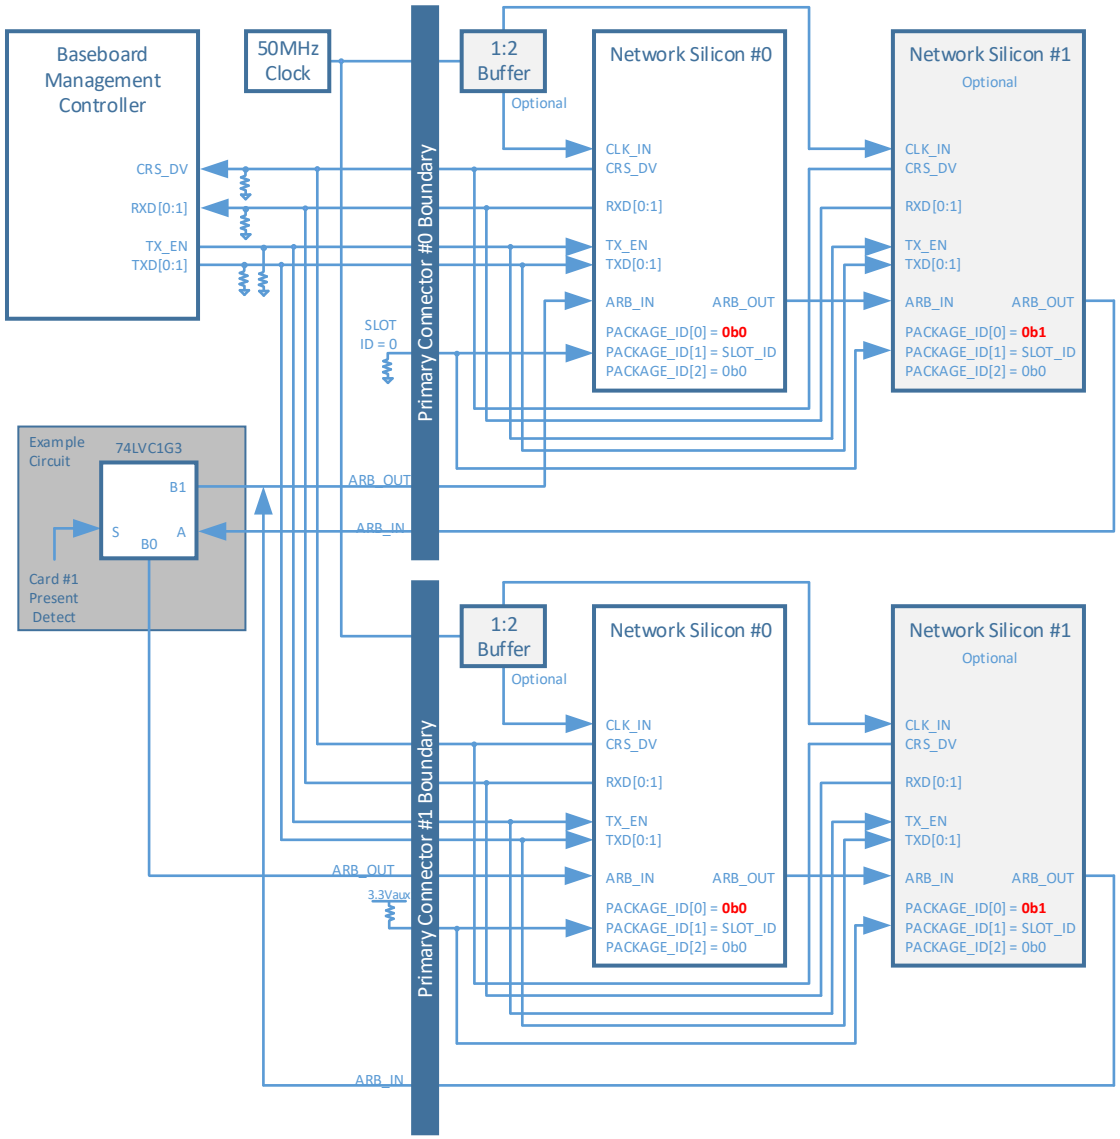

Note 1: For baseboard designs with a single Primary Connector, connect ARB_IN to ARB_OUT to complete the NC-SI hardware arbitration ring. For designs with multiple Primary Connectors, connect ARB_IN and ARB_OUT to an analog mux to complete the NC-SI arbitration ring based on the number of cards installed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

Figure 49.

Note 2: For add-in 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 AC/DC specifications are defined in Section XXX. An example connection diagram is shown in Figure 50.

Table 21: Pin Descriptions – Scan Chain

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, 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. 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 50, below. The CLK pin shall be pulled up to 3.3Vaux through a 1kOhm resistor.
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. 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. For NIC implementations, the DATA_OUT pin may be left floating if it is not used for add-in card configuration. The DATA_OUT pin shall be pulled up to 3.3Vaux through a 1kOhm resistor.
DATA_IN	OCP_B4	Input	Scan clock data input to the baseboard. This bit stream is used to shift out NIC status bits. For baseboard implementations, the DATA_IN pin shall be pulled up to 3.3Vaux 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. 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 50.
LD#	OCP_B3	Output	Scan clock shift register load. Used to latch configuration data on the add-in card. For baseboard implementations, the LD# pin shall be pulled up to 3.3Vaux through a 1kOhm resistor if the scan chain is not used to prevent the add-in card from erroneous data latching. 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 50. The LD# pin shall be pulled up to 3.3Vaux through a 1kOhm resistor.

The scan chain provides side band status indication between the add-in 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 3.3Vaux power domain.

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 and subsequent revisions of this specification.

Table 22: Pin Descriptions – Scan Chain DATA_OUT Bit Definition

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

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 (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 ob1 for unused bits for implementations using less than four shift registers.

Table 23: Pin Descriptions – Scan Bus DATA_IN Bit Definition

Byte.bit	DATA_OUT Field Name	Default Value	Description
0.0	PRSNTB[0]#	obX	PRSNTB[3:0]# bits shall reflect the same state as the signals on the Primary Connector.
0.1	PRSNTB[1]#	obX	
0.2	PRSNTB[2]#	obX	
0.3	PRSNTB[3]#	obX	
0.4	WAKE_N	obX	PCIe WAKE_N signal shall reflect the same state as the signal on the Primary Connector.
0.5	TEMP_WARN	obo	Temperature monitoring pin from the on-card thermal solution. This pin shall be asserted high when temperature sensor exceeds the temperature warning threshold.
0.6	TEMP_CRIT	obo	Temperature monitoring pin from the on-card thermal solution. This pin shall be asserted high when temperature sensor exceeds the temperature critical threshold.
0.7	FAN_ON_AUX	obo	When high, FAN_ON_AUX shall request the system fan to be enabled for extra cooling in the S5 state.
1.0	LINK_ACT0	ob1	Port 0..3 link/activity indication. Active low. obo – Link LED is illuminated on the host platform. ob1 – Link LED is not illuminated on the host platform. **Steady** = link is detected on the port. **Blinking** = activity is detected on the port. The blink rate should blink low for 50-500ms during activity periods. **Off** = the physical link is down or disabled
1.1	LINK_ACT1	ob1	
1.2	LINK_ACT2	ob1	
1.3	LINK_ACT3	ob1	
1.4	SPEED_A0	ob1	Port 0..3 speed A (max rate) indication. Active low. obo – Port is linked at maximum speed. ob1 – Port is not linked at the maximum speed or no link is present.
1.5	SPEED_A1	ob1	
1.6	SPEED_A2	ob1	
1.7	SPEED_A3	ob1	
2.0	ScanChainVer[0]	ob1	ScanChainVer[1:0] shall be used to indicate the scan chain bit definition version. The encoding shall be as follows:
2.1	ScanChainVer[1]	ob1	

			ob11 – Scan chain bit definition version 1 corresponding to OCP NIC 3.0 version 1.0. All other encoding values shall be reserved.
2.2	RSVD	ob1	Byte 2 bits [2:7] are reserved. These bits shall default to the value of ob1. These bits may be used in future versions of the scan chain.
2.3	RSVD	ob1	
2.4	RSVD	ob1	
2.5	RSVD	ob1	
2.6	RSVD	ob1	
2.7	RSVD	ob1	
3.0	LINK_ACT4	ob1	Port 4..7 link/activity indication. Active low. obo – Link LED is illuminated on the host platform. ob1 – Link LED is not illuminated on the host platform. **Steady** = link 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	LINK_ACT5	ob1	
3.2	LINK_ACT6	ob1	
3.3	LINK_ACT7	ob1	
3.4	SPEED_A4	ob1	Port 4..7 speed A (max rate) indication. Active low. obo – Port is linked at maximum speed. ob1 – Port is not linked at the maximum speed or no link is present.
3.5	SPEED_A5	ob1	
3.6	SPEED_A6	ob1	
3.7	SPEED_A7	ob1	

Figure 50: Scan Bus Connection Example

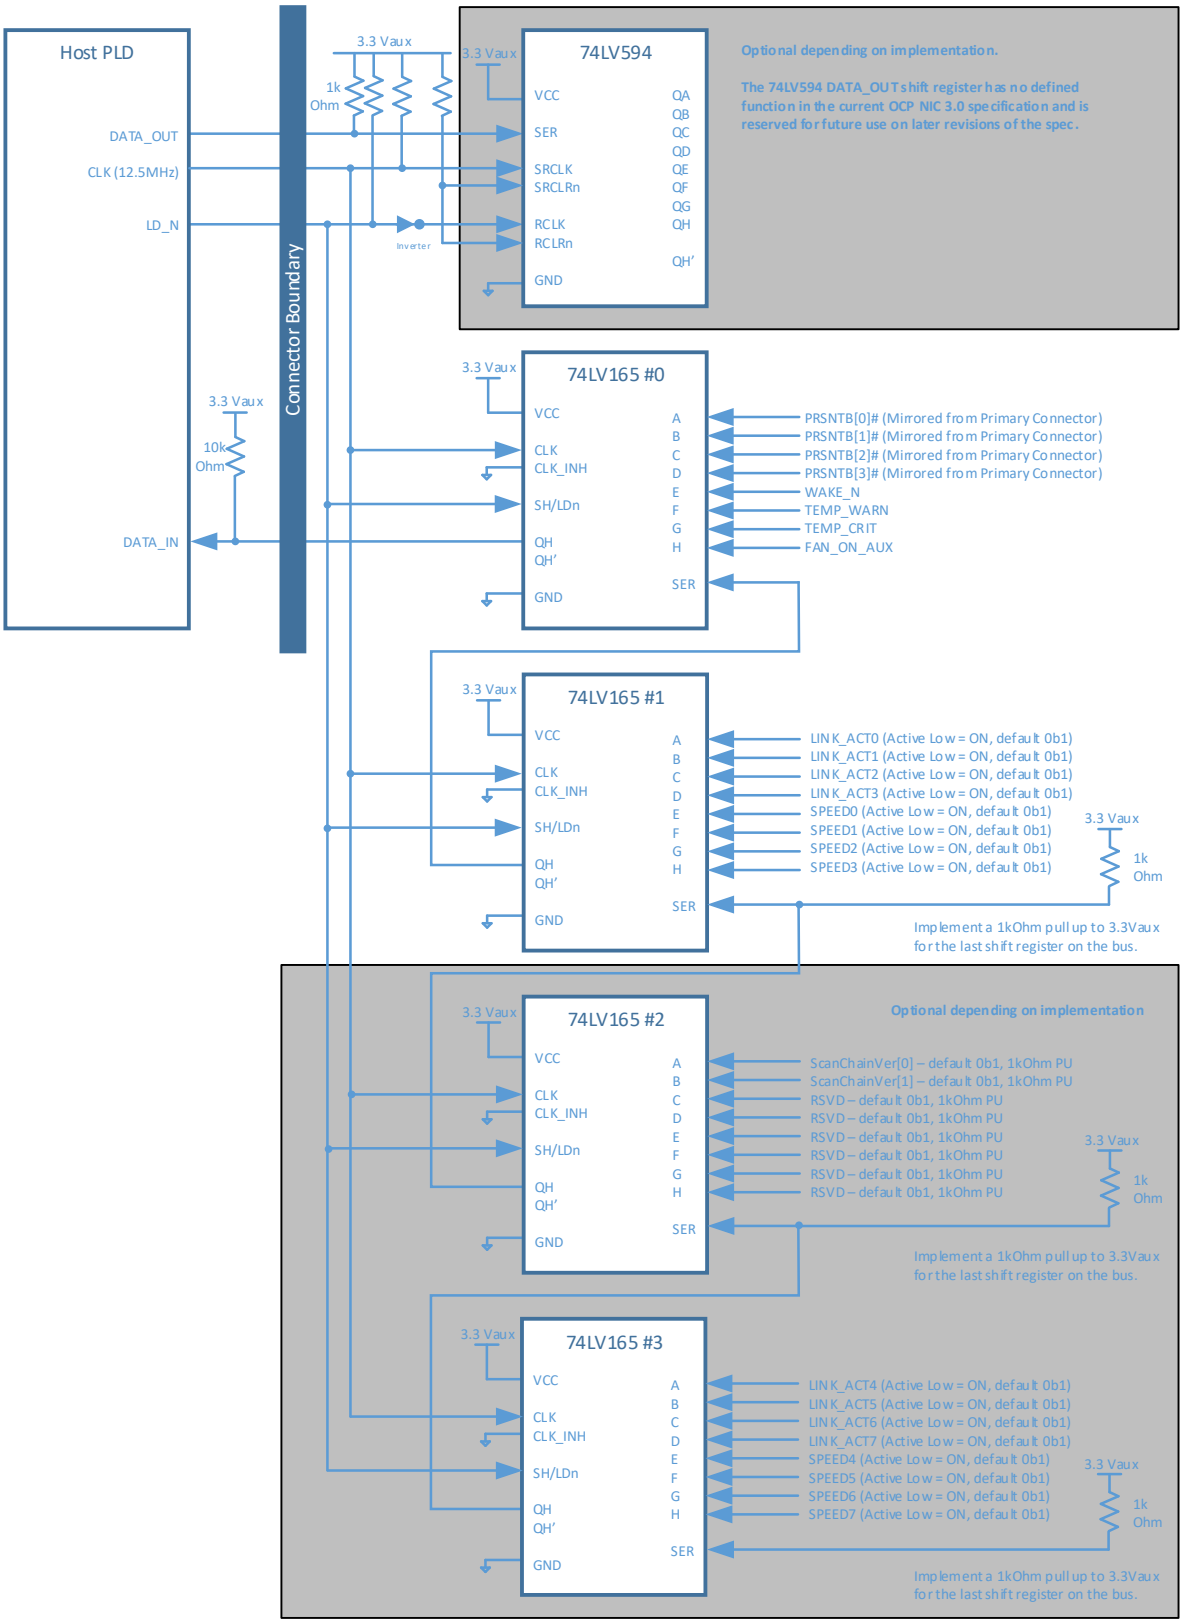

3.5.4 Primary Connector Miscellaneous Pins – OCP Bay (Primary Connector)

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 connection diagram is shown in Figure XXX.

