



OPEN
Compute Project

OCP NIC 3.0 Design Specification

Version 0.01

Author: OCP Server Workgroup, OCP NIC subgroup

Table of Contents

1	Overview	9
1.1	License	9
1.2	Background	10
1.3	Acknowledgements	11
1.4	Overview	11
1.4.1	Mechanical Form factor overview	11
1.4.2	Electrical overview	14
1.5	References	17
2	Card Form Factor	18
2.1	Overview	18
2.2	Form Factor Options	18
2.3	I/O bracket	20
2.4	Line Side I/O Implementations	20
2.5	LED Implementations	21
2.5.1	Baseboard LEDs Configuration Over the Scan Chain	21
2.5.2	Add-in Card LED Configuration	23
2.5.3	Add-in Card LED Ordering	24
2.6	Mechanical Keepout Zones	25
2.6.1	Baseboard Keep Out Zones	25
2.6.2	Add-in Card Keep Out Zones	25
2.7	Labeling Requirements	25
2.8	Insulation Requirements	25
2.9	NIC Implementation Examples	26
2.10	Non-NIC Use Cases	26
2.10.1	PCIe Retimer card	26
2.10.2	Accelerator card	26
2.10.3	Storage HBA / RAID card	26
3	Card Edge and Baseboard Connector Interface	26
3.1	Gold Finger Requirements	26

3.1.1	Gold Finger Mating Sequence.....	27
3.2	Baseboard Connector Requirements	29
3.3	Pin definition.....	30
3.4	Signal Descriptions – Common.....	36
3.4.1	PCIe Interface Pins.....	36
3.4.2	PCIe Present and Bifurcation Control Pins.....	40
3.4.3	SMBus Interface Pins	42
3.4.4	Power Supply Pins	43
3.4.5	Miscellaneous Pins.....	45
3.5	Signal Descriptions – OCP Bay (Primary Connector).....	45
3.5.1	PCIe Interface Pins – OCP Bay (Primary Connector)	46
3.5.2	NC-SI Over RBT Interface Pins – OCP Bay (Primary Connector).....	48
3.5.3	Scan Chain Pins – OCP Bay (Primary Connector).....	55
3.5.4	Primary Connector Miscellaneous Pins – OCP Bay (Primary Connector).....	62
3.6	PCIe Bifurcation Mechanism	63
3.6.1	PCIe Add-in Card to Baseboard Bifurcation Configuration (PRSNTA#, PRSNTB[3:0]#).....	64
3.6.2	PCIe Baseboard to Add-in Card Bifurcation Configuration (BIF[2:0]#)	64
3.6.3	PCIe Bifurcation Decoder.....	65
3.6.4	Bifurcation Detection Flow.....	67
3.6.5	PCIe Bifurcation Examples	68
3.7	PCIe Clocking Topology.....	74
3.8	PCIe Bifurcation Results and REFCLK Mapping.....	76
3.9	Power Capacity and Power Delivery.....	85
3.9.1	NIC Power Off.....	85
3.9.2	ID Mode.....	85
3.9.3	Aux Power Mode (S5)	86
3.9.4	Main Power Mode (S0).....	86
3.10	Power Supply Rail Requirements and Slot Power Envelopes	86
3.11	Hot Swap Considerations for 12V and 3.3V Rails.....	87
3.12	Power Sequence Timing Requirements	87
4	Management.....	89
4.1	Sideband Management Interface and Transport.....	89

4.2	NC-SI Traffic	89
4.3	Management Controller (MC) MAC Address Provisioning	90
4.4	Temperature Reporting	90
4.5	Power Consumption Reporting	91
4.6	Link Status/Speed Reporting	91
4.7	Pluggable Module Status Reporting	92
4.8	Out-Of-Band Firmware Update	93
4.9	NC-SI Over RBT Sideband Interface	93
4.9.1	NC-SI Over RBT Addressing	93
4.9.2	Arbitration Ring Connections	94
4.10	SMBus 2.0 Interface	94
4.10.1	SMBus Address Map	94
4.11	FRU EEPROM	95
4.11.1	FRU EEPROM Address, Size and Availability	95
4.11.2	FRU EEPROM Content Requirements	96
4.12	FW Requirements	98
4.12.1	Firmware Update	98
4.12.2	Secure Firmware	98
4.12.3	Firmware Queries	98
4.12.4	Multi-Host Firmware Queries	98
4.13	Thermal Reporting Interface	98
4.13.1	Emulated Thermal Reporting	99
4.13.2	Remote on-die sensing	102
4.13.3	PLDM Method	104
4.13.4	Thermal reporting accuracy	104
5	Data Network Requirements	105
5.1	Network Boot	105
6	Routing Guidelines and Signal Integrity Considerations	106
6.1	NC-SI Over RBT	106
6.2	PCIe	106
7	Thermal and Environmental	107
7.1	Environmental Requirements	107
7.1.1	Thermal Reporting interface	107

7.1.2	Thermal Simulation Boundary Example	107
7.2	Shock & Vibration	107
7.3	Regulatory	107
8	Revision History	108

List of Figures

Figure 1: Representative Small OCP NIC 3.0 Card with Quad SFP Ports	10
Figure 2: Representative Large OCP NIC 3.0 Card with Dual QSFP Ports and on-board DRAM.....	11
Figure 3: Small and Large Card Form-Factors (not to scale).....	12
Figure 4: Example Small Card Form Factor	18
Figure 5: Example Large Card Form Factor	19
Figure 6: Primary Connector (4C + OCP Bay) with 4C and 2C Add-in Cards.....	20
Figure 7: LED Ordering – Example Small Card Link/Activity and Speed LED Placement	24
Figure 8: Small Card Bottom Side Insulator and Mechanical Envelope	25
Figure 9: Large Card Bottom Side Insulator and Mechanical Envelope.....	26
Figure 10: Small Size Primary Connector Gold Finger Mating Card Dimensions – x16 – Top Side	27
Figure 11: Small Size Primary Connector Gold Finger Mating Card Dimensions – x16 – Bottom Side	27
Figure 12: Large Size Card Gold Finger Mating Card Dimensions – x32 – Top Side	27
Figure 13: Large Size Card Gold Finger Mating Card Dimensions – x32 – Bottom Side.....	27
Figure 14: 168-pin Base Board Primary Connector – Right Angle	30
Figure 15: 140-pin Base Board Secondary Connector – Right Angle	30
Figure 16: 168-pin Base Board Primary Connector – Straddle Mount.....	30
Figure 17: 140-pin Base Board Secondary Connector – Straddle Mount.....	30
Figure 18: Primary and Secondary Connector Locations for Large Card Support.....	30
Figure 19: PCIe Present and Bifurcation Control Pins.....	42
Figure 20: Example Power Supply Topology	44
Figure 21: NC-SI Over RBT Connection Example – Single Primary Connector.....	53
Figure 22: NC-SI Over RBT Connection Example – Dual Primary Connector.....	54
Figure 23: Scan Bus Connection Example	61
Figure 24: PCIe Bifurcation Pin Connections Support.....	64
Figure 25: Single Host (1 x16) and 1 x16 Add-in Card (Single Controller)	68
Figure 26: Single Host (2 x8) and 2 x8 Add-in Card (Dual Controllers).....	69
Figure 27: Four Hosts (4 x4) and 4 x4 Add-in Card (Single Controller).....	70
Figure 28: Four Hosts (4 x4) and 4 x4 Add-in Card (Four Controllers).....	71
Figure 29: Single Host with no Bifurcation (1 x16) and 2 x8 Add-in Card (Two Controllers).....	73
Figure 30: PCIe Interface Connections for 1 x16 and 2 x8 Add-in Cards.....	74
Figure 31: PCIe Interface Connections for a 4 x4 Add-in Card.....	75
Figure 32: Baseboard Power States.....	85
Figure 33: Power Sequencing.....	87
Figure 34: Block Diagram for Emulated Thermal Reporting.....	100
Figure 35: Block Diagram for Remote on-die Sensing.....	102

List of Tables

Table 1: OCP 3.0 Form Factor Dimensions.....	12
Table 2: Baseboard to OCP NIC Form factor Compatibility Chart.....	13
Table 3: OCP NIC 3.0 Card Definitions	20
Table 4: OCP 3.0 Line Side I/O Implementations.....	21
Table 5: Baseboard LED Configurations with Two Physical LEDs per Port.....	22
Table 6: Add-in Card LED Configuration with Two Physical LEDs per Port.....	23
Table 7: Contact Mating Positions for the Primary and Secondary Connectors	27
Table 8: Primary Connector Pin Definition (x16) (4C + OCP Bay).....	31
Table 9: Secondary Connector Pin Definition (x16) (4C).....	34
Table 10: Pin Descriptions – PCIe 1.....	36
Table 11: Pin Descriptions – PCIe Present and Bifurcation Control Pins.....	40
Table 12: Pin Descriptions – SMBus.....	42
Table 13: Pin Descriptions – Power.....	43
Table 14: Pin Descriptions – Miscellaneous 1	45
Table 15: Pin Descriptions – PCIe 2.....	46
Table 16: Pin Descriptions – NC-SI Over RBT	48
Table 17: Pin Descriptions – Scan Chain.....	55
Table 18: Pin Descriptions – Scan Chain DATA_OUT Bit Definition.....	57
Table 19: Pin Descriptions – Scan Bus DATA_IN Bit Definition	58
Table 20: Pin Descriptions – Miscellaneous 2	62
Table 21: PCIe Bifurcation Decoder for x16 and x8 Card Widths	66
Table 22: PCIe Clock Associations.....	74
Table 23: Bifurcation for Single Host, Single Socket and Single Upstream Link (BIF[2:0]#=0b000).....	77
Table 24: Bifurcation for Single Host, Single Socket and Single/Dual Upstream Links (BIF[2:0]#=0b000)	78
Table 25: Bifurcation for Single Host, Single Socket and Single/Dual/Quad Upstream Links (BIF[2:0]#=0b000)	79
Table 26: Bifurcation for Single Host, Dual Sockets and Dual Upstream Links (BIF[2:0]#=0b001)	80
Table 27: Bifurcation for Single Host, Four Sockets and Dual Upstream Links (BIF[2:0]#=0b010).....	81
Table 28: Bifurcation for Dual Host, Dual Sockets and Dual Upstream Links (BIF[2:0]#=0b101).....	82
Table 29: Bifurcation for Quad Host, Quad Sockets and Quad Upstream Links (BIF[2:0]#=0b110).....	83
Table 30: Bifurcation for Quad Host, Quad Sockets and Quad Upstream Links (BIF[2:0]#=0b110).....	84
Table 31: Power States.....	85
Table 32: Baseboard Power Supply Rail Requirements – Slot Power Envelopes.....	86
Table 33: Power Sequencing Parameters	88
Table 34: Sideband Management Interface and Transport Requirements	89
Table 35: NC-SI Traffic Requirements	89
Table 36: MC MAC Address Provisioning Requirements.....	90
Table 37: Temperature Reporting Requirements.....	90
Table 38: Power Consumption Reporting Requirements	91
Table 39: Link Status/Speed Reporting Requirements.....	92



Table 40: Pluggable Module Status Reporting Requirements.....	92
Table 41: Out-Of-Band Firmware Update Requirements.....	93
Table 42: Power Sequencing Parameters	94
Table 43: FRU EEPROM Record – OEM Record 0xC0, Offset 0x00.....	96
Table 44: FRU EEPROM Record – OEM Record 0xC0, Offset 0x01.....	96
Table 45: FRU EEPROM Record – OEM Record 0xC0, Offset 0x02.....	96
Table 46: FRU EEPROM Record – OEM Record 0xC0, Offset 0x03.....	97
Table 47: FRU EEPROM Record – OEM Record 0xC0, Offset 0x04.....	97
Table 48: FRU EEPROM Record – OEM Record 0xC0, Offset 0x05.....	97
Table 49: Implementation Requirement for TMP421 Registers	101

1 Overview

1.1 License

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

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

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

1.2 Background

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

- NICs with a higher TDP
- Support up to 80W of power delivery to a single connector (Small) card; and 150W to a dual connector (Large) card
- Support up to PCIe Gen5 on the system and add-in card
- Support for up to 32 lanes of PCIe per add-in card
- Support for single host, multi-root complex and multi-host environments
- Support a greater board area for more complex add-in card designs
- Support for Smart NIC implementations with on-board DRAM and accelerators
- Simplification of FRU installation and removal while reducing overall down time

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

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





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



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

This specification is created under OCP Server workgroup – OCP NIC subgroup. An electronic copy of this specification can be found on the Open Compute Project website:

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

1.3 Acknowledgements

Placeholder [Editor's note TN 20171219: I suggest adding a table of contributing companies and individuals to this section].

1.4 Overview

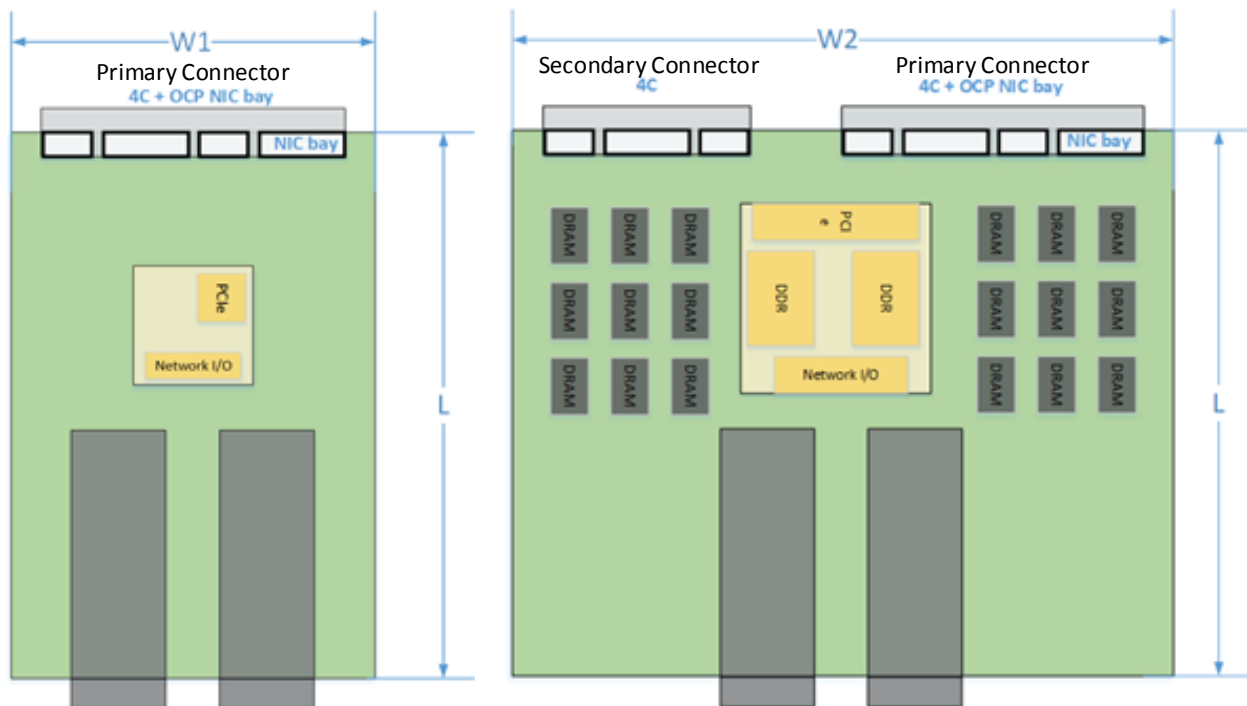
1.4.1 Mechanical Form factor overview

The OCP NIC 3.0 specification defines a third generation mechanical form factor that allows for interoperability between compliant baseboards and add-in cards.

OCP NIC 3.0 cards have two form factors – Small and Large. These cards are shown in Figure 3 below. The components shown in the figures are for illustrative purposes. The Small form factor card has one connector (Primary connector) on baseboard. The Large form factor card has one or two connectors (Primary Connector only and both the Primary and Secondary Connectors) on the baseboard.

Both the Primary and Secondary connectors are defined in and compliant to SFF-TA-1002. [Editor's note: plan to submit change back to SFF-TA-1002]. On the add-in card side, the card edge is implemented with gold fingers. The small card gold finger area only occupies the Primary Connector area for up to 16 PCIe lanes. The large card gold finger area may occupy both the Primary and Secondary connectors for up to 32 PCIe lanes, or optionally just the Primary connector for up to 16 PCIe lane implementations. The gold finger design follows SFF-TA-1002 as well.

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



The two form factor dimensions are shown in Table 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 NIC with a similar profile as an OCP NIC 2.0 add-in card; up to 16 PCIe lanes.
Large	W2 = 139 mm	L = 115 mm	4C + OCP sideband 168 pins	4C 140 pins	Larger PCB width to support additional NICs; up to 32 PCIe lanes.

The OCP NIC 3.0 design allows downward compatibility between the two card sizes. Table 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 16 PCIe lanes	Not Supported
Large	Up to 16 PCIe lanes	Up to 32 PCIe lanes

There are two baseboard connector options available for system designers: straddle mount and right angle (RA). The straddle mount connector option allows the OCP NIC and baseboard to exist in a co-planer position. To achieve this, a cutout exists on the baseboard and is defined in this specification. Alternatively, the right angle option allows the OCP NIC to be installed on top of the baseboard. A baseboard cutout is not required for the right angle connector. The right angle option allows the baseboard to use this area for additional routing or backside component placement. The straddle mount and right angle connectors are shown in Section 3.2.

For both the baseboard and OCP card, this specification defines the component and routing keep out areas. Refer to Section 2.6 for details.



Both the straddle mount and right angle implementations shall accept the same OCP add-in card and shall be supported in the baseboard chassis regardless of the baseboard connector selection (right angle or straddle mount) so long as the baseboard slot side and add-in card sizes are a supported combination as shown in Table 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 16 lanes of PCIe. The Secondary connector, when used in conjunction with the Primary connector, allows Large form-factor implementations and may support up to 32 lanes of PCIe.

1.4.2.1 Primary Connector

The Primary connector provides all OCP specific management functions as well as up to 16 lanes of PCIe between the OCP NIC and the system motherboard.

