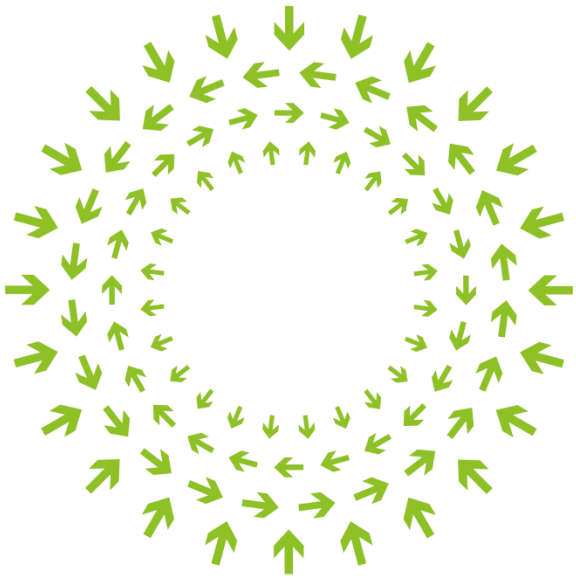

OPEN

Compute Project

OCP NIC 3.0 Design Specification

Version 0.01

Author: OCP Server Workgroup, OCP NIC subgroup

Table of Contents

List of Figures

List of Tables

1 Overview

1.1 License

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

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

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

1.2 Background

The OCP NIC 3.0 specification is a follow-on to the OCP 2.0 form-factor for PCIe add-in cards. The OCP NIC 3.0 specification supports two basic card sizes: Small Card, and Large Card. The Small Card allows for up to 16 PCIe lanes on the card edge while a Large Card supports up to 32 PCIe. Compared to the OCP NIC 2.0 release, the updated specification provides a broader solution space for NIC and system vendors to support the following use case scenarios:

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

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

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

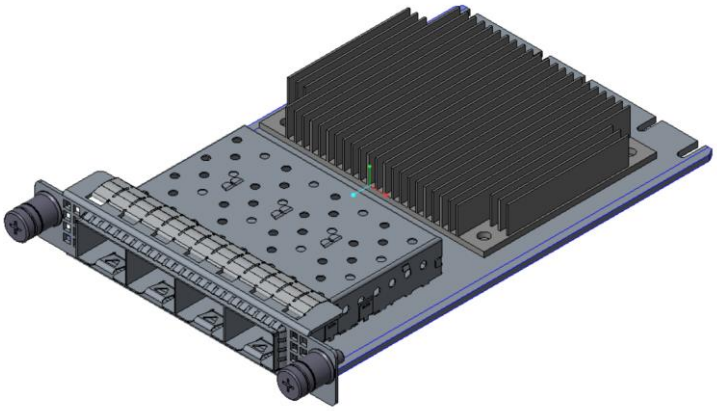

Figure 2: Representative Large OCP NIC 3.0 Card with Dual QSFP Ports and on-board DRAM

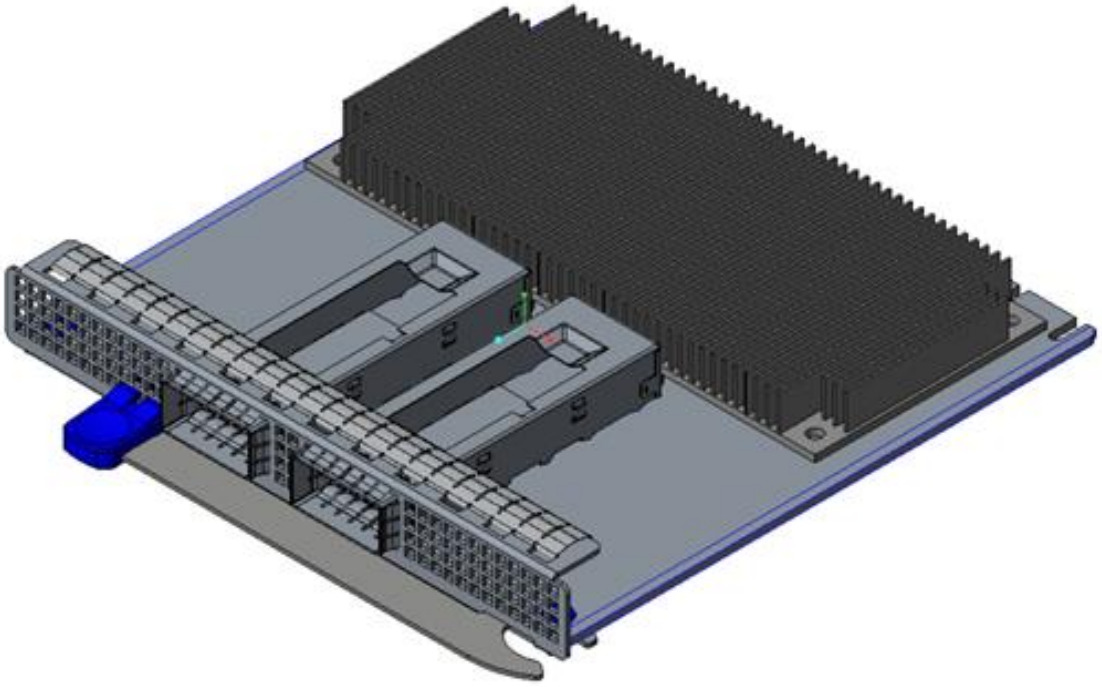

In order to achieve the features outlined in this specification, OCP 3.0 compliant cards are not backwards compatible to the 2.0 cards.

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

Placeholder

1.4 Overview

1.4.1 Mechanical Form factor overview

The OCP NIC 3.0 specification defines a third generation mechanical form factor that allows for interoperability between 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 two connectors (Primary Connector and Secondary Connector) on the baseboard. Both the Primary and Secondary connectors are defined in and compliant to SFF-TA-1002. [Editor's note: plan to submit change back to SFF-TA-1002]. On the NIC side, the card edge is implemented with gold fingers. The gold finger design follows SFF-TA-1002 as well.

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

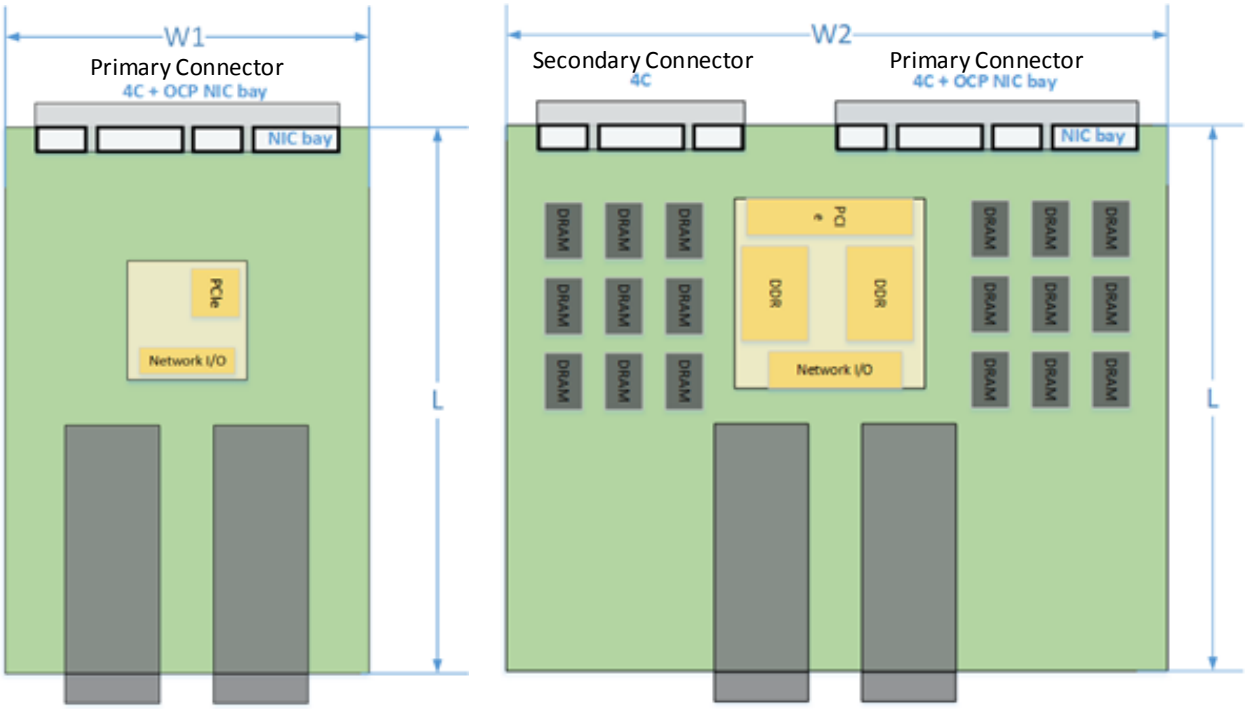

The two form factor dimensions are shown in Table 1.

Table 1: OCP 3.0 Form Factor Dimensions

Form Factor	Width	Depth	Primary Connector	Secondary Connector	Typical Use Case

Small	W1 = 76 mm	L = 115 mm	4C + OCP sideband 168 pins	N/A	Low profile and general NIC with a similar profile as an OCP NIC 2.0 add-in card; up to x16 PCIe.
Large	W2 = 139 mm	L = 115 mm	4C + OCP sideband 168 pins	4C 140 pins	Larger PCB width to support feature rich NICs; up to x32 PCIe.

The OCP NIC 3.0 design allows downward compatibility between the two card sizes. Table 2 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 2: Baseboard to OCP NIC Form factor Compatibility Chart

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

There are two baseboard connector options available for system designers: straddle mount and right angle (RA). The straddle mount connector option allows the OCP NIC and baseboard to exist in a co-planer position. To achieve this, a cutout exists on the baseboard and is defined in this specification. Alternatively, the right angle option allows the OCP NIC to be installed on top of the baseboard. A baseboard cutout is not required for the right angle connector. The right angle option allows the 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 2.

This specification defines the form factor at the add-in card level, including the front panel, latching mechanism and card guide features [TBD; pending on the Mechanical work across stakeholders].

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 x16 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

1.4.2.2 Secondary connector

1.5 References

- DMTF Standard. *DSP0222, Network Controller Sideband Interface (NC-SI) Specification.* Distributed Management Task Force, Inc, Rev 1.0.1, January 24th, 2013.
- EDSFF. *Enterprise and Datacenter SSD Form Factor Connector Specification*. Enterprise and Datacenter SSD Form Factor Working Group, Rev 0.9 (draft), August 2nd 2017.
- NXP Semiconductors. *I²C-bus specification and user manual.* NXP Semiconductors, Rev 6, April 4th, 2014.

- Open Compute Project. *OCP NIC Subgroup*. Online. http://www.opencompute.org/wiki/Server/Mezz
- PCIe Base Specification. *PCI Express Base Specification, Revision 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 3.0, December 20th, 2014.
- SNIA. *SFF-TA-1002, Specification for Protocol Agnostic Multi-Lane High Speed Connector*. SNIA SFF TWG Technology Affiliate, Rev 0.0.9.1, September 9th, 2017.

2 Card Form Factor

2.1 Overview

2.2 Form Factor Options

OCP3.0 provides two fundamental form factor options: a small card (76mm x 115mm) and a large card (139mm x 115mm).

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

The small card uses the 4C connector for up to a x16 PCIe interface. The small cards implement and additional 28-pin OCP bay for management functions and support for up to a four PCIe hosts. The small size card provides sufficient faceplate area to accommodate up to 2x QSFP modules, 4x SFP modules, or 4x RJ-45 for BASE-T operation. The small card form factor supports up to 79W of delivered power to the card edge.

Figure 4: Example Small Card Form Factor

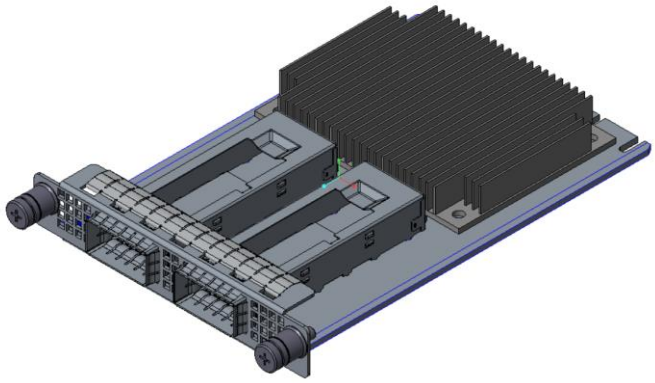

The large card uses provides the same functionality as the small card, but with support up to a x32 PCIe interface. The large card utilizes both the Primary and Secondary connectors. The large size card supports higher power envelopes and provides additional board area for more complex designs. The small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

Figure 5: Example Large Card Form Factor

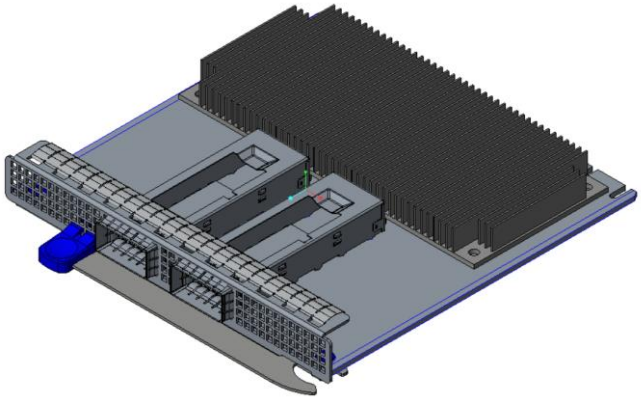

For both form-factors, an add-in card may optionally implement a subset of pins to support a x8 PCIe connection. This is implemented using a 2C card edge per SFF-TA-1002. The Primary Connector may support a 2C sized add-in card along with the 28 pin OCP bay. The following diagram from the SFF-TA-1002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

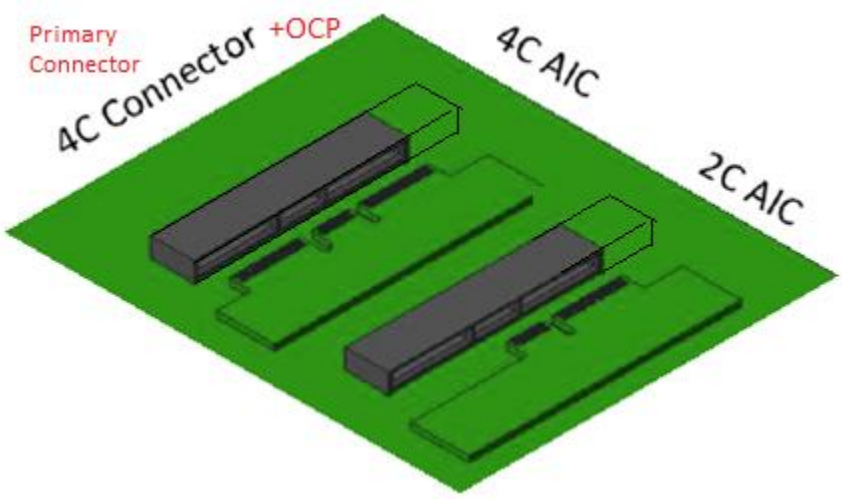

Table 3 summarizes the supported card form factors. Small form factors cards support the Primary Connector and up to 16 PCIe lanes. Large form factor cards support both the Primary and Secondary Connectors and up to 32 PCIe lanes.