Table 24: Pin Descriptions – Miscellaneous 2

Signal Name	Pin #	Baseboard Direction	Signal Description
PWRBRK#	OCP_B2	Output, OD	Power break. Active low, open drain. This signal shall be pulled up to 3.3Vaux on the add-in 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. When this signal is driven low by the baseboard, the Emergency Power Reduction State is requested. The add-in card shall move to a lower power consumption state.
NIC_PWR_GOOD	OCP_B1	Input	NIC Power Good. Active high. This signal is driven by the add-in card. When high, this signal shall indicate that all of the add-in card power rails are operating within nominal tolerances. When low, this signal shall indicate that the add-in card power supplies are not yet within nominal tolerances or are in a fault condition. 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 add-in card is present. For add-in cards this signal shall indicate the add-in card power is "good". This signal may be implemented by a cascaded power good or a discrete power good monitor output.
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 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 shall connect 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 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 51.

Figure 51: PCIe Bifurcation Pin Connections Support

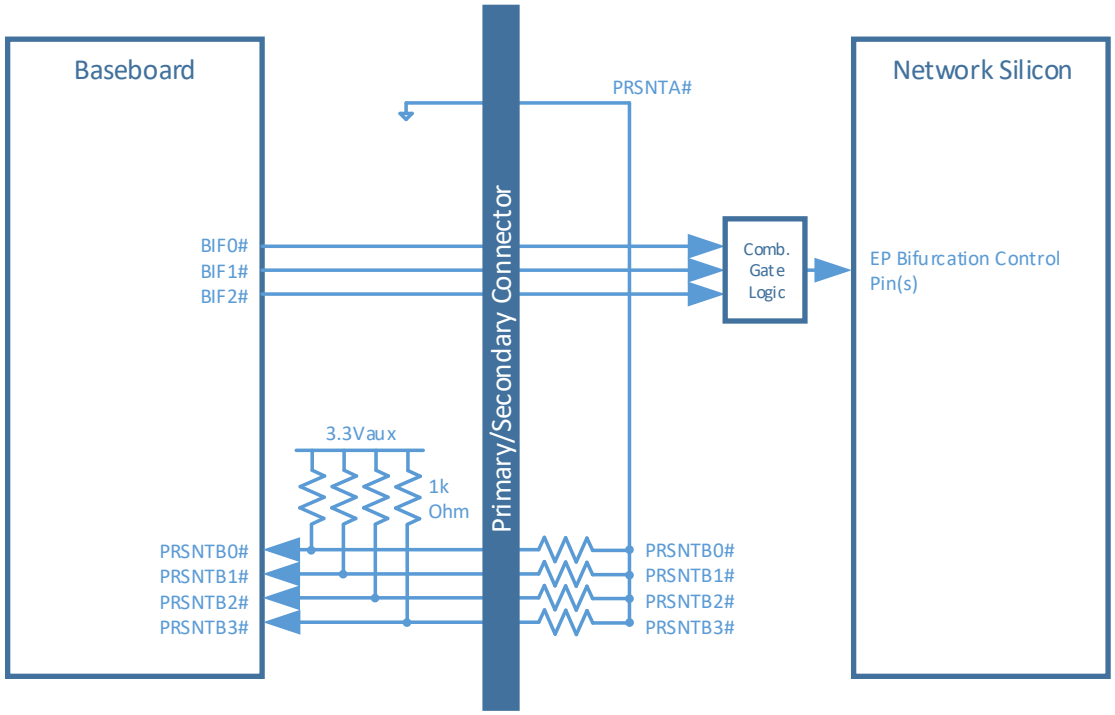

3.6.1 PCIe Add-in Card to Baseboard Bifurcation Configuration (PRSNTA#, PRSNTB[3:0]#)

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 3.3Vaux on the baseboard and are active low signals. A state of ob1111 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 25 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 25 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 combination 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 25 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 25: PCIe Bifurcation Decoder for x16 and x8 Card Widths

Minimum Required Card Edge	Network Card – Supported PCIe Configurations			Host	Single Host					RSVD	Dual Host		RSVD	Quad Host		QuadOct Host
					1 Host	1 Host	1 Host	1 Host		RSVD	2 Hosts		RSVD	4 Hosts		4 or 8 Hosts
				Host CPU Sockets	1 Upstream Socket	1 Upstream Socket	1 Upstream Socket	2 Upstream Sockets	4 Upstream Sockets	RSVD	2 Upstream Sockets (1 Socket per Host)	4 Upstream Sockets (1 Socket per Host)	RSVD	4 Upstream Sockets (1 Socket per Host)	4/8 Upstream Sockets (1 Socket per Host)	
				Total PCIe Links	1 Link (No Bifurcation)	1 or 2 Links	1, 2, or 4 Links	2 Links	4 Links	RSVD	2 Links	4 Links	RSVD	4 Links	4 or 8 x2 links	
				System Support	1x16, 1x8, 1x4, 1x2, 1x1	1x16, 1x8, 1x4, 1x2, 1x1	1x16, 1x8, 1x4, 1x2, 1x1	1x8, 1x4, 1x2, 1x1		RSVD			RSVD			
						2 x8, 2 x4, 2 x2, 2 x1	2 x8, 2 x4, 2 x2, 2 x1	2 x8, 2 x4, 2 x2, 2 x1			2 x8, 2 x4, 2 x2, 2 x1					
							4 x4, 4 x2, 4 x1	4 x4, 4 x2, 4 x1				4 x4, 4 x2, 4 x1		4 x2, 4 x1		
			System Encoding		0b000	0b000	0b000	0b001	0b010	0b011	0b101	0b110	0b100	0b110	0b111	
		Card Short Name	**Add-in Card Encoding PRSNTB[3:0]#**		-	-	-	-	-	-	-	-	-	-	-	-
n/a	Not Present	Card Not Present	0b1111	RSVD - Card not present in the system												
2C	1x8	1x8, 1x4, 1x2, 1x1	0b1110		1x8	1x8	1x8	1x8 (Socket 0 only)	1x4 (Socket 0 only)	-	1x8 (Host 0 only)		-	1x4 (Host 0 only)		1x2 (Host 0 only)
2C	1x4	1x4, 1x2, 1x1	0b1110		1x4	1x4	1x4	1x4 (Socket 0 only)	1x4 (Socket 0 only)	-	1x4 (Host 0 only)		-	1x4 (Host 0 only)		1x2 (Host 0 only)
2C	1x2	1x2, 1x1	0b1110		1x2	1x2	1x2	1x2 (Socket 0 only)	1x2 (Socket 0 only)	-	1x2 (Host 0 only)		-	1x2 (Host 0 only)		1x2 (Host 0 only)
2C	1x1	1x1	0b1110		1x1	1x1	1x1	1x1 (Socket 0 only)	1x1 (Socket 0 only)	-	1x1 (Host 0 only)		-	1x1 (Host 0 only)		1x1 (Host 0 only)
2C	1x8 Option B	1x8, 1x4, 1x2, 1x1 2 x4, 2 x2, 2 x1	0b1101		1x8	1x8	1x8	1x8 (Socket 0 only)	2 x4	-	1x8 (Host 0 only)		-	2 x4		2 x2 (Host 0 & 1 only)
4C	2 x8 Option B	2 x8, 2 x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	0b1101		1x8*	2 x8	2 x8	2 x8	4 x4	-	2 x8		-	4 x4		2 x2 (Host 0 & 1 only)
2C	1x8 Option D	1x8, 1x4 2 x4, 4 x2 (First 8 lanes), 4 x1	0b1100		1x8	2 x4	2 x4	1x8 (Socket 0 only)	2 x4	-	1x8 (Host 0 only)		-	2 x4		4 x2
4C	1x16 Option D	1x16, 1x8, 1x4 2 x8, 2 x4, 4 x4, 4 x2 (First 8 lanes), 4 x1	0b1100		1x16	1x16	1x16	2 x8	4 x4	-	2 x8		-	4 x4		4 x2
RSVD	RSVD	RSVD	0b1011	RSVD - The encoding of 0b1011 is reserved due to insufficient spacing between PRSNTA and PRSNTB2 pin to provide positive card identification.												
2C	2 x4	2 x4, 2 x2, 2 x1 1 x4, 1 x2, 1 x1	0b1010		1x4	2 x4	2 x4	1x4 (Socket 0 only)	2 x4	-	1x4 (Host 0 only)		-	2 x4		2 x2 (Host 0 & 1 only)
RSVD	RSVD	RSVD for future x8 encoding	0b1001		-	-	-	-	-	-	-		-	-		-
RSVD	RSVD	RSVD for future x8 encoding	0b1000		-	-	-	-	-	-	-		-	-		-
4C	1x16	1x16, 1x8, 1x4, 1x2, 1x1	0b0111		1x16	1x16	1x16	1x8 (Socket 0 only)	1x4 (Socket 0 only)	-	1x8 (Host 0 only)		-	1x4 (Host 0 only)		1x2 (Host 0 only)
4C	2 x8 Option A	2 x8, 2 x4, 2 x2, 2 x1	0b0110		1x8*	2 x8	2 x8	2 x8	2 x4 (Socket 0 & 2 only)	-	2 x8		-	2 x4 (Host 0 & 2 only)		1 x2 (Host 0 & 1 only)
4C	1x16 Option B	1x16, 1x8, 1x4, 1x2, 1x1 2 x8, 2 x4, 2 x2, 2 x1	0b0101		1x16	1x16	1x16	2 x8	2 x4 (Socket 0 & 2 only)	-	2 x8		-	2 x4 (Host 0 & 2 only)		2 x2 (Host 0 & 1 only)
4C	1x16 Option C	1x16, 1x8, 1x4 2 x8, 2 x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	0b0100		1x16	1x16	1x16	2 x8	4 x4	-	2 x8		-	4 x4		2 x2 (Host 0 & 1 only)
4C	4 x4	4 x4, 4 x2, 4 x1	0b0011		1x4*	2 x4*	4 x4	2 x4 (EP, 0 and 2 only)	4 x4	-	2 x4 (EP, 0 and 2 only)		-	4 x4		4 x2 (Host 0 & 1 only)
RSVD	RSVD	RSVD	0b0010		-	-	-	-	-	-	-		-	-		-
RSVD	RSVD	RSVD	0b0001		-	-	-	-	-	-	-		-	-		-
RSVD	RSVD	RSVD	0b0000		-	-	-	-	-	-	-		-	-		-

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 shall read the state of the PRSNTB[3:0]# pins. An add-in card is present in the system if the resulting value is not ob1111.
2. Firmware determines the add-in card PCIe lane width capabilities per Table 25 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 add-in card supports the requested link override.
5. PERST# shall be deasserted after the >100ms window as defined by the PCIe specification. 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.

Figure 52 illustrates a single host baseboard that supports x16 with a single controller add-in card that also supports x16. The PRSTNB[3:0]# state is ob0111. The BIF[2:0]# state is obooo 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 52: Single Host (1 x16) and 1 x16 Add-in Card (Single Controller)