Management Function Overview (OCP Bay):

- DSP0222 1.1 compliant Network Controller Sideband Interface (NC-SI) RMIII Based Transport (RBT) Physical Interface
- Power management and status reporting
 - Power disable
 - State change control
- SMBus 2.0
- Control / status serial bus

- NIC-to-Host status
 - Port LED Link/Activity
 - Environmental Indicators
- Host-to-NIC configuration Information
- Multi-host PCIe support signals (2x PCIe resets, 2x reference clocks).
 - The OCP bay provides PERST2#, PERST3#, REFCLK2 and REFCLK3. This enables support for up to four hosts when used in conjunction with PERST0#, PERST1#, REFCLK0 and REFCLK1 in the Primary 4C region.
- PCIe Wake signal

See Section 3.5 for a complete list of pin and function descriptions for the OCP Bay portion of the primary connector.

PCIe Interface Overview (4C Connector):

- 16x differential transmit/receive pairs
 - Up to PCIe Gen 5 support
- 2x 100 MHz differential reference clocks
- Control signals
 - 2x PCIe Resets
 - Link Bifurcation Control
 - Card power disable/enable
- Power
 - 12V /12V AUX
 - 3.3V AUX

See Section 3.4 for a complete list of pin and function descriptions for the 4C connector.

1.4.2.2 Secondary Connector

The secondary connector provides an additional 16 lanes of PCIe and their respective control signals.

PCIe Interface Overview (4C Connector):



- 16x differential transmit/receive pairs
 - Up to PCIe Gen 5 support
- 2x 100 MHz differential reference clocks
- Control signals
 - 2x PCIe Resets
 - Link Bifurcation Control
 - Card power disable/enable
- Power
 - 12V /12V AUX
- 3.3V AUX

See Section 3.4 for a complete list of pin and function descriptions for the 4C connector.

1.5 References

- DMTF Standard. *DSP0222, Network Controller Sideband Interface (NC-SI) Specification*. Distributed Management Task Force, Rev 1.1.0, September 23rd, 2015.
- DMTF Standard. *DSP0222, Network Controller Sideband Interface (NC-SI) Specification*. Distributed Management Task Force, Rev 1.2.0, Work-in-progress.
- EDSFF. *Enterprise and Datacenter SSD Form Factor Connector Specification*. Enterprise and Datacenter SSD Form Factor Working Group, Rev 0.9 (draft), August 2nd 2017.
- *IPMI Platform Management FRU Information Storage Definition*, v1.2, February 28th, 2013.
- NXP Semiconductors. *I²C-bus specification and user manual*. NXP Semiconductors, Rev 6, April 4th, 2014.
- Open Compute Project. *OCP NIC Subgroup*. Online.
<http://www.opencompute.org/wiki/Server/Mezz>
- PCIe Base Specification. PCI Express Base Specification, Revision 4.0 (draft).
- PCIe CEM Specification. PCI Express Card Electromechanical Specification, Revision 4.0 (draft).
- SMBus Management Interface Forum. *System Management Bus (SMBus) Specification*. System Management Interface Forum, Inc, Version 2.0, August 3rd, 2000.
- SNIA. *SFF-TA-1002, Specification for Protocol Agnostic Multi-Lane High Speed Connector*. SNIA SFF TWG Technology Affiliate, Rev 1.0, December 12th, 2017.

2 Card Form Factor

2.1 Overview


2.2 Form Factor Options

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

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


The small card uses the Primary 4C connector to provide up to a x16 PCIe interface to the host. The additional 28-pin OCP bay carries sideband management interfaces as well as OCP NIC 3.0 specific control signals for multi-host PCIe support. The small size card provides sufficient faceplate area to accommodate up to 2x QSFP modules, 4x SFP modules, or 4x RJ-45 for BASE-T operation. The small card form factor supports up to 80W of delivered power to the card edge.

Figure 4: Example Small Card Form Factor



The large card uses the Primary 4C + OCP bay connector to provide the same functionality as the small card along with an additional Secondary 4C connector to provide up to a x32 PCIe interface. The large card may utilize both the Primary and Secondary connectors, or just the Primary connector for lower PCIe lane count applications. Table 3 summarizes the large card permutations. The large size card supports higher power envelopes and provides additional board area for more complex designs. The large card form factor supports up to 150W of delivered power to the card edge across the two connectors.

Figure 5: Example Large Card Form Factor



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

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




Table 3 summarizes the supported card form factors. Small form factor cards support the Primary Connector and up to 16 PCIe lanes. Large form factor cards support implementations with both the Primary and Secondary Connectors and up to 32 PCIe lanes, or a Primary Connector only implementation with up to 16 PCIe lanes.

Table 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 (x8)			2C	OCP Bay
Large (x16)			4C	OCP Bay
Large (x24)		2C	4C	OCP Bay
Large (x32)	4C		4C	OCP Bay

2.3 I/O bracket

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 implementations have been optimized to 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	2x QSFP28
Large	4x SFP28
Large	4x RJ-45

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

2.5 LED Implementations

LEDs may be implemented on the card Scan Chain (as defined in Section 3.5.3) for remote link/activity indication on the baseboard or optionally be implemented on the OCP NIC 3.0 I/O bracket if there is sufficient space for local indication. These two cases are described below. In both cases, the actual link rate may be directly queried through the management interface.

2.5.1 Baseboard LEDs Configuration Over the Scan Chain

A small form-factor OCP NIC 3.0 with a fully populated I/O bracket (2x QSFP28, 4x SFP28 or 4x RJ-45) has insufficient space for discrete on-board (faceplate) LED indicators. In this case, the line side link and activity LED indicators are implemented on the baseboard system via the Scan Chain. The Scan Chain bit stream is defined in Section 3.5.3. This LED implementation is required for all add-in cards. The baseboard LED implementation uses two discrete LEDs (Link/Activity and Speed indication). Table 5 describes the baseboard LED configuration for baseboard implementations.

Table 5: Baseboard LED Configurations with Two Physical LEDs per Port

LED Pin	LED Color	Description
Link / Activity	Green	Active low. Multifunction LED. This LED shall be used to indicate link and link activity. When the link is up and no link activity is present, then this LED shall be lit and solid. This indicates that the link is established, there are no local or remote faults, and the link is ready for data packet transmission/reception. When the link is up and there is link activity, then this LED should blink at the interval of 50-500ms during link activity. The Link/Activity LED shall be located on the left hand side or located on top for each port when the baseboard is viewed in the horizontal plane.
	Off	
Speed	Green	Active low. Multifunction LED.
	Off	The LED is Green when the port is linked at its maximum speed. The LED is off when the device is linked at a speed lower than the highest capable speed, or no link is present. The bicolor speed LED shall be located on the right hand side or located on the bottom for each port when the baseboard is viewed in the horizontal plane.

At the time of this writing, the Scan Chain definition allows for up to one link/activity and one speed LED per port. A total of up to 8 ports are supported in the Scan Chain. The bit stream defines the LEDs to be active low (on). The Scan Chain LED implementation allows

the NIC LED indicators to be remotely located on the OCP 3.0 compliant chassis (e.g. front LED indicators with rear I/O cards).

2.5.2 Add-in Card LED Configuration

For low I/O count small form-factor cards without built in light pipes (such as 1x QSFP28, 2x SFP28, or 2x RJ-45), or a large form-factor OCP 3.0 NIC, where additional I/O bracket area is available, the card may optionally implement on-board link/activity indications in place of the Scan Chain LED stream. The recommended local (on-card) LED implementation uses two physical LEDs (a discrete Link/Activity LED and a bi-colored Speed A/Speed B LED). Table 6 describes the add-in card LED implementations.

Table 6: Add-in Card LED Configuration with Two Physical LEDs per Port

LED Pin	LED Color	Description
Link / Activity	Green	Active low. Multifunction LED.
		<p>This LED shall be used to indicate link and link activity.</p> <p>When the link is up and no link activity is present, then this LED shall be lit and solid. This indicates that the link is established, there are no local or remote faults, and the link is ready for data packet transmission/reception.</p> <p>When the link is up and there is link activity, then this LED should blink at the interval of 50-500ms during link activity.</p> <p>The Link/Activity LED shall be located on the left hand side or located on the top for each port when the add-in card is viewed in the horizontal plane.</p>
Speed	Green	Active low. Bicolor multifunction LED.
	Amber	
	Off	
		The LED is Green when the port is linked at its maximum speed.

		<p>The LED is Amber when the port is linked at its second highest speed.</p> <p>The LED is off when the device is linked at a speed lower than the second highest capable speed, or no link is present.</p> <p>The Amber Speed LED indicator may be used for port identification through vendor specific link diagnostic software.</p> <p>The bicolor speed LED shall be located on the right hand side or located on the bottom for each port when the add-in card is viewed in the horizontal plane.</p>
--	--	--

2.5.3 Add-in Card LED Ordering

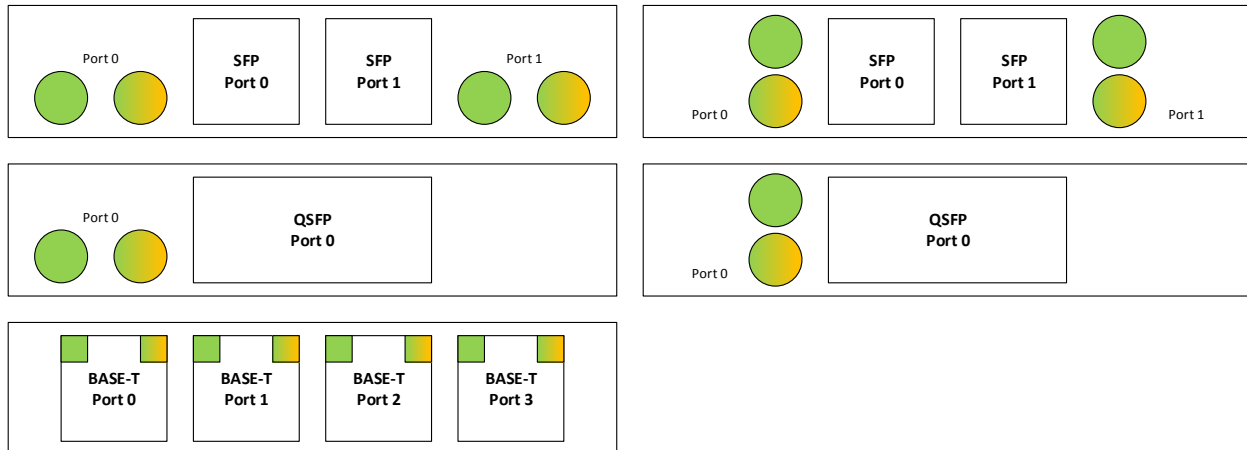
For all add-in card use cases, each port shall implement the green Link/Activity LED and a bicolor green/amber speed A/B LED. For all baseboards, each port shall implement the green Link/Activity LED and a green speed A LED.

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

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

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

Figure 7: LED Ordering – Example Small Card Link/Activity and Speed LED Placement



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

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

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

2.8 Insulation Requirements

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

Figure 8: Small Card Bottom Side Insulator and Mechanical Envelope

TBD <need 2D drawings>



Figure 9: Large Card Bottom Side Insulator and Mechanical Envelope

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 Requirements

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

Small Size cards fit in the Primary Connector. Primary Connector compliant cards are 76mm x 115mm and may implement the full 168-pins. The Primary Connector cards may optionally implement a subset of gold finger pins if there is a reduced PCIe width requirement (such as 1 x8 and below). In this case, the card edge gold finger may implement a 2C design. The overall board thickness is 1.60mm. The gold finger dimensions for the Primary Connector compliant cards are shown below.

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

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

TBD

Figure 11: Small Size Primary Connector Gold Finger Mating Card Dimensions – x16 – Bottom Side

TBD

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

TBD

Figure 13: 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 7 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 7: Contact Mating Positions for the Primary and Secondary Connectors

Side B	Side A
--------	--------

		AIC Plug (Free)		Receptacle (Fixed)			AIC Plug (Free)		Receptacle (Fixed)
		2 nd Mate	1 st Mate				2 nd Mate	1 st Mate	
OCP B1					OCP A1				
OCP B2					OCP A2				
OCP B3					OCP A3				
OCP B4					OCP A4				
OCP B5					OCP A5				
OCP B6					OCP A6				
OCP B7					OCP A7				
OCP B8					OCP A8				
OCP B9					OCP A9				
OCP B10					OCP A10				
OCP B11					OCP A11				
OCP B12					OCP A12				
OCP B13					OCP A13				
OCP B14					OCP A14				
Mechanical Key									
B1					A1				
B2					A2				
B3					A3				
B4					A4				
B5					A5				
B6					A6				
B7					A7				
B8					A8				
B9					A9				
B10					A10				
B11					A11				
B12					A12				
B13					A13				
B14					A14				
B15					A15				
B16					A16				
B17					A17				
B18					A18				
B19					A19				
B20					A20				
B21					A21				
B22					A22				
B23					A23				
B24					A24				
B25					A25				
B26					A26				
B27					A27				
B28					A28				
Mechanical Key									
B29					A29				
B30					A30				
B31					A31				
B32					A32				
B33					A33				



3.2 Baseboard Connector Requirements

The OCP NIC 3.0 connectors are compliant to the “4C connector” as defined in the SFF-TA-1002 specification for a right angle or straddle mount form-factor. The 4C connector is 140-pins in width and includes support for up to 32 differential pairs to support a x16 PCIe connection. The connector also provides 6 pins of 12V for payload power. This implementation is common between both the Primary and Secondary Connectors. In addition, the Primary Connector has a 28-pin OCP Bay to the right of pin 1. These pins are



used for management and support for up to a 4 x2 and 4 x4 multi-host configuration on the Primary Connector. The Primary and Secondary Connector drawings are shown in Figure 14, Figure 15, Figure 16 and Figure 17 below.

Figure 14: 168-pin Base Board Primary Connector – Right Angle

TBD

Figure 15: 140-pin Base Board Secondary Connector – Right Angle

TBD

Figure 16: 168-pin Base Board Primary Connector – Straddle Mount

TBD

Figure 17: 140-pin Base Board Secondary Connector – Straddle Mount

TBD

In order to support the large form factor, systems must locate the Primary and Secondary Connectors per the mechanical drawing shown in Figure 18.

Figure 18: 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 8 and Table 9. All signal directions are shown from the perspective of the baseboard.

A baseboard system may provide a combination of Primary Connectors only, or Primary and Secondary Connectors to support multiple sizes of add-in cards. Both connectors share common functionality with power, SMBus 2.0, x16 PCIe and bifurcation control. The Primary Connector has an additional OCP Bay (pins OCP_A[1:14], OCP_B[1:14]) with additional REFCLKs for supporting up to four PCIe hosts, NC-SI connectivity and a Scan Chain for information exchange between the host and card. The NIC is required to implement the

Scan Chain, while the baseboard may choose to optionally implement it. Depending on the baseboard form-factor, multiple OCP NIC 3.0 compliant cards may be designed into the system.

The pins common to the Primary and Secondary Connectors are shown in Section 3.4. The OCP Bay pins on the Primary Connector only are shown in Section 3.5.

Cards or systems that do not require the use of a PCIe x16 connection may optionally implement a subset electrical connections as applicable to the design. For example, a x8 (or smaller) card using the first 8 PCIe lanes that is compliant with the Primary Connector pinout. Refer to Sections 3.1 and 3.2 for mechanical details. For these cases, the Primary Connector matches the 2C dimensions as defined in SFF-TA-1002.

In all cases, the physical baseboard connectors shall support x16 PCIe widths and must be implemented with the Primary (4C + OCP Bay) and Secondary (4C) connectors.

Table 8: Primary Connector Pin Definition (x16) (4C + OCP Bay)

Side B		Side A		Primary Connector (x16, 168-pin add-in card with OCP Bay)	Primary Connector (x8, 112-pin add-in card with OCP bay)
OCP_B1	NIC_PWR_GOOD	PERST2#	OCP_A1		
OCP_B2	PWRBRK#	PERST3#	OCP_A2		
OCP_B3	LD#	WAKE_N#	OCP_A3		
OCP_B4	DATA_IN	RBT_ARB_IN	OCP_A4		
OCP_B5	DATA_OUT	RBT_ARB_OUT	OCP_A5		
OCP_B6	CLK	GND	OCP_A6		
OCP_B7	SLOT_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 9: Secondary Connector Pin Definition (x16) (4C)

Side B		Side A		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	PRSNB2#	A12		
B13	GND	GND	A13		
B14	REFCLKn0	REFCLKn1	A14		
B15	REFCLKp0	REFCLKp1	A15		
B16	GND	GND	A16		
B17	PETn0	PERn0	A17		
B18	PETp0	PERp0	A18		
B19	GND	GND	A19		
B20	PETn1	PERn1	A20		
B21	PETp1	PERp1	A21		
B22	GND	GND	A22		
B23	PETn2	PERn2	A23		
B24	PETp2	PERp2	A24		
B25	GND	GND	A25		
B26	PETn3	PERn3	A26		
B27	PETp3	PERp3	A27		
B28	GND	GND	A28		
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 31.