Table 3: 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 (x24)		2C	4C		OCP Bay
Large (x32)	4C		4C		OCP Bay

2.3 I/O bracket

TBD <need input from OCP mechanical groups>

2.4 Line Side I/O Implementations

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

Table 4: OCP 3.0 Line Side I/O Implementations

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

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

2.5 LED Implementation

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

For small form-factor low I/O count cards (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 addition to the Scan Chain.

For both cases, the OCP NIC 3.0 specification recommends the following LED definitions:

Table 5: Default LED Configuration

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

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

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

2.6 Mechanical Keepout Zones

2.6.1 Baseboard Keep Out Zones

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

2.6.2 Add-in Card Keep Out Zones

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

2.7 Labeling Requirements

TBD

2.8 Insulation Requirements

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

TBD <need 2D drawings>

2.9 NIC Implementation Examples

TBD

2.10 Non-NIC Use Cases

"PCIe interface with extra management sideband"

2.10.1 PCIe Retimer card

2.10.2 Accelerator card

2.10.3 Storage HBA / RAID card

3 Card Edge and Baseboard Connector Interface

3.1 Gold Finger Requirement

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 1.60mm. The gold finger dimensions for the Primary Connector compliant cards are shown below.

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

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

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

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

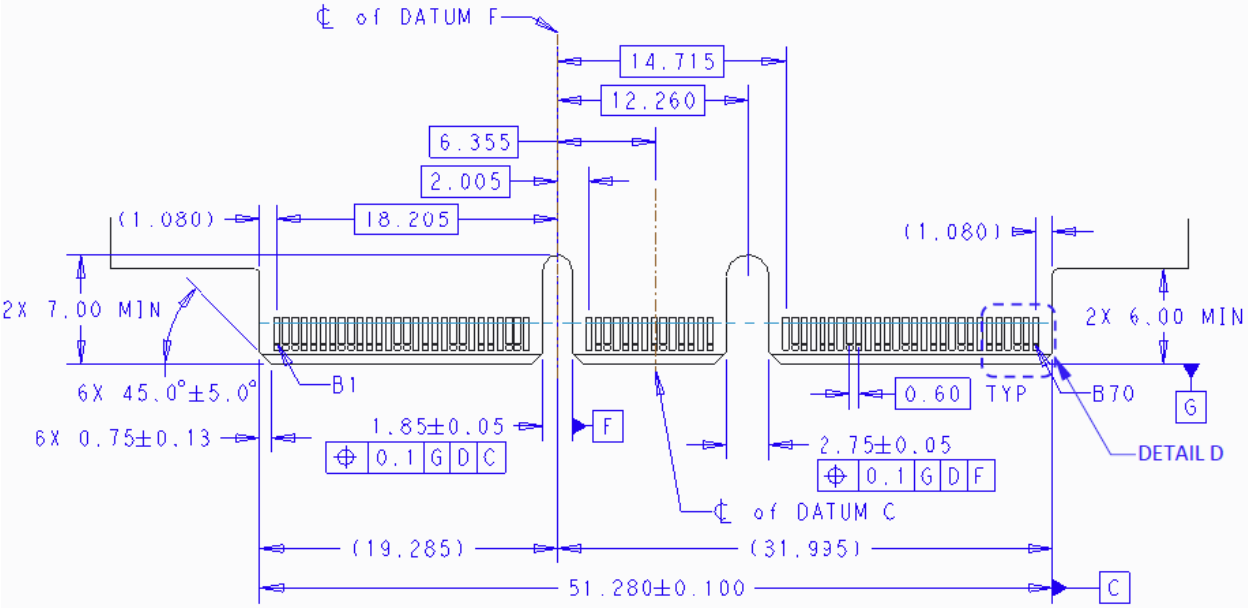

Figure 8: Small Size Primary Connector Gold Finger Mating Card Dimensions – x16 – Bottom Side

TBD

Figure 9: Large Size Card Gold Finger Mating Card Dimensions – x32 – Top Side

TBD

Figure 10: Large Size Card Gold Finger Mating Card Dimensions – x32 – Bottom Side

TBD

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 6 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 6: Contact Mating Positions for the Primary and Secondary Connectors

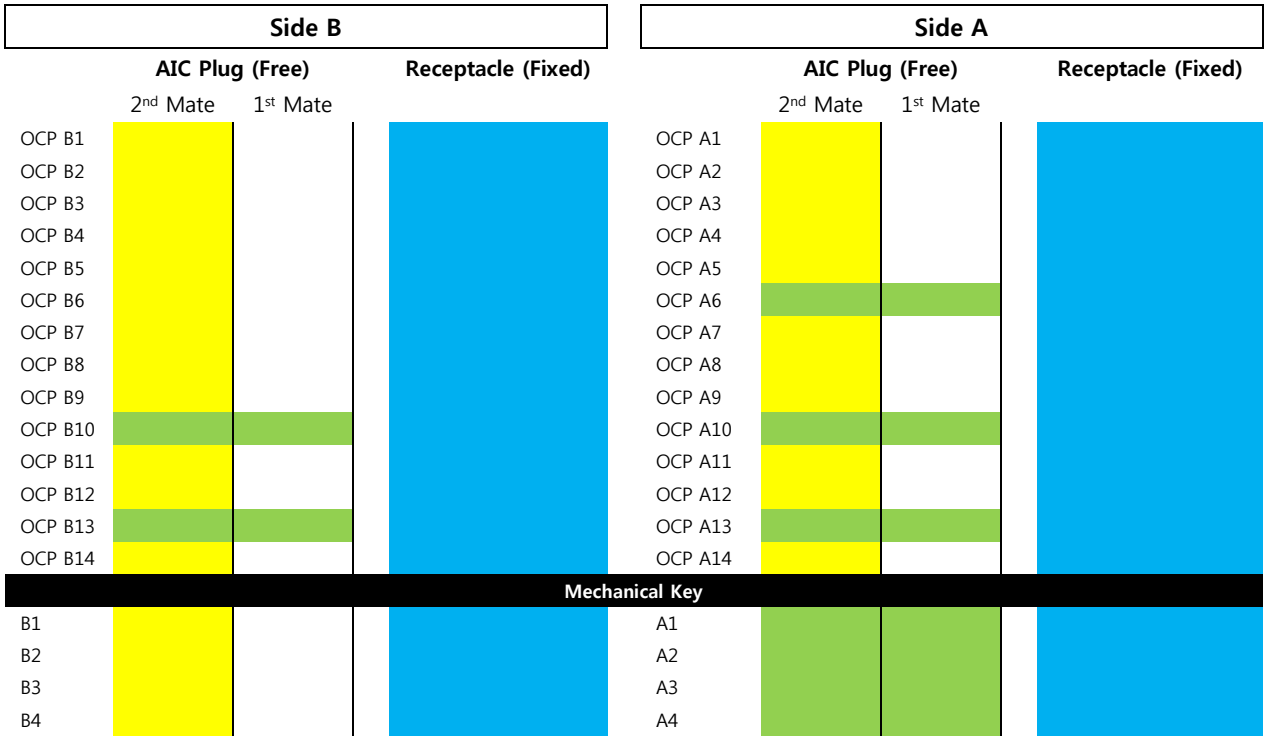

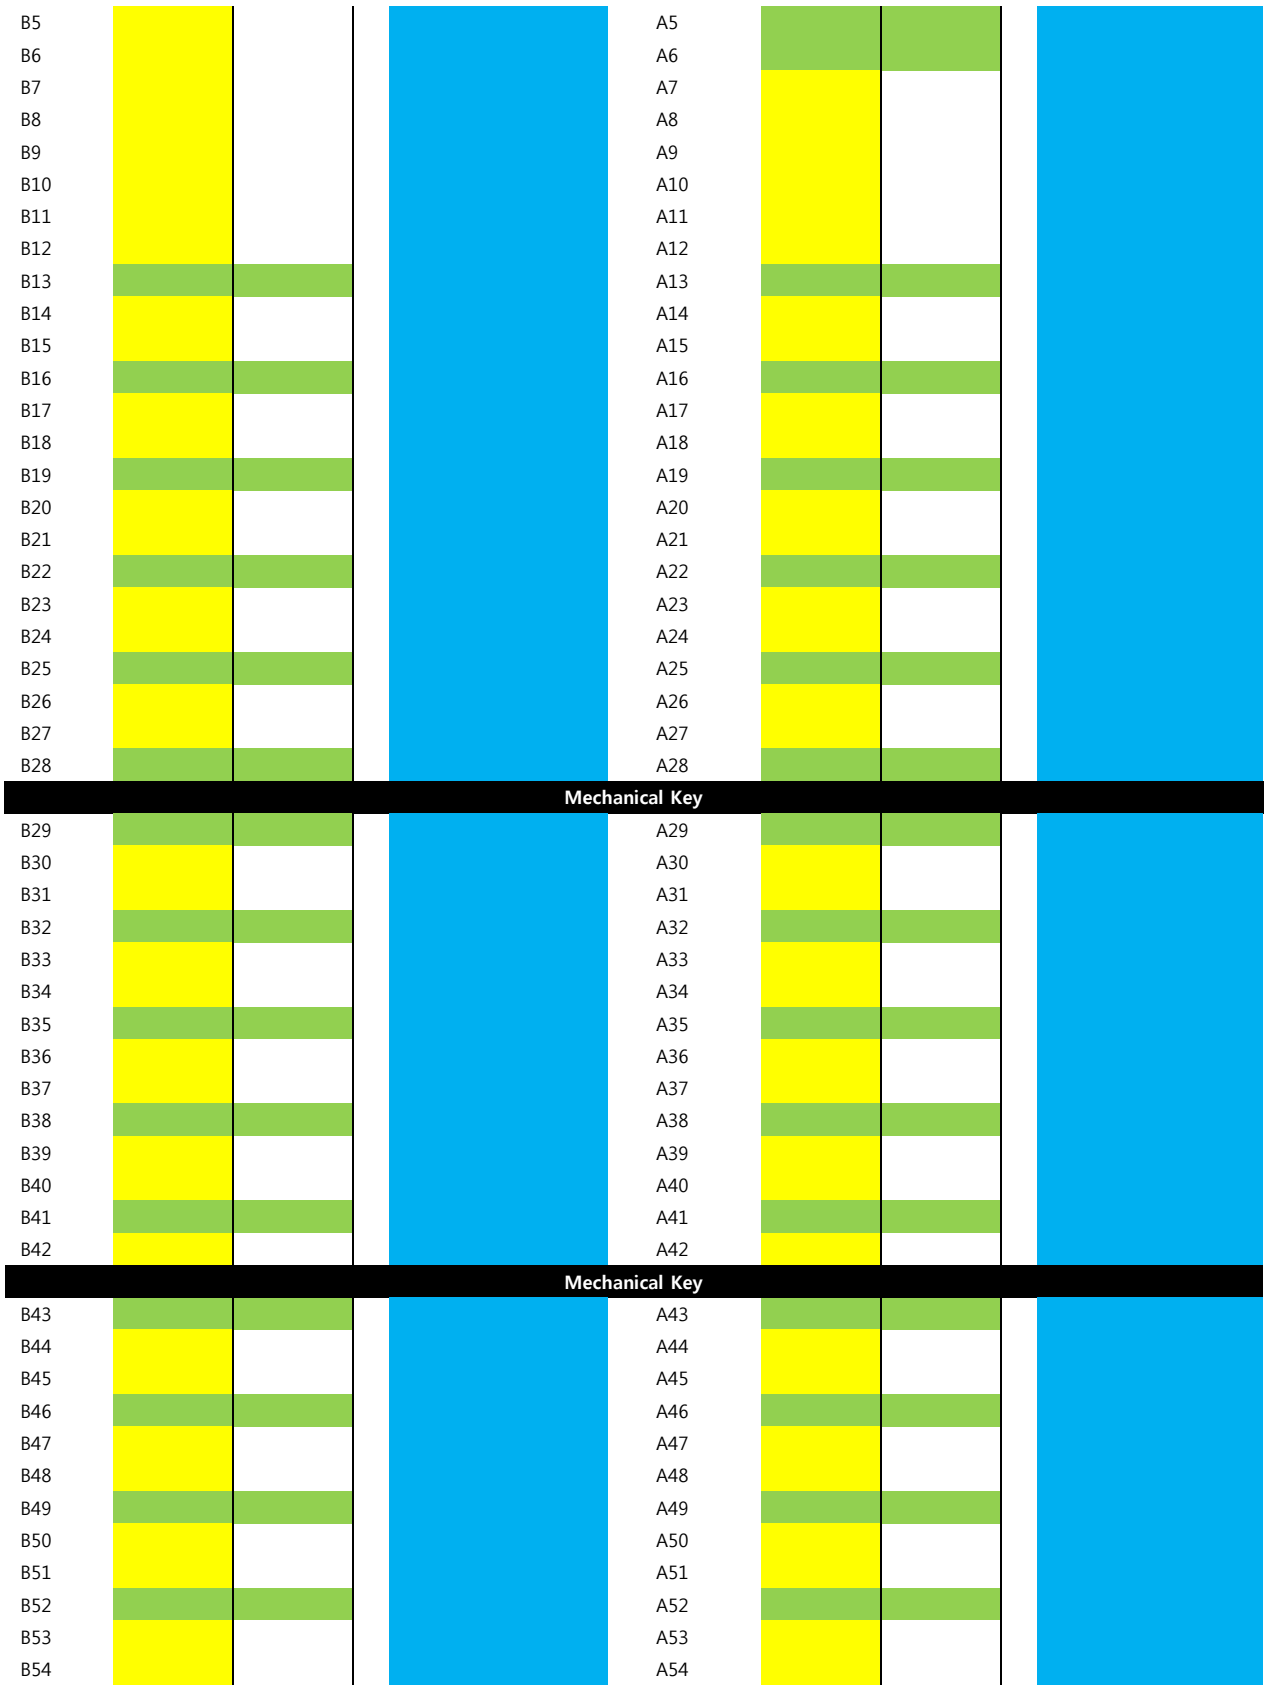
B5
B6
B7
B8
B9
B10
B11
B12
B13
B14
B15
B16
B17
B18
B19
B20
B21
B22
B23
B24
B25
B26
B27
B28
A5
A6
A7
A8
A9
A10
A11
A12
A13
A14
A15
A16
A17
A18
A19
A20
A21
A22
A23
A24
A25
A26
A27
A28
Mechanical Key
B29
B30
B31
B32
B33
B34
B35
B36
B37
B38
B39
B40
B41
B42
A29
A30
A31
A32
A33
A34
A35
A36
A37
A38
A39
A40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

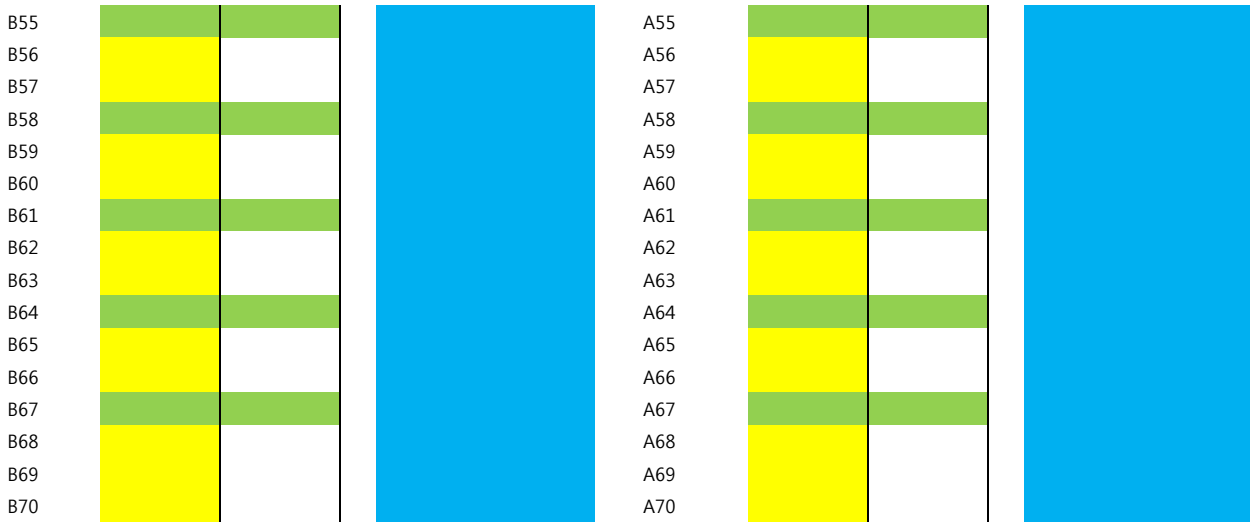

3.2 Baseboard Connector Requirement

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 x4 multi-host configuration on the Primary Connector. The Primary and Secondary Connector drawings are shown in Figure 11, Figure 12, Figure 13 and Figure 14 below.