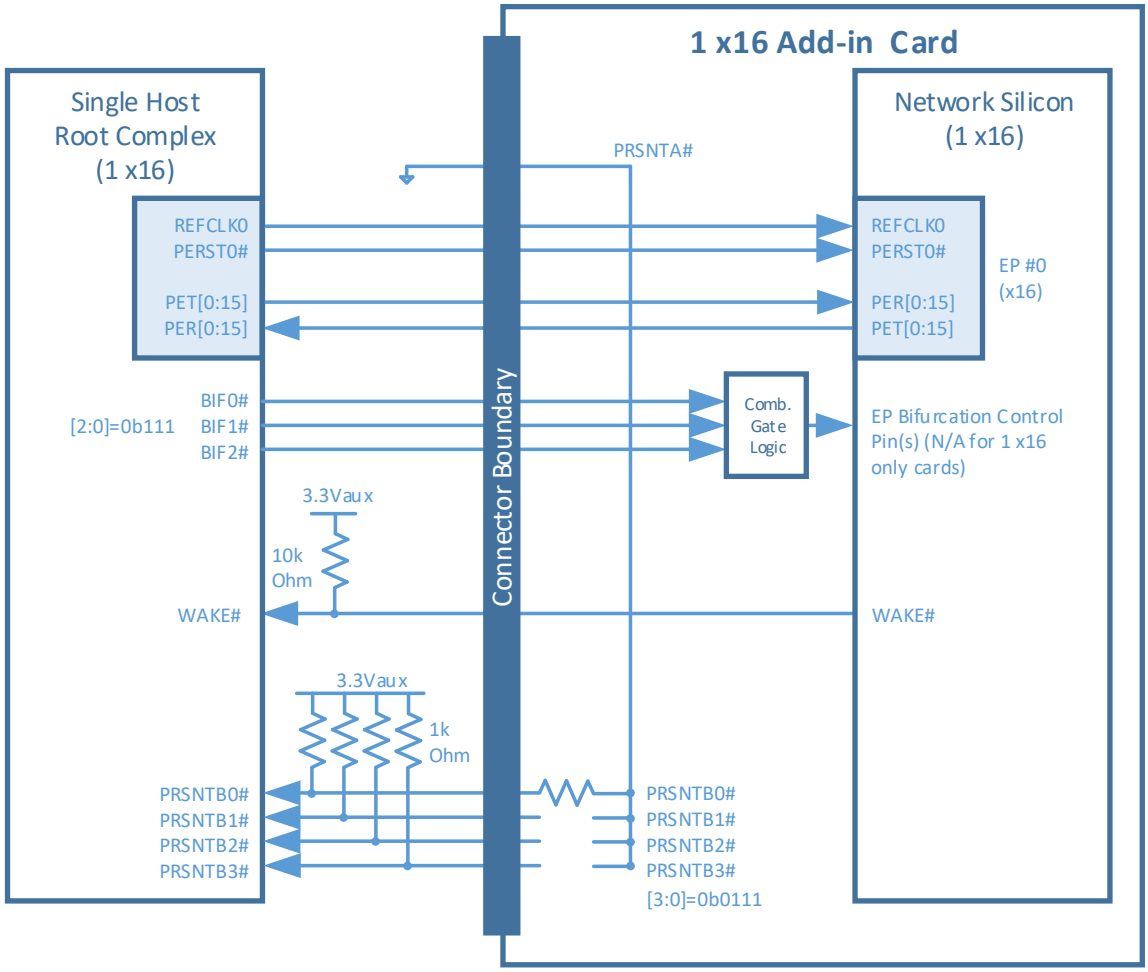

Figure 53 illustrates a single host baseboard that supports 2 x8 with a single controller add-in card that also supports 2 x8. The PRSTNB[3:0]# state is ob0110. The BIF[2:0]# state is ob111 as there is no need to instruct the end-point network controllers to a specific bifurcation. The PRSNTB encoding 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 53: Single Host (2 x8) and 2 x8 Add-in Card (Dual Controllers)

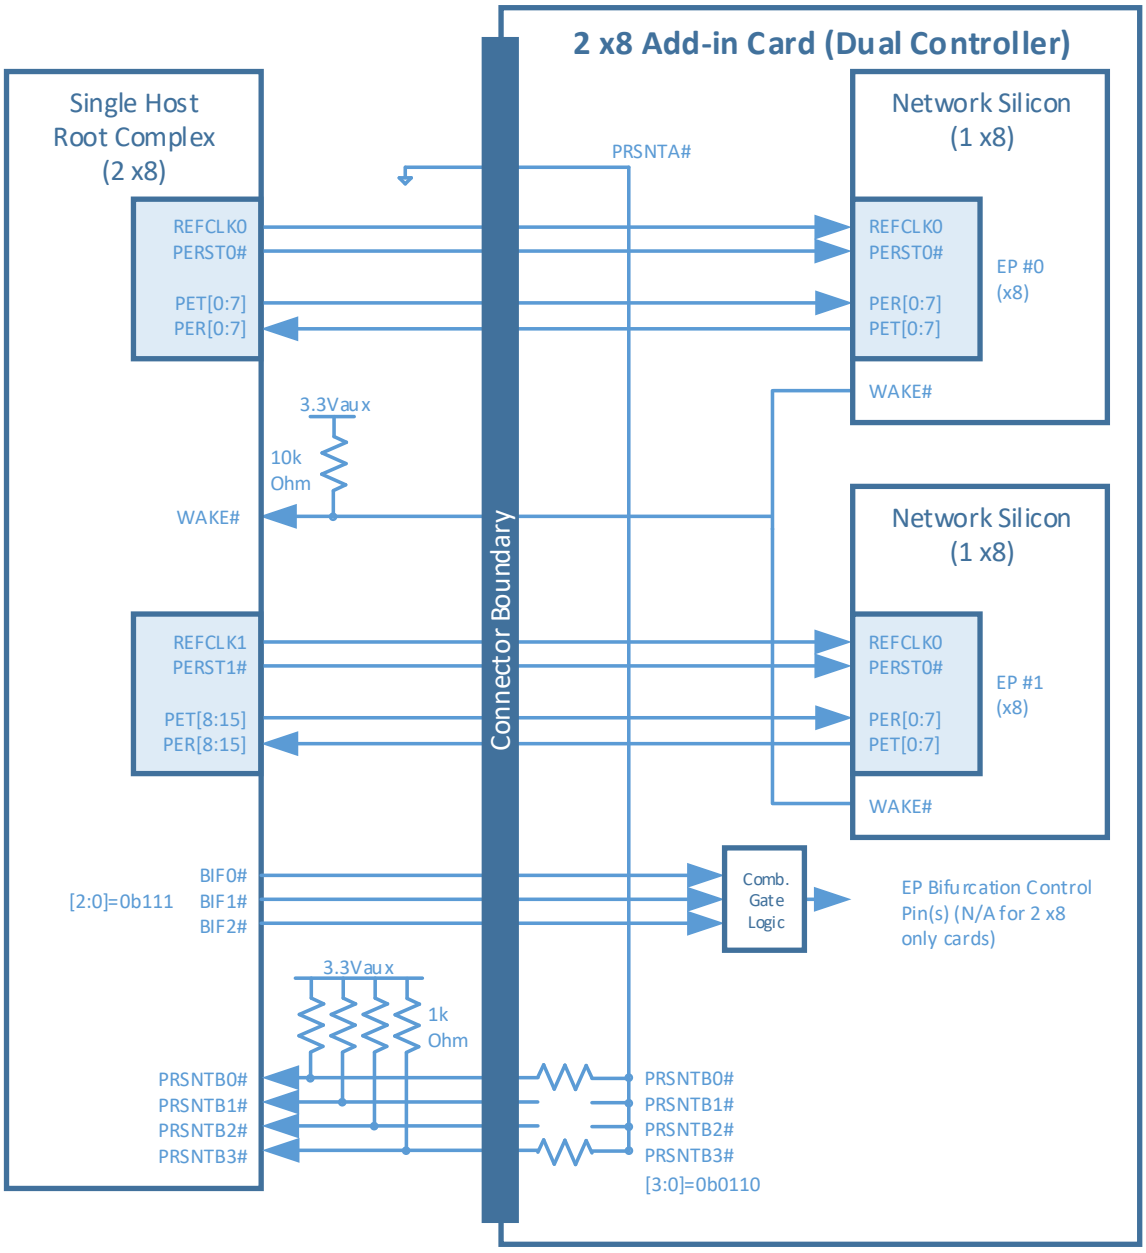

Figure 54 illustrates a four host baseboard that supports 4 x4 with a single controller add-in card that supports 1 x16, 2 x8 and 4 x4. The PRSTNB[3:0]# state is ob0011. The BIF[2:0]# state is ob101 as the end point network controller is forced to bifurcate to 4 x4. The PRSNTB encoding notifies the baseboard 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 54: Four Hosts (4 x4) and 4 x4 Add-in Card (Single Controller)

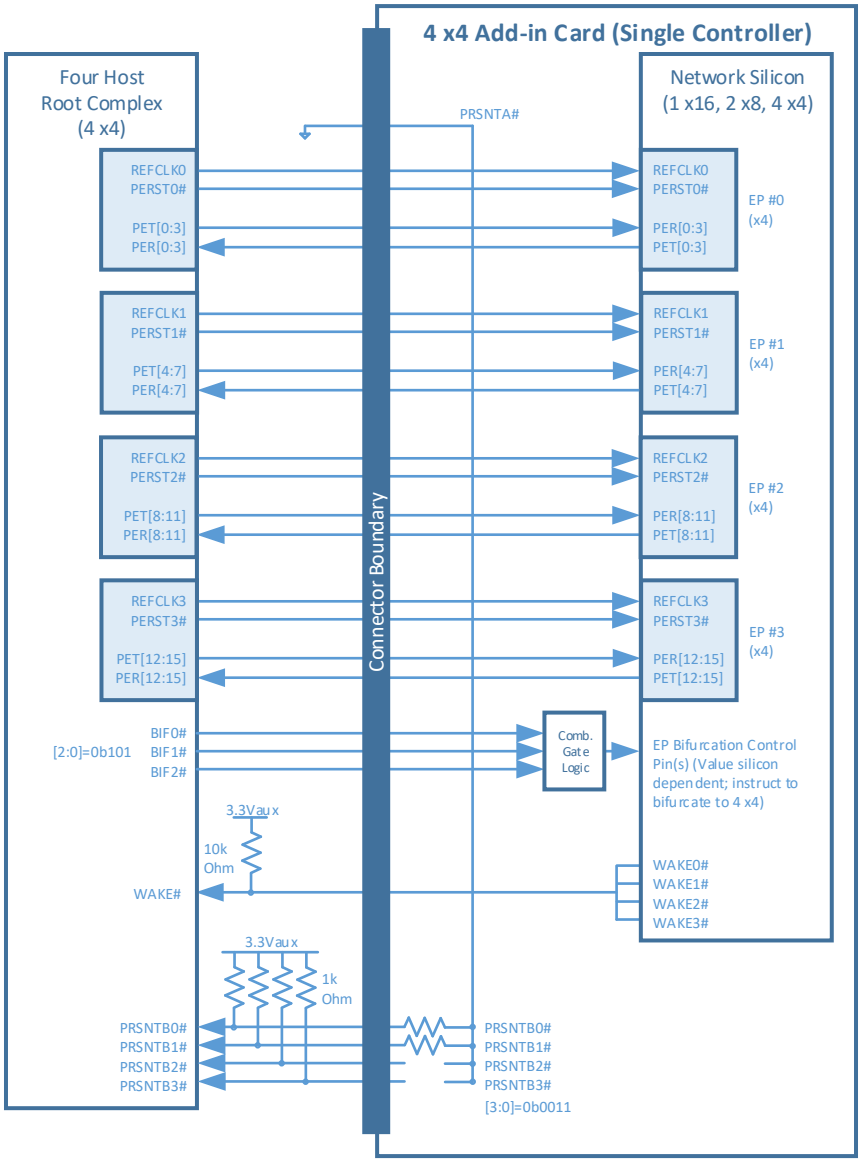

Figure 55 illustrates a four host baseboard that supports 4 x4 with a four controller add-in card that supports 4 x4. The PRSTNB[3:0]# state is ob0011. The BIF[2:0]# state is ob111 as there is no need to instruct the end-point network controllers to a specific bifurcation. The PRSNTB encoding notifies 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 55: Four Hosts (4 x4) and 4 x4 Add-in Card (Four Controllers)

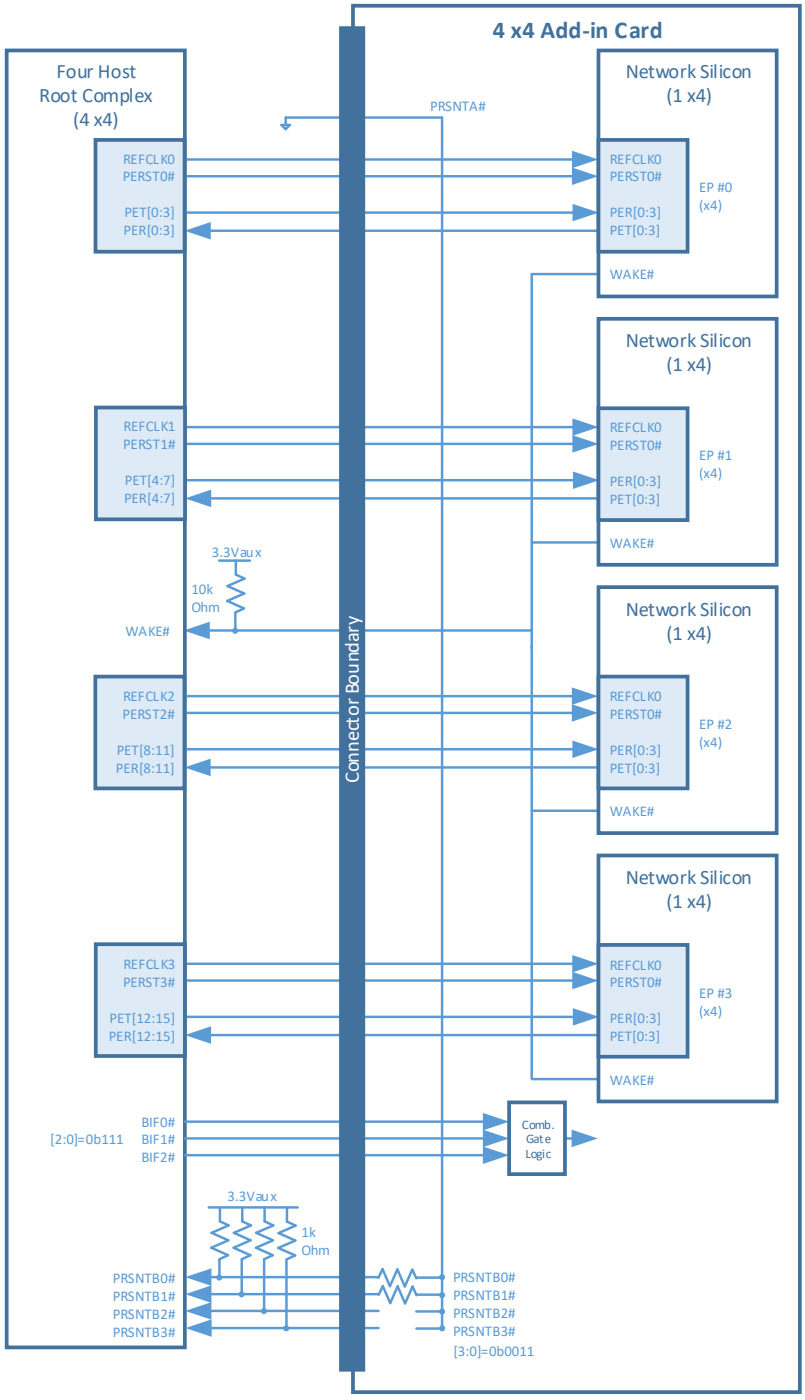
4 x4 Add-in Card
Four Host Root Complex (4 x4)
PRSNTA#
Network Silicon (1 x4)
REFCLK0
PERST0#
PET[0:3]
PER[0:3]
REFCLK0
PERST0#
PER[0:3]
PET[0:3]
EP #0 (x4)
WAKE#
Network Silicon (1 x4)
REFCLK1
PERST1#
PET[4:7]
PER[4:7]
REFCLK0
PERST0#
PER[0:3]
PET[0:3]
EP #1 (x4)
WAKE#
3.3Vaux
10k Ohm
WAKE#
Network Silicon (1 x4)
REFCLK2
PERST2#
PET[8:11]
PER[8:11]
REFCLK0
PERST0#
PER[0:3]
PET[0:3]
EP #2 (x4)
WAKE#
Connector Boundary
Network Silicon (1 x4)
REFCLK3
PERST3#
PET[12:15]
PER[12:15]
REFCLK0
PERST0#
PER[0:3]
PET[0:3]
EP #3 (x4)
WAKE#
BIF0#
BIF1#
BIF2#
[2:0]=0b111
Comb. Gate Logic
3.3Vaux
1k Ohm
PRSNTB0#
PRSNTB1#
PRSNTB2#
PRSNTB3#
PRSNTB0#
PRSNTB1#
PRSNTB2#
PRSNTB3#
[3:0]=0b0011