Table 10: Pin Descriptions – PCIe 1

Signal Name	Pin #	Baseboard Direction	Signal Description
REFCLKn0 REFCLKp0	B14 B15	Output	PCIe compliant differential reference clock #0, and #1. 100MHz reference clocks are used for the add-in card PCIe core logic. For baseboards, the REFCLK0 and REFCLK1 signals shall be available at the connector. For add-in cards, the required REFCLKs shall be connected per the endpoint datasheet. 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	<p>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.</p> <p>The PCIe transmit pins shall be AC coupled on the baseboard with capacitors. The capacitors shall be placed next to the baseboard transmitters. The AC coupling capacitor value shall be between 176nF (min) and 265nF (max).</p> <p>For baseboards, the PET[0:15] signals are required at the connector.</p> <p>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.</p> <p>Refer to Section 6.1 in the PCIe CEM Specification, Rev 4.0 for details.</p>
PETn1 PETp1	B20 B21	Output	
PETn2 PETp2	B23 B24	Output	
PETn3 PETp3	B26 B27	Output	
PETn4 PETp4	B30 B31	Output	
PETn5 PETp5	B33 B34	Output	
PETn6 PETp6	B36 B37	Output	
PETn7 PETp7	B39 B40	Output	
PETn8 PETp8	B44 B45	Output	
PETn9 PETp9	B47 B48	Output	
PETn10 PETp10	B50 B51	Output	
PETn11 PETp11	B53 B54	Output	
PETn12 PETp12	B56 B57	Output	
PETn13 PETp13	B59 B60	Output	

PETn14	B62	Output	
PETp14	B63		
PETn15	B65	Output	
PETp15	B66		
PERn0	A17	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.
PERp0	A18		
PERn1	A20	Input	The PCIe receive pins shall be AC coupled on the add-in card with capacitors. The capacitors shall be placed next to the add-in card transmitters. The AC coupling capacitor value shall be between 176nF (min) and 265nF (max).
PERp1	A21		
PERn2	A23	Input	For baseboards, the PER[0:15] signals are required at the connector.
PERp2	A24		
PERn3	A26	Input	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.
PERp3	A27		
PERn4	A30	Input	Refer to Section 6.1 in the PCIe CEM Specification, Rev 4.0 for details.
PERp4	A31		
PERn5	A33	Input	
PERp5	A34		
PERn6	A36	Input	
PERp6	A37		
PERn7	A39	Input	
PERp7	A40		
PERn8	A44	Input	
PERp8	A45		
PERn9	A47	Input	
PERp9	A48		
PERn10	A50	Input	
PERp10	A51		
PERn11	A53	Input	
PERp11	A54		
PERn12	A56	Input	
PERp12	A57		
PERn13	A59	Input	
PERp13	A60		

PERn14 PERp14	A62 A63	Input	
PERn15 PERp15	A65 A66	Input	
PERST0# PERST1#	B10 A11	Output	<p>PCIe Reset #0, #1. Active low.</p> <p>When PERSTn# is deasserted, the signal shall indicate the applied power is within tolerance and stable for the add-in card.</p> <p>PERST# shall be deasserted at least 100ms after the power rails are within the operating limits per the PCIe CEM Specification. The PCIe REFCLKs shall also become stable within this period of time.</p> <p>PERST shall be pulled high to 3.3Vaux on the baseboard.</p> <p>For OCP NIC 3.0, PERST deassertion shall also indicate the full card power envelope is available to the add-in card.</p> <p>For baseboards, the PERST[0:1]# signals are required at the connector.</p> <p>For add-in cards, the required PERST[0:1]# signals shall be connected to the endpoint silicon.</p> <p>Note: For cards that only support 1 x16, PERST0# is used. For cards that support 2 x8, PERST0# is used for the first eight PCIe lanes,</p>



			<p>and PERST1# is used for the second eight PCIe lanes.</p> <p>Refer to Section 2.2 in the PCIe CEM Specification, Rev 4.0 for details.</p>
--	--	--	---

3.4.2 PCIe Present and Bifurcation Control Pins

This section provides the pin assignments for the PCIe present and bifurcation control signals. The AC/DC specifications are defined in Section XXX. An example connection diagram is shown in Figure 19.

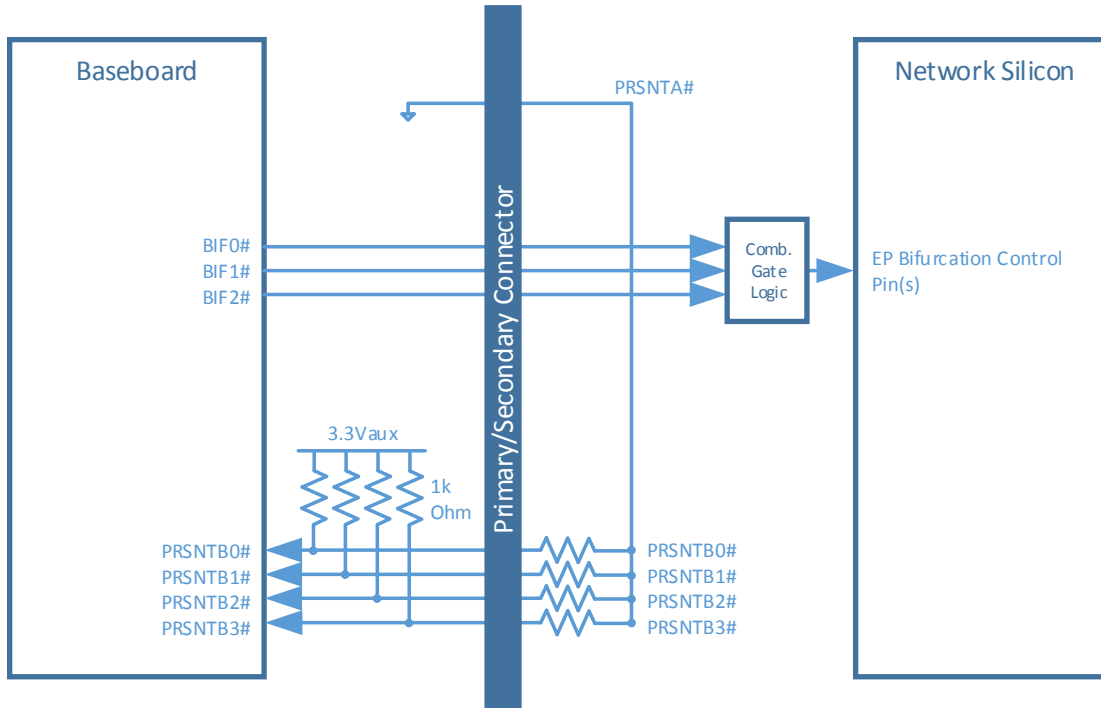
The PRSNTA#/PRSNTB[0:3]# state shall be used to determine if a card has been physically plugged in. The BIF[0:2]# pins shall be latched before PWRDIS deassertion to ensure the correct values are detected by the system. Changing the pin states after this timing window is not allowed. Refer to the AC timing diagram in Section XXX for details.

Table 11: Pin Descriptions – PCIe Present and Bifurcation Control Pins

Signal Name	Pin #	Baseboard Direction	Signal Description
PRSNTA#	A12	Output	<p>Present A is used for add-in card presence and PCIe capabilities detection.</p> <p>For baseboards, this pin shall be directly connected to GND.</p> <p>For add-in cards, this pin shall be directly connected to the PRSNTB[3:0]# pins.</p>
PRSNTB0#	B42	Input	Present B [0:3]# are used for add-in card presence and PCIe capabilities detection.
PRSNTB1#	A42		
PRSNTB2#	A10		
PRSNTB3#	B70		

			<p>For baseboards, these pins shall be connected to the I/O hub and pulled up to +3.3Vaux using 1kOhm resistors.</p> <p>For add-in cards, these pins shall be strapped to PRSNTA# per the encoding definitions described in Section 3.6.</p> <p>Note: PRSNTB3# is located at the bottom of the 4C connector and is only applicable for add-in cards with a PCIe width of x16 (or greater). Add-in cards that implement a 2C card edge do not use the PRSNTB3# pin for capabilities or present detection.</p>
<p>BIF0# BIF1# BIF2#</p>	<p>A7 A8 A9</p>	<p>Output</p>	<p>Bifurcation [0:2]# pins allow the baseboard to force configure the add-in card bifurcation.</p> <p>For baseboards, these pins shall be outputs driven from the baseboard I/O hub and allow the system to force configure the add-in card bifurcation. The baseboard may optionally tie the BIF[0:2]# signals to 3.3Vaux or to ground per the definitions are described in Section 3.6 if no dynamic bifurcation configuration is required.</p> <p>For add-in cards, these signals shall connect to the endpoint bifurcation pins if it is supported.</p> <p>Note: the required combinatorial logic output for endpoint bifurcation is dependent on the specific silicon and is not defined in this specification.</p>

Figure 19: PCIe Present and Bifurcation Control Pins



3.4.3 SMBus Interface Pins

This section provides the pin assignments for the SMBus interface signals. The AC/DC specifications are defined in the SMBus 2.0 and I²C bus specifications. An example connection diagram is shown in Figure XXX.

Table 12: Pin Descriptions – SMBus

Signal Name	Pin #	Baseboard Direction	Signal Description
SMCLK	A7	Output, OD	<p>SMBus clock. Open drain, pulled up to 3.3Vaux on the baseboard.</p> <p>For baseboards, the SMCLK from the platform SMBus master shall be connected to the connector.</p>

			For add-in cards, the SMCLK from the endpoint silicon shall be connected to the card edge gold fingers.
SMDAT	A8	Input / Output, OD	<p>SMBus Data. Open drain, pulled up to 3.3Vaux on the baseboard.</p> <p>For baseboards, the SMDAT from the platform SMBus master shall be connected to the connector.</p> <p>For add-in cards, the SMDAT from the endpoint silicon shall be connected to the card edge gold fingers.</p>
SMRST#	A9	Output, OD	<p>SMBus reset. Open drain.</p> <p>For baseboards, this pin shall be pulled up to 3.3Vaux. The SMRST pin may be used to reset optional downstream SMBus devices (such as temperature sensors). The SMRST# implementation shall be mandatory for baseboard implementations.</p> <p>For add-in cards, SMRST# is optional and is dependent on the add-in card implementation.</p>

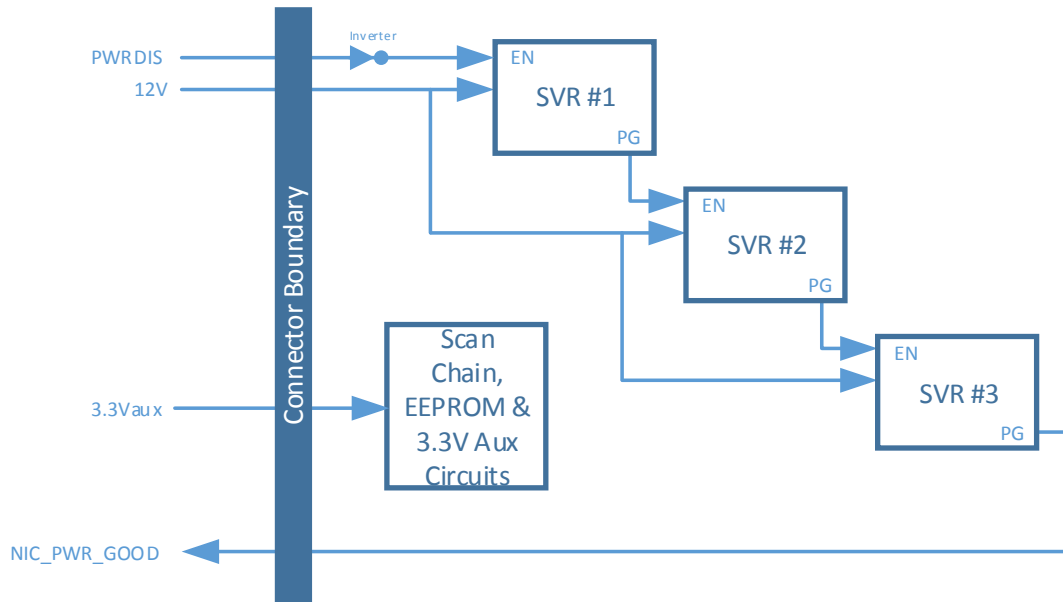
3.4.4 Power Supply Pins

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 20.

Table 13: 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	<p>12V main or 12V Aux power; total of 6 pins per connector. The 12V pins shall be rated to 1.1A per pin with a maximum derated power delivery of 80W.</p> <p>The +12V power pins shall be within the rail tolerances as defined in Section 3.10 when the PWRDIS pin is driven low by the baseboard.</p>
+3.3V/3.3V_AUX	B11	Power	<p>3.3V main or 3.3V Aux power; total of 1 pin per connector. The 3.3V pin shall be rated to 1.1A for a maximum derated power delivery of 3.63W.</p> <p>The 3.3Vaux/main power pin shall be within the rail tolerances as defined in Section 3.10 when the PWRDIS pin is driven low by the baseboard.</p>
PWRDIS	B12	Output, O/D	<p>Power disable. Active high. Open-drain</p> <p>This signal shall be pulled up to 3.3V through a 10kOhm resistor on the baseboard.</p> <p>When high, all add-in card supplies shall be disabled.</p> <p>When low, add-in card supplies shall be enabled.</p>

Figure 20: Example Power Supply Topology



3.4.5 Miscellaneous Pins

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

Table 14: Pin Descriptions – Miscellaneous 1

Signal Name	Pin #	Baseboard Direction	Signal Description
RFU, N/C	B68, B69, A68, A69, A70	Input / Output	Reserved future use pins. These pins shall be left as no connect.

3.5 Signal Descriptions – OCP Bay (Primary Connector)

The following section describes the functions in the Primary Connector 28-pin OCP bay. This 28-pin bay is shown in Section 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 15: Pin Descriptions – PCIe 2

Signal Name	Pin #	Baseboard Direction	Signal Description
REFCLKn2 REFCLKp2	OCP_B11 OCP_B12	Output	PCIe compliant differential reference clock #2, and #3. 100MHz reference clocks are used for the add-in card PCIe core logic. For baseboards, the REFCLK2 and REFCLK3 signals are required at the Primary connector. For add-in cards, the required REFCLKs shall be connected per the endpoint datasheet. Note: REFCLK2 and REFCLK3 are not used for cards that only support a 1 x16, 1 x8 or 2 x8 connection. Refer to Section 2.1 in the PCIe CEM Specification, Rev 4.0 for details.
REFCLKn3 REFCLKp3	OCP_A11 OCP_A12	Output	
PERST2# PERST3#	OCP_A1 OCP_A2	Output	PCIe Reset #2, #3. Active low.

			<p>When PERSTn# is deasserted, the signal shall indicate the applied power is within tolerance and stable for the add-in card.</p> <p>PERST# shall be deasserted at least 100ms after the power rails are within the operating limits per the PCIe CEM Specification. The PCIe REFCLKs shall also become stable within this period of time.</p> <p>PERST shall be pulled high to 3.3Vaux on the baseboard.</p> <p>For OCP NIC 3.0, PERST deassertion shall also indicate the full card power envelope is available to the add-in card.</p> <p>For baseboards, the PERST[0:1]# signals are required at the connector.</p> <p>For add-in cards, the required PERST[0:1]# signals shall be connected to the endpoint silicon.</p> <p>Note: PERST2# and PERST3# are not used for cards that only support a 1 x16 or 2 x8 connection.</p> <p>Refer to Section 2.2 in the PCIe CEM Specification, Rev 4.0 for details.</p>
WAKE#	OCP_A3	Input, OD	WAKE#. Open drain. Active low.

			<p>This signal shall be driven by the add-in card to notify the baseboard to restore PCIe link. For add-in cards that support multiple WAKE# signals, their respective WAKE# pins may be tied together as the signal is open-drain to form a wired-OR.</p> <p>For baseboards, this signal shall be pulled up to +3.3V on the baseboard with a 10kOhm resistor. This signals shall be connected to the system WAKE# signal.</p> <p>For add-in cards, this signal shall be directly connected to the endpoint silicon WAKE# pin(s).</p> <p>Refer to Section 2.3 in the PCIe CEM Specification, Rev 4.0 for details.</p>
--	--	--	---

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

This section provides the pin assignments for the NC-SI over RBT interface signals on the Primary Connector OCP bay. The AC/DC specifications are defined in the NC-SI specification. An example connection diagram is shown in Figure 21.

Table 16: Pin Descriptions – NC-SI Over RBT

Signal Name	Pin #	Baseboard Direction	Signal Description
RBT_CLK_IN	OCP_A14	Output	Reference clock input. Synchronous clock reference for receive, transmit and control interface. The clock shall have a nominal frequency of 50MHz \pm 100ppm.

			<p>For baseboards, this pin shall be connected between the baseboard NC-SI over RBT PHY and the Primary connector OCP bay. This signal requires a 100kOhm pull down resistor on the baseboard. If the baseboard does not support NC-SI over RBT, then this signal shall be terminated to ground through a 100kOhm pull down resistor.</p> <p>For add-in cards, this pin shall be connected between the gold finger to the endpoint silicon. This pin shall be left as a no connect if NC-SI is not supported.</p>
RBT_CRS_DV	OCP_B14	Input	<p>Carrier sense/receive data valid. This signal is used to indicate to the baseboard that the carrier sense/receive data is valid.</p> <p>For baseboards, this pin shall be connected between the baseboard NC-SI over RBT PHY and the connector. This signal requires a 100kOhm pull down resistor on the baseboard. If the baseboard does not support NC-SI over RBT, then this signal shall be terminated to ground through a 100kOhm pull down resistor.</p> <p>For add-in cards, this pin shall be connected between the gold finger to the endpoint silicon. This pin shall be left as a no connect if NC-SI is not supported.</p>
RBT_RXD0 RBT_RXD1	OCP_B9 OCP_B8	Input	<p>Receive data. Data signals from the network controller to the BMC.</p>

			<p>For baseboards, this pin shall be connected between the baseboard NC-SI over RBT PHY and the connector. This signal requires a 100kOhm pull-up resistor to 3.3Vaux on the baseboard. If the baseboard does not support NC-SI over RBT, then this signal shall be terminated to 3.3Vaux through a 100kOhm pull-up.</p> <p>For add-in cards, this pin shall be connected between the gold finger and the RBT_RXD[0:1] pins on endpoint silicon. This pin shall be left as a no connect if NC-SI is not supported.</p>
RBT_TX_EN	OCP_A7	Output	<p>Transmit enable.</p> <p>For baseboards, this pin shall be connected between the baseboard NC-SI over RBT PHY and the connector. This signal requires a 100kOhm pull down resistor to ground on the baseboard. If the baseboard does not support NC-SI over RBT, then this signal shall be terminated to ground through a 100kOhm pull down.</p> <p>For add-in cards, this pin shall be connected between the gold finger to the endpoint silicon. This pin shall be left as a no connect if NC-SI is not supported.</p>
RBT_TXD0 RBT_TXD1	OCP_A9 OCP_A8	Output	<p>Transmit data. Data signals from the BMC to the network controller.</p> <p>For baseboards, this pin shall be connected between the baseboard NC-SI over RBT PHY</p>