Figure 11: 168-pin Base Board Primary Connector – Right Angle

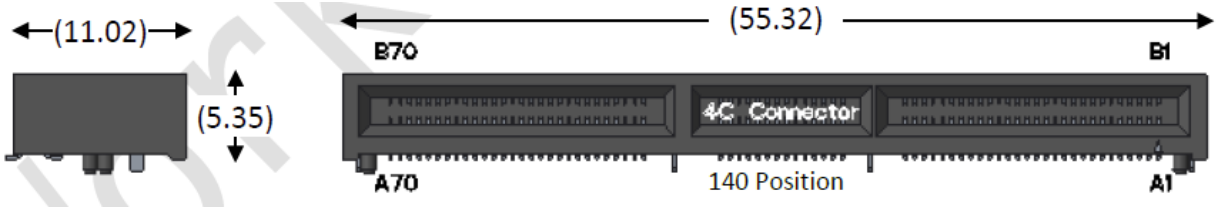

Figure 12: 140-pin Base Board Secondary Connector – Right Angle

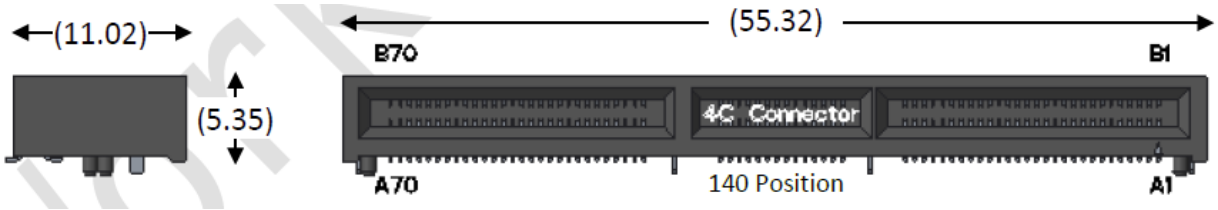

Figure 13: 168-pin Base Board Primary Connector – Straddle Mount

TBD

Figure 14: 140-pin Base Board Secondary Connector – Straddle Mount

TBD

In order to the support the large form factor, systems must locate the Primary and Secondary Connectors per the mechanical drawing shown in Figure 15.

Figure 15: Primary and Secondary Connector Locations for Large Card Support

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 7 and Table 8. 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, x16 PCIe Gen4 and bifurcation control. The Primary Connector has an additional OCP Bay (pins OCP_A[1:14], OCP_B[1:14]) with additional REFCLKs for supporting up to four PCIe hosts, NC-SI connectivity and a Scan Chain for information exchange between the host and card. The NIC is required to implement the Scan Chain, while the baseboard may choose to optionally implement it. Depending on the baseboard form-factor, multiple OCP NIC 3.0 compliant cards may be designed into the system.

The pins common to the Primary and Secondary 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 7: Primary Connector Pin Definition (x16) (4C + OCP Bay)

Side B		Side A			
OCP_B1	NIC_PWR_GOOD	WAKE_N	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#	PERST2#	OCP_A2		
OCP_B3	LD#	PERST3#	OCP_A3		
OCP_B4	DATA_IN	RBT_ARB_IN	OCP_A4		
OCP_B5	DATA_OUT	RBT_ARB_OUT	OCP_A5		
OCP_B6	CLK	GND	OCP_A6		
OCP_B7	SLOT_ID0	RBT_TX_EN	OCP_A7		
OCP_B8	RBT_RXD1	RBT_TXD1	OCP_A8		
OCP_B9	RBT_RXD0	RBT_TXD0	OCP_A9		
OCP_B10	GND	GND	OCP_A10		
OCP_B11	REFCLKn2	REFCLKn3	OCP_A11		
OCP_B12	REFCLKp2	REFCLKp3	OCP_A12		
OCP_B13	GND	GND	OCP_A13		
OCP_B14	RBT_CRS_DV	RBT_CLK_IN	OCP_A14		
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 8: Secondary Connector Pin Definition (x16) (4C)

Side B				Secondary Connector (x16, 140-pin add-in card)	Secondary Connector (x8, 84-pin add-in card)
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		

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 26.

Table 9: 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 are required at the connector. For add-in cards, the required REFCLKs shall be connected per the endpoint datasheet. **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,
REFCLKn1 REFCLKp1	A14 A15	Output	

			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 are AC coupled on the baseboard with capacitors and are placed next to the baseboard transmitters. The AC coupling capacitor must be between 176nF (min) and 265nF (max). 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 are AC coupled on the add-in card with capacitors and are placed next to the add-in card transmitters. The AC coupling capacitor must be between 176nF (min) and 265nF (max). 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.
PERn1 PERp1	A20 A21	Input	
PERn2 PERp2	A23 A24	Input	
PERn3 PERp3	A26 A27	Input	
PERn4 PERp4	A30 A31	Input	
PERn5 PERp5	A33 A34	Input	
PERn6 PERp6	A36 A37	Input	
PERn7 PERp7	A39 A40	Input	
PERn8 PERp8	A44 A45	Input	
PERn9 PERp9	A47 A48	Input	
PERn10 PERp10	A50 A51	Input	
PERn11 PERp11	A53 A54	Input	
PERn12 PERp12	A56 A57	Input	
PERn13 PERp13	A59 A60	Input	

PERn14 PERp14	A62 A63	Input	
PERn15 PERp15	A65 A66	Input	
PERST0# PERST1#	B10 A11	Output	PCIe Reset #0, #1. Active low. Indicates when the applied power is within tolerance and stable for the add-in card. PERST# goes high at least 100ms after the power rails are within operating limits per the PCIe CEM Specification. The PCIe REFCLKs also become stable within this period of time. PERST is pulled high on the baseboard. For OCP NIC 3.0, PERST deassertion also indicates 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 16.

The PRSNTA#/PRSNTB[0:3]#/BIF[0:2]# pins much be latched at least 1 ms of the system AC power on 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 10: 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 is directly connected to GND. For add-in cards, this pin is 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 are connected to the I/O hub and are pulled up to +3.3Vaux using 1kOhm resistors. For add-in cards, these pins are strapped to PRSNTA#. The encoding definitions are described in Section 3.6.

			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 are outputs driven from the baseboard I/O hub and allows 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 if no dynamic bifurcation configuration is required. For add-in cards, these signals connect to the endpoint bifurcation pins if it is supported. The BIF[0:2]# encoding definitions are described in Section 3.6. Note: the required combinatorial logic output for endpoint bifurcation is dependent on the specific silicon and is not defined in this specification.

Figure 16: PCIe Present and Bifurcation Control Pins

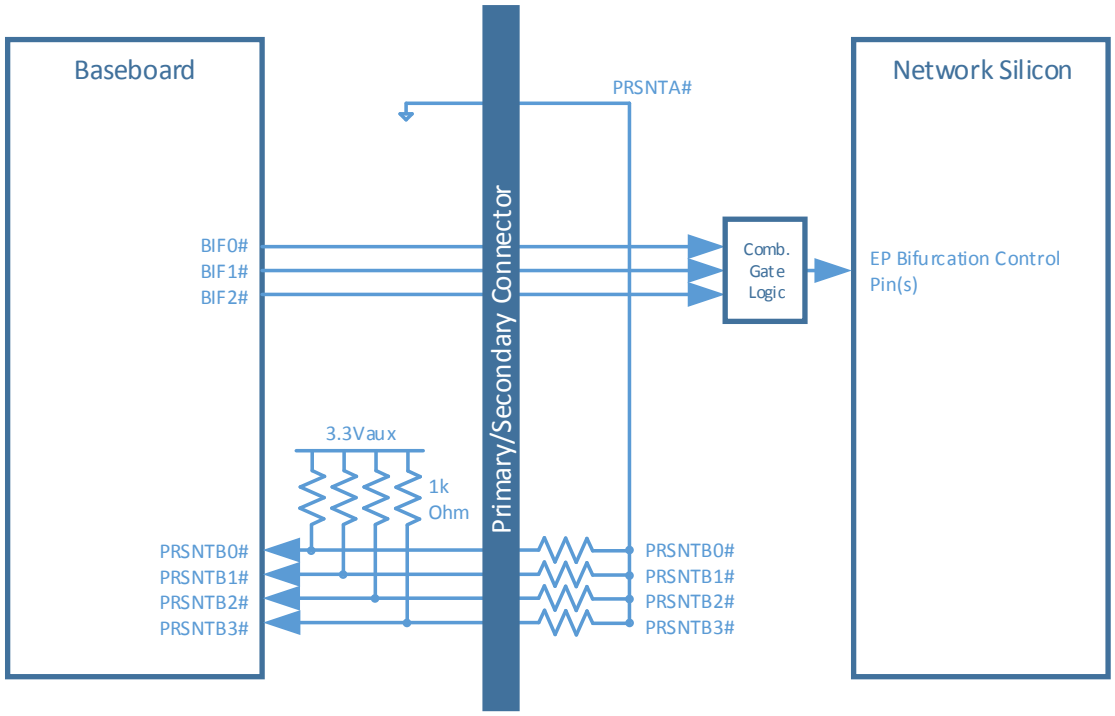

3.4.3 SMBus Interface Pins

This section provides the pin assignments for the SMBus interface signals. The AC/DC specifications are defined in the SMBus and I^2C bus specifications. An example connection diagram is shown in Figure XXX.

Table 11: 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, connect the SMCLK from the platform SMBus master to the connector. For add-in cards, connect the SMCLK from the endpoint silicon 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, connect the SMDAT from the platform SMBus master to the connector. For add-in cards, connect the SMDAT from the endpoint silicon to the card edge gold fingers.
SMRST#	A9	Output, OD	SMBus reset. Open drain. For baseboards, this pin is pulled up to 3.3Vauxand is used to reset optional downstream SMBus devices (such as temperature sensors). SMRST# is a mandatory signal for baseboard implementations. For add-in cards, SMRST# is optional.

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 XXX.

Table 12: Pin Descriptions – Power

Signal Name	**Pin #**	**Baseboard Direction**	**Signal Description**
GND	Various	GND	Ground return; a total of 46 ground pins are on the main 140-pin connector area.
+12V/+12V_AUX	B1, B2, B3, B4, B5, B6	Power	12V main or 12V Aux power; total of 6 pins per connector. The 12V pins are rated to 1.1A per pin with a maximum derated power delivery of 79.2W.

			The +12V power pins must be within the rail tolerances (TBD tolerance for Aux) when the PWRDIS pin is driven low by the baseboard.
+3.3V/3.3V_AUX	B11	Power	3.3V main or 3.3V Aux power; total of 1 pin per connector. The 3.3V pin is rated to 1.1A for a maximum derated power delivery of 3.63W. The 3.3Vaux/main power pin must be within the rail tolerances when the PWRDIS pin is driven low by the baseboard.
PWRDIS	B12	Output	Power disable. Active high. This signal is driven by the baseboard. When high, this signal notifies the add-in card to turn off all supplies connected to +12V power. When low, this signal notifies the add-in card to enable the on-card power supplies.

3.4.5 Miscellaneous Pins

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

Table 13: Pin Descriptions – Miscellaneous 1

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

	A69, A70		

3.5 Signal Descriptions – OCP Bay (Primary Connector)

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

The 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 3.7.

Table 14: 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.
REFCLKn3 REFCLKp3	OCP_A11 OCP_A12	Output	

			Note: REFCLK2 and REFCLK3 are not used for cards that only support a 1 x16 or 2 x8 connection. Refer to Section 2.1 in the PCIe CEM Specification, Rev 4.0 for details.
PERST2# PERST3#	OCP_A2 OCP_A3	Output	PCIe Reset #2, #3. Active low. Indicates when the applied power is within tolerance and stable for the add-in card. PERST# goes high at least 100ms after the power rails are within operating limits per the PCIe CEM Specification. The PCIe REFCLKs also become stable within this period of time. PERST is pulled high on the baseboard. For OCP NIC 3.0, PERST deassertion also indicates 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_A1	Input, OD	WAKE#. Open drain. Active low. This signal is driven by the add-in card to notify the baseboard restore the PCIe link. For add-in cards that support multiple WAKE# signals, their respective WAKE# pins may be tied together as the signal is open-drain to form a wired-OR. For baseboards, this signal is pulled up to +3.3V on the baseboard with a 10kOhm resistor and is connected to the system WAKE# signal. For add-in cards, this signal is connected directly 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 17.

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.