Figure 56 illustrates a single host baseboard that supports 1 x16 with a dual controller add-in 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 four 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 56: Single Host with no Bifurcation (1 x16) and 2 x8 Add-in Card (Two Controllers)

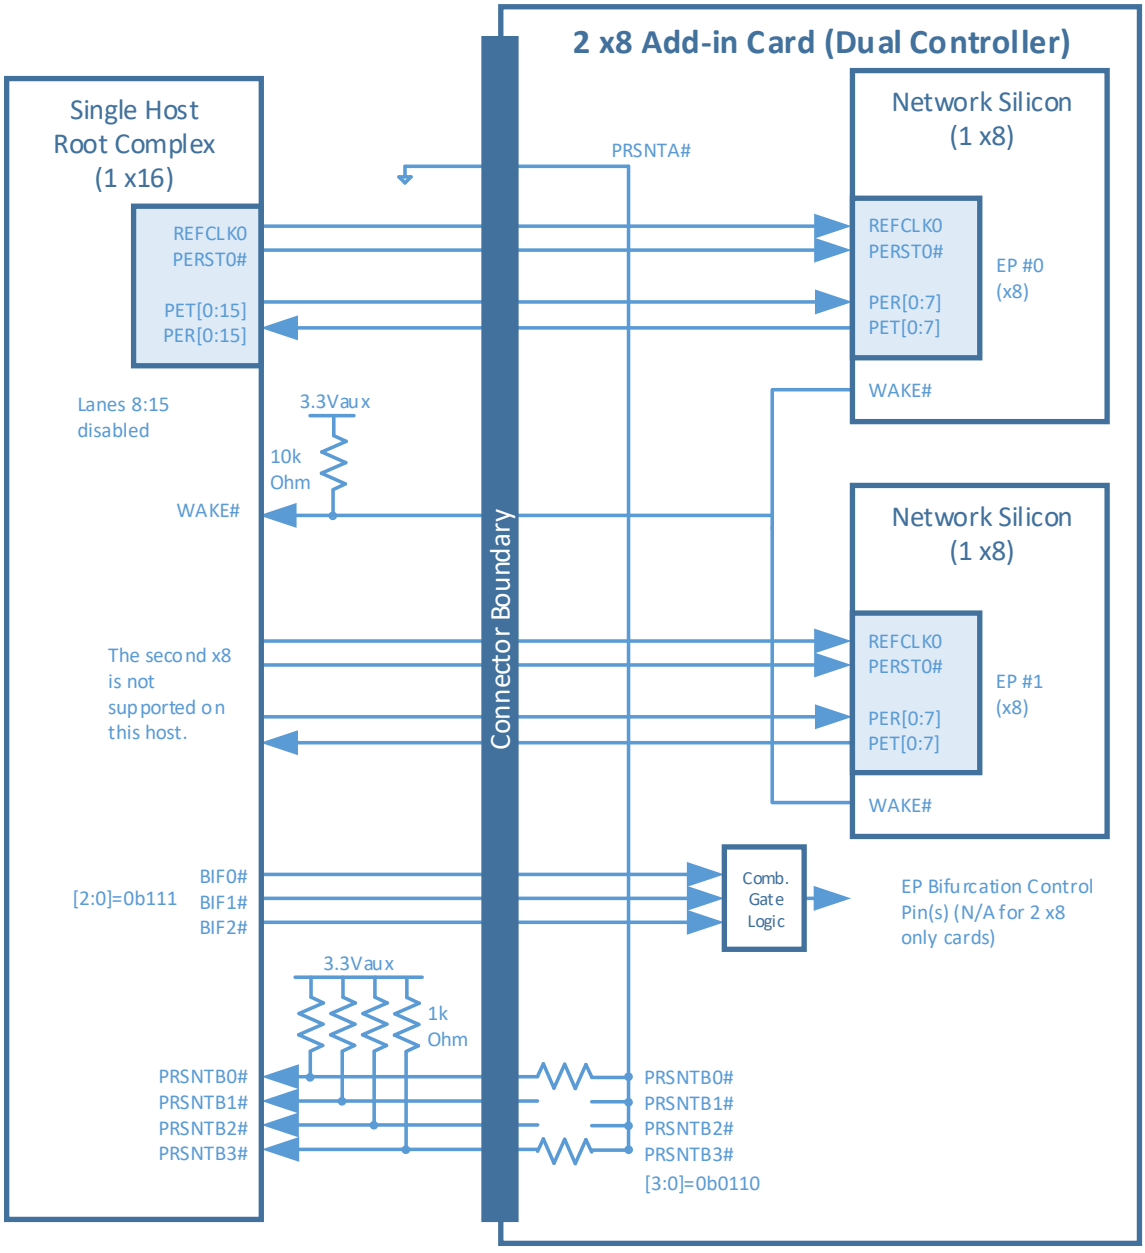

3.7 PCIe Clocking Topology

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 26. Cards that implement both the Primary and Secondary connectors have a total of up to 6 REFCLKs.

Table 26: 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 add-in card, the following REFCLK connection rules must be followed:

- For a 1 x16 capable add-in card, REFCLK0 shall be used for lanes [0:15].
- For a 2 x8 capable add-in card, REFCLK0 shall be used for lanes [0:7] and REFCLK1 shall be used for lanes [8:15].
- For a 4 x4 capable add-in 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 57: PCIe Interface Connections for 1 x16 and 2 x8 Add-in Cards

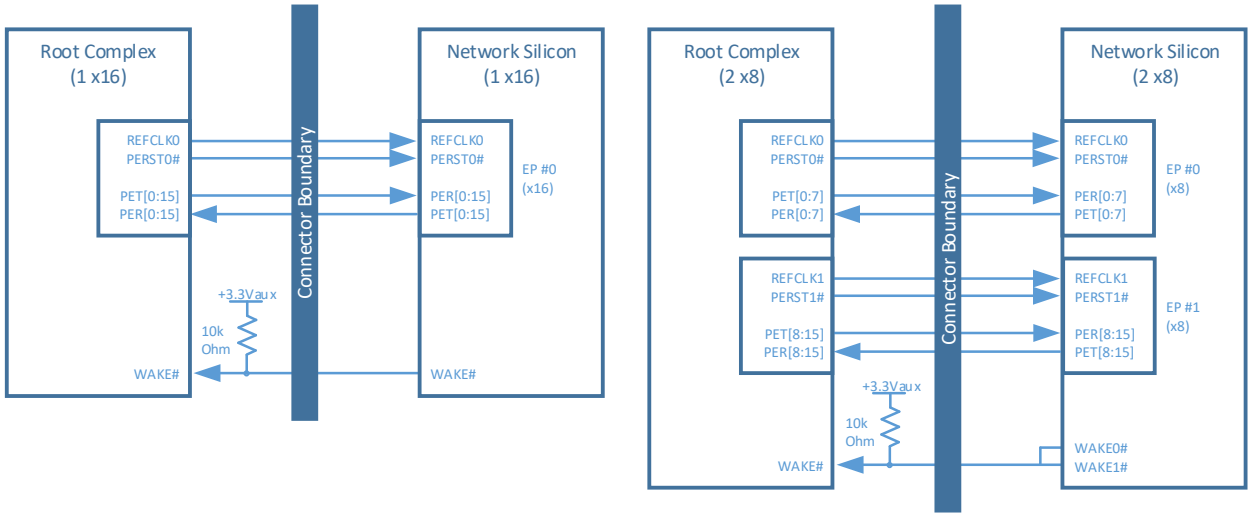

Figure 58: PCIe Interface Connections for a 4 x4 Add-in Card

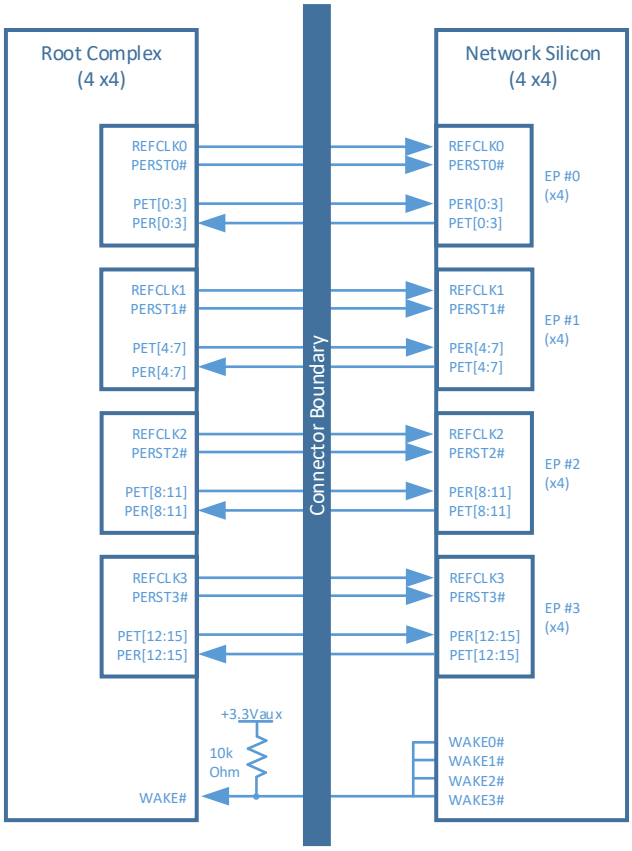

3.8 PCIe Bifurcation Results and REFCLK Mapping

For the cases where the baseboard and add-in 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 27: 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 PRSNT#[3:0]#	Host	1:16, 1:8, 1:4, 1:2, 1	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
n/a	Not Present	Card Not Present	0b1111	1 Host	1 Upstream Socket	1 Link	0b000	-																
2C	1x8	1x8, 1x4, 1x2, 1x1	0b1110	1 Host	1 Upstream Socket	1 Link	0b000	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	0b000	1x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	1x2	1x2, 1x1	0b1110	1 Host	1 Upstream Socket	1 Link	0b000	1x2	Link 0, Lane 0	Link 0, Lane 1														
2C	1x1	1x1	0b1110	1 Host	1 Upstream Socket	1 Link	0b000	1x1	Link 0, Lane 0															
2C	1x8 Option B	1x8, 1x4, 1x2, 1x1 2 x4, 2 x2, 2 x1	0b1101	1 Host	1 Upstream Socket	1 Link	0b000	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	2 x8 Option B	2 x8, 2 x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	0b1101	1 Host	1 Upstream Socket	1 Link	0b000	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	1x8, 1x4 2 x4, 4 x2 (First 8 lanes), 4 x1	0b1100	1 Host	1 Upstream Socket	1 Link	0b000	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	1x16 Option D	1x16, 1x8, 1x4 2 x8, 2 x4, 4 x4, 4 x2 (First 8 lanes), 4 x1	0b1100	1 Host	1 Upstream Socket	1 Link	0b000	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
RSVD	RSVD	RSVD	0b1011	1 Host	1 Upstream Socket	1 Link	0b000	-																
2C	2 x4	2 x4, 2 x2, 2 x1 1 x4, 1 x2, 1 x1	0b1010	1 Host	1 Upstream Socket	1 Link	0b000	1x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
RSVD	RSVD	RSVD for future x8 encoding	0b1001	1 Host	1 Upstream Socket	1 Link	0b000	-																
RSVD	RSVD	RSVD for future x8 encoding	0b1000	1 Host	1 Upstream Socket	1 Link	0b000	-																
4C	1x16	1x16, 1x8, 1x4, 1x2, 1x1 2 x8, 2 x4, 2 x2, 2 x1	0b0111	1 Host	1 Upstream Socket	1 Link	0b000	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	2 x8, 2 x4, 2 x2, 2 x1	0b0110	1 Host	1 Upstream Socket	1 Link	0b000	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	1x16 Option B	1x16, 1x8, 1x4, 1x2, 1x1 2 x8, 2 x4, 2 x2, 2 x1	0b0101	1 Host	1 Upstream Socket	1 Link	0b000	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	1x16, 1x8, 1x4 2 x8, 2 x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	0b0100	1 Host	1 Upstream Socket	1 Link	0b000	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	4 x4	4 x4, 4 x2, 4 x1	0b0011	1 Host	1 Upstream Socket	1 Link	0b000	1x4*	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
RSVD	RSVD	RSVD	0b0010	1 Host	1 Upstream Socket	1 Link	0b000	-																
RSVD	RSVD	RSVD	0b0001	1 Host	1 Upstream Socket	1 Link	0b000	-																
RSVD	RSVD	RSVD	0b0000	1 Host	1 Upstream Socket	1 Link	0b000	-																