			<p>and the connector. This signal requires a 100kOhm pull-up resistor to 3.3Vaux on the baseboard. If the baseboard does not support NC-SI over RBT, then this signal shall be terminated to 3.3Vaux through a 100kOhm pull-up.</p> <p>For add-in cards, this pin shall be connected between the gold finger to the RBT_TXD[0:1] pins on the endpoint silicon. This pin shall be left as a no connect if NC-SI is not supported.</p>
RBT_ARB_OUT	OCP_A5	Output	<p>NC-SI hardware arbitration output. This pin shall only be used if the endpoint silicon supports hardware arbitration. This pin shall be connected to the RBT_ARB_IN signal of an adjacent device in the hardware arbitration ring.</p> <p>The baseboard shall implement a multiplexing implementation that directs the RBT_ARB_OUT to the RBT_ARB_IN pin of the next NC-SI capable device in the ring, or back to the RBT_ARB_IN pin of the source device if there is a single device on the ring.</p> <p>For baseboards, this pin shall be connected between the baseboard OCP connector(s) to complete the hardware arbitration ring. If the baseboard does not support NC-SI over RBT, this signal shall be directly connected to the RBT_ARB_IN pin.</p>

			<p>For add-in cards, this pin shall be connected from the gold finger to the RBT_ARB_IN pin on the endpoint silicon. This pin shall be left as a no connect if NC-SI is not supported.</p>
RBT_ARB_IN	OCP_A4	Input	<p>NC-SI hardware arbitration input. This pin shall only be used if the endpoint silicon supports hardware arbitration. This pin shall be connected to the RBT_ARB_OUT signal of an adjacent device in the hardware arbitration ring.</p> <p>The baseboard shall implement a multiplexing implementation that directs the RBT_ARB_IN to the RBT_ARB_OUT pin of the next NC-SI capable device in the ring, or back to the RBT_ARB_OUT pin of the source device if there is a single device on the ring.</p> <p>For baseboards, this pin shall be connected between the baseboard OCP connector(s) to complete the hardware arbitration ring. If the baseboard does not support NC-SI over RBT, this signal shall be directly connected to the RBT_ARB_OUT pin.</p> <p>For add-in cards, this pin shall be connected between the gold finger to the RBT_ARB_OUT pin on the endpoint silicon. Leave this pin as a no connect if NC-SI is not supported.</p>
SLOT_ID	OCP_B7	Output	<p>NC-SI Address pin. This pin shall only be used if the end point silicon supports package identification.</p>

		<p>For baseboards, this pin shall be used to identify the slot ID value. This pin shall be directly to GND for SlotID = 0. This pin shall be pulled up to 3.3Vaux for SlotID = 1.</p> <p>For add-in cards, this pin shall be connected to the endpoint device GPIO associated with the Package ID[1] field. Refer to Section 4.9.1 and the device datasheet for details.</p> <p>For add-in cards with multiple endpoint devices, the SLOT_ID pin may be used to configure a different Package ID value so long as the resulting combination does not cause addressing interferences.</p> <p>For endpoint devices without NC-SI support, this pin shall be left as a no connect on the add-in card.</p>
--	--	--

Figure 21: NC-SI Over RBT Connection Example – Single Primary Connector

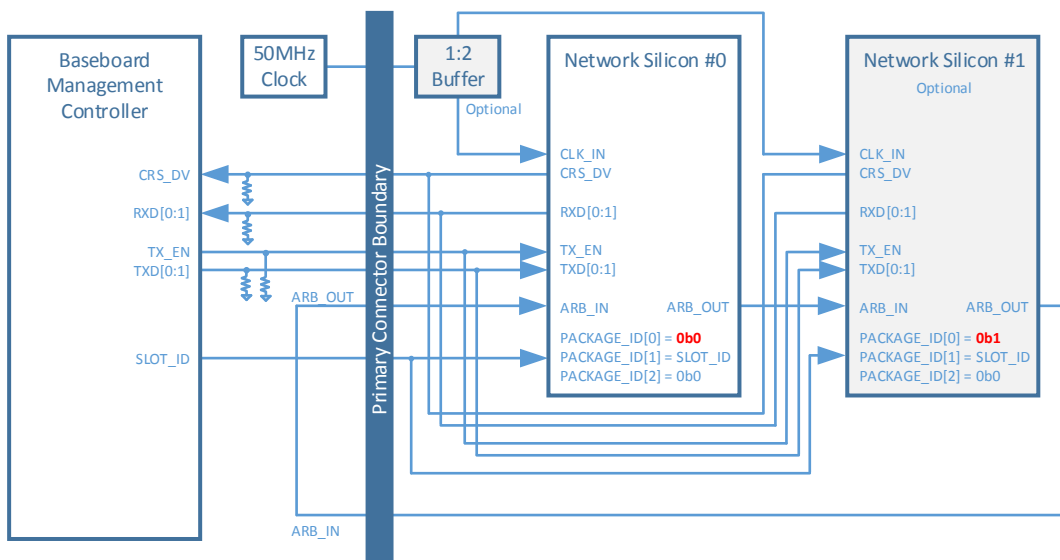
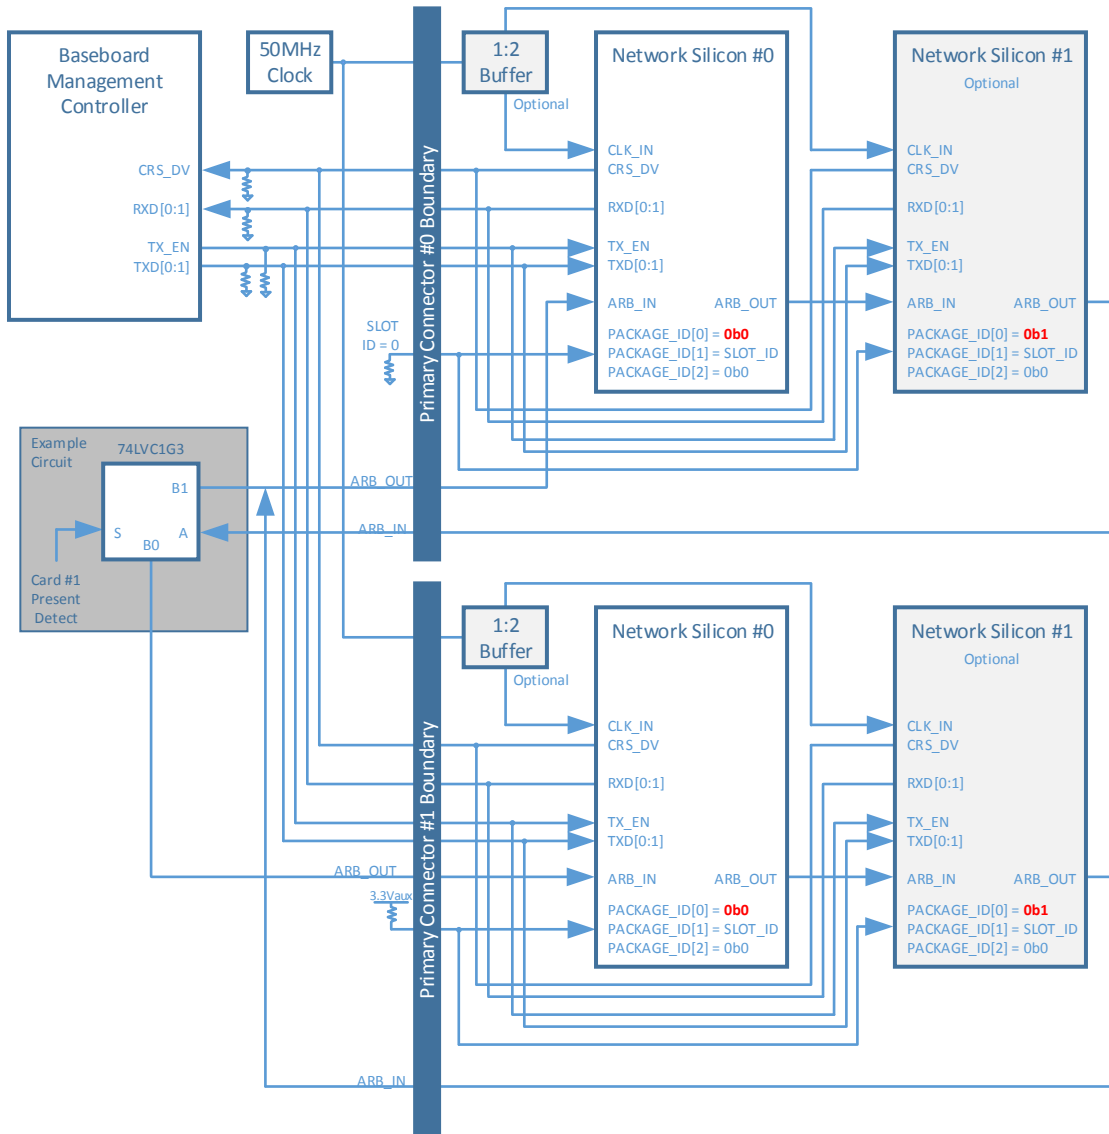


Figure 22: NC-SI Over RBT Connection Example – Dual Primary Connector



Note 1: For baseboard designs with a single Primary Connector, connect ARB_IN to ARB_OUT to complete the NC-SI hardware arbitration ring. For designs with multiple Primary Connectors, connect ARB_IN and ARB_OUT to an analog mux to complete the NC-SI arbitration ring based on the number of cards installed in the system. An example dual Primary Connector implementation is shown in Figure 22.

Note 2: For add-in cards with two discrete endpoint silicon, the Package ID[0] bit shall be statically set based on its silicon instance. For example, the figure above shows Network

Silicon #0 and Network Silicon #1. Network Silicon #0 has Package ID[0] = 0b0, Network Silicon #1 has Package ID[0] = 0b1.

3.5.3 Scan Chain Pins – OCP Bay (Primary Connector)

This section provides the pin assignments for the Scan Bus interface signals on the Primary Connector OCP Bay. The AC/DC specifications are defined in Section XXX. An example connection diagram is shown in Figure 23.

Table 17: Pin Descriptions – Scan Chain

Signal Name	Pin #	Baseboard Direction	Signal Description
CLK	OCP_B6	Output	<p>Scan clock. The CLK is an output pin from the baseboard to the add-in card. The CLK may run up to 12.5MHz.</p> <p>For baseboard implementations, the CLK pin shall be connected to the Primary Connector. The CLK pin shall be tied directly to GND if the scan chain is not used.</p> <p>For NIC implementations, the CLK pin shall be connected to Shift Registers 0 & 1, and optionally connected to Shift Registers 2 & 3 (if implemented) as defined in the text and Figure 23, below. The CLK pin shall be pulled up to 3.3Vaux through a 1kOhm resistor.</p>
DATA_OUT	OCP_B5	Output	<p>Scan clock data output from the baseboard to the add-in card. This bit stream is used to shift in NIC configuration data.</p> <p>For baseboard implementations, the DATA_OUT pin shall be connected to the Primary</p>

			<p>Connector. The DATA_OUT pin shall be tied directly to GND if the scan chain is not used.</p> <p>For NIC implementations, the DATA_OUT pin may be left floating if it is not used for add-in card configuration. The DATA_OUT pin shall be pulled up to 3.3Vaux through a 1kOhm resistor.</p>
DATA_IN	OCP_B4	Input	<p>Scan clock data input to the baseboard. This bit stream is used to shift out NIC status bits.</p> <p>For baseboard implementations, the DATA_IN pin shall be pulled up to 3.3Vaux 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.</p> <p>For NIC implementations, the DATA_IN scan chain is required. The DATA_IN pin shall be connected to Shift Registers 0 & 1, as defined in the text and Figure 23.</p>
LD#	OCP_B3	Output	<p>Scan clock shift register load. Used to latch configuration data on the add-in card.</p> <p>For baseboard implementations, the LD# pin shall be pulled up to 3.3Vaux through a 1kOhm resistor if the scan chain is not used to prevent the add-in card from erroneous data latching.</p> <p>For NIC implementations, the LD# pin implementation is required. The LD# pin shall be connected to Shift Registers 0 & 1 as defined in the text and Figure 23. The LD# pin</p>

			shall be pulled up to 3.3Vaux through a 1kOhm resistor.
--	--	--	---

The scan chain provides side band status indication between the add-in card and the baseboard. The scan chain bit definition is defined in the two tables below. The scan chain data stream is 32-bits in length for both the DATA_OUT and the DATA_IN streams. The scan chain implementation is optional on the host, but is mandatory on all OCP NIC 3.0 cards. The scan chain components operates on the 3.3Vaux power domain.

The DATA_OUT bus is an output from the host. The DATA_OUT bus provides initial configuration options to the add-in card. At the time of this writing, the default implementation does not use the DATA_OUT stream and is not implemented on the NIC. However, all baseboard systems that implement the Scan Chain shall connect DATA_OUT between the platform and the Primary Connector for future-proofing NIC implementations and subsequent revisions of this specification.

Table 18: Pin Descriptions – Scan Chain DATA_OUT Bit Definition

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

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

DATA_IN shift registers 0 & 1 shall be mandatory for scan chain implementations. DATA_IN shift registers 2 & 3 are optional depending on the card type and fields being reported to the host. DATA_IN shift register 2 may be used to indicate future definitions of the scan



chain bit stream. DATA_IN shift registers 3 (in conjunction with shift register 2) are required for reporting link/activity indication on card implementations with 5-8 ports.

The host should read the DATA_IN bus multiple (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.

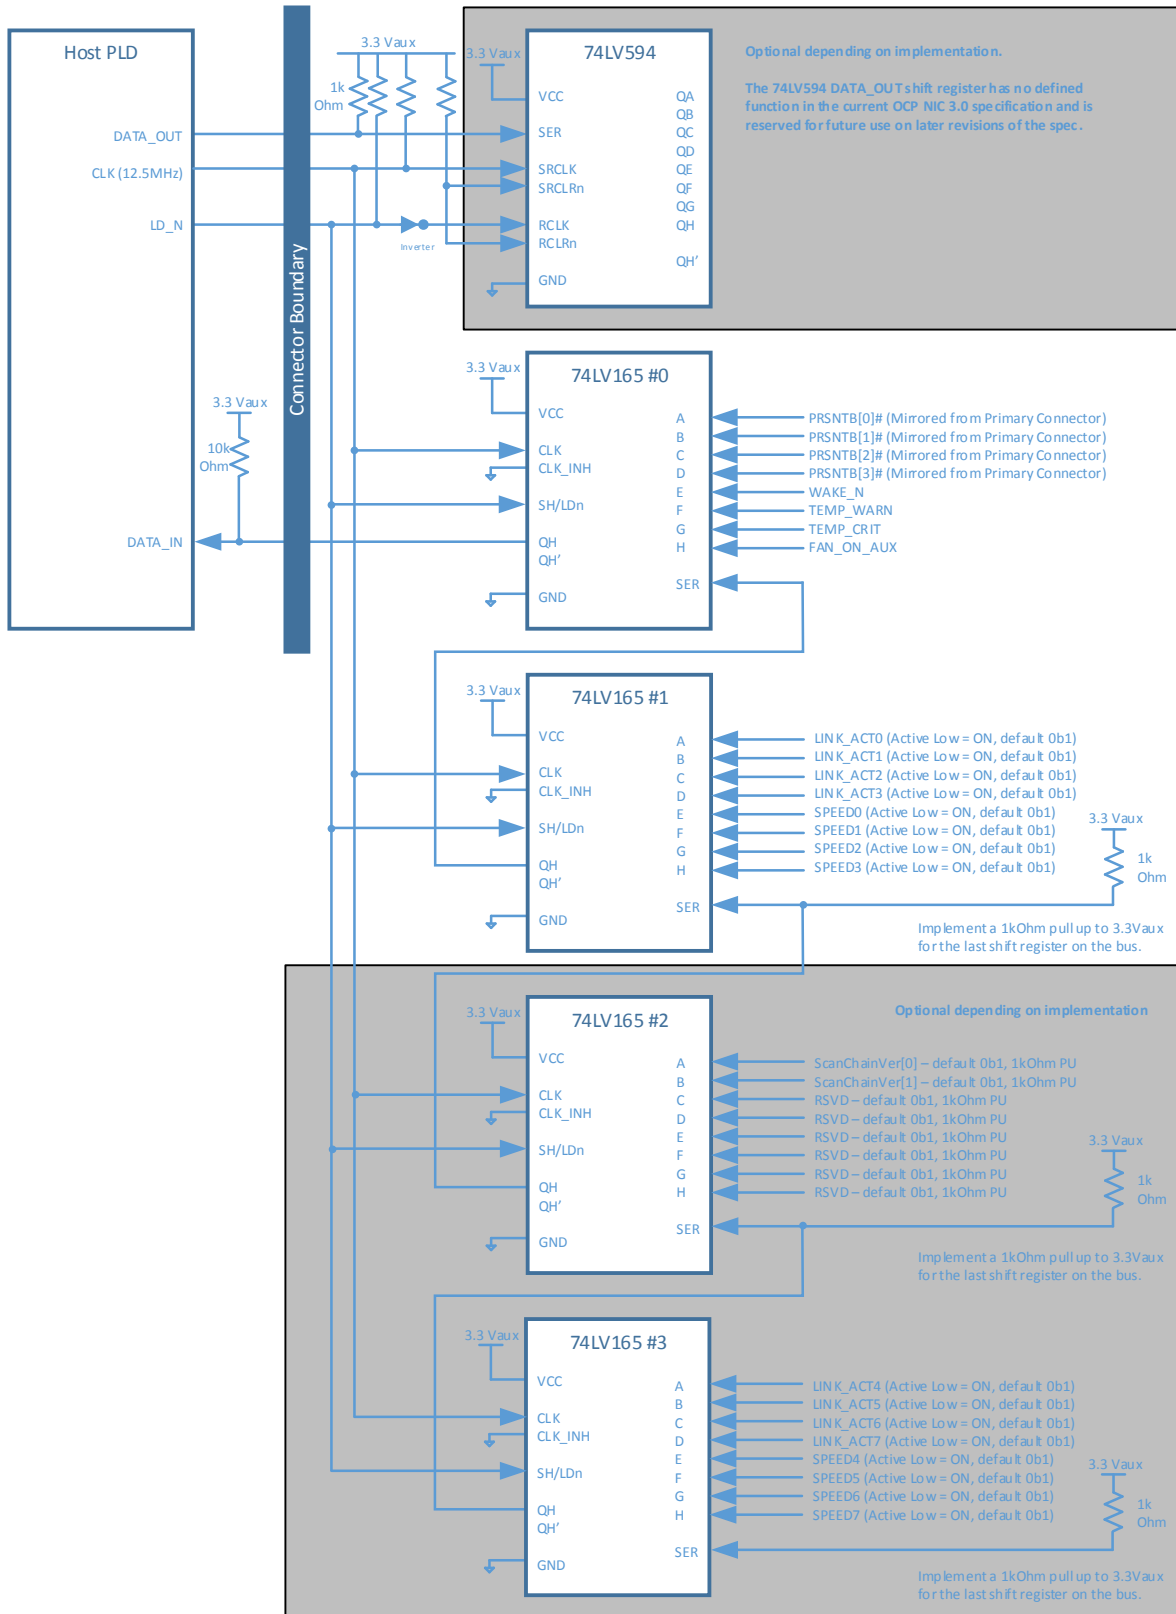
Table 19: Pin Descriptions – Scan Bus DATA_IN Bit Definition