Table 15: Pin Descriptions – NC-SI Over RBT

Signal Name	Pin #	Baseboard Direction	Signal Description
RBT_CLK_IN	OCP_A14	Output	Reference clock input. Synchronous clock reference for receive, transmit and control interface. The clock has a nominal frequency of 50MHz ±100ppm. For baseboards, connect this pin 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 terminate this signal to ground through a 100kOhm pull down resistor. For add-in cards, connect this pin from the gold finger to the endpoint silicon. Leave this pin 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, connect this pin 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 terminate this signal to ground through a 100kOhm pull down resistor. For add-in cards, connect this pin from the gold finger to the endpoint silicon. Leave this pin 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, connect this pin 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 terminate this signal to 3.3Vaux through a 100kOhm pull-up. For add-in cards, connect this pin from the gold finger to the endpoint silicon. Leave this pin as a no connect if NC-SI is not supported.
RBT_TX_EN	OCP_A7	Output	Transmit enable. For baseboards, connect this pin 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 terminate this signal to ground through a 100kOhm pull down. For add-in cards, connect this pin from the gold finger to the endpoint silicon. Leave this pin 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, connect this pin 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 terminate this signal to 3.3Vaux through a 100kOhm pull-up. For add-in cards, connect this pin from the gold finger to the endpoint silicon. Leave this pin as a no connect if NC-SI is not supported.
RBT_ARB_OUT	OCP_A5	Output	NC-SI hardware arbitration output. Used only if the end point silicon supports hardware arbitration. Connects to the ARB_IN signal of an adjacent device. The ARB_IN pin is also routed to the card edge to allow multiple devices and OCP slots on the baseboard to share the NC-SI ring. 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, connect this pin between the baseboard OCP connector(s) to complete the hardware arbitration ring. If the baseboard does not support NC-SI over RBT, connect this signal directly to the RBT_ARB_IN pin. For add-in cards, connect this pin from the gold finger to the endpoint silicon. Leave this pin as a no connect if NC-SI is not supported.

RBT_ARB_IN	OCP_A4	Input	NC-SI hardware arbitration input. Used only if the end point silicon supports hardware arbitration. Connects to the ARB_OUT signal of an adjacent device. The ARB_OUT pin is also routed to the card edge to allow multiple devices and OCP slots on the baseboard to share the NC-SI ring. 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, connect this pin between the baseboard OCP connector(s) to complete the hardware arbitration ring. If the baseboard does not support NC-SI over RBT, connect this signal directly to the RBT_ARB_OUT pin. For add-in cards, connect this pin from the gold finger to 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. Used only if the end point silicon supports package identification. For baseboards, this pin is used to identify the slot ID value. Connect this pin directly to GND for SlotID = 0; or pull this pin up to 3.3Vaux for SlotID = 1.

			For add-in cards, connect this pin to the endpoint SLOT_ID pin for device address selection. For add-in cards with multiple end point devices, the SLOT_ID pin may be used to configure a different PHYAD bit so long as the resulting combination does not cause addressing interferences. For end point devices without NC-SI support, leave this pin as a no connect.

Figure 17: NC-SI Over RBT Connection Example

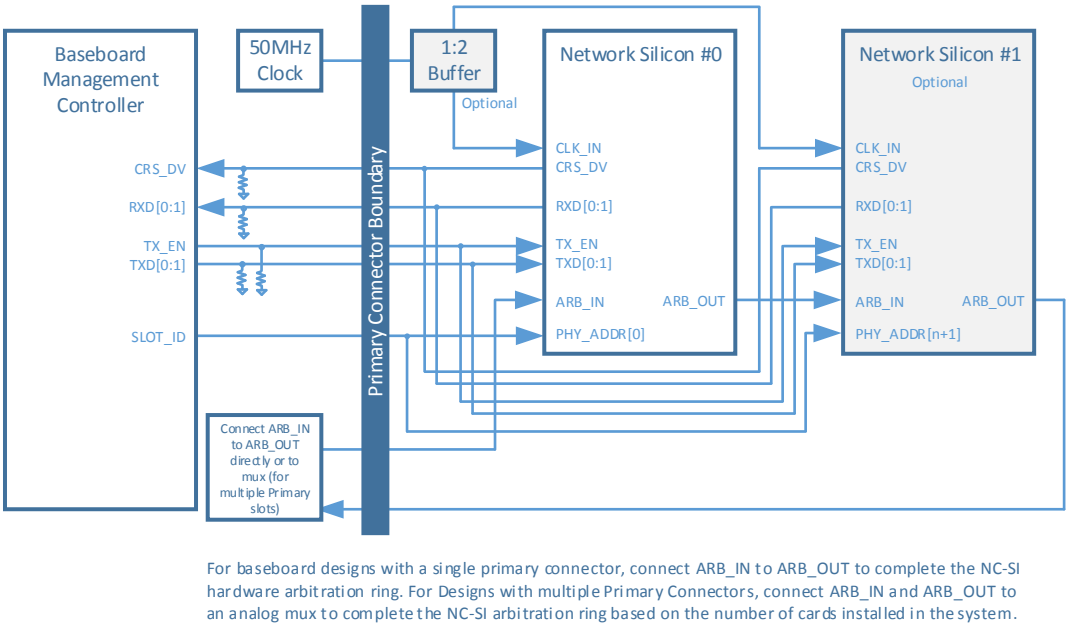

For baseboard designs with a single primary connector, connect ARB_IN to ARB_OUT to complete the NC-SI hardware arbitration ring. For Designs with multiple Primary Connectors, connect ARB_IN and ARB_OUT to an analog mux to complete the NC-SI arbitration ring based on the number of cards installed in the system.

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 18.

Table 16: 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, connect the CLK pin to the Primary Connector. Tie the CLK pin directly to GND if the scan chain is not used. For NIC implementations, the CLK pin must be connected to Shift Registers 0 & 1, and optionally to Shift Registers 2 & 3 (if implemented) as defined in the text and Figure 18, below. Pull the CLK pin 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, connect the DATA_OUT pin to the Primary Connector. Tie the DATA_OUT pin 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. Pull the DATA_OUT pin 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 connection to Shift Registers 0 & 1, as defined in the text and Figure 18, are required.
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 must be connected to Shift Registers 0 & 1 as defined in the text and Figure 18. Pull the LD# pin 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 17: Pin Descriptions – Scan Chain DATA_OUT Bit Definition

Byte.bit	**DATA_OUT Field Name**	**Default Value**	**Description**
0.[0..7]	RSVD	0b000000	Reserved. Byte 0 value is 0h00.
1.[0..7]	RSVD	0h00	Reserved. Byte 1 value is 0h00.
2.[0..7]	RSVD	0h00	Reserved. Byte 2 value is 0h00.
3.[0..7]	RSVD	0h00	Reserved. Byte 3 value is 0h00.

The DATA_IN bus is an input to the host and provides NIC status indication. The default implementation is completed with two 8-bit 74LV165 parallel in to serial out shift registers in a cascaded implementation. Up to four shift registers may be implemented to provide additional NIC status indication to the host platform.

DATA_IN shift registers 0 & 1 are mandatory for all cards. 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 defintions 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 0b1 for unused bits for implementations using less than four shift registers.

Table 18: Pin Descriptions – Scan Bus DATA_IN Bit Definition

Byte.bit	DATA_OUT Field Name	Default Value	Description
0.0	PRSNTB[0]#	0bX	PRSNTB[3:0]# value is mirrored from the Primary Connector.
0.1	PRSNTB[1]#	0bX	
0.2	PRSNTB[2]#	0bX	
0.3	PRSNTB[3]#	0bX	
0.4	WAKE_N	0bX	PCIe WAKE_N signal is mirrored from the Primary Connector.
0.5	TEMP_WARN	0b0	Temperature monitoring pin from the on-card thermal solution. This pin is asserted high when temperature sensor exceeds the temperature warning threshold.
0.6	TEMP_CRIT	0b0	Temperature monitoring pin from the on-card thermal solution. This pin is asserted high when temperature sensor exceeds the temperature critical threshold.
0.7	FAN_ON_AUX	0b0	When high, FAN_ON_AUX requests the system fan to be enabled for extra cooling in the S5 state.
1.0	LINK0	0b1	Port 0..3 link indication. Active low. 0b0 – Link LED is illuminated on the host platform. 0b1 – Link LED is not illuminated on the host platform. Steady = link is detected on the port. Off = no link is detected on the port.
1.1	LINK1	0b1	
1.2	LINK2	0b1	
1.3	LINK3	0b1	
1.4	ACT0	0b1	Port 0..3 activity indication. Active low. 0b0 – Link LED is illuminated on the host platform.
1.5	ACT1	0b1	
1.6	ACT2	0b1	
1.7	ACT3	0b1	

			0b1 – Link LED is not illuminated on the host platform. Steady = no activity is detected on the port Blink = activity is detected on the port. Off = no link, see also LINK[3:0] LED bits. The LED blink duty cycle is dependent on the add-in card implementation. The recommended duty cycle is 50%.
2.0	ScanChainVer[0]	0b1	ScanChainVer[1:0] is used to indicate the scan chain bit definitions. The encoding is as follows: 0b11 – Scan chain bit definitions version 1 corresponding to OCP 3.0 spec version 1.0. All other encodings are reserved.
2.1	ScanChainVer[1]	0b1	
2.2	RSVD	0b1	Byte 2 bits [2:7] are reserved. These bits shall default to the value of 0b1. These bits may be used in future versions of the scan chain.
2.3	RSVD	0b1	
2.4	RSVD	0b1	
2.5	RSVD	0b1	
2.6	RSVD	0b1	
2.7	RSVD	0b1	
3.0	LINK4	0b1	Port 4..7 link indication. Active low. 0b0 – Link LED is illuminated on the host platform. 0b1 – Link LED is not illuminated on the host platform. Steady = link is detected on the port. Off = no link is detected on the port.
3.1	LINK5	0b1	
3.2	LINK6	0b1	
3.3	LINK7	0b1	

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

Figure 18: Scan Bus Connection Example

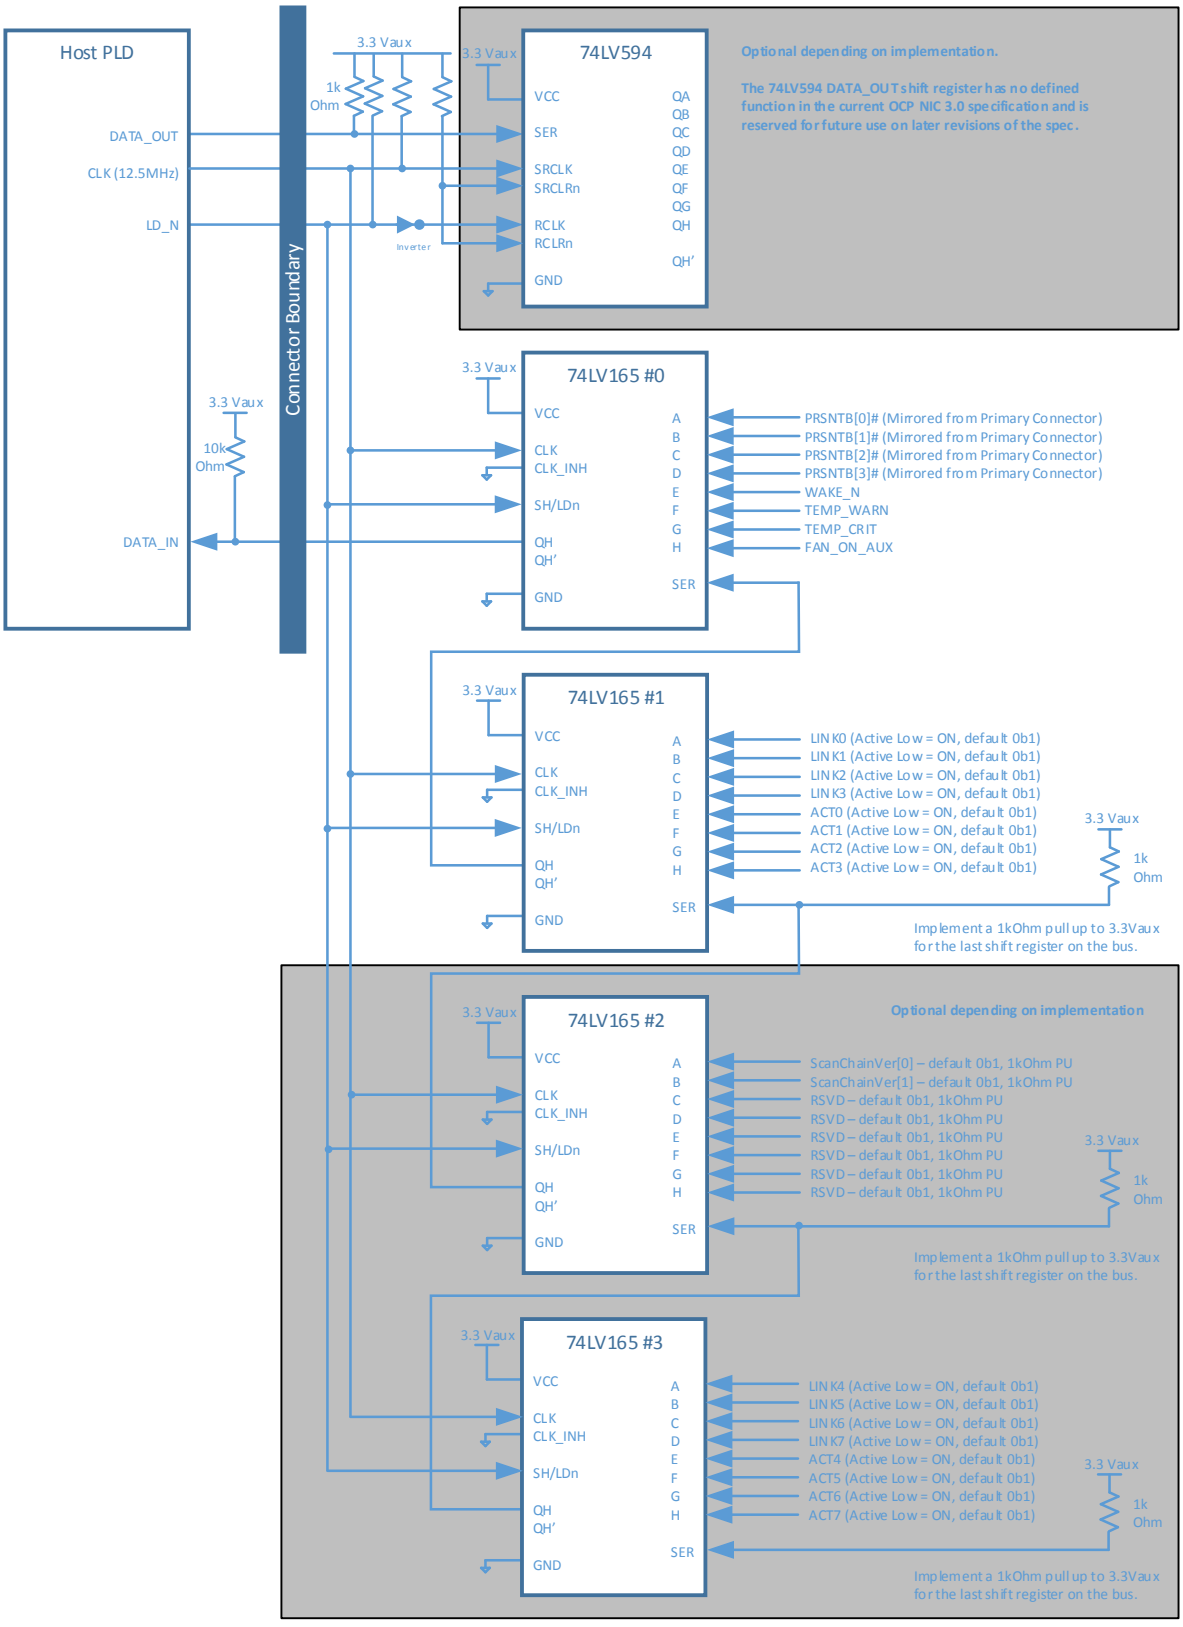

3.5.4 Primary Connector Miscellaneous Pins – OCP Bay (Primary Connector)

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 19: Pin Descriptions – Miscellaneous 2

Signal Name	Pin #	Baseboard Direction	Signal Description
PWRBRK#	OCP_B2	Output, OD	Power break. Active low, open drain. This signal is pulled up to 3.3Vaux on the add-in card with a minimum of 95kOhm and the baseboard with a stiffer resistance in-order to meet the timing specs as shown in the PCIe CEM Specification. This signal is driven low by the baseboard and is used to notify that an 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 indicates that all of the add-in card power rails are operating within nominal tolerances. When low the add-in card power supplies are not yet ready 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 is 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 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 connects to the PRSNTB# pins as a hard coded value on the add-in card. The encoding of the PRSNTB[3:0]# pins allows the baseboard to determine the PCIe Links available on the add-in card.
- BIF[3:0]#. The BIF# pin states are controlled by the baseboard and allows the baseboard to override the default end point bifurcation for silicon that support bifurcation. Additional combinatorial logic is required and is specific to the card silicon. The combinatorial logic is not covered in this specification.