Table 28: 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							1:x16, 1:x8, 1:x4, 1:x2, 1:x1 2:x8, 2:x4, 2:x2, 2:x1																			
Min Card Width	**Card Short Name**	**Supported Bifurcation Modes**	**Addr-in-Card Encoding PRSNT[B: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**
n/a	Not Present	Card Not Present	0b1111	1Host	1Upstream Socket	1 or 2 Links	0b000	-																
2C	1x8	1x8, 1x4, 1x2, 1x1	0b1110	1Host	1Upstream Socket	1 or 2 Links	0b000	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	1Host	1Upstream Socket	1 or 2 Links	0b000	1x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	1x2	1x2, 1x1	0b1110	1Host	1Upstream Socket	1 or 2 Links	0b000	1x2	Link 0, Lane 0	Link 0, Lane 1														
2C	1x1	1x1	0b1110	1Host	1Upstream Socket	1 or 2 Links	0b000	1x1	Link 0, Lane 0															
2C	1x8 Option B	1x8, 1x4, 1x2, 1x1 2x4, 2x2, 2x1	0b1101	1Host	1Upstream Socket	1 or 2 Links	0b000	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	2x8, 2x4, 2x2, 2x1 4x4, 4x2, 4x1	0b1101	1Host	1Upstream Socket	1 or 2 Links	0b000	2x8	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 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 1, Lane 4	Link 1, Lane 5	Link 1, Lane 6	Link 1, Lane 7
2C	1x8 Option D	1x8, 1x4 2x4, 4 x2 (First 8 lanes), 4 x1	0b1100	1Host	1Upstream Socket	1 or 2 Links	0b000	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 D	1x16, 1x8, 1x4 2x8, 2x4, 4 x4, 4 x2 (First 8 lanes), 4 x1	0b1100	1Host	1Upstream Socket	1 or 2 Links	0b000	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
RSVD	RSVD	RSVD	0b1011	1Host	1Upstream Socket	1 or 2 Links	0b000	-																
2C	2 x4	2 x4, 2 x2, 2 x1 1 x4, 1 x2, 1 x1	0b1010	1Host	1Upstream Socket	1 or 2 Links	0b000	1x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
RSVD	RSVD	RSVD for future c.8 encoding	0b1001	1Host	1Upstream Socket	1 or 2 Links	0b000	-																
RSVD	RSVD	RSVD for future c.8 encoding	0b1000	1Host	1Upstream Socket	1 or 2 Links	0b000	-																
4C	1x16	1x16, 1x8, 1x4, 1x2, 1x1	0b0111	1Host	1Upstream Socket	1 or 2 Links	0b000	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, 2 x2, 2 x1	0b0110	1Host	1Upstream Socket	1 or 2 Links	0b000	2x8	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 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 1, Lane 4	Link 1, Lane 5	Link 1, Lane 6	Link 1, Lane 7
4C	1x16 Option B	1x16, 1x8, 1x4, 1x2, 1x1 2x8, 2x4, 2 x2, 2 x1	0b0101	1Host	1Upstream Socket	1 or 2 Links	0b000	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	1x16, 1x8, 1x4 2x8, 2x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	0b0100	1Host	1Upstream Socket	1 or 2 Links	0b000	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	4 x4	4 x4, 4 x2, 4 x1	0b0011	1Host	1Upstream Socket	1 or 2 Links	0b000	2 x4*	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3					Link 2, Lane 0	Link 2, Lane 1	Link 2, Lane 2	Link 2, Lane 3				
RSVD	RSVD	RSVD	0b0010	1Host	1Upstream Socket	1 or 2 Links	0b000	-																
RSVD	RSVD	RSVD	0b0001	1Host	1Upstream Socket	1 or 2 Links	0b000	-																
RSVD	RSVD	RSVD	0b0000	1Host	1Upstream Socket	1 or 2 Links	0b000	-																

Table 29: Bifurcation for Single Host, Single Socket and Single/Dual/Quad Upstream Links (BIF[2:0]#=obooo)

Single Host, Single Upstream Socket, One, Two or Four Upstream Links				1:x16, 1:x8, 1:x4, 1:x2,1 2:x8, 2:x4, 2:x2, 2:x1 4:x4, 4:x2, 4:x1																				
Min Card Width	**Card Short Name**	**Supported Bifurcation Modes**	**Add-in-Card Encoding** PRSNT#[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**
n/a	Not Present	Card Not Present	0b1111	1Host	1Upstream Socket	1,2, or 4 Links	0b000	-	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	1:x8, 1:x4, 1:x2, 1:x1	0b1110	1Host	1Upstream Socket	1,2, or 4 Links	0b000	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	1x4	1:x4, 1:x2, 1:x1	0b1110	1Host	1Upstream Socket	1,2, or 4 Links	0b000	1:x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	1x2	1:x2, 1:x1	0b1110	1Host	1Upstream Socket	1,2, or 4 Links	0b000	1:x2	Link 0, Lane 0	Link 0, Lane 1														
2C	1x1	1:x1	0b1110	1Host	1Upstream Socket	1,2, or 4 Links	0b000	1:x1	Link 0, Lane 0															
2C	1:x8 Option B	1:x8, 1:x4, 1:x2, 1:x1 2:x4, 2:x2, 2:x1	0b1101	1Host	1Upstream Socket	1,2, or 4 Links	0b000	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	2:x8 Option B	2:x8, 2:x4, 2:x2, 2:x1 4:x4, 4:x2, 4:x1	0b1101	1Host	1Upstream Socket	1,2, or 4 Links	0b000	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 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 1, Lane 4	Link 1, Lane 5	Link 1, Lane 6	Link 1, Lane 7
2C	1:x8 Option D	1:x8, 1:x4 2:x4, 4:x2 (First 8 lanes), 4:x1	0b1100	1Host	1Upstream Socket	1,2, or 4 Links	0b000	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	0b1100	1Host	1Upstream Socket	1,2, or 4 Links	0b000	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
RSVD	RSVD	RSVD	0b1011	1Host	1Upstream Socket	1,2, or 4 Links	0b000	-																
2C	2:x4	2:x4, 2:x2, 2:x1 1:x4, 1:x2, 1:x1	0b1010	1Host	1Upstream Socket	1,2, or 4 Links	0b000	2:x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3								
RSVD	RSVD	RSVD for future x8 cascading	0b1001	1Host	1Upstream Socket	1,2, or 4 Links	0b000	-																
RSVD	RSVD	RSVD for future x8 cascading	0b1000	1Host	1Upstream Socket	1,2, or 4 Links	0b000	-																
4C	1:x16	1:x16, 1:x8, 1:x4, 1:x2, 1:x1 2:x8, 2:x4, 2:x2, 2:x1 4:x4, 4:x2, 4:x1	0b0111	1Host	1Upstream Socket	1,2, or 4 Links	0b000	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 4:x4, 4:x2, 4:x1	0b0110	1Host	1Upstream Socket	1,2, or 4 Links	0b000	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 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 1, Lane 4	Link 1, Lane 5	Link 1, Lane 6	Link 1, Lane 7
4C	1:x16 Option B	1:x16, 1:x8, 1:x4, 1:x2, 1:x1 2:x8, 2:x4, 2:x2, 2:x1 4:x4, 4:x2, 4:x1	0b0101	1Host	1Upstream Socket	1,2, or 4 Links	0b000	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 4:x4, 4:x2, 4:x1	0b0100	1Host	1Upstream Socket	1,2, or 4 Links	0b000	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	0b0011	1Host	1Upstream Socket	1,2, or 4 Links	0b000	4:x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 2, Lane 0	Link 2, Lane 1	Link 2, Lane 2	Link 2, Lane 3	Link 3, Lane 0	Link 3, Lane 1	Link 3, Lane 2	Link 3, Lane 3
RSVD	RSVD	RSVD	0b0010	1Host	1Upstream Socket	1,2, or 4 Links	0b000	-																
RSVD	RSVD	RSVD	0b0001	1Host	1Upstream Socket	1,2, or 4 Links	0b000	-																
RSVD	RSVD	RSVD	0b0000	1Host	1Upstream Socket	1,2, or 4 Links	0b000	-																

Table 30: Bifurcation for Single Host, Dual Sockets and Dual Upstream Links (BIF[2:0]#=0b001)

Single Host, Two Upstream Sockets, Two Upstream Links

Min Card Width	Card Short Name	Supported Bifercation Modes	Add-in-Card Encoding PRSNT#[3:0]#	Host	1x8, 1x4, 1x2, 1x1 2x8, 2x4, 2 x2, 2x1 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
n/a	Not Present	Card Not Present	0b1111	1 Host	2 Upstream Sockets	2 Links	0b001	-																
2C	1x8	1x8, 1x4, 1x2, 1x1	0b1110	1 Host	2 Upstream Sockets	2 Links	0b001	1x8 (Socket 0 only)	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	2 Upstream Sockets	2 Links	0b001	1x4 (Socket 0 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	1x2	1x2, 1x1	0b1110	1 Host	2 Upstream Sockets	2 Links	0b001	1x2 (Socket 0 only)	Link 0, Lane 0	Link 0, Lane 1														
2C	1x1	1x1	0b1110	1 Host	2 Upstream Sockets	2 Links	0b001	1x1 (Socket 0 only)	Link 0, Lane 0															
2C	1x8 Option B	1x8, 1x4, 1x2, 1x1 2 x4, 2 x2, 2 x1	0b1101	1 Host	2 Upstream Sockets	2 Links	0b001	1x8 (Socket 0 only)	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	2x8, 2 x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	0b1101	1 Host	2 Upstream Sockets	2 Links	0b001	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 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 1, Lane 4	Link 1, Lane 5	Link 1, Lane 6	Link 1, Lane 7
2C	1x8 Option D	1x8, 1x4 2 x4, 4 x2 (First 8 lanes), 4 x1	0b1100	1 Host	2 Upstream Sockets	2 Links	0b001	1x8 (Socket 0 only)	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	1x16 Option D	1x16, 1x8, 1x4 2x8, 2 x4, 4 x4, 4 x2 (First 8 lanes), 4 x1	0b1100	1 Host	2 Upstream Sockets	2 Links	0b001	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 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 1, Lane 4	Link 1, Lane 5	Link 1, Lane 6	Link 1, Lane 7
RSVD	RSVD	RSVD	0b1011	1 Host	2 Upstream Sockets	2 Links	0b001	-																
2C	2 x4	2 x4, 2 x2, 2 x1 1 x4, 1 x2, 1 x1	0b1010	1 Host	2 Upstream Sockets	2 Links	0b001	1 x4 (Socket 0 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
RSVD	RSVD	RSVD for future, 8 encoding	0b1001	1 Host	2 Upstream Sockets	2 Links	0b001	-																
RSVD	RSVD	RSVD for future, 8 encoding	0b1000	1 Host	2 Upstream Sockets	2 Links	0b001	-																
4C	1x16	1x16, 1x8, 1x4, 1x2, 1x1	0b0111	1 Host	2 Upstream Sockets	2 Links	0b001	1x8 (Socket 0 only)	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 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 1, Lane 4	Link 1, Lane 5	Link 1, Lane 6	Link 1, Lane 7
4C	2x8 Option A	2 x8, 2 x4, 2 x2, 2 x1	0b0110	1 Host	2 Upstream Sockets	2 Links	0b001	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 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 1, Lane 4	Link 1, Lane 5	Link 1, Lane 6	Link 1, Lane 7
4C	1x16 Option B	1x16, 1x8, 1x4, 1x2, 1x1 2 x8, 2 x4, 2 x2, 2 x1	0b0101	1 Host	2 Upstream Sockets	2 Links	0b001	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 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 1, Lane 4	Link 1, Lane 5	Link 1, Lane 6	Link 1, Lane 7
4C	1x16 Option C	1x16, 1x8, 1x4 2 x8, 2 x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	0b0100	1 Host	2 Upstream Sockets	2 Links	0b001	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 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 1, Lane 4	Link 1, Lane 5	Link 1, Lane 6	Link 1, Lane 7
4C	4 x4	4 x4, 4 x2, 4 x1	0b0011	1 Host	2 Upstream Sockets	2 Links	0b001	2 x4 (EP 0 and 2 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3					Link 2, Lane 0	Link 2, Lane 1	Link 2, Lane 2	Link 2, Lane 3				
RSVD	RSVD	RSVD	0b0010	1 Host	2 Upstream Sockets	2 Links	0b001	-																
RSVD	RSVD	RSVD	0b0001	1 Host	2 Upstream Sockets	2 Links	0b001	-																
RSVD	RSVD	RSVD	0b0000	1 Host	2 Upstream Sockets	2 Links	0b001	-																

Table 31: Bifurcation for Single Host, Four Sockets and Dual Upstream Links (BIF[2:0]#=0b010)