Byte.bit	DATA_OUT Field Name	Default Value	Description
0.0	PRSNTB[0]#	0bX	PRSNTB[3:0]# bits shall reflect the same state as the signals on the Primary Connector.
0.1	PRSNTB[1]#	0bX	
0.2	PRSNTB[2]#	0bX	
0.3	PRSNTB[3]#	0bX	
0.4	WAKE_N	0bX	PCIe WAKE_N signal shall reflect the same state as the signal on the Primary Connector.
0.5	TEMP_WARN	0b0	Temperature monitoring pin from the on-card thermal solution. This pin shall be asserted high when temperature sensor exceeds the temperature warning threshold.
0.6	TEMP_CRIT	0b0	Temperature monitoring pin from the on-card thermal solution. This pin shall be asserted high when temperature sensor exceeds the temperature critical threshold.
0.7	FAN_ON_AUX	0b0	When high, FAN_ON_AUX shall request the system fan to be enabled for extra cooling in the S5 state.
1.0	LINK_ACT0	0b1	Port 0..3 link/activity indication. Active low.
1.1	LINK_ACT1	0b1	

1.2	LINK_ACT2	0b1	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. Blinking = activity is detected on the port. The blink rate should blink low for 50-500ms during activity periods. Off = the physical link is down or disabled
1.3	LINK_ACT3	0b1	
1.4	SPEED_A0	0b1	Port 0..3 speed A (max rate) indication. Active low. 0b0 – Port is linked at maximum speed. 0b1 – Port is not linked at the maximum speed or no link is present.
1.5	SPEED_A1	0b1	
1.6	SPEED_A2	0b1	
1.7	SPEED_A3	0b1	
2.0	ScanChainVer[0]	0b1	ScanChainVer[1:0] shall be used to indicate the scan chain bit definition version. The encoding shall be as follows: 0b11 – Scan chain bit definitions version 1 corresponding to OCP NIC 3.0 version 1.0. All other encoding values shall be 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	LINK_ACT4	0b1	Port 4..7 link/activity indication. Active low.
3.1	LINK_ACT5	0b1	
3.2	LINK_ACT6	0b1	

3.3	LINK_ACT7	0b1	<p>0b0 – Link LED is illuminated on the host platform.</p> <p>0b1 – Link LED is not illuminated on the host platform.</p> <p>Steady = link is detected on the port.</p> <p>Blinking = activity is detected on the port. The blink rate should blink low for 50-500ms during activity periods.</p> <p>Off = the physical link is down or disabled</p>
3.4	SPEED_A4	0b1	<p>Port 4..7 speed A (max rate) indication. Active low.</p> <p>0b0 – Port is linked at maximum speed.</p> <p>0b1 – Port is not linked at the maximum speed or no link is present.</p>
3.5	SPEED_A5	0b1	
3.6	SPEED_A6	0b1	
3.7	SPEED_A7	0b1	

Figure 23: 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 20: Pin Descriptions – Miscellaneous 2

Signal Name	Pin #	Baseboard Direction	Signal Description
PWRBRK#	OCP_B2	Output, OD	<p>Power break. Active low, open drain.</p> <p>This signal shall be pulled up to 3.3V_{aux} on the add-in card with a minimum of 95kΩ. The pull up on the baseboard shall be a stiffer resistance in-order to meet the timing specs as shown in the PCIe CEM Specification.</p> <p>When this signal is driven low by the baseboard, the Emergency Power Reduction State is requested. The add-in card shall move to a lower power consumption state.</p>
NIC_PWR_GOOD	OCP_B1	Input	<p>NIC Power Good. Active high. This signal is driven by the add-in card.</p> <p>When high, this signal shall indicate that all of the add-in card power rails are operating within nominal tolerances.</p> <p>When low, this signal shall indicate that the add-in card power supplies are not yet within nominal tolerances or are in a fault condition.</p>

			<p>For baseboards, this pin may be connected to the platform I/O hub as a NIC power health status indication. This signal shall be pulled down to ground with a 100kOhm resistor on the baseboard to prevent a false power good indication if no add-in card is present.</p> <p>For add-in cards this signal shall indicate the add-in card power is “good”. This signal may be implemented by a cascaded power good or a discrete power good monitor output.</p>
GND	OCP_A6 OCP_A10 OCP_A13 OCP_B10 OCP_B13	GND	Ground return; a total of 5 ground pins are on the OCP bay area.

3.6 PCIe Bifurcation Mechanism

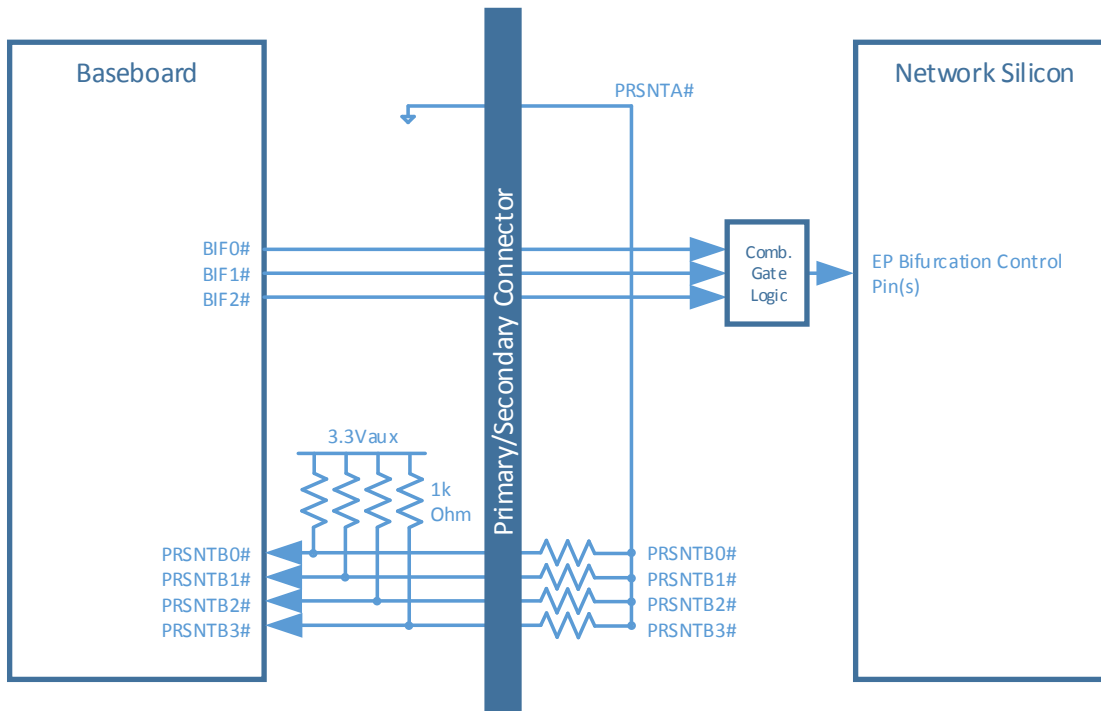
OCP3.0 baseboards and add-in cards support multiple bifurcation combinations. Single socket baseboards with a single or multiple root ports, as well as a multi-socket baseboards with a single or multiple root ports are supported. The bifurcation mechanism also supports add-in cards with a single or multiple end points. These features are accomplished via I/O pins on the Primary and Secondary connector:

- PRSNTA#, PRSNTB[3:0]#. The PRSNTA# pin shall connect to the PRSNTB# pins as a hard coded value on the add-in card. The encoding of the PRSNTB[3:0]# pins allows the baseboard to determine the PCIe Links available on the add-in card.
- BIF[3:0]#. The BIF# pin states shall be controlled by the baseboard to allow the baseboard to override the default end point bifurcation for silicon that support bifurcation. Additional combinatorial logic is required and is specific to the card silicon. The combinatorial logic is not covered in this specification. The BIF[3:0]# pins

may optionally be hardcoded for baseboards that do not require a dynamic bifurcation override.

A high level bifurcation connection diagram is shown in Figure 24.

Figure 24: 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 21 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 21 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 21 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.

3.6.4 Bifurcation Detection Flow

[Need input and clarification from system vendors]

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

1. The baseboard shall read the state of the PRSNTB[3:0]# pins. An add-in card is present in the system if the resulting value is not 0b1111.
2. Firmware determines the add-in card PCIe lane width capabilities per Table 21 by reading the PRSNTB[3:0]# pins.
3. The baseboard reconfigures the PCIe bifurcation on its ports to match the highest common lane width and lowest common link count on the card.
4. For cases where the baseboard request a link count override (such as requesting a 4-host baseboard requesting 4 x4 operation on a supported card that would otherwise default to a 2 x8 case), the BIF[0:2]# pins shall be asserted as appropriate. Asserting the BIF[0:2]# pins assumes the add-in card supports the requested link override.
5. PERST# shall be deasserted after the >100ms window as defined by the PCIe specification. Refer to Section 3.12 for timing details.

3.6.5 PCIe Bifurcation Examples

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

Figure 25 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 25: Single Host (1 x16) and 1 x16 Add-in Card (Single Controller)

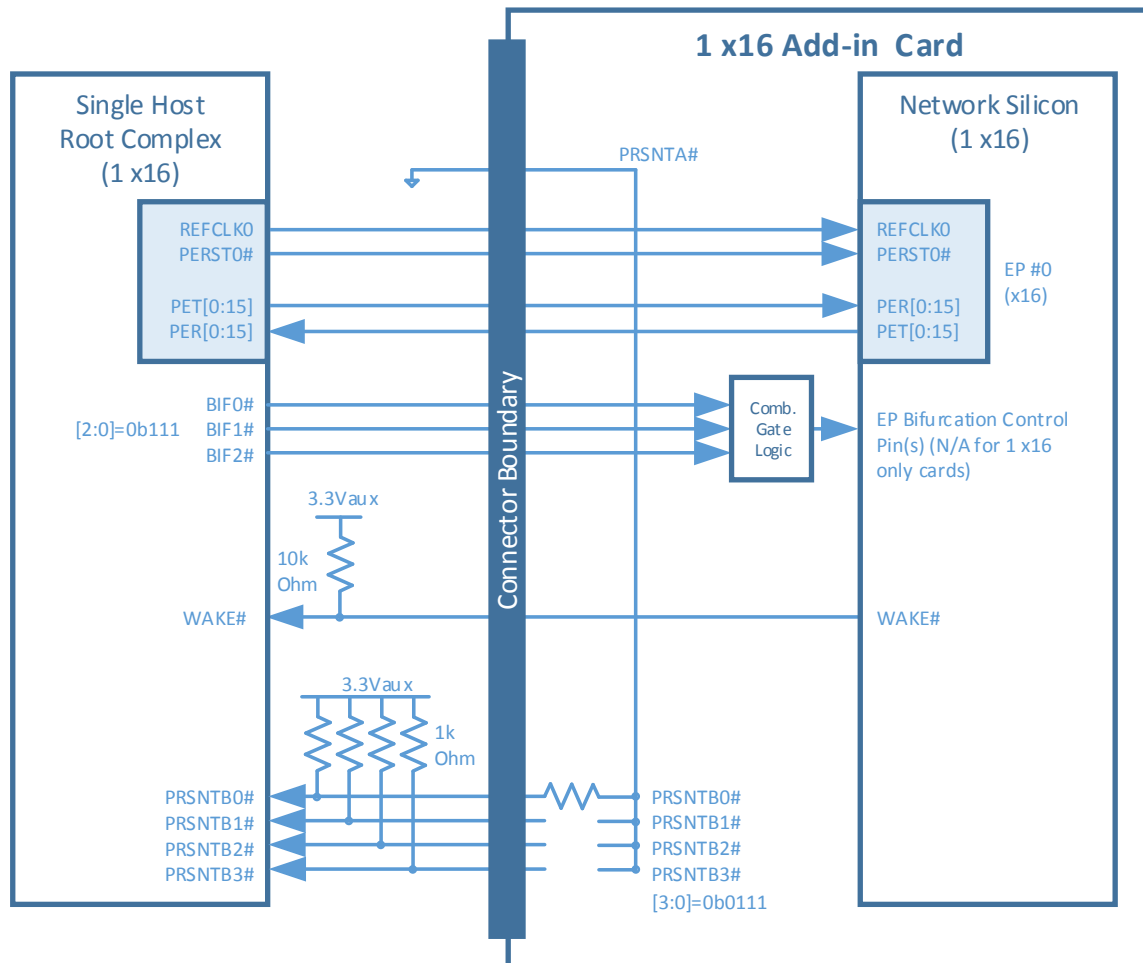


Figure 26 illustrates a single host baseboard that supports 2 x8 with a single controller add-in card that also supports 2 x8. The PRSTNB[3:0]# state is 0b0110. The BIF[2:0]# state is

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

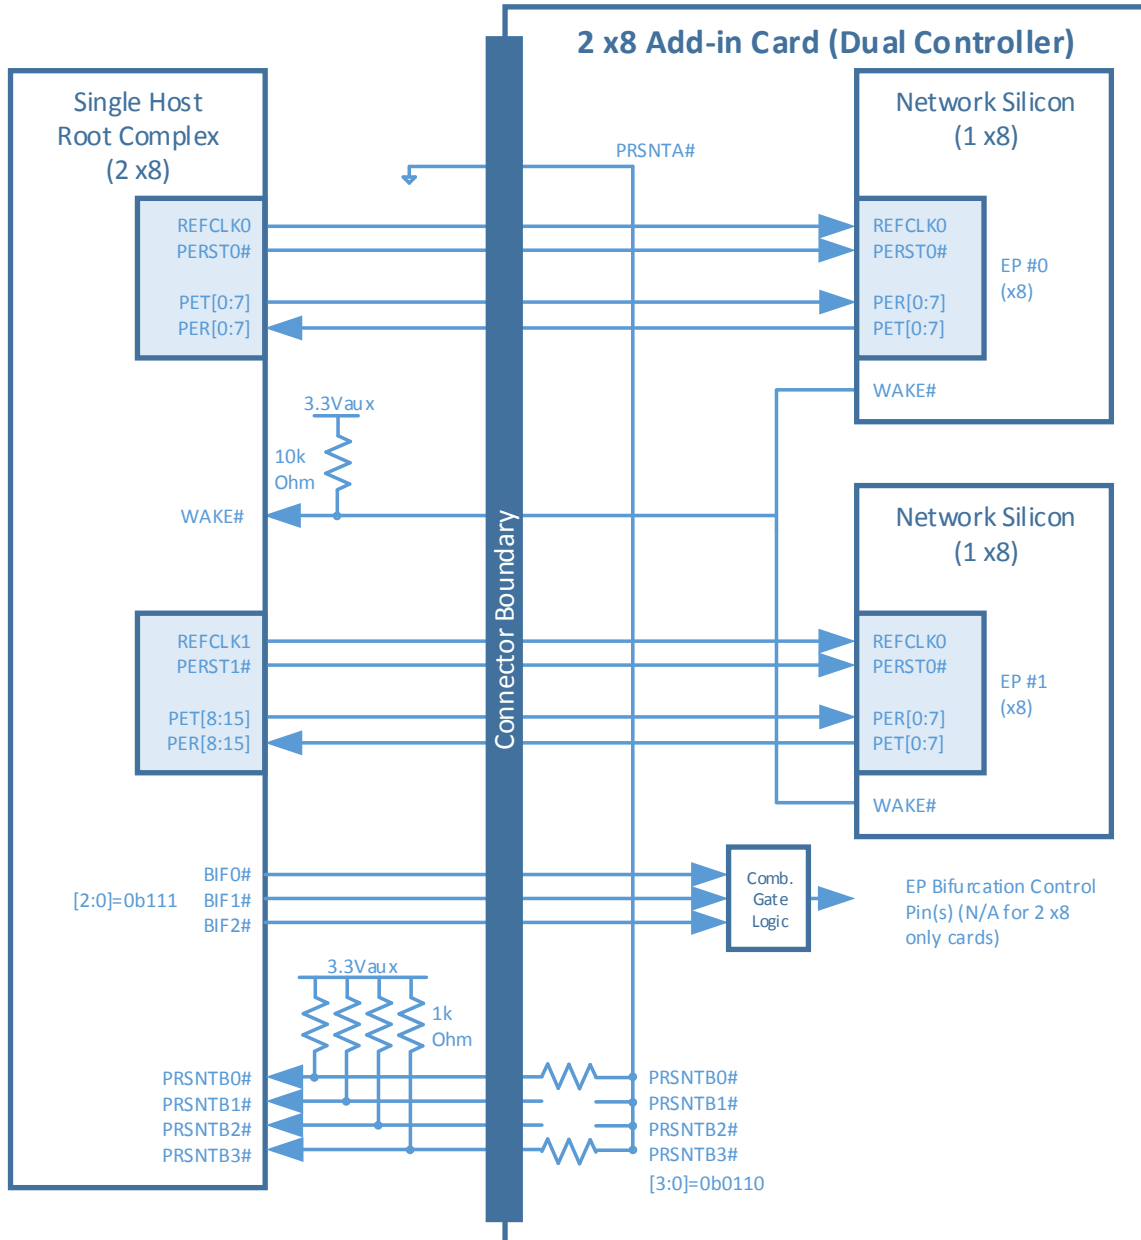


Figure 27 illustrates a four host baseboard that supports 4 x4 with a single controller add-in card that supports 1 x16, 2 x8 and 4 x4. The PRSTNB[3:0]# state is 0b0011. The BIF[2:0]# state is 0b101 as the end point network controller is forced to bifurcate to 4 x4. The PRSNTB encoding notifies the baseboard 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 27: Four Hosts (4 x4) and 4 x4 Add-in Card (Single Controller)




Figure 28 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 28: Four Hosts (4 x4) and 4 x4 Add-in Card (Four Controllers)