The high level bifurcation connections are shown in Figure 19.

Figure 19: PCIe Bifurcation Pin Connections Support

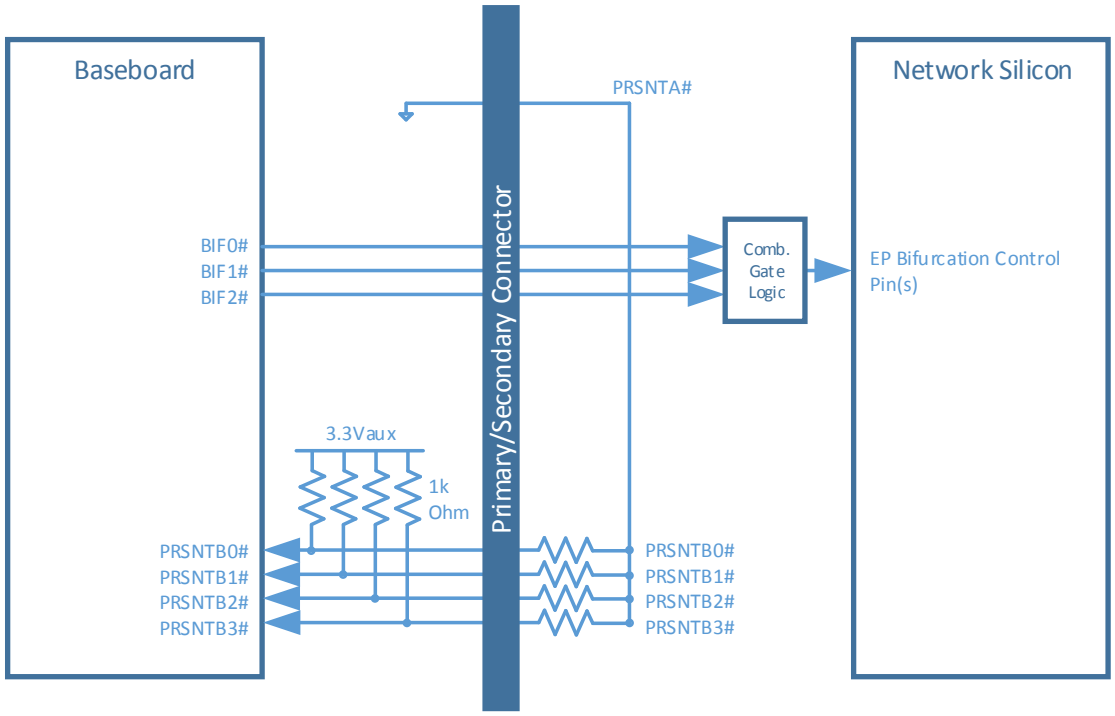

3.6.1 PCIe Add-in Card to Baseboard Bifurcation Configuration (PRSNTA#, PRSNTB[3:0]#)

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 0b1111 indicates that no card is present in the system. Depending on the capabilities of the add-in card, a selection of PRSNTB[3:0]# signals may be strapped to the PRSNTA# signal and is pulled low by the baseboard. The encoding of the PRSTNB[3:0]# bits is shown in Table 20 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 20 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 20 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 20: PCIe Bifurcation Decoder for x16 and x8 Card Widths

				Single Host					RSVD	RSVD	Dual Host	Quad Host	Quad/Oct Host
	Network Card – Supported PCIe Configurations		Host	1 Host	1 Host	1 Host	1 Host	1 Host	RSVD	RSVD	2 Hosts	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	RSVD	2 Upstream Sockets (1 Socket per Host)	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	RSVD	2 Links	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, 2x1				2 x8, 2 x4, 2 x2, 2 x1		
						4 x4, 4 x2, 4 x1		4 x4, 4 x2, 4x1				4 x4, 4 x2, 4 x1	4 x2, 4 x1
Minimum Required Card Edge	Card Short Name	x16 Cards	System Encoding	0b000	0b000	0b000	0b001	0b010	0b011	0b100	0b101	0b110	0b111
			Add-in-Card Encoding PRSNTB[3:0]#	-	-	-	-	-	-	-	-	-	-
n/a	Not Present	Card Not Present	0b**1111**	RSVD – Card not present in the system.									
2C	1x8	1x8, 1x4, 1x2, 1x1	0b**1110**	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	0b**1110**	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	0b**1110**	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	0b**1110**	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	0b**1101**	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	0b**1101**	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	0b**1100**	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	0b**1100**	1x16	1x16	1x16	2 x8	4 x4	-	-	2 x8	4 x4	4 x2
RSVD	RSVD	RSVD	0b**1011**	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 1x4, 1x2, 1x1	0b**1010**	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	0b**1001**	-	-	-	-	-	-	-	-	-	-
RSVD	RSVD	RSVD for future x8 encoding	0b**1000**	-	-	-	-	-	-	-	-	-	-
4C	1x16	1x16, 1x8, 1x4, 1x2, 1x1	0b**0111**	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	0b**0110**	1x8*	2 x8	2 x8	2 x8	2 x4 (Socket 0 & 2 only)	-	-	2 x8	2 x4 (Host 0 & 2 only)	1x2 (Host 0 & 1 only)
4C	1x16 Option B	1x16, 1x8, 1x4, 1x2, 1x1 2 x8, 2 x4, 2 x2, 2 x1	0b**0101**	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	0b**0100**	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	0b**0011**	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	0b**0010**	-	-	-	-	-	-	-	-	-	-
RSVD	RSVD	RSVD	0b**0001**	-	-	-	-	-	-	-	-	-	-
RSVD	RSVD	RSVD	0b**0000**	-	-	-	-	-	-	-	-	-	-

3.6.4 Bifurcation Detection Flow

[Need input and clarification from system vendors]

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

1. The baseboard reads the state of the PRSNTB[3:0]# pins. If the resulting value is not 0b1111, an add-in card is present.
2. Firmware determines the add-in card PCIe lane width capabilities per Table 20 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 would be asserted as appropriate. Doing so assumes the add-in card supports the requested link override.
5. PERST# is deasserted after the >100ms window as defined by the PCIe specification. Refer to Section 3.12 for timing details.

3.6.5 PCIe Bifurcation Examples

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

Figure 20 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 0b0111. The BIF[2:0]# state is 0b000 as there is no need to instruct the end-point network controller to a specific bifurcation. The PRSNTB encoding notifies the baseboard that this card is only capable of 1 x16. The single host baseboard determines that it is also capable of supporting 1 x16. The resulting link width is 1 x16.

Figure 20: Single Host (1 x16) and 1 x16 Add-in Card (Single Controller)

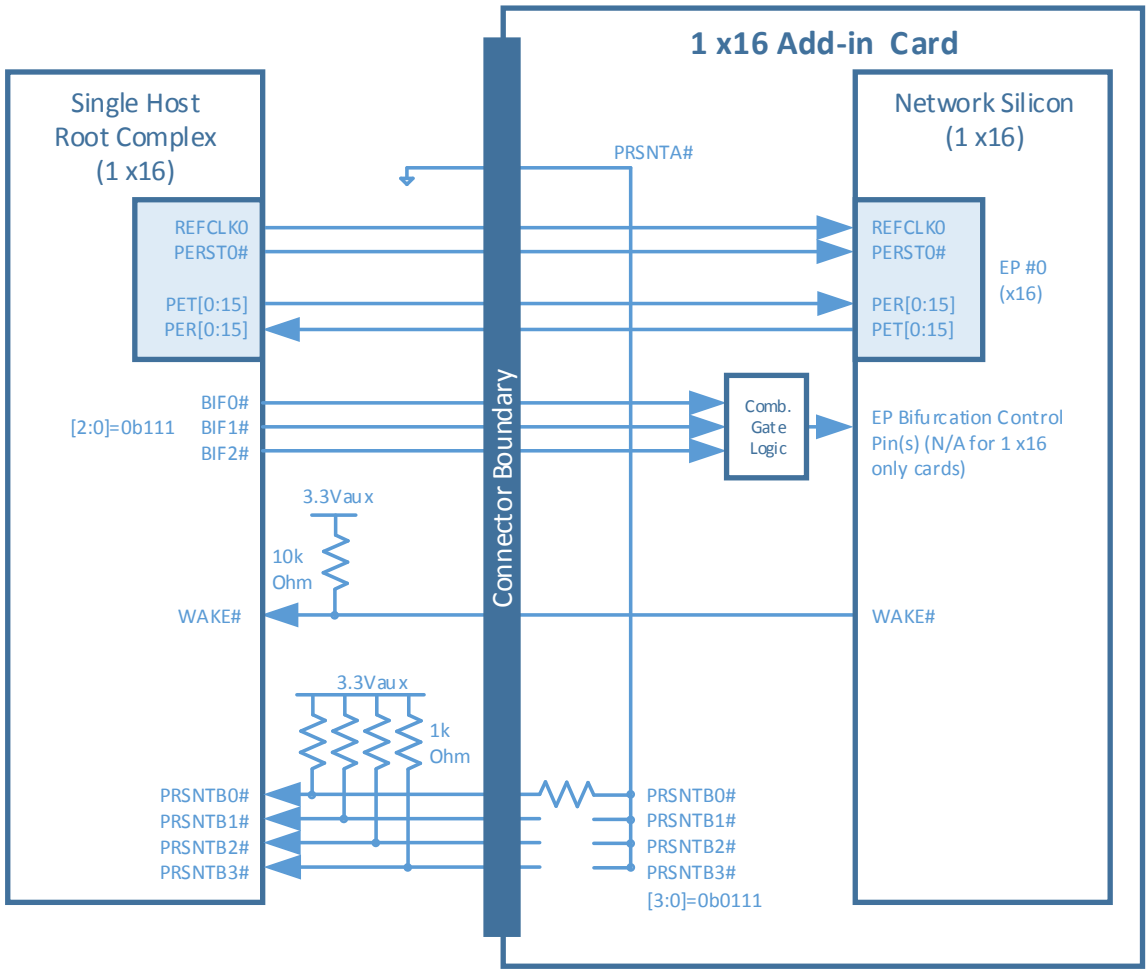

Figure 21 illustrates a single host baseboard that supports 2 x8 with a single controller add-in card that also supports 2 x8. The PRSTNB[3:0]# state is 0b0110. The BIF[2:0]# state is

0b111 as there is no need to instruct the end-point network controllers to a specific bifurcation. The PRSNTB encoding notifies the baseboard that this card is only capable of 2 x8. The single host baseboard determines that it is also capable of supporting 2 x8. The resulting link width is 2 x8.

Figure 21: Single Host (2 x8) and 2 x8 Add-in Card (Dual Controllers)

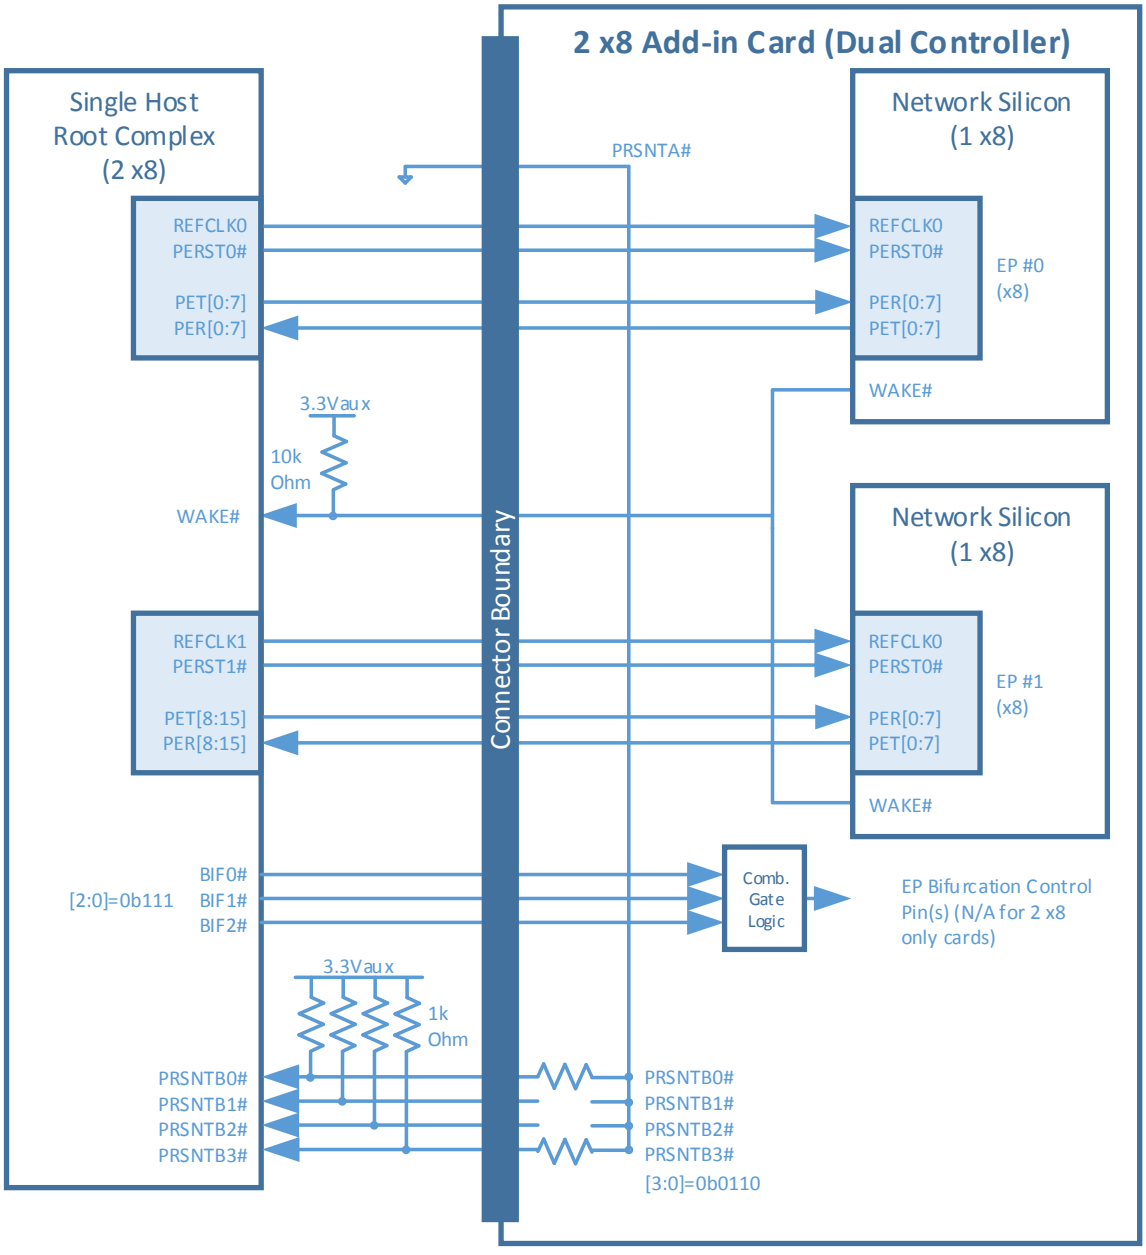

Figure 22 illustrates a four host baseboard that supports 4 x4 with a single controller add-in card that supports 1 x16, 2 x8 and 4 x4. The PRSTNB[3:0]# state is 0b0011. The BIF[2:0]# state is 0b101 as the end point network controller is forced to bifurcate to 4 x4. The PRSNTB encoding notifies the baseboard that this card is only capable of 1 x16, 2 x8 and 4 x4. The four host baseboard determines that it is also capable of supporting 4 x4. The resulting link width is 4 x4.