Single Host, Four Upstream Sockets, Four Upstream Links

Min Card Width	Card Short Name	Supported Bifurcation Modes	Add-in-Card Encoding PRSNTB[3:0]#	Host	4 x4, 4 x2, 4 x1 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
n/a	Not Present	Card Not Present	0b1111	1 Host	4 Upstream Sockets	4 Links	0b010	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	1x8	1x8, 1x4, 1x2, 1x1	0b1110	1 Host	4 Upstream Sockets	4 Links	0b010	1x4 (Socket 0 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	1x4	1x4, 1x2, 1x1	0b1110	1 Host	4 Upstream Sockets	4 Links	0b010	1x4 (Socket 0 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	1x2	1x2, 1x1	0b1110	1 Host	4 Upstream Sockets	4 Links	0b010	1x2 (Socket 0 only)	Link 0, Lane 0	Link 0, Lane 1														
2C	1x1	1x1	0b1110	1 Host	4 Upstream Sockets	4 Links	0b010	1x1 (Socket 0 only)	Link 0, Lane 0															
2C	1x8 Option B	1x8, 1x4, 1x2, 1x1 2 x4, 2 x2, 2 x1	0b1101	1 Host	4 Upstream Sockets	4 Links	0b010	2 x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3								
4C	2 x8 Option B	2 x8, 2 x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	0b1101	1 Host	4 Upstream Sockets	4 Links	0b010	4 x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 2, Lane 0	Link 2, Lane 1	Link 2, Lane 2	Link 2, Lane 3	Link 3, Lane 0	Link 3, Lane 1	Link 3, Lane 2	Link 3, Lane 3
2C	1x8 Option D	1x8, 1x4 2 x4, 4 x2 (First 8 lanes), 4 x1	0b1100	1 Host	4 Upstream Sockets	4 Links	0b010	2 x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3		Link 2, Lane 1	Link 2, Lane 2	Link 2, Lane 3	Link 3, Lane 0	Link 3, Lane 1	Link 3, Lane 2	Link 3, Lane 3
4C	1x16 Option D	1x16, 1x8, 1x4 2 x8, 2 x4, 4 x4, 4 x2 (First 8 lanes), 4 x1	0b1100	1 Host	4 Upstream Sockets	4 Links	0b010	4 x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 2, Lane 0	Link 2, Lane 1	Link 2, Lane 2	Link 2, Lane 3	Link 3, Lane 0	Link 3, Lane 1	Link 3, Lane 2	Link 3, Lane 3
RSVD	RSVD	RSVD	0b1011	1 Host	4 Upstream Sockets	4 Links	0b010	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3								
2C	2 x4	2 x4, 2 x2, 2 x1 1 x4, 1 x2, 1 x1	0b1010	1 Host	4 Upstream Sockets	4 Links	0b010	2 x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3								
RSVD	RSVD	RSVD for future x8 encoding	0b1001	1 Host	4 Upstream Sockets	4 Links	0b010	-																
RSVD	RSVD	RSVD for future x8 encoding	0b1000	1 Host	4 Upstream Sockets	4 Links	0b010	-																
4C	1x16	1x16, 1x8, 1x4, 1x2, 1x1 2 x8, 2 x4, 2 x2, 2 x1	0b0111	1 Host	4 Upstream Sockets	4 Links	0b010	1 x4 (Socket 0 only) 2 x4 (Socket 0 & 2 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
4C	2 x8 Option A	1x16, 1x8, 1x4, 1x2, 1x1 2 x8, 2 x4, 2 x2, 2 x1	0b0110	1 Host	4 Upstream Sockets	4 Links	0b010	2 x4 (Socket 0 & 2 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3					Link 2, Lane 0	Link 2, Lane 1	Link 2, Lane 2	Link 2, Lane 3				
4C	1x16 Option B	1x16, 1x8, 1x4, 1x2, 1x1 2 x8, 2 x4, 2 x2, 2 x1	0b0101	1 Host	4 Upstream Sockets	4 Links	0b010	2 x4 (Socket 0 & 2 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3					Link 2, Lane 0	Link 2, Lane 1	Link 2, Lane 2	Link 2, Lane 3				
4C	1x16 Option C	1x16, 1x8, 1x4 2 x8, 2 x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	0b0100	1 Host	4 Upstream Sockets	4 Links	0b010	4 x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 2, Lane 0	Link 2, Lane 1	Link 2, Lane 2	Link 2, Lane 3	Link 3, Lane 0	Link 3, Lane 1	Link 3, Lane 2	Link 3, Lane 3
4C	4 x4	4 x4, 4 x2, 4 x1	0b0011	1 Host	4 Upstream Sockets	4 Links	0b010	4 x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 2, Lane 0	Link 2, Lane 1	Link 2, Lane 2	Link 2, Lane 3	Link 3, Lane 0	Link 3, Lane 1	Link 3, Lane 2	Link 3, Lane 3
RSVD	RSVD	RSVD	0b0010	1 Host	4 Upstream Sockets	4 Links	0b010	-																
RSVD	RSVD	RSVD	0b0001	1 Host	4 Upstream Sockets	4 Links	0b010	-																
RSVD	RSVD	RSVD	0b0000	1 Host	4 Upstream Sockets	4 Links	0b010	-																

Table 32: Bifurcation for Dual Host, Dual Sockets and Dual Upstream Links (BIF[2:0]#=0b101)

Dual Host, Two Upstream Sockets, Two Upstream Links

Min Card Width	Card Short Name	Supported Bifercation Modes	Add-in-Card Encoding PRSNT[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
n/a	Not Present	Card Not Present	0b1111	2 Host	2 Upstream Sockets	2 Links	0b101																	
2C	1x8	1x8, 1x4, 1x2, 1x1	0b1110	2 Host	2 Upstream Sockets	2 Links	0b101	1x8 (Host 0 only)	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	2 Host	2 Upstream Sockets	2 Links	0b101	1x4 (Host 0 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	1x2	1x2, 1x1	0b1110	2 Host	2 Upstream Sockets	2 Links	0b101	1x2 (Host 0 only)	Link 0, Lane 0	Link 0, Lane 1														
2C	1x1	1x1	0b1110	2 Host	2 Upstream Sockets	2 Links	0b101	1x1 (Host 0 only)	Link 0, Lane 0															
2C	1x8 Option B	1x8, 1x4, 1x2, 1x1 2 x4, 2 x2, 2 x1	0b1101	2 Host	2 Upstream Sockets	2 Links	0b101	1x8 (Host 0 only)	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	2 x8 Option B	2 x8, 2 x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	0b1101	2 Host	2 Upstream Sockets	2 Links	0b101	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 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 1, Lane 4	Link 1, Lane 5	Link 1, Lane 6	Link 1, Lane 7
2C	1x8 Option D	1x8, 1x4, 1x2, 1x1 2 x4, 2 x2, 2 x1 4 x4, 4 x2 (First 8 lanes), 4 x1	0b1100	2 Host	2 Upstream Sockets	2 Links	0b101	1x8 (Host 0 only)	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 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 1, Lane 4	Link 1, Lane 5	Link 1, Lane 6	Link 1, Lane 7
4C	1x16 Option D	1x16, 1x8, 1x4 2 x8, 2 x4, 4 x4, 4 x2 (First 8 lanes), 4 x1	0b1100	2 Host	2 Upstream Sockets	2 Links	0b101	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 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 1, Lane 4	Link 1, Lane 5	Link 1, Lane 6	Link 1, Lane 7
RSVD	RSVD	RSVD	0b1011	2 Host	2 Upstream Sockets	2 Links	0b101	-																
2C	2 x4	2 x4, 2 x2, 2 x1 1 x4, 1 x2, 1 x1	0b1010	2 Host	2 Upstream Sockets	2 Links	0b101	1 x4 (Host 0 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
RSVD	RSVD	RSVD for future x8 encoding	0b1001	2 Host	2 Upstream Sockets	2 Links	0b101	-																
RSVD	RSVD	RSVD for future x8 encoding	0b1000	2 Host	2 Upstream Sockets	2 Links	0b101	-																
4C	1 x16	1 x16, 1 x8, 1 x4, 1 x2, 1 x1 2 x8, 2 x4, 2 x2, 2 x1	0b0111	2 Host	2 Upstream Sockets	2 Links	0b101	1 x8 (Host 0 only)	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	2 x8 Option A	2 x8, 2 x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	0b0110	2 Host	2 Upstream Sockets	2 Links	0b101	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 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 1, Lane 4	Link 1, Lane 5	Link 1, Lane 6	Link 1, Lane 7
4C	1 x16 Option B	1 x16, 1 x8, 1 x4, 1 x2, 1 x1 2 x8, 2 x4, 2 x2, 2 x1	0b0101	2 Host	2 Upstream Sockets	2 Links	0b101	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 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 1, Lane 4	Link 1, Lane 5	Link 1, Lane 6	Link 1, Lane 7
4C	1 x16 Option C	1 x16, 1 x8, 1 x4 2 x8, 2 x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	0b0100	2 Host	2 Upstream Sockets	2 Links	0b101	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 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 1, Lane 4	Link 1, Lane 5	Link 1, Lane 6	Link 1, Lane 7
4C	4 x4	4 x4, 4 x2, 4 x1	0b0011	2 Host	2 Upstream Sockets	2 Links	0b101	2 x4 (EP 0 and 2 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3					Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3				
RSVD	RSVD	RSVD	0b0010	2 Host	2 Upstream Sockets	2 Links	0b101	-																
RSVD	RSVD	RSVD	0b0001	2 Host	2 Upstream Sockets	2 Links	0b101	-																
RSVD	RSVD	RSVD	0b0000	2 Host	2 Upstream Sockets	2 Links	0b101	-																

Table 33: Bifurcation for Quad Host, Quad Sockets and Quad Upstream Links (BIF[2:0]#=0b110)

Quad Host, Four Upstream Sockets, Four Upstream Links

Min Card Width	Card Short Name	Supported Bifurcation Modes	Add-in-Card Encoding PRSNTB[3:0]#	Host	4 x4, 4 x2, 4 x1 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
n/a	Not Present	Card Not Present	0b1111	4 Host	4 Upstream Sockets	4 Links	0b110	-																
2C	1x8	1x8, 1x4, 1x2, 1x1	0b1110	4 Host	4 Upstream Sockets	4 Links	0b110	1x4 (Host 0 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	1x4	1x4, 1x2, 1x1	0b1110	4 Host	4 Upstream Sockets	4 Links	0b110	1x4 (Host 0 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	1x2	1x2, 1x1	0b1110	4 Host	4 Upstream Sockets	4 Links	0b110	1x2 (Host 0 only)	Link 0, Lane 0	Link 0, Lane 1														
2C	1x1	1x1	0b1110	4 Host	4 Upstream Sockets	4 Links	0b110	1x1 (Host 0 only)	Link 0, Lane 0															
2C	1x8 Option B	1x8, 1x4, 1x2, 1x1 2x4, 2 x2, 2 x1	0b1101	4 Host	4 Upstream Sockets	4 Links	0b110	2 x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3								
4C	2 x8 Option B	2x8, 2 x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	0b1101	4 Host	4 Upstream Sockets	4 Links	0b110	4 x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 2, Lane 0	Link 2, Lane 1	Link 2, Lane 2	Link 2, Lane 3	Link 3, Lane 0	Link 3, Lane 1	Link 3, Lane 2	Link 3, Lane 3
2C	1x8 Option D	1x8, 1x4 2x4, 4 x2 (First 8 lanes), 4 x1	0b1100	4 Host	4 Upstream Sockets	4 Links	0b110	2 x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 2, Lane 0	Link 2, Lane 1	Link 2, Lane 2	Link 2, Lane 3	Link 3, Lane 0	Link 3, Lane 1	Link 3, Lane 2	Link 3, Lane 3
4C	1x16 Option D	1x16, 1x8, 1x4 2x8, 2 x4, 4 x4, 4 x2 (First 8 lanes), 4 x1	0b1100	4 Host	4 Upstream Sockets	4 Links	0b110	4 x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3								
RSVD	RSVD	RSVD	0b1011	4 Host	4 Upstream Sockets	4 Links	0b110	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3								
2C	2x4	2 x4, 2 x2, 2 x1 1 x4, 1 x2, 1 x1	0b1010	4 Host	4 Upstream Sockets	4 Links	0b110	2 x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
RSVD	RSVD	RSVD for future x8 encoding	0b1001	4 Host	4 Upstream Sockets	4 Links	0b110	-																
RSVD	RSVD	RSVD for future x8 encoding	0b1000	4 Host	4 Upstream Sockets	4 Links	0b110	-																
4C	1x16	1x16, 1x8, 1x4, 1x2, 1x1	0b0111	4 Host	4 Upstream Sockets	4 Links	0b110	1 x4 (Host 0 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
4C	2 x8 Option A	2x8, 2 x4, 2 x2, 2 x1	0b0110	4 Host	4 Upstream Sockets	4 Links	0b110	2 x4, 1 x4 (Host 0 & 2 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3					Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3				
4C	1x16 Option B	1x16, 1x8, 1x4, 1x2, 1x1 2 x8, 2 x4, 2 x2, 2 x1	0b0101	4 Host	4 Upstream Sockets	4 Links	0b110	2 x4, 2 x4 (Host 0 & 2 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3					Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3				
4C	1x16 Option C	1x16, 1x8, 1x4 2 x8, 2 x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	0b0100	4 Host	4 Upstream Sockets	4 Links	0b110	4 x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 2, Lane 0	Link 2, Lane 1	Link 2, Lane 2	Link 2, Lane 3	Link 3, Lane 0	Link 3, Lane 1	Link 3, Lane 2	Link 3, Lane 3
4C	4 x4	4 x4, 4 x2, 4 x1	0b0011	4 Host	4 Upstream Sockets	4 Links	0b110	4 x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3	Link 2, Lane 0	Link 2, Lane 1	Link 2, Lane 2	Link 2, Lane 3	Link 3, Lane 0	Link 3, Lane 1	Link 3, Lane 2	Link 3, Lane 3
RSVD	RSVD	RSVD	0b0010	4 Host	4 Upstream Sockets	4 Links	0b110	-																
RSVD	RSVD	RSVD	0b0001	4 Host	4 Upstream Sockets	4 Links	0b110	-																
RSVD	RSVD	RSVD	0b0000	4 Host	4 Upstream Sockets	4 Links	0b110	-																

Table 34: Bifurcation for Quad Host, Quad Sockets and Quad Upstream Links (BIF[2:0]#=0b110)