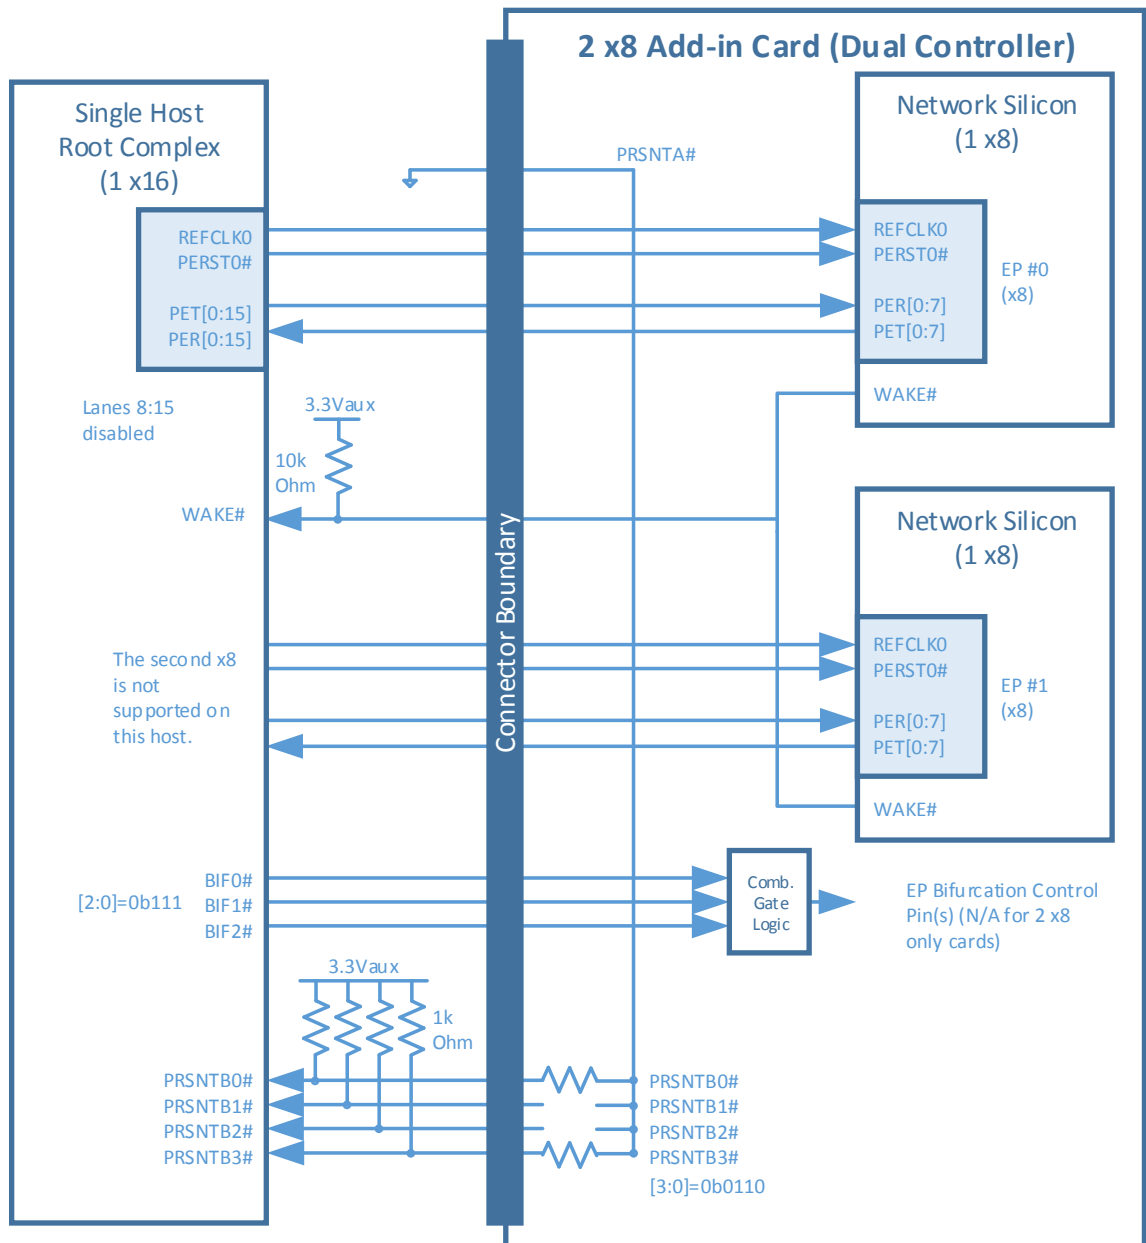


Figure 29 illustrates a single host baseboard that supports 1 x16 with a dual controller add-in card that supports 2 x8. The PRSTNB[3:0]# state is 0b0110. The BIF[2:0]# state is 0b111 as there is no need to instruct the end-point network controllers to a specific bifurcation. The PRSNTB encoding notifies 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 29: 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 22. Cards that implement both the Primary and Secondary connectors have a total of up to 6 REFCLKs.

Table 22: PCIe Clock Associations

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

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

- For a 1 x16 capable add-in card, REFCLK0 shall be used for lanes [0:15].
- For a 2 x8 capable add-in card, REFCLK0 shall be used for lanes [0:7] and REFCLK1 shall be used for lanes [8:15].
- For a 4 x4 capable add-in card, REFCLK0 shall be used for lanes [0:3], REFCLK1 shall be used for lanes [4:7], REFCLK2 shall be used for lanes [8:11] and REFCLK3 shall be used for lanes [12:15]. Pins for REFCLK2 and REFCLK3 are described in Section 3.5.1 and are located on the 28-pin OCP bay.

Figure 30: PCIe Interface Connections for 1 x16 and 2 x8 Add-in Cards

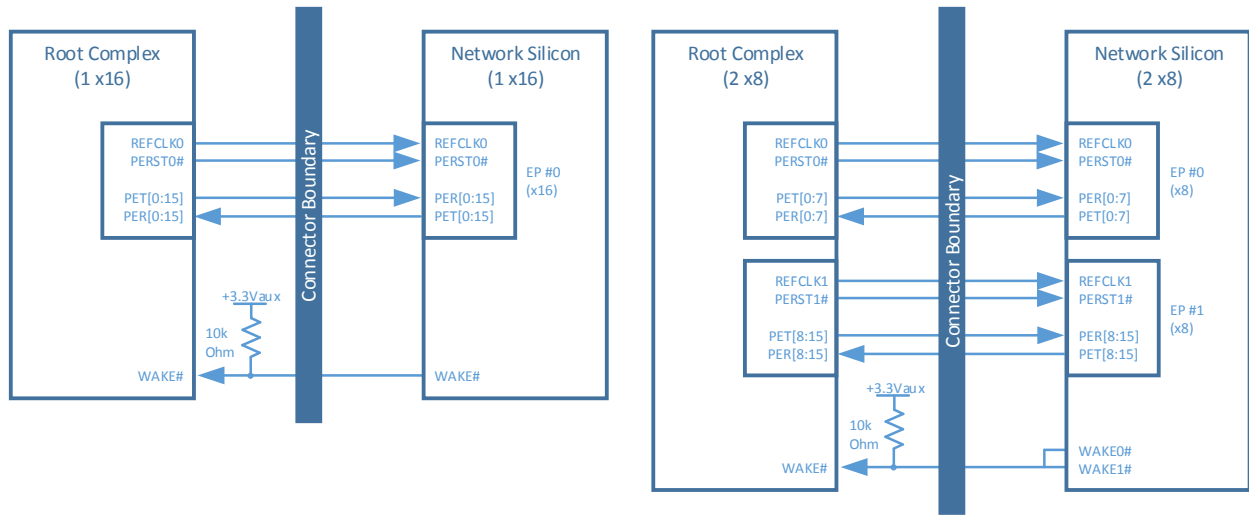
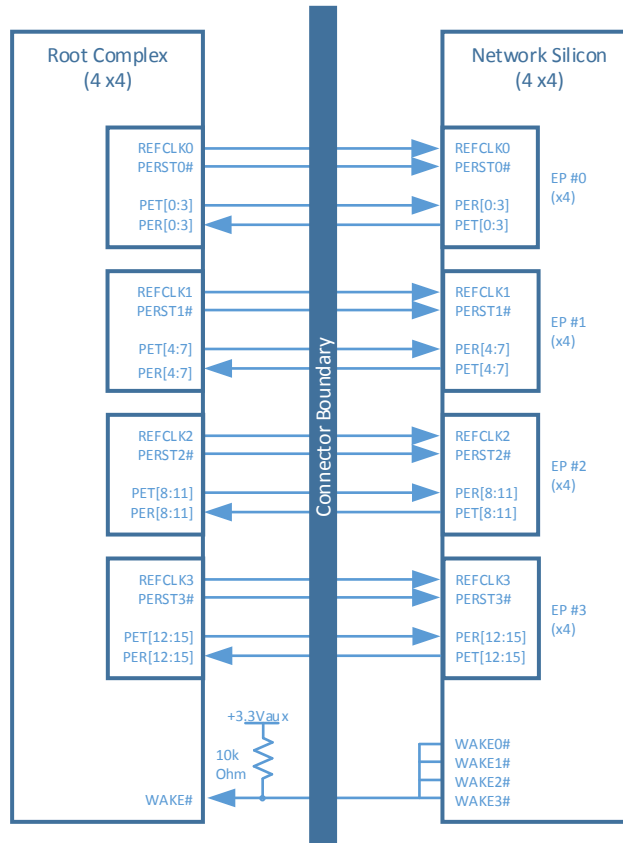


Figure 31: PCIe Interface Connections for a 4 x4 Add-in Card





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

Card Width	Card Name	Card Short Name	Supported Bifurcation	Address-Card Encoding	Host	Upstream Definer	Upstream Link	BIF[2:0] #	Resulting Link	1x16, 1x8, 1x4, 1x2, 1															
										Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	1x8	1x8	Card Not Present	0b1111	1Host	1Upstream Socket	1Link	0b0000	-	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	1x4	1x4, 1x2, 1x1	Card Not Present	0b1110	1Host	1Upstream Socket	1Link	0b0000	1x8	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	1x4	1x4, 1x2, 1x1	Card Not Present	0b1110	1Host	1Upstream Socket	1Link	0b0000	1x4	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	1x2	1x2, 1x1	Card Not Present	0b1110	1Host	1Upstream Socket	1Link	0b0000	1x2	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	1x1	1x1	Card Not Present	0b1110	1Host	1Upstream Socket	1Link	0b0000	1x1	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	1x8 Option B	1x8, 1x4, 1x2, 1x1	Card Not Present	0b1101	1Host	1Upstream Socket	1Link	0b0000	1x8	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
4C	2x8 Option B	2x8, 2x4, 2x2, 2x1	Card Not Present	0b1101	1Host	1Upstream Socket	1Link	0b0000	1x8*	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	1x8 Option D	1x8, 1x4, 1x2, 1x1	Card Not Present	0b1100	1Host	1Upstream Socket	1Link	0b0000	1x8	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
4C	1x16 Option D	1x16, 1x8, 1x4, 4x4, 4x2 (First 8 lanes), 4x1	Card Not Present	0b1100	1Host	1Upstream Socket	1Link	0b0000	1x16	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	RSVD	RSVD	RSVD	0b1011	1Host	1Upstream Socket	1Link	0b0000	-	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	RSVD	RSVD	RSVD	0b1010	1Host	1Upstream Socket	1Link	0b0000	1x4	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
RSVD	RSVD	RSVD	RSVD for lanes x8 ascending	0b1001	1Host	1Upstream Socket	1Link	0b0000	-	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
RSVD	RSVD	RSVD	RSVD for lanes x8 ascending	0b1000	1Host	1Upstream Socket	1Link	0b0000	-	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
4C	1x16	1x16	Card Not Present	0b0111	1Host	1Upstream Socket	1Link	0b0000	1x16	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
4C	2x8 Option A	2x8, 2x4, 2x2, 2x1	Card Not Present	0b0110	1Host	1Upstream Socket	1Link	0b0000	1x8*	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
4C	1x16 Option B	1x16, 1x8, 1x4, 1x2, 1x1	Card Not Present	0b0101	1Host	1Upstream Socket	1Link	0b0000	1x16	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
4C	1x16 Option C	1x16, 1x8, 1x4, 1x2, 1x1	Card Not Present	0b0100	1Host	1Upstream Socket	1Link	0b0000	1x16	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
4C	1x16 Option C	1x16, 1x8, 1x4, 1x2, 1x1	Card Not Present	0b0100	1Host	1Upstream Socket	1Link	0b0000	1x16	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
4C	4x4	4x4	Card Not Present	0b0011	1Host	1Upstream Socket	1Link	0b0000	1x4*	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
RSVD	RSVD	RSVD	RSVD	0b0010	1Host	1Upstream Socket	1Link	0b0000	-	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
RSVD	RSVD	RSVD	RSVD	0b0001	1Host	1Upstream Socket	1Link	0b0000	-	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
RSVD	RSVD	RSVD	RSVD	0b0000	1Host	1Upstream Socket	1Link	0b0000	-	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



Table 24: 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 Modes	Supported Bifurcations	Add-in-Card Encoding	Host	Upstream Devices	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
1x8, 1x4, 1x2, 1x1	1Host	Card Not Present	1x8, 1x4, 1x2, 1x1	0b1111	1Host	1 Upstream Socket	0b0000	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	1x8	1x4, 1x2, 1x1	1x8, 1x4, 1x2, 1x1	0b1110	1Host	1 Upstream Socket	0b0000	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	1x4	1x2, 1x1	1x4, 1x2, 1x1	0b1110	1Host	1 Upstream Socket	0b0000	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	1x2	1x1	1x2, 1x1	0b1110	1Host	1 Upstream Socket	0b0000	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	1x1	1x1	1x1	0b1110	1Host	1 Upstream Socket	0b0000	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	1x8 Option B	2x4, 2x2, 2x1	1x8, 1x4, 1x2, 1x1	0b1101	1Host	1 Upstream Socket	0b0000	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
4C	2x8 Option B	2x8, 2x4, 2x2, 2x1	2x8, 2x4, 2x2, 2x1	0b1101	1Host	1 Upstream Socket	0b0000	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	1x8 Option D	4x2 (Firt 8 lanes), 4x1	1x8, 1x4, 4x1	0b1100	1Host	1 Upstream Socket	0b0000	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
4C	1x16 Option D	2x8, 2x4, 4x2, 4x1	1x16, 1x8, 1x4, 2x8, 2x4, 4x2, 4x1	0b1100	1Host	1 Upstream Socket	0b0000	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
RSVD	RSVD	RSVD	RSVD	0b1011	1Host	1 Upstream Socket	0b0000	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	2x4	1x4, 1x2, 2x1	1x4, 1x2, 2x1	0b1010	1Host	1 Upstream Socket	0b0000	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
RSVD	RSVD	RSVD	RSVD	0b1001	1Host	1 Upstream Socket	0b0000	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
RSVD	RSVD	RSVD	RSVD	0b1000	1Host	1 Upstream Socket	0b0000	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
4C	1x16	2x8, 2x4, 2x2, 2x1	1x16, 1x8, 1x4, 1x2, 1x1	0b0111	1Host	1 Upstream Socket	0b0000	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
4C	2x8 Option A	1x16, 1x8, 1x4, 1x2, 1x1	2x8, 2x4, 2x2, 2x1	0b0110	1Host	1 Upstream Socket	0b0000	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
4C	1x16 Option B	2x8, 2x4, 2x2, 2x1	1x16, 1x8, 1x4, 1x2, 1x1	0b0101	1Host	1 Upstream Socket	0b0000	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
4C	1x16 Option C	4x4, 4x2, 4x1	1x16, 1x8, 1x4, 2x2, 2x1	0b0100	1Host	1 Upstream Socket	0b0000	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
4C	4x4	4x4, 4x2, 4x1	4x4, 4x2, 4x1	0b0011	1Host	1 Upstream Socket	0b0000	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
RSVD	RSVD	RSVD	RSVD	0b0010	1Host	1 Upstream Socket	0b0000	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
RSVD	RSVD	RSVD	RSVD	0b0001	1Host	1 Upstream Socket	0b0000	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
RSVD	RSVD	RSVD	RSVD	0b0000	1Host	1 Upstream Socket	0b0000	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

Table 27: Bifurcation for Single Host, Four Sockets and Dual Upstream Links
(BIF[2:0]#=#0b010)

Single Host, Four Upstream Sockets, Four Upstream Links										4 x4, 4x2, 4x1															
Card Width	Card Short Name	Supported Modes	Address-Card Encoding PFSM[B:0]#	Host	Upstream Devices	Upstream Links	BIF[2:0]	Resulting Link	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15	
2C	1x8	Card Not Present 1x8, 1x4, 1x2, 1x1	0b1111	1 Host	4 Upstream Sockets	4 Links	0b010	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3													
2C	1x4	1x4, 1x2, 1x1	0b1110	1 Host	4 Upstream Sockets	4 Links	0b010	{Socket 0 only}	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3													
2C	1x4	1x2, 1x1	0b1110	1 Host	4 Upstream Sockets	4 Links	0b010	{Socket 0 only}	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3													
2C	1x2	1x1	0b1110	1 Host	4 Upstream Sockets	4 Links	0b010	{Socket 0 only}	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3													
2C	1x1	1x1	0b1101	1 Host	4 Upstream Sockets	4 Links	0b010	{Socket 0 only}	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3													
2C	1x8 Option B	2x8, 2x4, 2x2, 2x1	0b1101	1 Host	4 Upstream Sockets	4 Links	0b010	4x4	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	2x8 Option B	1x8, 1x4	0b1100	1 Host	4 Upstream Sockets	4 Links	0b010	2x4	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	1x8 Option D	4x2 (First 8 lanes), 4x1	0b1100	1 Host	4 Upstream Sockets	4 Links	0b010	4x4	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	2x8, 2x4	0b1100	1 Host	4 Upstream Sockets	4 Links	0b010	2x4	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3	Link 1, Lane 0	Link 1, Lane 1	Link 1, Lane 2	Link 1, Lane 3									
2C	2x4	2x4, 2x2, 2x1	0b1011	1 Host	4 Upstream Sockets	4 Links	0b010	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3													
RSVD	RSVD	1x4, 1x2, 1x1	0b1010	1 Host	4 Upstream Sockets	4 Links	0b010	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3													
RSVD	RSVD	RSVD for lanes x8 encoding	0b1001	1 Host	4 Upstream Sockets	4 Links	0b010	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3													
RSVD	RSVD	RSVD for lanes x8 encoding	0b1000	1 Host	4 Upstream Sockets	4 Links	0b010	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3													
4C	1x16	1x16, 1x8, 1x4, 1x2, 1x1	0b0111	1 Host	4 Upstream Sockets	4 Links	0b010	{Socket 0 only}	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3													
4C	2x8 Option A	2x8, 2x4, 2x2, 2x1	0b0110	1 Host	4 Upstream Sockets	4 Links	0b010	{Socket 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	0b0101	1 Host	4 Upstream Sockets	4 Links	0b010	{Socket 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	2x8, 2x4, 2x2, 2x1	0b0100	1 Host	4 Upstream Sockets	4 Links	0b010	4x4	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	4x4	4x4, 4x2, 4x1	0b0011	1 Host	4 Upstream Sockets	4 Links	0b010	4x4	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	0b0100	1 Host	4 Upstream Sockets	4 Links	0b010	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3													
RSVD	RSVD	RSVD	0b0011	1 Host	4 Upstream Sockets	4 Links	0b010	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3													
RSVD	RSVD	RSVD	0b0000	1 Host	4 Upstream Sockets	4 Links	0b010	-	Link 0, Lane 0	Link 0, Lane 1	Link 0, Lane 2	Link 0, Lane 3													