Figure 22: Four Hosts (4 x4) and 4 x4 Add-in Card (Single Controller)

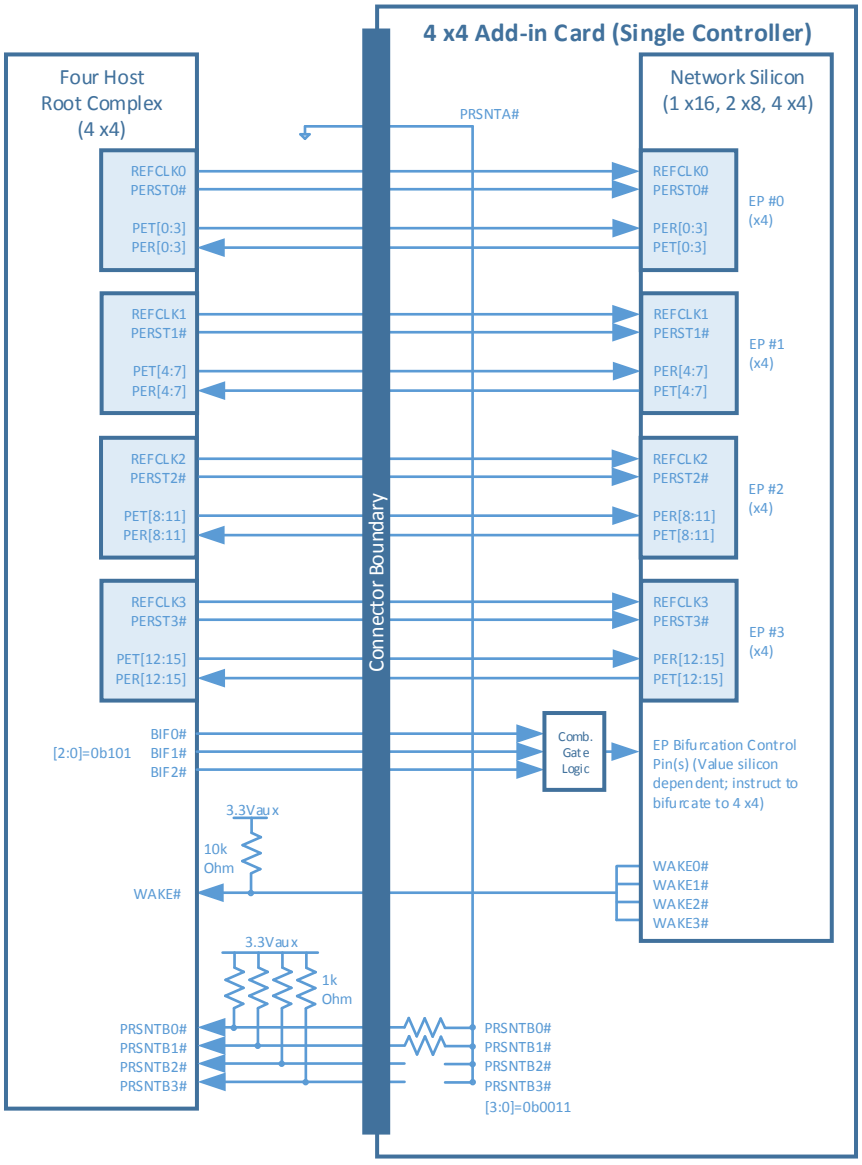

Figure 23 illustrates a four host baseboard that supports 4 x4 with a four controller add-in card that supports 4 x4. The PRSTNB[3:0]# state is 0b0011. The BIF[2:0]# state is 0b111 as there is no need to instruct the end-point network controllers to a specific bifurcation. The PRSNTB encoding notifies the baseboard that this card is only capable of 4 x4. The four host baseboard determines that it is also capable of supporting 4 x4. The resulting link width is 4 x4.

Figure 23: Four Hosts (4 x4) and 4 x4 Add-in Card (Four Controllers)

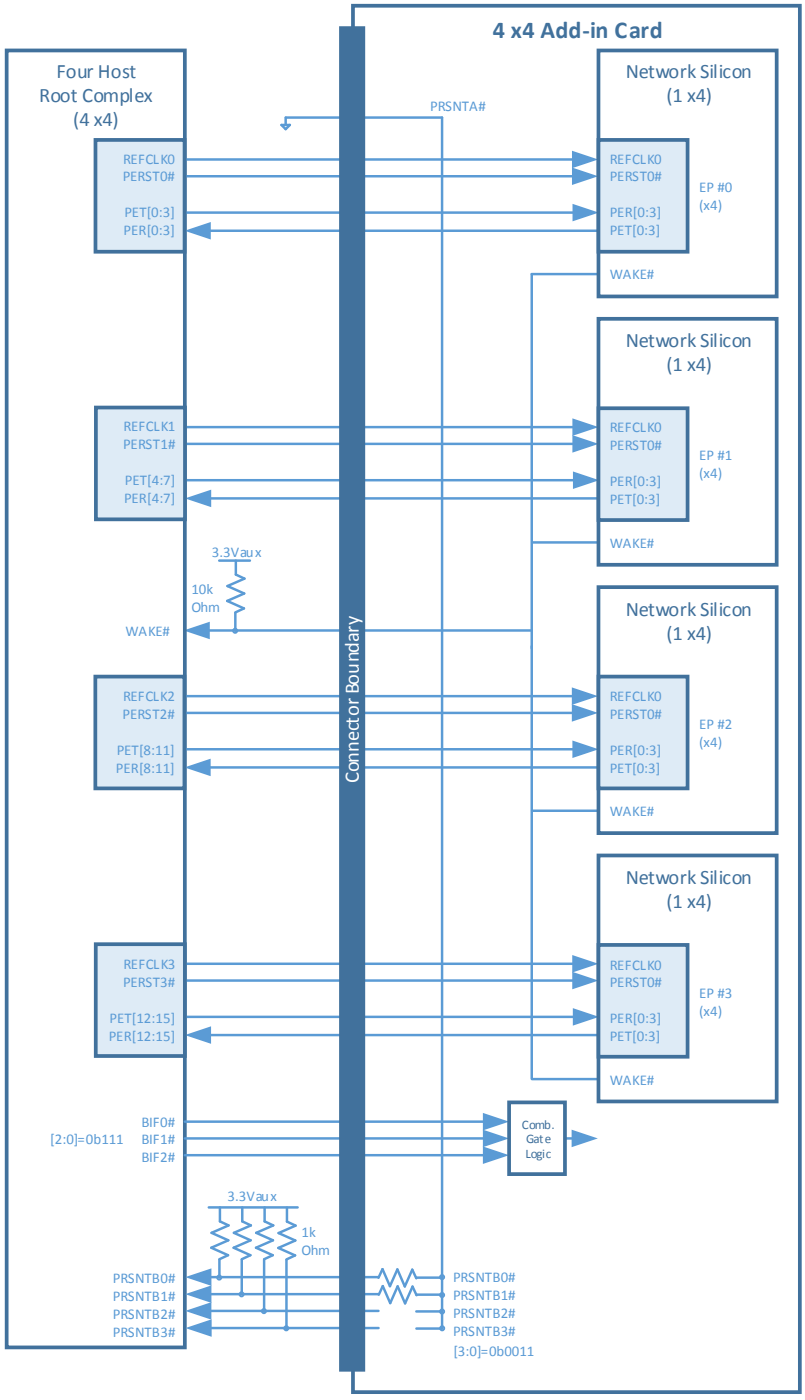
4 x4 Add-in Card
Four Host Root Complex (4 x4)
Connector Boundary
PRSNTA#
REFCLK0
PERST0#
PET[0:3]
PER[0:3]
REFCLK1
PERST1#
PET[4:7]
PER[4:7]
3.3Vaux
10k Ohm
WAKE#
REFCLK2
PERST2#
PET[8:11]
PER[8:11]
REFCLK3
PERST3#
PET[12:15]
PER[12:15]
BIF0#
BIF1#
BIF2#
[2:0]=0b111
3.3Vaux
1k Ohm
PRSNTB0#
PRSNTB1#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

Figure 24 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 24: Single Host with no Bifurcation (1 x16) and 2 x8 Add-in Card (Two Controllers)

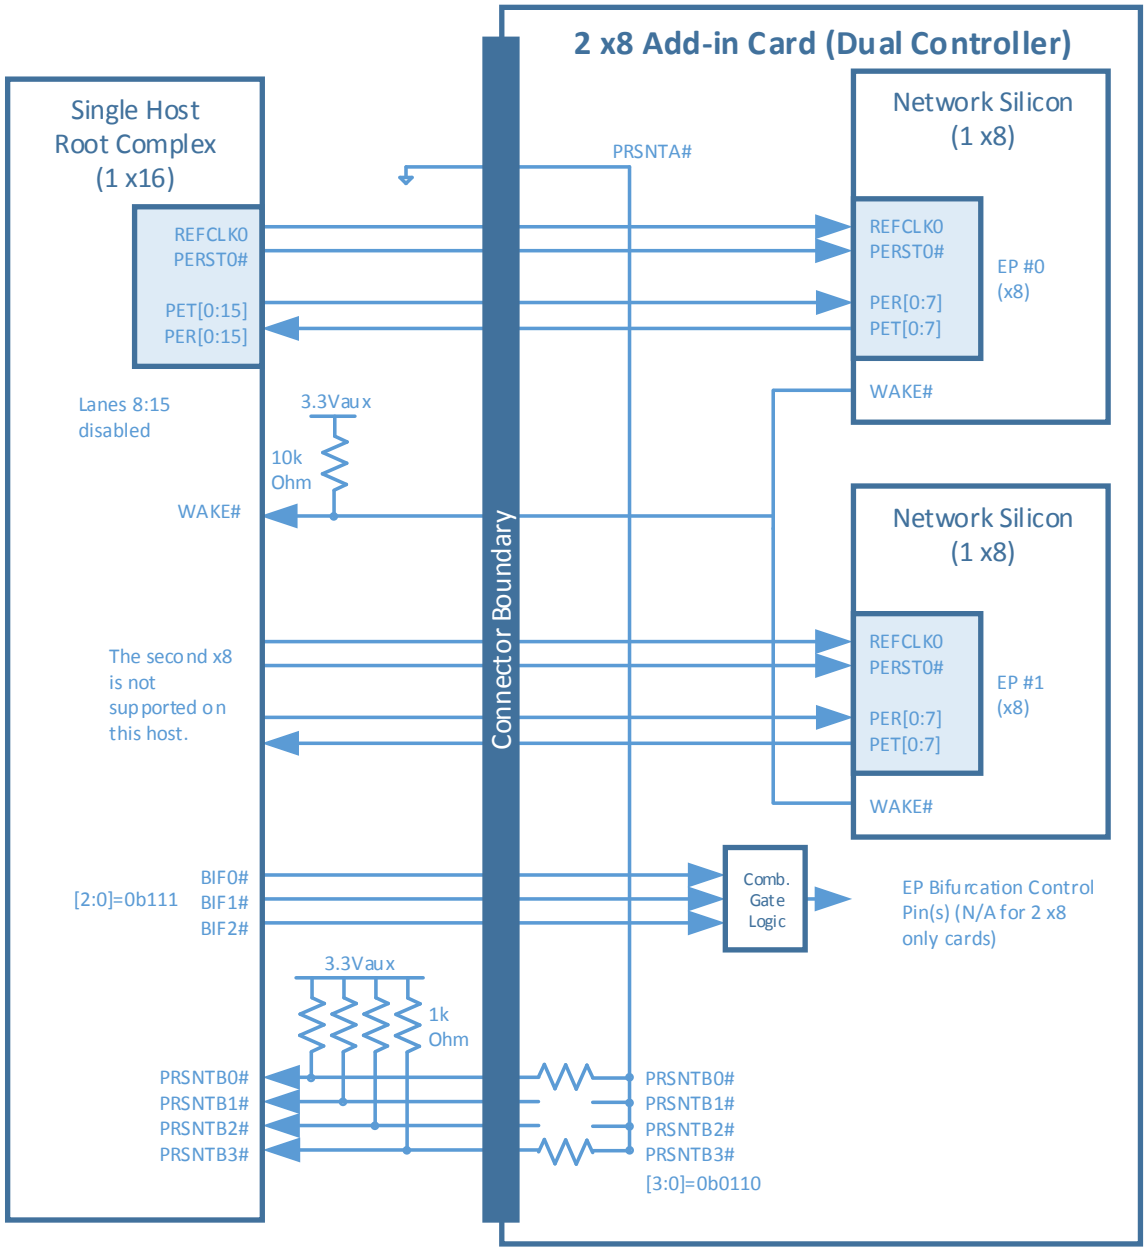

3.7 PCIe Clocking Topology

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

Table 21: 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 is used for lanes [0:15].
- For a 2 x8 capable add-in card, REFCLK0 is used for lanes [0:7] and REFCLK1 is used for lanes [8:15].
- For a 4 x4 capable add-in card, REFCLK0 is used for lanes [0:3], REFCLK1 is used for lanes [4:7], REFCLK2 is used for lanes [8:11] and REFCLK3 is used for lanes [12:15]. Pins for REFCLK2 and REFCLK3 are described in Section 3.5.1 and are located on the 28-pin OCP bay.