Quad/Oct Host, Four/Eight Upstream Sockets, Four/Eight Upstream links

Mia 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
					4 x2, 4 x1																			
n/a	Not Present	Card Not Present	0b1111	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	-																
2C	1 x8	1 x8, 1 x4, 1 x2, 1 x1	0b1110	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	1 x2 (Host 0 only)	Link 0, Lane 0	Link 0, Lane 1														
2C	1 x4	1 x4, 1 x2, 1 x1	0b1110	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	1 x2 (Host 0 only)	Link 0, Lane 0	Link 0, Lane 1														
2C	1 x2	1 x2, 1 x1	0b1110	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	1 x2 (Host 0 only)	Link 0, Lane 0	Link 0, Lane 1														
2C	1 x1	1 x1	0b1110	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	1x1 (Host 0 only)	Link 0, Lane 0															
2C	1 x8 Option B	1 x8, 1 x4, 1 x2, 1 x1 2 x4, 2 x2, 2 x1	0b1101	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	2 x2 (Host 0 & 1 only)	Link 0, Lane 0	Link 0, Lane 1	Link 1, Lane 0	Link 1, Lane 1												
4C	2 x8 Option B	2 x8, 2 x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	0b1101	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	2 x2 (Host 0 & 1 only)	Link 0, Lane 0	Link 0, Lane 1	Link 1, Lane 0	Link 1, Lane 1												
2C	1 x8 Option D	1 x8, 1 x4 2 x4, 4 x2 (First 8 lanes), 4 x1	0b1100	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	4 x2	Link 0, Lane 0	Link 0, Lane 1	Link 1, Lane 0	Link 1, Lane 1	Link 2, Lane 0	Link 2, Lane 1	Link 3, Lane 0	Link 3, Lane 1								
4C	1 x16 Option D	1 x16, 1 x8, 1 x4 2 x8, 2 x4, 4 x4, 4 x2 (First 8 lanes), 4 x1	0b1100	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	4 x2	Link 0, Lane 0	Link 0, Lane 1	Link 1, Lane 0	Link 1, Lane 1	Link 2, Lane 0	Link 2, Lane 1	Link 3, Lane 0	Link 3, Lane 1								
RSVD	RSVD	RSVD	0b1011	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	-																
2C	2 x4	2 x4, 2 x2, 2 x1 1 x4, 1 x2, 1 x1	0b1010	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	2 x2 (Host 0 & 1 only)	Link 0, Lane 0	Link 0, Lane 1	Link 1, Lane 0	Link 1, Lane 1												
RSVD	RSVD	RSVD for future x8 encoding	0b1001	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	-																
RSVD	RSVD	RSVD for future x8 encoding	0b1000	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	-																
4C	1 x16	1 x16, 1 x8, 1 x4, 1 x2, 1 x1 2 x8, 2 x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	0b0111	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	1 x2 (Host 0 only)	Link 0, Lane 0	Link 0, Lane 1														
4C	2 x8 Option A	2 x8, 2 x4, 2 x2, 2 x1	0b0110	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	1 x2 (Host 0 & 1 only)	Link 0, Lane 0	Link 0, Lane 1							Link 1, Lane 0	Link 1, Lane 1						
4C	1 x16 Option B	1 x16, 1 x8, 1 x4, 1 x2, 1 x1 2 x8, 2 x4, 2 x2, 2 x1	0b0101	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	2 x2 (Host 0 & 1 only)	Link 0, Lane 0	Link 0, Lane 1							Link 1, Lane 0	Link 1, Lane 1						
4C	1 x16 Option C	1 x16, 1 x8, 1 x4 2 x8, 2 x4, 2 x2, 2 x1 4 x4, 4 x2, 4 x1	0b0100	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	2 x2 (Host 0 & 1 only)	Link 0, Lane 0	Link 0, Lane 1							Link 2, Lane 0	Link 2, Lane 1						
4C	4 x4	4 x4, 4 x2, 4 x1	0b0011	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	4 x2 (Host 0 & 1 only)	Link 0, Lane 0	Link 0, Lane 1			Link 1, Lane 0	Link 1, Lane 1										
RSVD	RSVD	RSVD	0b0010	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	-																
RSVD	RSVD	RSVD	0b0001	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	-																
RSVD	RSVD	RSVD	0b0000	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	-																

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 59. The max available power envelopes for each of these states are defined in Table 35.

Figure 59: Baseboard Power States

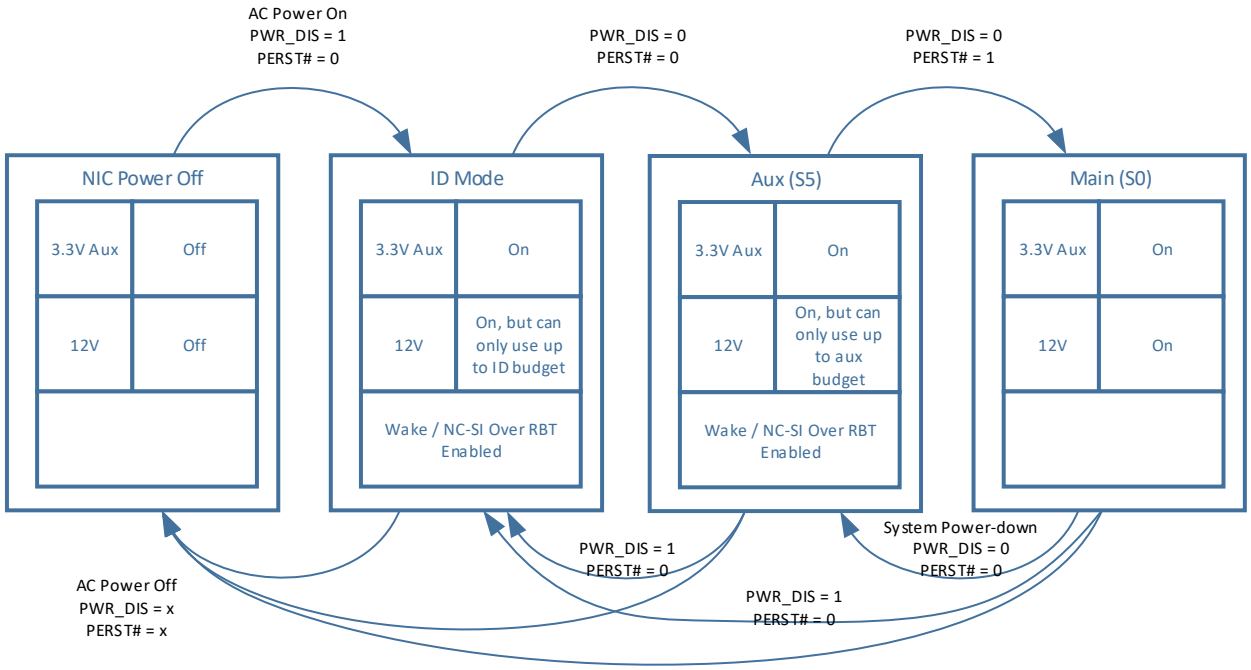

Table 35: Power States

Power State	PWRDIS	PERSTn	FRU	Scan Chain	RBT Link	3.3V	12V
NIC Power Off	Low	Low					
ID Mode	High	Low	X	X		X	X
Aux Power Mode (S5)	Low	Low	X	X	X	X	X
Main Power Mode (S0)	Low	High	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 Aux is available for powering up management only functions. FRU accesses are only allowed in this mode. An add-in card shall transition to this mode when PWRDIS=1 and PERST#=0.

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. An add-in card shall transition to this mode when PWRDIS=0 and PERST#=0.

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 80W may be delivered on 12V, and 3.63W on the 3.3V pins. An add-in card shall transition to this mode when PWRDIS=0 and PERST#=1.

3.10 Power Supply Rail Requirements and Slot Power Envelopes

The baseboard provides 3.3Vaux and 12Vaux/main 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 36: 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					
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					
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.

3.11 Hot Swap Considerations for 12V and 3.3V Rails

For baseboards that support system hot (powered on) add-in card insertions and extractions, the system implementer shall consider the use of hotswap controllers on both the 12Vmain/aux and 3.3Vaux pins to prevent damage to the baseboard or the add-in 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 12Vmain/aux and 3.3Vmain/aux 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 ob1111. 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 XXX), 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]#, 3.3Vaux, 12Vaux/12Vmain relative to PWRDIS, BIF[2:0]#, PERSTn*, the add-in card power ramp and NIC_PWR_GOOD.

Figure 60: Power Sequencing

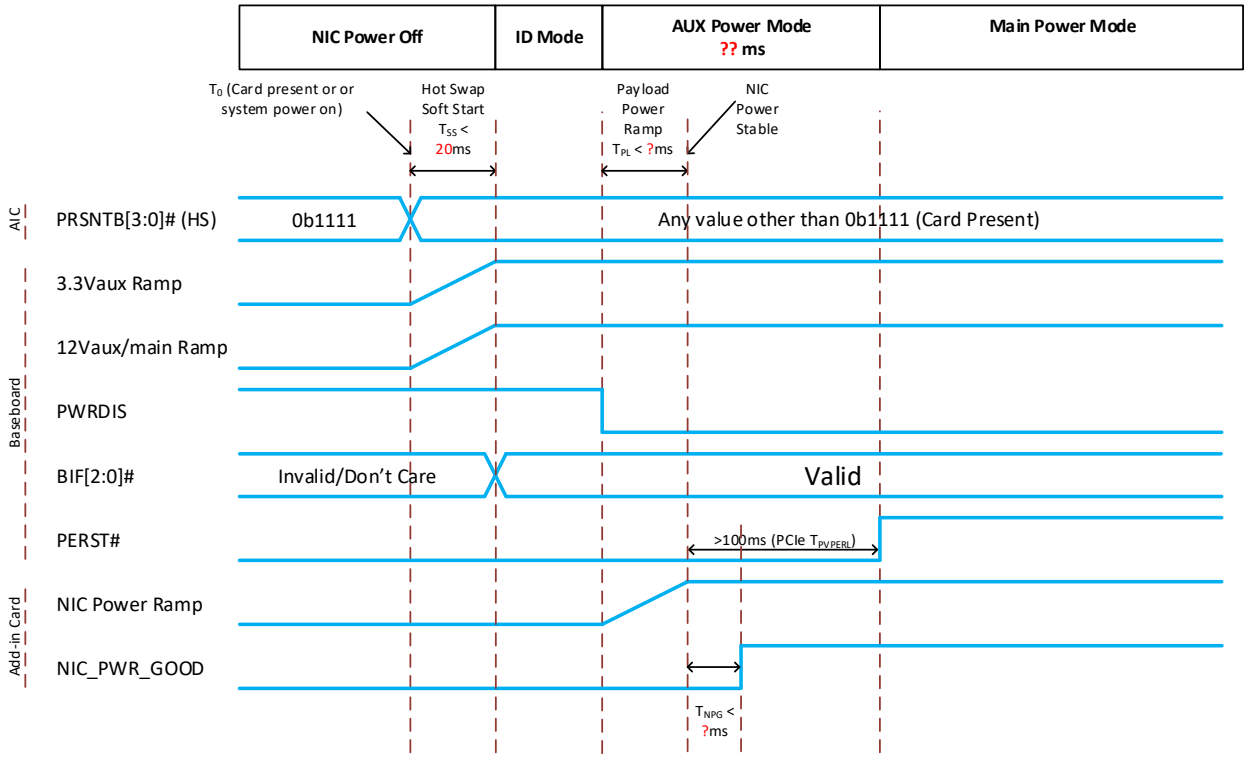

Table 37: Power Sequencing Parameters

Parameter	Value	Units	Description
T_{ss}	20	ms	Max time between system 3.3Vaux and 12Vaux/main ramp to power stable.
T_{PL}	<?	ms	Max time between the NIC payload power ramp to NIC_PWR_GOOD assertion
T_{NPG}	<?	ms	Max time between NIC power stable and NIC_PWR_GOOD assertion.
T_{PVPERL}	>100	ms	Max time between PWRDIS deassertion and PERST# deassertion. This value is from the PCIe CEM Specification, Rev 4.0.

4 Management

OCP NIC 3.0 card management is an important aspect to overall system management. This section specifies a common set of management requirements for OCP NIC 3.0 implementations. There are three types of implementations (No Management Type, RBT Type, and MCTP Type) depending on the physical sideband management interfaces, transports, and traffic supported over different transports. An OCP NIC 3.0 implementation shall support at least one type of implementation for card management. A No Management implementation should not be used for an Ethernet add-in card

4.1 Sideband Management Interface and Transport

OCP NIC 3.0 sideband management interfaces are used by a Management Controller (MC) or Baseboard Management Controller (BMC) to communicate with the NIC. Table 38 summarizes the sideband management interface and transport requirements.

Table 38: Sideband Management Interface and Transport Requirements

Requirement	RBT Type	MCTP Type
NC-SI 1.1 or later compliant RMII Based Transport (RBT) including physical interface defined in Section 10 of DSP0222.	Required	
Both of the following physical sideband interfaces: • SMBus 2.0 • PCIe VDM		Required
Management Component Transport Protocol (MCTP) Base 1.3 on both of the following physical bindings: • MCTP/SMBus (DSP0237 1.1) • MCTP/PCIe VDM (DSP0238 1.1)		Required

4.2 NC-SI Traffic

DSP0222 defines two types of NC-SI traffic: Pass-Through and Control. Table 39 summarizes the NC-SI traffic requirements.

Table 39: NC-SI Traffic Requirements

Requirement	RBT Type	MCTP Type
NC-SI Control over RBT (DSP0222 1.1 or later compliant)	Required	
NC-SI Control over MCTP (DSP0261 1.2 or later compliant)		Required
NC-SI Pass-Through over RBT (DSP0222 1.1 or later compliant)	Required	
NC-SI Pass-Through over MCTP (DSP0261 1.2 or later compliant)		Recommended

4.3 Management Controller (MC) MAC Address Provisioning

An OCP NIC 3.0 add-in card 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 40 summarizes MC MAC address provisioning requirements.

Table 40: MC MAC Address Provisioning Requirements

Requirement	RBT Type	MCTP Type
One or more MAC Addresses shall be provisioned for the MC.	Required	
Support at least one of the following mechanism for provisioned MC MAC Address retrieval: • NC-SI Control/RBT (DSP0222 1.1 or later compliant) • NC-SI Control/MCTP (DSP0261 1.2 or later compliant)	Required	

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 optical modules providing physical network media connectivity. It is important for the system management that temperatures of these components can be retrieved over sideband interfaces. Table 41 summarizes temperature reporting requirements.