Table 29: Bifurcation for Quad Host, Quad Sockets and Quad Upstream Links
(BIF[2:0]#=0b110)

Quad Host, Four Upstream Sockets, Four Upstream Links				BIF[2:0]#=0b110																				
Min Card Width	Card Short Modes	Supported Bifurcation	Add-in-Card Encoding	Host	Upstream Devices	Upstream Links	BIF[2:0]	Resulting Link	Lane 0	Lane 1	Lane 2	Lane 3	Lane 4	Lane 5	Lane 6	Lane 7	Lane 8	Lane 9	Lane 10	Lane 11	Lane 12	Lane 13	Lane 14	Lane 15
2C	1x8	Card Not Present	0b1111	4 Host	4 Upstream Sockets	4 Links	0b110	-	Link 0	Link 0	Link 0	Link 0												
2C	1x4	1x8, 1x4, 1x2, 1x1	0b1110	4 Host	4 Upstream Sockets	4 Links	0b110	1x4 (Host 0 only)	Link 0	Link 0	Link 0	Link 0												
2C	1x4	1x4, 1x2, 1x1	0b1110	4 Host	4 Upstream Sockets	4 Links	0b110	1x4 (Host 0 only)	Link 0	Link 0	Link 0	Link 0												
2C	1x2	1x2, 1x1	0b1110	4 Host	4 Upstream Sockets	4 Links	0b110	1x2 (Host 0 only)	Link 0	Link 0	Link 0	Link 0												
2C	1x1	1x1	0b1110	4 Host	4 Upstream Sockets	4 Links	0b110	1x1 (Host 0 only)	Link 0	Link 0	Link 0	Link 0												
2C	1x8 Option B	1x8, 1x4, 1x2, 1x1	0b1101	4 Host	4 Upstream Sockets	4 Links	0b110	2x4	Link 0	Link 0	Link 0	Link 0												
2C	2x8 Option B	2x8, 2x2, 2x1, 2x1	0b1101	4 Host	4 Upstream Sockets	4 Links	0b110	4x4	Link 0	Link 0	Link 0	Link 0												
4C	2x8 Option B	4x4, 4x2, 4x1	0b1100	4 Host	4 Upstream Sockets	4 Links	0b110	2x4	Link 0	Link 0	Link 0	Link 0												
2C	1x8 Option D	1x8, 1x4, 4x2 (First 8 lanes), 4x1	0b1100	4 Host	4 Upstream Sockets	4 Links	0b110	4x4	Link 0	Link 0	Link 0	Link 0												
4C	1x16 Option D	1x16, 1x8, 1x4, 2x8, 2x4, 4x4, 4x2 (First 8 lanes), 4x1	0b1011	4 Host	4 Upstream Sockets	4 Links	0b110	2x4	Link 0	Link 0	Link 0	Link 0												
RSVD	RSVD	RSVD	0b0110	4 Host	4 Upstream Sockets	4 Links	0b110	-	Link 0	Link 0	Link 0	Link 0												
2C	2x4	1x4, 1x2, 1x1	0b0110	4 Host	4 Upstream Sockets	4 Links	0b110	2x4	Link 0	Link 0	Link 0	Link 0												
RSVD	RSVD	RSVD for lanes x8 encoding	0b0001	4 Host	4 Upstream Sockets	4 Links	0b110	-	Link 0	Link 0	Link 0	Link 0												
RSVD	RSVD	RSVD for lanes x8 encoding	0b0111	4 Host	4 Upstream Sockets	4 Links	0b110	-	Link 0	Link 0	Link 0	Link 0												
4C	1x16	2x8, 2x4, 2x2, 2x1	0b0110	4 Host	4 Upstream Sockets	4 Links	0b110	1x4 (Host 0 only)	Link 0	Link 0	Link 0	Link 0												
4C	2x8 Option A	1x8, 1x4, 1x2, 1x1	0b0110	4 Host	4 Upstream Sockets	4 Links	0b110	2x4	Link 0	Link 0	Link 0	Link 0												
4C	1x16 Option B	2x8, 2x4, 2x1, 2x1	0b0101	4 Host	4 Upstream Sockets	4 Links	0b110	2x8	Link 0	Link 0	Link 0	Link 0												
4C	1x8	1x8, 1x4	0b0100	4 Host	4 Upstream Sockets	4 Links	0b110	4x4	Link 0	Link 0	Link 0	Link 0												
4C	1x16 Option C	2x8, 2x4, 4x2, 4x1	0b0100	4 Host	4 Upstream Sockets	4 Links	0b110	4x4	Link 0	Link 0	Link 0	Link 0												
4C	4x4	4x4, 4x2, 4x1	0b0011	4 Host	4 Upstream Sockets	4 Links	0b110	4x4	Link 0	Link 0	Link 0	Link 0												
RSVD	RSVD	RSVD	0b0010	4 Host	4 Upstream Sockets	4 Links	0b110	-	Link 0	Link 0	Link 0	Link 0												
RSVD	RSVD	RSVD	0b0001	4 Host	4 Upstream Sockets	4 Links	0b110	-	Link 0	Link 0	Link 0	Link 0												
RSVD	RSVD	RSVD	0b0000	4 Host	4 Upstream Sockets	4 Links	0b110	-	Link 0	Link 0	Link 0	Link 0												

3.9 Power Capacity and Power Delivery

There are four permissible power states: NIC Power Off, ID Mode, Aux Power Mode (S5), and Main Power Mode (S0). The transition of these states is shown in Figure 32. The max available power envelopes for each of these states are defined in Table 31.

Figure 32: Baseboard Power States

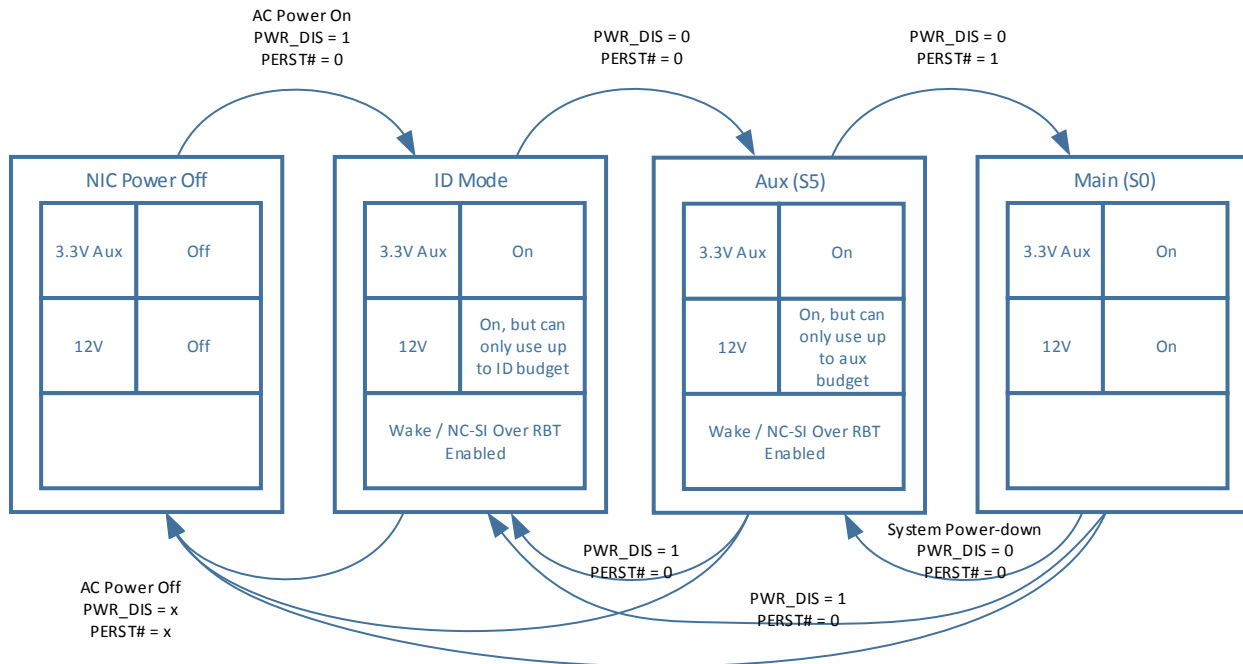


Table 31: Power States

Power State	PWRDIS	PERSTn	FRU	Scan Chain	RBT Link	3.3V	12V
NIC Power Off	Low	Low					
ID Mode	High	Low	X	X		X	X
Aux Power Mode (S5)	Low	Low	X	X	X	X	X
Main Power Mode (S0)	Low	High	X	X	X	X	X

3.9.1 NIC Power Off

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

3.9.2 ID Mode



In the ID Mode, only 3.3V Aux is available for powering up management only functions. FRU accesses are only allowed in this mode. An add-in card shall transition to this mode when PWRDIS=1 and PERST#=0.

3.9.3 Aux Power Mode (S5)

In Aux Power Mode provides both 3.3V Aux as well as 12V Aux is available. 12V Aux may be used to deliver power to the add-in card, but only up to the Aux budget of 35W. An add-in card shall transition to this mode when PWRDIS=0 and PERST#=0.

3.9.4 Main Power Mode (S0)

In Main Power Mode provides both 3.3V and 12V (Main) across the OCP connector. The add-in card operates in full capacity. Up to 80W may be delivered on 12V, and 3.63W on the 3.3V pins. An add-in card shall transition to this mode when PWRDIS=0 and PERST#=1.

3.10 Power Supply Rail Requirements and Slot Power Envelopes

The baseboard provides 3.3Vaux and 12Vaux/main to both the Primary and Secondary connectors. The rail requirements are leveraged from the PCIe CEM 4.0 specification. For OCP NIC 3.0 cards, the requirements are as follows:

Table 32: Baseboard Power Supply Rail Requirements – Slot Power Envelopes

Power Rail	15W Slot Small Card Hot Aisle	25W Slot Small Card Hot Aisle	35W Slot Small Card Hot Aisle	80W Slot Small Card Cold Aisle	150W Large Card Cold Aisle
3.3V					
Voltage Tolerance	±9% (max)	±9% (max)	±9% (max)	±9% (max)	±9% (max)
Supply Current					
ID Mode	375mA (max)	375mA (max)	375mA (max)	375mA (max)	375mA (max)
Aux Mode	1.1A (max)	1.1A (max)	1.1A (max)	1.1A (max)	2.2A (max)
Main Mode	1.1A (max)	1.1A (max)	1.1A (max)	1.1A (max)	2.2A (max)
Capacitive Load	150µF (max)	150µF (max)	150µF (max)	150µF (max)	300µF (max)
12V					
Voltage Tolerance	±8% (max)	±8% (max)	±8% (max)	±8% (max)	±8% (max)
Supply Current					
ID Mode	100mA (max)	100mA (max)	100mA (max)	100mA (max)	100mA (max)
Aux Mode	0.7A (max)	1.1A (max)	1.5A (max)	3.3A (max)	6.3A (max)
Main Mode	1.25A (max)	2.1A (max)	2.9A (max)	6.6A (max)	12.5A (max)
Capacitive Load	500µF (max)	500µF (max)	1000µF (max)	1000µF (max)	2000µF (max)

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

3.11 Hot Swap Considerations for 12V and 3.3V Rails

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

The hotswap controller may gate the 12Vmain/aux and 3.3Vmain/aux based on the PRSNTB[3:0]# value. Per Section 3.6.3, a card is present in the system when the encoded value is not 0b1111. The PRSNTB[3:0]# may be AND'ed together and connected to the hotswap controller to accomplish this result. Per the OCP NIC 3.0 mechanical definition (Section 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 33: Power Sequencing

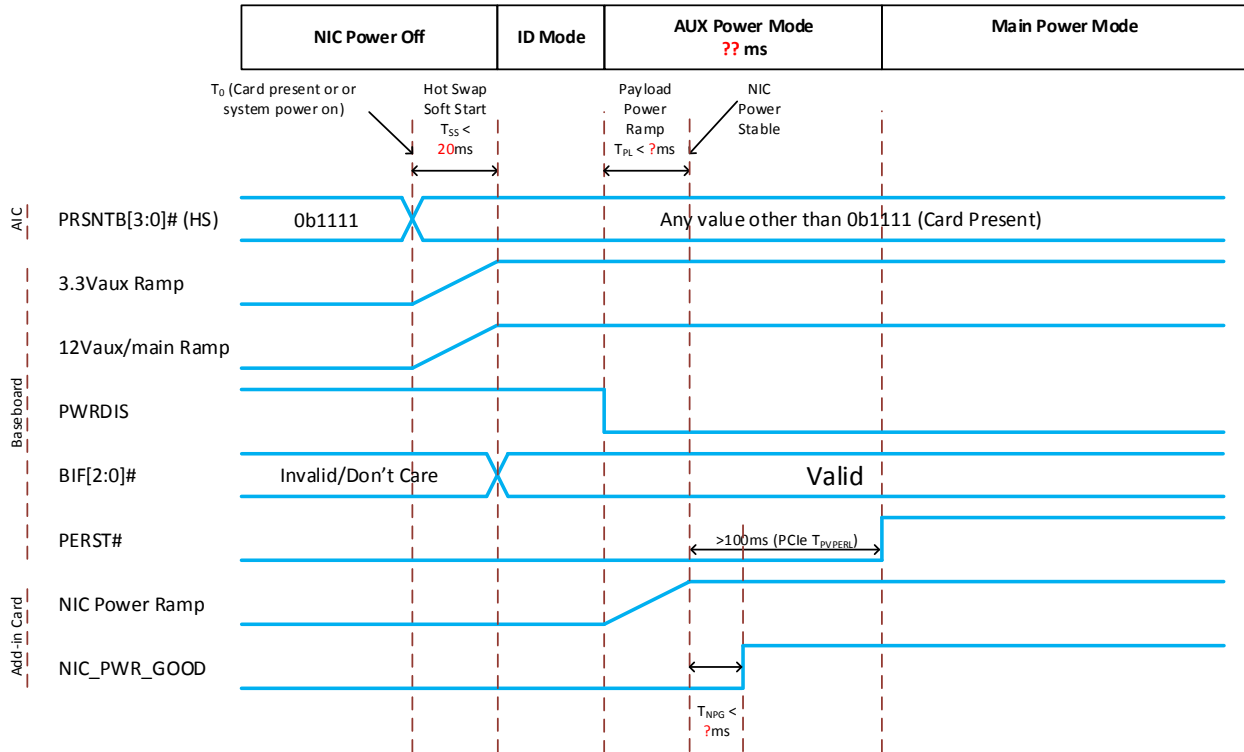


Table 33: Power Sequencing Parameters

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

4 Management

OCP NIC 3.0 card management is an important aspect to overall system management. This section specifies a common set of management requirements for OCP NIC 3.0 implementations. There are two types of implementations (Type 1 and Type 2) depending on the physical sideband management interfaces, transports, and traffic supported over different transports. An OCP NIC 3.0 implementation shall support at least one type of implementation for card management.

4.1 Sideband Management Interface and Transport

OCP NIC 3.0 sideband management interfaces are used by a Management Controller (MC) or Baseboard Management Controller (BMC) to communicate with the NIC. Table 34 summarizes the sideband management interface and transport requirements.

Table 34: Sideband Management Interface and Transport Requirements

Requirement	Type 1	Type 2
NC-SI 1.1 or later compliant RMI Based Transport (RBT) including physical interface defined in Section 10 of DSP0222.	Required	Optional
At least one of the following physical sideband interfaces: <ol style="list-style-type: none"> SMBus 2.0 PCIe VDM 	Required	Required
Management Component Transport Protocol (MCTP) Base 1.3 on at least one of the following physical bindings: <ol style="list-style-type: none"> MCTP/SMBus (DSP0237 1.1) MCTP/PCIe VDM (DSP0238 1.1) 	Required	Required

4.2 NC-SI Traffic

DSP0222 defines two types of NC-SI traffic: Pass-Through and Control. Table 35 summarizes the NC-SI traffic requirements.

Table 35: NC-SI Traffic Requirements

Requirement	Type 1	Type 2
-------------	--------	--------



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

4.3 Management Controller (MC) MAC Address Provisioning

An OCP NIC 3.0 add-in card shall provision one or more MAC addresses for Out-Of-Band (OOB) management traffic. The number of MC MAC addresses provisioned is implementation dependent. These MAC addresses are not exposed to the host(s) as available MAC addresses. The MC is not required to use these provisioned MAC addresses. Table 36 summarizes MC MAC address provisioning requirements.

Table 36: MC MAC Address Provisioning Requirements

Requirement	Type 1	Type 2
One or more MAC Addresses shall be provisioned for the MC.	Required	Optional
Support at least one of the following mechanism for provisioned MC MAC Address retrieval: <ol style="list-style-type: none"> 1. NC-SI Control/RBT (DSP0222 1.1 or later compliant) 2. NC-SI Control/MCTP (DSP0261 1.2 or later compliant) 	Required	Optional

4.4 Temperature Reporting

An OCP NIC 3.0 implementation can have several silicon components including one or more ASICs implementing NIC functions and one or more optical modules providing physical network media connectivity. It is important for the system management that temperatures of these components can be retrieved over sideband interfaces. Table 37 summarizes temperature reporting requirements.