Figure 25: PCIe Interface Connections for 1 x16 and 2 x8 Add-in Cards

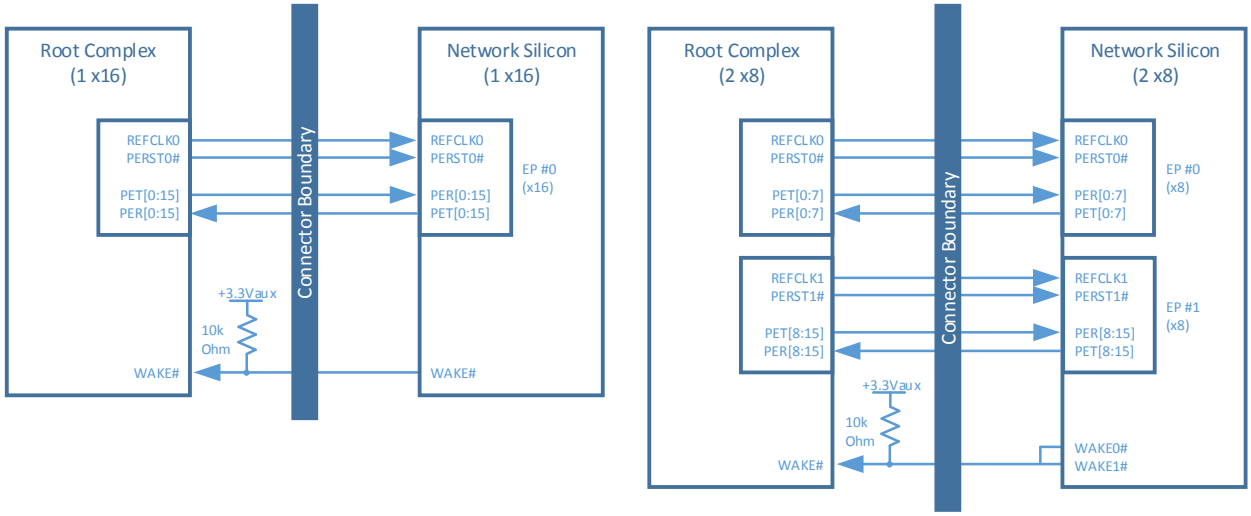

Figure 26: PCIe Interface Connections for a 4 x4 Add-in Card

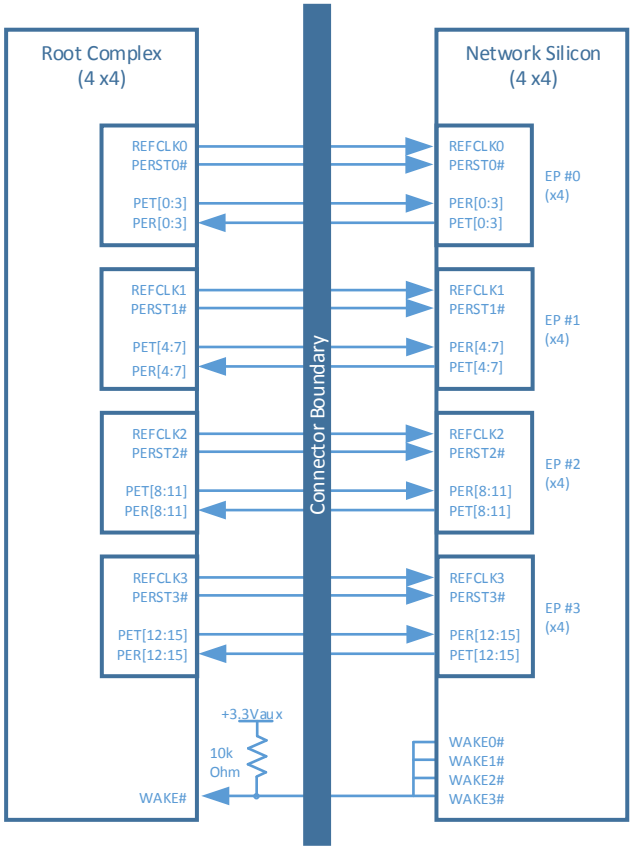

3.8 PCIe Bifurcation Results and REFCLK Mapping

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

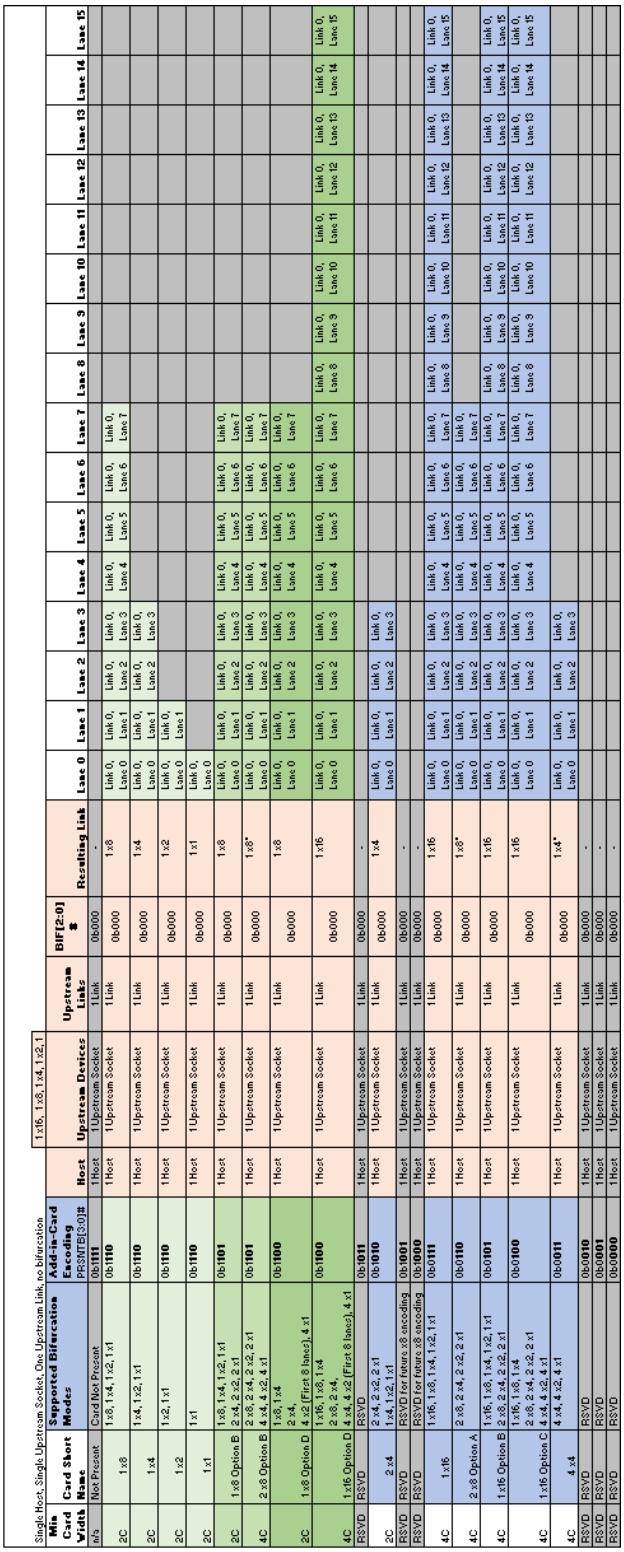

Single Host, Single Upstream Socket, One Upstream Link, no bifurcation

Min Card Width	Card Short Name	Supported Bifurcation Modes	Add-in-Card Encoding PRSNTB[3:0]#	Host	Upstream Devices	1:x16, 1:x8, 1:x4, 1:x2, 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	-	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	1 Host	1 Upstream Socket	1 Link	0b000	1:x8	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	1:x4	1:x4, 1:x2, 1:x1	0b1110	1 Host	1 Upstream Socket	1 Link	0b000	1:x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	1:x2	1:x2, 1:x1	0b1110	1 Host	1 Upstream Socket	1 Link	0b000	1:x2	Link 0, Lane 0	Link 0, Lane 1														
2C	1:x1	1:x1	0b1110	1 Host	1 Upstream Socket	1 Link	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	1 Host	1 Upstream Socket	1 Link	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	1 Host	1 Upstream Socket	1 Link	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	1:x8 Option D	1:x8, 1:x4 2:x4, 4:x2 (First 8 lanes), 4:x1	0b1100	1 Host	1 Upstream Socket	1 Link	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	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 D	1:x16, 1:x8, 1:x4 2:x8, 2:x4, 4:x4, 4:x2 (First 8 lanes), 4:x1	0b1100	1 Host	1 Upstream Socket	1 Link	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	1 Host	1 Upstream Socket	1 Link	0b000	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	2:x4	2:x4, 2:x2, 2:x1 1:x4, 1:x2, 1:x1	0b1010	1 Host	1 Upstream Socket	1 Link	0b000	1:x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
RSVD	RSVD	RSVD for future x8 encoding	0b0001	1 Host	1 Upstream Socket	1 Link	0b000	-																
RSVD	RSVD	RSVD for future x8 encoding	0b0000	1 Host	1 Upstream Socket	1 Link	0b000	-																
4C	1:x16	1:x16, 1:x8, 1:x4, 1:x2, 1:x1	0b0111	1 Host	1 Upstream Socket	1 Link	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	0b0110	1 Host	1 Upstream Socket	1 Link	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 B	1:x16, 1:x8, 1:x4, 1:x2, 1:x1 2:x8, 2:x4, 2:x2, 2:x1	0b0101	1 Host	1 Upstream Socket	1 Link	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	1 Host	1 Upstream Socket	1 Link	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	1 Host	1 Upstream Socket	1 Link	0b000	1:x4*	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 23: 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

Min Card Width	Card Short Name	Supported Bifurcation Modes	Addin-Card Encoding PRSNTB[3:0]#	Host	Upstream Devices	1:x16, 1:x8, 1:x4, 1:x2, 1 2:x8, 2:x4, 2:x2, 2:x1 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	1or 2 Links	0b000	-																	
2C	1x8	1x8, 1x4, 1x2, 1x1	0b1110	1Host	1Upstream Socket	1or 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	1or 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	1or 2 Links	0b000	1x2	Link 0, Lane 0	Link 0, Lane 1														
2C	1x1	1x1	0b1110	1Host	1Upstream Socket	1or 2 Links	0b000	1x1	Link 0, Lane 0															
2C	1x8 Option B	1x8, 1x4, 1x2, 1x1 2x4, 2x2, 2x1	0b1101	1Host	1Upstream Socket	1or 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	1or 2 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	1x8 Option D	1x8, 1x4 2x4 4 x2 (First 8 lanes), 4 x1	0b1100	1Host	1Upstream Socket	1or 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	1or 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								
RSVD	RSVD	RSVD	0b1011	1Host	1Upstream Socket	1or 2 Links	0b000	-																
2C	2 x4	2 x4, 2 x2, 2 x1 1 x4, 1 x2, 1 x1	0b1010	1Host	1Upstream Socket	1or 2 Links	0b000	1 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	1Host	1Upstream Socket	1or 2 Links	0b000	-																
RSVD	RSVD	RSVD for future x8 encoding	0b1000	1Host	1Upstream Socket	1or 2 Links	0b000	-																
4C	1x16	1x16, 1x8, 1x4, 1x2, 1x1 2x8, 2x4, 2x2, 2x1 4x4, 4x2, 4x1	0b0111	1Host	1Upstream Socket	1or 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, 2x2, 2x1 4x4, 4x2, 4x1	0b0110	1Host	1Upstream Socket	1or 2 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	1x16 Option B	1x16, 1x8, 1x4, 1x2, 1x1 2x8, 2x4, 2x2, 2x1 4x4, 4x2, 4x1	0b0101	1Host	1Upstream Socket	1or 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, 2x2, 2x1 4x4, 4x2, 4x1	0b0100	1Host	1Upstream Socket	1or 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	1or 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	1or 2 Links	0b000	-																
RSVD	RSVD	RSVD	0b0001	1Host	1Upstream Socket	1or 2 Links	0b000	-																
RSVD	RSVD	RSVD	0b0000	1Host	1Upstream Socket	1or 2 Links	0b000	-																

Table 24: Bifurcation for Single Host, Single Socket and Single/Dual/Quad Upstream Links (BIF[2:0]#=0b000)

Single Host, Single Upstream Socket, One, Two or Four Upstream Links

Min Card Width	Card Short Name	Supported Bifurcation Modes	Add-in-Card Encoding PRSNTB[3:0]#	Host	Upstream Devices (1x16, 1x8, 1x4, 1x2, 1 / 2x8, 2x4, 2x2, 2x1 / 4x4, 4x2, 4x1)	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								
x1	x8	1x8, 1x4, 1x2, 1x1	0b1110	1Host	1Upstream Socket	1, 2, or 4 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								
x1	x4	1x4, 1x2, 1x1	0b1110	1Host	1Upstream Socket	1, 2, or 4 Links	0b000	1x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
x1	x2	1x2, 1x1	0b1110	1Host	1Upstream Socket	1, 2, or 4 Links	0b000	1x2	Link 0, Lane 0	Link 0, Lane 1														
x1	x1	1x1	0b1110	1Host	1Upstream Socket	1, 2, or 4 Links	0b000	1x1	Link 0, Lane 0															
x1	x8 Option B	1x8, 1x4, 1x2, 1x1 2x4, 2x2, 2x1	0b1101	1Host	1Upstream Socket	1, 2, or 4 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, 2, or 4 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 4x2 (First 8 lanes), 4x1	0b1100	1Host	1Upstream Socket	1, 2, or 4 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, 4x4, 4x2 (First 8 lanes), 4x1	0b1100	1Host	1Upstream Socket	1, 2, or 4 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, 2, or 4 Links	0b000	-																
2C	2 x4	2x4, 2x2, 2x1 1x4, 1x2, 1x1	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 encoding	0b1001	1Host	1Upstream Socket	1, 2, or 4 Links	0b000	-																
RSVD	RSVD	RSVD for future x8 encoding	0b1000	1Host	1Upstream Socket	1, 2, or 4 Links	0b000	-																
4C	1x16	1x16, 1x8, 1x4, 1x2, 1x1 2x8, 2x4, 2x2, 2x1 4x4, 4x2, 4x1	0b0111	1Host	1Upstream Socket	1, 2, or 4 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, 2x2, 2x1 4x4, 4x2, 4x1	0b0110	1Host	1Upstream Socket	1, 2, or 4 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, 2x2, 2x1 4x4, 4x2, 4x1	0b0101	1Host	1Upstream Socket	1, 2, or 4 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, 2x2, 2x1 4x4, 4x2, 4x1	0b0100	1Host	1Upstream Socket	1, 2, or 4 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	4x4, 4x2, 4x1	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 25: 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 Bifurcation Modes	Add-in-Card Encoding PRSNTB[3:0]#	Host	Upstream Devices (1x8, 1x4, 1x2, 1x1 / 2x8, 2x4, 2 x2, 2x1)	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	2 Upstream Sockets	2 Links	0b001	-																
2C	1x8	1x8, 1x4, 1x2, 1x1	0b1110	1Host	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	1Host	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	1Host	2 Upstream Sockets	2 Links	0b001	1x2 (Socket 0 only)	Link 0, Lane 0	Link 0, Lane 1														
2C	1x1	1x1	0b1110	1Host	2 Upstream Sockets	2 Links	0b001	1x1 (Socket 0 only)	Link 0, Lane 0															
2C	1x8 Option B	1x8, 1x4, 1x2, 1x1 2x4, 2 x2, 2 x1	0b1101	1Host	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, 2x4, 2 x2, 2 x1 4x4, 4 x2, 4 x1	0b1101	1Host	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 2x4, 4 x2 (First 8 lanes), 4 x1	0b1100	1Host	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	1Host	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	1Host	2 Upstream Sockets	2 Links	0b001	-																
2C	2 x4	2 x4, 2 x2, 2 x1 1 x4, 1 x2, 1 x1	0b1010	1Host	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, x8 encoding	0b1001	1Host	2 Upstream Sockets	2 Links	0b001	-																
RSVD	RSVD	RSVD for future, x8 encoding	0b1000	1Host	2 Upstream Sockets	2 Links	0b001	-																
4C	1x16	1x16, 1x8, 1x4, 1x2, 1x1 2x8, 2 x4, 2 x2, 2 x1 4x4, 4 x2, 4 x1	0b0111	1Host	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 A	2x8, 2x4, 2 x2, 2 x1 4x4, 4 x2, 4 x1	0b0110	1Host	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 2x8, 2 x4, 2 x2, 2 x1 4x4, 4 x2, 4 x1	0b0101	1Host	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 2x8, 2 x4, 2 x2, 2 x1 4x4, 4 x2, 4 x1	0b0100	1Host	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	1Host	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	1Host	2 Upstream Sockets	2 Links	0b001	-																
RSVD	RSVD	RSVD	0b0001	1Host	2 Upstream Sockets	2 Links	0b001	-																
RSVD	RSVD	RSVD	0b0000	1Host	2 Upstream Sockets	2 Links	0b001	-																