Table 41: Temperature Reporting Requirements

Requirement	RBT Type	MCTP Type
ASIC Temperature Reporting	Recommended	Recommended
Optical Modules Temperature Reporting	Recommended	Recommended
Support at least one of the following mechanisms for ASIC temperature reporting: • NC-SI Control (DSP0222 1.1 or later compliant) • PLDM for Platform Monitoring and Control (DSP0248 1.1 compliant)	Recommended	Recommended
Support at least one of the following mechanisms for optical modules temperature reporting: • NC-SI Control (DSP0222 1.1 or later compliant) • PLDM for Platform Monitoring and Control (DSP0248 1.1 compliant)	Recommended	Recommended
Where the temperature sensor reporting function is implemented, the temperature reporting accuracy on the card shall be within ±3°C	Required	Required

4.5 Power Consumption Reporting

An OCP NIC 3.0 implementation may be able to report the power consumed by ASICs implementing NIC functions. It is important for the system management that the information about the power consumption can be retrieved over sideband interfaces. Table 42 summarizes power consumption reporting requirements.

Table 42: Power Consumption Reporting Requirements

Requirement	RBT Type	MCTP Type
ASIC Power Consumption Reporting	Optional	Optional

Support at least one of the following mechanisms for ASIC power consumption reporting: • NC-SI Control (DSP0222 1.1 or later compliant) • PLDM for Platform Monitoring and Control (DSP0248 1.1 compliant)	Optional	Optional

4.6 Link Status/Speed Reporting

Link status/speed reporting is important for network operations and link management. Table 43 summarizes link status and speed reporting requirements.

Table 43: Link Status/Speed Reporting Requirements

Requirement	RBT Type	MCTP Type
Link Status Reporting	Required	Required
Support at least one of the following mechanisms for reporting the link status: • NC-SI Control (DSP0222 1.1 compliant) • PLDM for Platform Monitoring and Control (DSP0248 1.1 compliant)	Required	Required
Link Speed Reporting	Required	Required
Support at least one of the following mechanisms for reporting the link speed: • NC-SI Control (DSP0222 1.1 compliant) • PLDM for Platform Monitoring and Control (DSP0248 1.1 compliant)	Required	Required

4.7 Pluggable Module Status Reporting

Pluggable modules like optical modules or direct attach cables are used to connect OCP NIC to physical media. It is important to know the presence of pluggable modules and information about insertion/deletion of pluggable modules. Table 44 summarizes pluggable module status reporting requirements.

Table 44: Pluggable Module Status Reporting Requirements

Requirement	RBT Type	MCTP Type
Pluggable Module Presence Reporting	Recommended	Recommended
Support at least one of the following mechanisms for reporting the pluggable module presence status: • NC-SI Control (DSP0222 1.1 or later compliant) • PLDM for Platform Monitoring and Control (DSP0248 1.1 compliant)	Recommended	Recommended
Pluggable Module Insertions/Deletions Reporting	Recommended	Recommended
Support at least one of the following mechanisms for reporting the pluggable module insertions/deletions:	Recommended	Recommended

• NC-SI Control (DSP0222 1.1 or later compliant) • PLDM for Platform Monitoring and Control (DSP0248 1.1 compliant)		

4.8 Out-Of-Band Firmware Update

An OCP 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 Out-Of-Band mechanism for the firmware update. Table 45 summarizes out-of-band firmware update requirements.

Table 45: Out-Of-Band Firmware Update Requirements

Requirement	RBT Type	MCTP Type
Support PLDM for Firmware Update (DSP0267 1.0 compliant)	Optional	Optional

4.9 NC-SI Over RBT Sideband Interface

NC-SI Over RBT provides a low speed management path for the add-in card. This is implemented via RMII pins between the BMC and the add-in card. NC-SI Over RBT is the recommended management method for OCP NIC 3.0 cards. Protocol and implementation details can be found in the DMTF DSP0222 standard.

4.9.1 NC-SI Over RBT Addressing

NC-SI Over RBT capable devices must use a unique Package ID 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 3.3Vaux (Slot ID = 1).

For add-in cards, 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] is defaults to 0b0 in the NC-SI specification, but is optionally configurable if the target silicon supports configuring this bit. Package ID[1] is connected to the SLOT_ID pin and is directly connected to the Slot_ID pin. Package ID[0] is set to 0b0 for Network Silicon #0. For OCP NIC 3.0 add-in cards with two discrete silicon instances, Package ID[1] is set to 0b1. Refer to the 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.

4.9.2 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 must support operation with a one card, or both cards installed.

Figure 49 shows an example connection with dual Primary Connectors.

4.10 SMBus 2.0 Interface

The SMBus provides a low speed management bus for the add-in card. The FRU EEPROM and on-board temperature sensors are connected on this bus. Additionally, network controllers may utilize the SMBus interface for NC-SI over MCTP. Proper power domain isolation shall be implemented on the NIC.

Note: The preferred network controller management path is NC-SI Over RBT, as described in Section 4.9. Silicon devices may offer both a SMBus and NC-SI Over RBT interface. In such instances, the add-in card developer shall choose NC-SI Over RBT.

4.10.1 SMBus Address Map

All predefined SMBus addresses for OCP NIC 3.0 are shown in Table 46. Baseboard and add-in 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 46: Power Sequencing Parameters

Address (8-bit)	Device	Notes
0xTBD	Network Controller IC	Value dependent on NIC vendors.
0x98 / 0x99	Temperature Sensor	TMP422/423 Temperature sensor Optional. Used for remote on-die thermal sensing. Optional. Powered from Aux power domain.
0x9E / 0x9F	Temperature Sensor	TMP421 Temperature sensor. Optional. Used for remote on-die thermal sensing. Optional. Powered from Aux power domain.
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 add-in card gold finger to allow two NIC add-in cards to exist on the same I²C bus.

4.11 FRU EEPROM

4.11.1 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 add-in 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 0xA2/0xA3 for the write/read pair in 8-bit format. The size of EEPROM is 4Kbits for the base EEPROM map. Add-in 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(S_5) mode, and MAIN(S_0) mode.

4.11.2 FRU EEPROM Content Requirements

The FRU EEPROM shall follow the data format specified in the IPMI Platform Management FRU Information Storage Definition v1.2. Use OEM record 0xC0, offset 0x01 through 0x05 to store specific records for the OCP NIC.

Table 47: FRU EEPROM Record – OEM Record 0xC0, Offset 0x00

Offset 0	Description
	Manufacturer ID, LS Byte first (3 bytes total)

Table 48: FRU EEPROM Record – OEM Record 0xC0, Offset 0x01

Offset 1	Primary Connector PRSNTB [3:0]#
0b1110 (0x0E)	Follows Pinout; to be filled after the pinout table is fixed
0b1101 (0x0D)	
0b1100 (0x0C)	
0b1010 (0x0A)	
0b0111 (0x07)	
0b0110 (0x06)	
0b0101 (0x05)	
0b0100 (0x04)	
0b0011 (0x03)	
0b1011 (0x0B)	Not a valid reading – Wrong EEPROM programming
0b1111 (0x0F)	Not a valid reading – Wrong EEPROM programming
All others	RFU
No FRU device detected	No NIC connected / bad connection

Table 49: FRU EEPROM Record – OEM Record 0xC0, Offset 0x02

Offset 2	Secondary Connector PRSNTB [3:0]#
0b1110 (0x0E)	Follows Pinout; to be filled after the pinout table is fixed
0b1101 (0x0D)	
0b1100 (0x0C)	
0b1010 (0x0A)	
0b0111 (0x07)	
0b0110 (0x06)	
0b0101 (0x05)	
0b0100 (0x04)	
0b0011 (0x03)	
0b1011 (0x0B)	Not a valid reading – Wrong EEPROM programming

ob1111 (ox0F)	Not a valid reading – Wrong EEPROM programming
All others	RFU
No FRU device detected	No NIC connected / bad connection

Table 50: FRU EEPROM Record – OEM Record oxC0, Offset ox03

Offset 3	**Card max power in Aux(S5)**
ox01 - oxFE	Hex format in Watts when NIC is in AUX(S5) mode; LSB = 1x Watt; roundup to the nearest Watt for fractional values.
oxFF	Invalid entry
ox00	Invalid entry

Table 51: FRU EEPROM Record – OEM Record oxC0, Offset ox04

Offset 4	**Card max power in Main(S0)**
ox01 - oxFE	Hex format in Watts when NIC is in Main (S0) mode; LSB = 1x Watt; roundup to the nearest Watt for fractional values.
oxFF	Invalid entry
ox00	Invalid entry

Table 52: FRU EEPROM Record – OEM Record oxC0, Offset ox05

Offset 5	**Thermal Reporting Interface**
ox01	Emulated thermal reporting on SMBus
ox02	Remote on-die sensor with TMP421 on SMBus
ox04	PLDM thermal reporting via NC-SI over RBT

4.12 FW Requirements

(Editors note (Jia): Tentative list; collecting feedback)

4.12.1 Firmware Update

- The OCP NIC 3.0 add-in card shall support device firmware upgrades from the BMC controller.

4.12.2 Secure Firmware

- The OCP NIC 3.0 add-in card shall support secured firmware.
- Where the secured firmware feature is enabled, the OCP NIC 3.0 add-in card shall allow only update and execute signed firmware.

4.12.3 Firmware Queries

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

4.12.4 Multi-Host Firmware Queries

- A multi-host capable OCP NIC 3.0 add-in card shall gracefully handle concurrent in-band queries from multiple hosts and out-of-band access from the BMC the management status and firmware, device model, and device ID information.
- A multi-host capable OCP NIC 3.0 add-in 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 add-in card shall gracefully handle exceptions when more than one entity attempts to perform concurrent NIC firmware writes.

4.13 Thermal Reporting Interface

The OCP NIC 3.0 thermal reporting interface is defined on the primary connector SMBus or via NC-SI over RBT depending on the implementation method described in this section.

This requirement improves the system thermal management and allows the baseboard management device to access key component temperatures on an OCP NIC.

There are three defined methods to implement thermal reporting described in this section: emulated thermal reporting, remote on-die sensing and via Platform Level Data Model (PLDM).

Emulated thermal reporting and remote on-die sensing the two methods defined in OCP NIC 2.0 and is used in the current 3.0 spec release. In both cases, the BMC treats the temperature sensors as a TI/TMP421 (or equivalent) device. The emulated temperature sensor is accessible at slave address 0x3E/0x3F as the read/write pair in 8-bit format over the SMBus.

The third reporting method is using PLDM. This is the recommend implementation for OCP NIC 3.0. PLDM uses the NC-SI over RBT as the underlying protocol.

A thermal reporting interface is required for all OCP NIC 3.0 complaint cards with a TDP > 10W.

Thermal reporting interface shall be accessible in AUX(S5) mode, and MAIN(S0) mode.

4.13.1 Remote on-die sensing

Alternatively, one TMP421 sensor can be used to do on die temperature sensing for IC with thermal diode interface with TMP421 remote sensing channel; Connection diagram is shown in Figure 61. For add-in cards that require more than one remote on-die sense point, a TMP422/TMP423 can be used and slave address is 0x98/0x99 (8-bit).

Figure 61: Block Diagram for Remote on-die Sensing


Mezz Pin 75/76
Mezz 2.0 Pin A75/A76
TMP421
0x3E(8bit)

5 Data Network Requirements

5.1 Network Boot

OCP NIC 3.0 shall support network booting in uEFI system environment with both IPv4 and IPv6 network booting.

For UEFI booting, below features are required (tentative list; collecting feedback)

- EFI_DRIVER_BINDING_PROTOCOL (for starting and stopping the driver)
- EFI_DEVICE_PATH_PROTOCOL (provides location of the device)
- EFI_MANAGED_NETWORK_SERVICE_BINDING_PROTOCOL (asynchronous network packet I/O services)
- EFI_DRIVER_DIAGNOSTICS2_PROTOCOL & EFI_DRIVER_DIAGNOSTICS_PROTOCOL (driver will allow the UEFI shell command *drvdiag* to perform a cursory check of the connections managed by the driver)
- Human Interface Infrastructure (HII) protocols
- EFI_DRIVER_HEALTH_PROTOCOL
- EFI_FIRMWARE_UPDATE_PROTOCOL

6 Routing Guidelines and Signal Integrity Considerations

6.1 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.

6.2 PCIe

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

Add-in card suppliers shall follow the PCIe routing specifications. At this time, the OCP NIC subgroup is working to identify and agree to the channel budget for an add-in 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.

7 Thermal and Environmental

7.1 Environmental Requirements

Specifics are not included to permit adoption of OCP 3.0 NIC in systems with varying thermal requirements and boundary conditions. The system adopting OCP NIC should define air flow direction, local approach air temperature and speed to the NIC, operational altitude and relative humidity.

For example, a system configured with I/O facing the cold aisle, can specify approach air temperature and speed of 35°C and 200 LFM respectively, with airflow impinging on the I/O modules first and an operational altitude of 6000 feet.

7.1.1 Thermal Reporting interface

[Link to 4.6; this session can be incorporated into Chapter 4.6] The NIC should support temperature reporting for key components on the card (including readings from active I/O modules). To improve thermal efficiency, the recommended accuracy for temperature sensors on the card is ±3°C. A power reporting interface is optional for low-power NICs (less than 10W), but strongly recommended for higher-powered NICs.

7.1.2 Thermal Simulation Boundary Example

Placeholder for the link to upcoming test fixture documentation (under development).

7.2 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 add-in card and baseboard vendor to decide how the shock and vibration tests shall be done.

7.3 Regulatory

An OCP NIC 3.0 add-in card shall meet the following compliance requirements:

- **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.
- **REACH Regulation (EC) No 1907/2006** addresses the production and use of chemical substances and their potential impact on human health and the environment.
- 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.
- CE
- FCC Class A

An OCP NIC 3.0 add-in card is recommended to meet below compliance requirements:

- **Halogen Free:** IEC 61249-2-21 Definition of halogen free: 900ppm for Br or CI, or 1500ppm combined.
- **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.

8 Revision History

Author	Description	Revision	Date
Thomas Ng Intel Corporation	Initial draft	0.5	12/22/2017