Table 37: Temperature Reporting Requirements

Requirement	Type 1	Type 2
ASIC Temperature Reporting	Required	Required
Optical Modules Temperature Reporting	Optional	Optional

Support at least one of the following mechanisms for ASIC temperature reporting: <ol style="list-style-type: none"> 1. NC-SI Control (DSP0222 1.1 or later compliant) 2. PLDM for Platform Monitoring and Control (DSP0248 1.1 compliant) 3. TMP421 emulation over SMBus 2.0 	Required	Required
Support at least one of the following mechanisms for optical modules temperature reporting: <ol style="list-style-type: none"> 4. NC-SI Control (DSP0222 1.1 or later compliant) 5. PLDM for Platform Monitoring and Control (DSP0248 1.1 compliant) 	Optional	Optional

4.5 Power Consumption Reporting

An OCP NIC 3.0 implementation may be able to report the power consumed by ASICs implementing NIC functions. It is important for the system management that the information about the power consumption can be retrieved over sideband interfaces. Table 38 summarizes power consumption reporting requirements.

Table 38: Power Consumption Reporting Requirements

Requirement	Type 1	Type 2
ASIC Power Consumption Reporting	Optional	Optional
Support at least one of the following mechanisms for ASIC power consumption reporting: <ol style="list-style-type: none"> 1. NC-SI Control (DSP0222 1.1 or later compliant) 2. PLDM for Platform Monitoring and Control (DSP0248 1.1 compliant) 	Optional	Optional

4.6 Link Status/Speed Reporting

Link status/speed reporting is important for network operations and link management. Table 39 summarizes link status and speed reporting requirements.

Table 39: Link Status/Speed Reporting Requirements

Requirement	Type 1	Type 2
Link Status Reporting	Required	Required
Support at least one of the following mechanisms for reporting the link status: <ol style="list-style-type: none"> 1. NC-SI Control (DSP0222 1.1 compliant) 2. PLDM for Platform Monitoring and Control (DSP0248 1.1 compliant) 	Required	Required
Link Speed Reporting	Required	Required
Support at least one of the following mechanisms for reporting the link speed: <ol style="list-style-type: none"> 3. NC-SI Control (DSP0222 1.1 compliant) 4. PLDM for Platform Monitoring and Control (DSP0248 1.1 compliant) 	Required	Required

4.7 Pluggable Module Status Reporting

Pluggable modules like optical modules or direct attach cables are used to connect OCP NIC to physical media. It is important to know the presence of pluggable modules and information about insertion/deletion of pluggable modules. Table 40 summarizes pluggable module status reporting requirements.

Table 40: Pluggable Module Status Reporting Requirements

Requirement	Type 1	Type 2
Pluggable Module Presence Reporting	Optional	Optional
Support at least one of the following mechanisms for reporting the pluggable module presence status: <ol style="list-style-type: none"> 1. NC-SI Control (DSP0222 1.1 or later compliant) 2. PLDM for Platform Monitoring and Control (DSP0248 1.1 compliant) 	Optional	Optional

Pluggable Module Insertions/Deletions Reporting	Optional	Optional
Support at least one of the following mechanisms for reporting the pluggable module insertions/deletions: <ul style="list-style-type: none"> 3. NC-SI Control (DSP0222 1.1 or later compliant) 4. PLDM for Platform Monitoring and Control (DSP0248 1.1 compliant) 	Optional	Optional

4.8 Out-Of-Band Firmware Update

An OCP implementation can have different types of firmware components for data path, control path, and management path operations. It is desirable that OCP NIC 3.0 implementations support an OS-independent Out-Of-Band mechanism for the firmware update. Table 41 summarizes out-of-band firmware update requirements.

Table 41: Out-Of-Band Firmware Update Requirements

Requirement	Type 1	Type 2
Support PLDM for Firmware Update (DSP0267 1.0 compliant)	Optional	Optional

4.9 NC-SI Over RBT Sideband Interface

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

4.9.1 NC-SI Over RBT Addressing

NC-SI Over RBT capable devices must use a unique Package ID to ensure there are no addressing conflicts.

Baseboards use the Slot_ID pin on the Primary Connector for this identification. The Slot_ID value may be directly connected to GND (Slot ID = 0), or pulled up to 3.3Vaux (Slot ID = 1).

For add-in cards, Package ID[2:0] is a 3-bit field and is encoded in the NC-SI Channel ID as bits [7:5]. Package ID[2] defaults to 0b0 in the NC-SI specification, but is optionally



configurable if the target silicon supports configuring this bit. Package ID[1] is connected to the SLOT_ID pin and is directly connected to the Slot_ID pin. Package ID[0] is set to 0b0 for Network Silicon #0. For OCP NIC 3.0 add-in cards with two discrete silicon instances, Package ID[1] is set to 0b1. Refer to the endpoint device datasheet for details on the Package ID configuration options.

Up to four silicon devices are supported on the bus if only Package ID[1:0] is configurable (e.g. Package ID[2] is statically set to 0b0). Up to eight silicon devices are supported on the NC-SI bus if Package ID[2:0] are all configurable. Refer to the DMTF DSP0222 standard for more information.

4.9.2 Arbitration Ring Connections

For baseboards that implement two or more Primary Connectors, the NC-SI Over RBT arbitration ring may be connected to each other. The arbitration ring must support operation with a one card, or both cards installed. Figure 22 shows an example connection with dual Primary Connectors.

4.10 SMBus 2.0 Interface

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

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

4.10.1 SMBus Address Map

All predefined SMBus addresses for OCP NIC 3.0 are shown in Table 42. Baseboard and add-in card designers must ensure additional devices do not conflict. The addresses shown are in 8-bit format and represent the read/write address pair.

Table 42: Power Sequencing Parameters

Address (8-bit)	Device	Notes
-----------------	--------	-------

0xTBD	Network Controller IC	Value dependent on NIC vendors.
0x98 / 0x99	Temperature Sensor	TMP422/423 Temperature sensor Optional. Used for remote on-die thermal sensing. Optional. Powered from Aux power domain.
0x9E / 0x9F	Temperature Sensor	Emulated TMP421 Temperature sensor. Optional. Thermal reporting is emulated from the target device. The communication interface is over SMBus and is compliant to the TMP421 register definition.
0x9E / 0x9F	Temperature Sensor	TMP421 Temperature sensor. Optional. Used for remote on-die thermal sensing. Optional. Powered from Aux power domain.
0xA2 / 0xA3	EEPROM	On-board FRU EEPROM. Mandatory. Powered from Aux power domain.

4.11 FRU EEPROM

4.11.1 FRU EEPROM Address, Size and Availability

The FRU EEPROM provided for the baseboard to determine the card type and is directly connected to the SMBus on the card edge. Only one EEPROM is required for a single physical add-in card regardless of the PCIe width or number of physical card edge connectors it occupies. The FRU EEPROM shall be connected to the Primary connector SMBus.

The EEPROM is addressable at 0xA2/0xA3 for the write/read pair in 8-bit format. The size of EEPROM is **1Kbits** for the base EEPROM map. Add-in card suppliers may use larger size EEPROM if needed to store vendor specific information.

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



4.11.2 FRU EEPROM Content Requirements

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

Table 43: FRU EEPROM Record – OEM Record 0xC0, Offset 0x00

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

Table 44: FRU EEPROM Record – OEM Record 0xC0, Offset 0x01

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

Table 45: FRU EEPROM Record – OEM Record 0xC0, Offset 0x02

Offset 2	Secondary Connector PRSNTB [3:0]#
0b1110 (0x0E)	Follows Pinout; to be filled after the pinout table is fixed

0b1101 (0x0D)	
0b1100 (0x0C)	
0b1010 (0x0A)	
0b0111 (0x07)	
0b0110 (0x06)	
0b0101 (0x05)	
0b0100 (0x04)	
0b0011 (0x03)	
0b1011 (0x0B)	Not a valid reading – Wrong EEPROM programming
0b1111 (0x0F)	Not a valid reading – Wrong EEPROM programming
All others	RFU
No FRU device detected	No NIC connected / bad connection

Table 46: FRU EEPROM Record – OEM Record 0xC0, Offset 0x03

Offset 3	Card max power in Aux(S5)
0x01 - 0xFE	Hex format in Watts when NIC is in AUX(S5) mode; LSB = 1x Watt; roundup to the nearest Watt for fractional values.
0xFF	Invalid entry
0x00	Invalid entry

Table 47: FRU EEPROM Record – OEM Record 0xC0, Offset 0x04

Offset 4	Card max power in Main(S0)
0x01 - 0xFE	Hex format in Watts when NIC is in Main (S0) mode; LSB = 1x Watt; roundup to the nearest Watt for fractional values.
0xFF	Invalid entry
0x00	Invalid entry

Table 48: FRU EEPROM Record – OEM Record 0xC0, Offset 0x05

Offset 5	Thermal Reporting Interface
0x01	Emulated thermal reporting on SMBus



0x02	Remote on-die sensor with TMP421 on SMBus
0x04	PLDM thermal reporting via NC-SI over RBT

4.12 FW Requirements

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

4.12.1 Firmware Update

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

4.12.2 Secure Firmware

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

4.12.3 Firmware Queries

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

4.12.4 Multi-Host Firmware Queries

- A multi-host capable OCP NIC 3.0 add-in card shall gracefully handle concurrent in-band queries from multiple hosts and out-of-band access from the BMC the management status and firmware, device model, and device ID information.
- A multi-host capable OCP NIC 3.0 add-in card shall only permit one entity to perform write accesses to NIC firmware at a time, without creating contention.
- A multi-host capable OCP NIC 3.0 add-in card shall gracefully handle exceptions when more than one entity attempts to perform concurrent NIC firmware writes.

4.13 Thermal Reporting Interface

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

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

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

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

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

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

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

4.13.1 Emulated Thermal Reporting

Emulated Thermal Reporting requires each OCP NIC 3.0 compliant device to emulate its key temperatures following the TI/TMP421 (or equivalent) register mapping¹. The slave address 0x3E/0x3F as the read/write pair in 8-bit format over the Primary Connector SMBus.

The baseboard will threat the thermal sensor as a TMP421. The baseboard BMC controller must use two separate reads to obtain the MSB and LSB of temperature data. This information is used for system thermal monitoring and fan speed control.

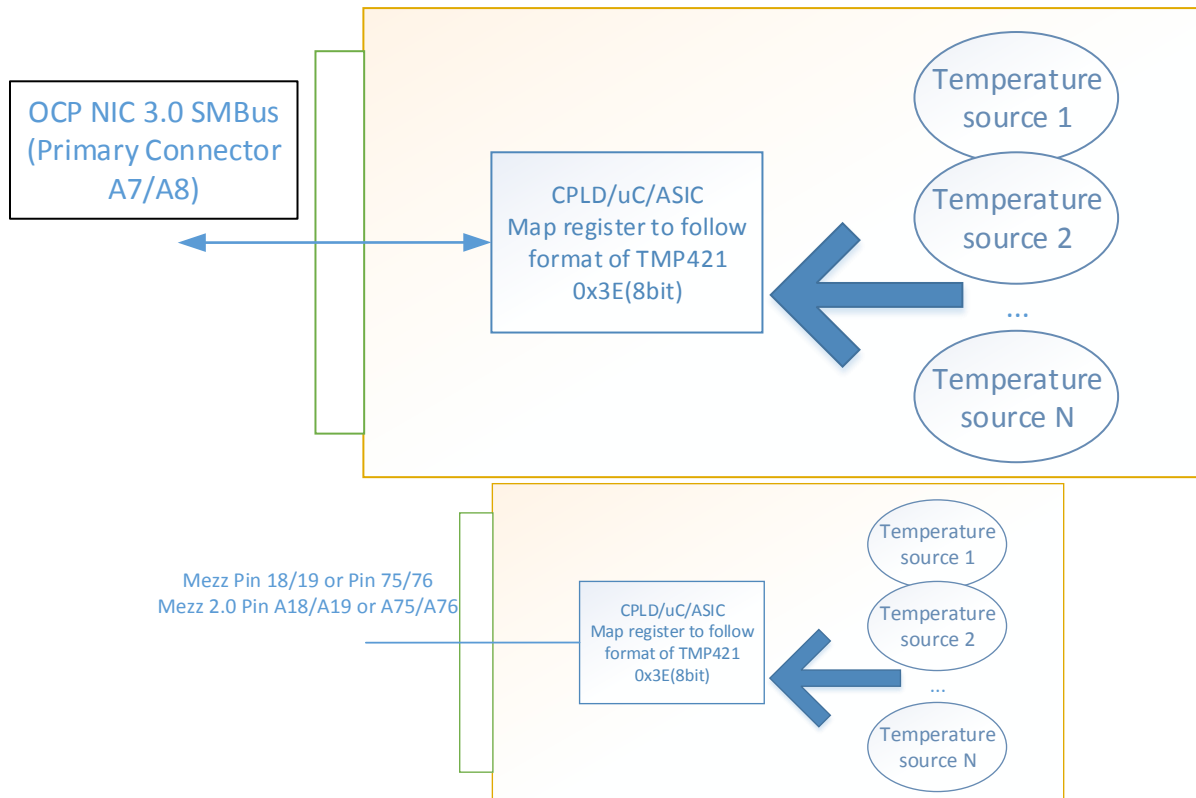
There are two temperatures for TMP421 register mapping – remote channel 1 and local. Remote channel 1 is typically used to represent key controller temperature of the card. This measures the temperature using a remote diode. The local channel is typically used to

¹ TMP421 specification: <http://www.ti.com/lit/ds/sbos398c/sbos398c.pdf>

represent highest of other key components temperature on the card, such as highest temperature of active cable module.

An implementation block diagram is shown in **Figure 34**.

Figure 34: Block Diagram for Emulated Thermal Reporting



If an additional temperature sensitive device needs monitoring, the emulated controller can be changed to a TMP422/TMP423 in addition to the register map. The TMP422/TMP423 slave address of emulated device is always 0x3E/0x3F as a read/write 8-bit address pair.

The vendor ID and device ID are mapped to offset 0xFE and 0xFF for the BMC to detect card types.

Power reporting and power capping are mapped to offset 0xF2 and 0xF3 as an optional feature to achieve device power monitoring and power capping level setting.

Table 49: describes the register implementation requirement for emulated method.

Table 49: Implementation Requirement for TMP421 Registers


Offset	Description	Original TMP offset	Implementation requirement for emulated method
0x0	Local Temperature (High Byte)	Y	Represents highest temperature of all other key components Required if any of the other key components or modules are critical for thermal design Otherwise it is an optional offset and return 0x00 if not used
0x1	Remote Temperature 1 (High Byte)	Y	Required; represent temperature of main controller
0x2	Remote Temperature 2 (High Byte)	Y	Optional; represent temperature of key component 1; return 0x00 if not used
0x3	Remote Temperature 3 (High Byte)	Y	Optional; represent temperature of key component 2; return 0x00 if not used
0x8	Status Register	Y	Not required
0x9	Configuration Register 1	Y	Not required; Emulated behavior follows SD=0, Temperature Range=0
0x0A	Configuration Register 2	Y	Required; follow TMP423 datasheet to declare the channel supported; RC=1
0x0B	Conversion Rate Register	Y	Not required; Equivalent emulated conversion rate should be >2 sample/s
0x0F	One-Shot Start	Y	Not required
0x10	Local Temperature (Low Byte)	Y	Optional; return 0x00 if not used
0x11	Remote Temperature 1 (Low Byte)	Y	Optional; return 0x00 if not used

0x12	Remote Temperature 2 (Low Byte)	Y	Optional; return 0x00 if not used
0x13	Remote Temperature 3 (Low Byte)	Y	Optional; return 0x00 if not used
0x21	N Correction 1	Y	Not required
0x22	N Correction 2	Y	Not required
0x23	N Correction 3	Y	Not required
0xF0	Manufacturer ID(High Byte)	N	High byte of PCIe vendor ID, if using emulated temperature sensor method
0xF1	Device ID(High Byte)	N	High byte of PCIe device ID, if using emulated temperature sensor method
0xF2	Power reporting	N	Optional; card power reporting; 1LSB=1W; Read only
0xF3	Power capping	N	Optional; card power capping; 1LSB=1W; Read/Write
0xFC	Software Reset	Y	Not required
0xFE	Manufacturer ID	Y(redefined)	Low byte of PCIe vendor ID, if using emulated temperature sensor method
0xFF	Device ID	Y(redefined)	Low byte of PCIe device ID, if using emulated temperature sensor method

4.13.2 Remote on-die sensing

Alternatively, one TMP421 sensor can be used to do on die temperature sensing for IC with thermal diode interface with TMP421 remote sensing channel; Connection diagram is shown in Figure 35. For add-in cards that require more than one remote on-die sense point, a TMP422/TMP423 can be used and slave address is 0x98/0x99 (8-bit).

Figure 35: Block Diagram for Remote on-die Sensing





4.13.3 PLDM Method

Placeholder; needs other editors' help.

4.13.4 Thermal reporting accuracy

The recommended accuracy for temperature sensors on the card is $\pm 3^{\circ}\text{C}$

5 Data Network Requirements

5.1 Network Boot

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

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

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

EFI_FIRMWARE_UPDATE_PROTOCOL



6 Routing Guidelines and Signal Integrity Considerations

6.1 NC-SI Over RBT

For the purposes of this specification, the min and max electrical trace length of the NC-SI signals shall be between 2 inches and 4 inches. The traces shall be implemented as 50 Ohm impedance controlled nets.

6.2 PCIe

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

Add-in card suppliers shall follow the PCIe routing specifications.

At this time, the OCP NIC subgroup is working to identify and agree to the channel budget for an add-in card and leave sufficient margin for the baseboard.

Refer to the PCIe CEM and PCIe Base specifications for end-to-end channel signal integrity considerations.

7 Thermal and Environmental

7.1 Environmental Requirements

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

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

7.1.1 Thermal Reporting interface

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

7.1.2 Thermal Simulation Boundary Example

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

7.2 Shock & Vibration

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

7.3 Regulatory

An OCP NIC 3.0 add-in card shall meet CE, CB, FCC Class A, WEEE, and RoHS requirements.