Table 26: 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	Upstream Devices	4 x4, 4 x2, 4x1	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	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	1Host	4 Upstream Sockets		4 Links	0b010	1 x4 (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	1Host	4 Upstream Sockets		4 Links	0b010	1 x4 (Socket 0 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
2C	1x2	1x2, 1x1	0b1110	1Host	4 Upstream Sockets		4 Links	0b010	1 x2 (Socket 0 only)	Link 0, Lane 0	Link 0, Lane 1														
2C	1x1	1x1	0b1110	1Host	4 Upstream Sockets		4 Links	0b010	1 x1 (Socket 0 only)	Link 0, Lane 0															
2C	1x8 Option B	1x8, 1x4, 1x2, 1x1 2 x4, 2 x2, 2 x1	0b1101	1Host	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	1Host	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	1Host	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	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	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	1Host	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	1Host	4 Upstream Sockets		4 Links	0b010	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	1Host	4 Upstream Sockets		4 Links	0b010	-																
RSVD	RSVD	RSVD for future x8 encoding	0b1000	1Host	4 Upstream Sockets		4 Links	0b010	-																
4C	1 x16	1 x16, 1 x8, 1 x4, 1 x2, 1 x1	0b0111	1Host	4 Upstream Sockets		4 Links	0b010	1 x4 (Socket 0 only)	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3												
4C	2 x8 Option A	2 x8, 2 x4, 2 x2, 2 x1	0b0110	1Host	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	1 x16 Option B	1 x16, 1 x8, 1 x4, 1 x2, 1 x1 2 x8, 2 x4, 2 x2, 2 x1	0b0101	1Host	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	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	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	1Host	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	1Host	4 Upstream Sockets		4 Links	0b010	-																
RSVD	RSVD	RSVD	0b0001	1Host	4 Upstream Sockets		4 Links	0b010	-																
RSVD	RSVD	RSVD	0b0000	1Host	4 Upstream Sockets		4 Links	0b010	-																

Table 27: 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 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	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	2x8 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 2 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	2x8 Option A	2 x8, 2 x4, 2 x2, 2 x1 1 x16, 1 x8, 1 x4, 1 x2, 1 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	1x16 Option B	2 x8, 2 x4, 2 x2, 2 x1 1 x16, 1 x8, 1 x4	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	1x16 Option C	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 28: 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]#	4 x4, 4 x2, 4 x1																				
				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	4 Host	4 Upstream Sockets	4 Links	0b110	-																
2C	1x8	1x8, 1x4, 1x2, 1x1	0b1110	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												
2C	1x4	1x4, 1x2, 1x1	0b1110	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												
2C	1x2	1x2, 1x1	0b1110	4 Host	4 Upstream Sockets	4 Links	0b110	1 x2 (Host 0 only)	Link 0, Lane 0	Link 0, Lane 1														
2C	1x1	1x1	0b1110	4 Host	4 Upstream Sockets	4 Links	0b110	1 x1 (Host 0 only)	Link 0, Lane 0															
2C	1x8 Option B	1x8, 1x4, 1x2, 1x1 2 x4, 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	2 x8, 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 2 x4, 4 x2 (First 8 lanes), 4 x1	0b1100	4 Host	4 Upstream Sockets	2 x4	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	1x16 Option D	1x16, 1x8, 1x4 2 x8, 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	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	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	2 x4	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	2 x8, 2 x4, 2 x2, 2 x1	0b0110	4 Host	4 Upstream Sockets	4 Links	0b110	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 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 (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 29: 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

					4 x2, 4 x1																			
Min Card Width	**Card Short Name**	**Supported Bifurcation Modes**	**Add-in-Card Encoding** PRSNTB[3:0]#	**Host**	**Upstream Devices**	**Upstream Links**	**BIF[2:0] #**	**Resulting Link #**	**Lane 0**	**Lane 1**	**Lane 2**	**Lane 3**	**Lane 4**	**Lane 5**	**Lane 6**	**Lane 7**	**Lane 8**	**Lane 9**	**Lane 10**	**Lane 11**	**Lane 12**	**Lane 13**	**Lane 14**	**Lane 15**
n/a	Not Present	Card Not Present	0b1111	4/8 Host	4/8 Upstream Sockets	4 or 8 x2 Links	0b111	-	Link 0, Lane 0	Link 0, Lane 1														
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	-	Link 0, Lane 0	Link 0, Lane 1	Link 1, Lane 0	Link 1, Lane 1												
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	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: AC Power Off, Management (FRU Only Mode), Aux Power Mode (S5), and Main Power Mode (S0). The transition of these states is shown in . The max available power envelopes for each of these states are defined in Table 30.

Figure 27: Baseboard Power Sequencing

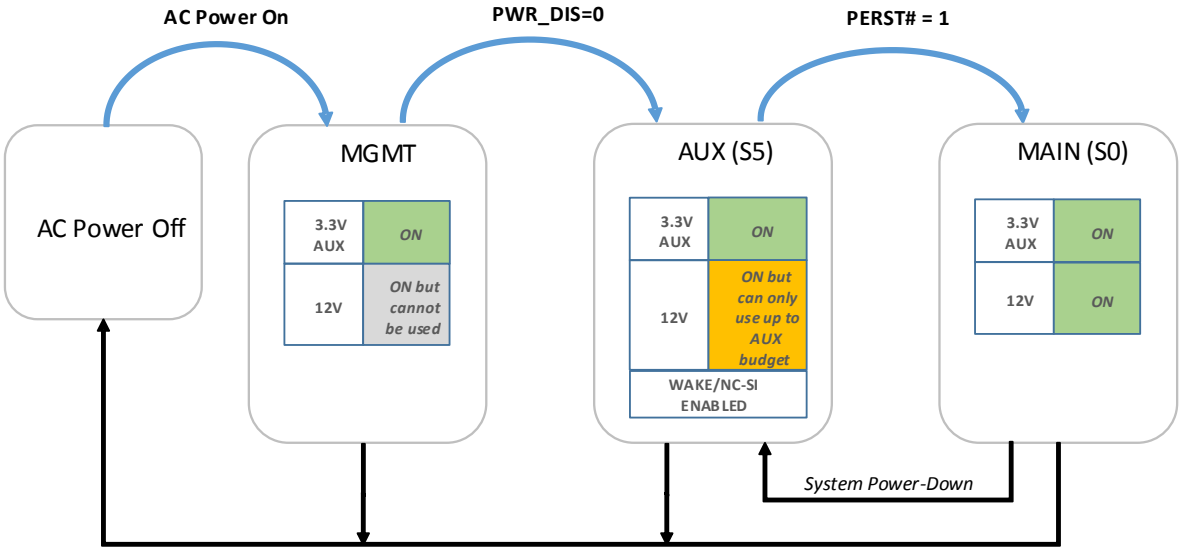

Table 30: Power Envelopes

Power State	Max Power	Notes
AC Power Off	0W	AC power removed; board off
Management (FRU only mode)	1W	Used only for board identification purposes.
Aux Power Mode (S5)	35W	
Main Power Mode (S0)	79.2W	Add-in card may use up to the 79.2W limit per connector. The connector is derated 1.1A of current per pin (6 pins total) for a 30degC rise in the thermoplastic connector shell.

3.9.1 AC Power Off

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

3.9.2 Management (FRU Only Mode)

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

3.9.3 Aux Power Mode (S5)

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

3.9.4 Main Power Mode (S0)

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

3.10 Add-in Card Input Capacitance

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 31: Power Supply Rail Requirements

Power Rail	35W (Aux Only Mode)	79.2W (Main Power Mode)
3.3Vaux		
Voltage Tolerance	±9% (max)	±9% (max)
Supply Current	1.1A (max)	1.1A (max)
Capacitive Load	150µF (max)	150µF (max)
12V		
Voltage Tolerance	±8% (max)	±8% (max)
Supply Current	2.92A (max)	6.6A (max)
Capacitive Load	2000µF (max)	2000µF (max)

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 0b1111. The PRSNTB[3:0]# can 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 28: Power Sequencing

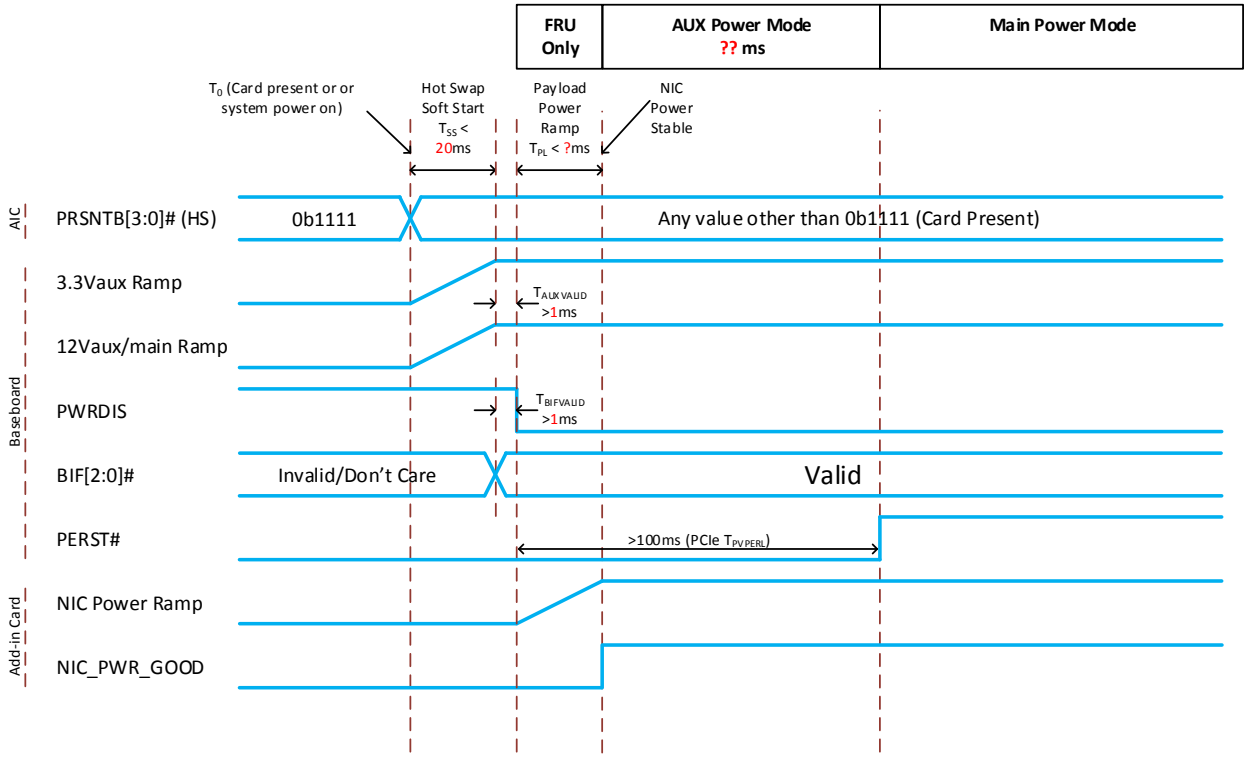

Table 32: Power Sequencing Parameters

Parameter	Value	Units	Description
T_{ss}	20	ms	Max time between system 3.3Vaux and 12Vaux/main ramp to power stable.
$T_{AUXVALID}$	>1	ms	Min time between 3.3Vaux valid to PWRDIS assertion.
$T_{BIFVALID}$	>1	ms	Min time between BIF[2:0]# valid to PWRDIS assertion. The BIF[2:0]# value sets the add-in card bifurcation mode (if applicable)
T_{PL}	<?	ms	Max time between the NIC payload power ramp to 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

4.1 SMBus Interface

The SMBus provides manageability of the add-in card.

4.2 NC-SI Sideband Interface

4.2.1 NC-SI addressing and Arb#

4.3 MAC Address Requirement

4.4 FRU EEPROM

4.4.1 Minimum EEPROM Size

4.4.2 EEPROM Map Definition

Editor's note [TN, 20171208]: the EEPROM map definition should include the card power class (e.g. be able to identify the max power required for 12V to the baseboard). This allows

the baseboard to intelligently enable/disable the card based on the available baseboard power budget.

4.4.3 EEPROM Address

4.5 FW Requirement (TBD)

4.6 Thermal Reporting Interface

5 Data Network Requirement

5.1 Network Booting (collect view from OEMs and hyperscale)

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

6 Routing Guidelines and Signal Integrity Considerations

6.1 NC-SI Over RBT

Min Length: 2"

Max length: 4"

Impedance: 50 Ohm single ended

7 Thermal and Environmental

7.1 Environmental Requirements

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

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

7.1.1 Thermal Simulation Boundary Example

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

7.2 Shock & Vibration

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

7.3 Regulation

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

8 Revision History

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