

Open. Together.



OCP
SUMMIT




Executive Track

Project Olympus 3U PCIe Expansion Server Deep Dive

Raymond Miles, VP Architecture, ZT Systems

Mark D. Chubb, Director Architecture ZT Systems



OPEN
GOLD[®]



Open. Together.

Agenda

ZT Systems Overview

Project Olympus Building Blocks

ZT 3U PCIe Expansion Server Products

Chassis

Flexible Riser Card Options

Power Distribution Board


Slide Rails

ZT Systems Overview



- **Leading provider of optimized servers and storage servers for hyperscale data centers**
 - Founded in 1994
 - HQ in Secaucus, NJ; global sites & capabilities
 - Private company
- **Built for Hyperscale**
 - Maximum design flexibility
 - Continuity of Supply
 - Industry leading achievement to aggressive delivery SLAs
 - Full range of flexible service and support options
 - Broad range of component options

A solution design and manufacturing partner with a record of success serving the world's largest data center customers




OPEN
Compute Project
SOLUTION PROVIDER®



Open. Together.

ZT OCP Engagement

- History
 - Engaged since 2012
 - Gold Member since 2014
 - Engineering projects
 - Solution Provider
 - Sponsored 8 summits
- Products
 - 2U Open Storage (2014)
 - 1U XPO200 (2017)
 - XPO200-3UN PCIe Expansion System (2019)



Project Olympus HW Building Blocks

- **Motherboards:** Flexible expansion
- **Power Supply:** 1KW three-phase
- **Universal PMDU:** AC power cord adaption
- **Rack:** 19” EIA-310 standard
- **Rack Manager:** Restful API or Redfish via Ethernet

Project Olympus
is an
OCP Accepted™
Specification




OPEN
ACCEPTED™

3U PCIe Expansion Server



SERVER

Project Olympus Rack Fully Compatible



3 x Load Balanced 1kW PSUs

Supports Project Olympus motherboards


Power Distribution Board (PDB)

PSUs Blind mate into Rack PMDU

6 x 60mm Dual Rotor Fans (room for 8 Fans)

2.5"/3.5" SSD or HDD

Supports Slide Rail Kit for smoother travel



5 x FHFL, Double-wide 300W x16 PCIe Cards + 1 x FHHL Single-wide 75W x16 PCIe Card
AMD® SKU

OR

6 x FHFL, Double-wide 300W x16 PCIe Cards + 1 x FHHL Single-wide 75W x16 PCIe Card

12 x FHFL, Single-wide 75W x16 PCIe Cards + 1 x FHHL Single-wide 75W x16 PCIe Card

Intel® SKU



Open. Together.

ZT Solution: 3U PCIe Expansion Server

(AMD® EPYC™ Based; ZT-XPO200-3UA)

Feature	Qty	Description
Chassis	1	3U, 19" EIA310-D Compliant supporting Project Olympus PMDU connections
Motherboard	1	Project Olympus 2-Slot AMD® EPYC™ Motherboard (OCP leveraged) from ZT Systems
Processor	2	AMD® EPYC™ 7551 (Naples), 180W, 32C, 2GHz
Memory	16 (32max)	32GB DDR4, DR (1DPC), 2667 R-DIMM; Total System Memory: 512GB
PCIe Riser 3 & 5	2	2-Slot Active PCIe x16 Riser Card with 48-lane PCIe Switch
PCIe Riser 4	1	2-Slot Passive PCIe x32 Riser Card
GPU Card	4 (5 max)	AMD® MI25 FHFL, Double-Wide, 300W PCIe x16 Card
Network Card	1	10G Single Port SFP+ PCIe2.0 x8 5GT/s
HDD/SSD	5 (8 max)	M.2 960GB NVMe SSD, PCIe x4 110mm
M.2 Riser	1 (2 max)	2-Slot M.2 Riser (in PCIe Slot 1)
Security	1	TPM2.0 SPI Module
System Fans	6	60mmx56mm Dual Rotor Fan
Power Supply	3	Project Olympus 1020W 3-Phase, non-LES PSU
Power Distribution	1	PDB and Cable Harnesses to support 12V Power to MB, Risers, GPUs, and System Fans

Items in Orange are being provided to OCP by ZT Systems



Open. Together.

ZT Solution: 3U PCIe Expansion Server

(Intel® Xeon®-SP Based; ZT-XPO200-3UN)

Feature	Qty	Description
Chassis	1	3U, 19" EIA310-D Compliant supporting Project Olympus PMDU connections
Motherboard	1	Project Olympus 2-Slot Intel® Xeon®-SP Motherboard
Processor	2	Intel® Xeon® Platinum 8168 (Skylake), 205W, 24C, 2.7GHz
Memory	12 (24max)	32GB DDR4, DR (1DPC), 2667 R-DIMM; Total System Memory: 384GB
PCIe Riser 3 & 5	2	4-Slot Active PCIe x16 Riser Card with 96-lane PCIe Switch
PCIe Riser 4	1	5-Slot Active PCIe x16 Riser Card with 96-lane PCIe Switch & Sliver to Edge Conn x8 Cables
GPU Card	12	Nvidia® P4 Low Profile, Single-Wide, 75W PCIe x16 Card
Network Card	1	10G Single Port SFP+ PCIe2.0 x8 5GT/s
HDD/SSD	1 (4 max)	M.2 960GB NVMe SSD, PCIe x4 110mm
Security	1	TPM2.0 SPI Module
System Fans	6	60mmx56mm Dual Rotor Fan
Power Supply	3	Project Olympus 1020W 3-Phase, non-LES PSU
Power Distribution	1	PDB and Cable Harnesses to support 12V Power to MB, Risers, GPUs, and System Fans

Items in Orange are being provided to OCP by ZT Systems




Open. Together.

3U PCIe Expansion System Options



ZT Systems: ZT-XPO200-3UA



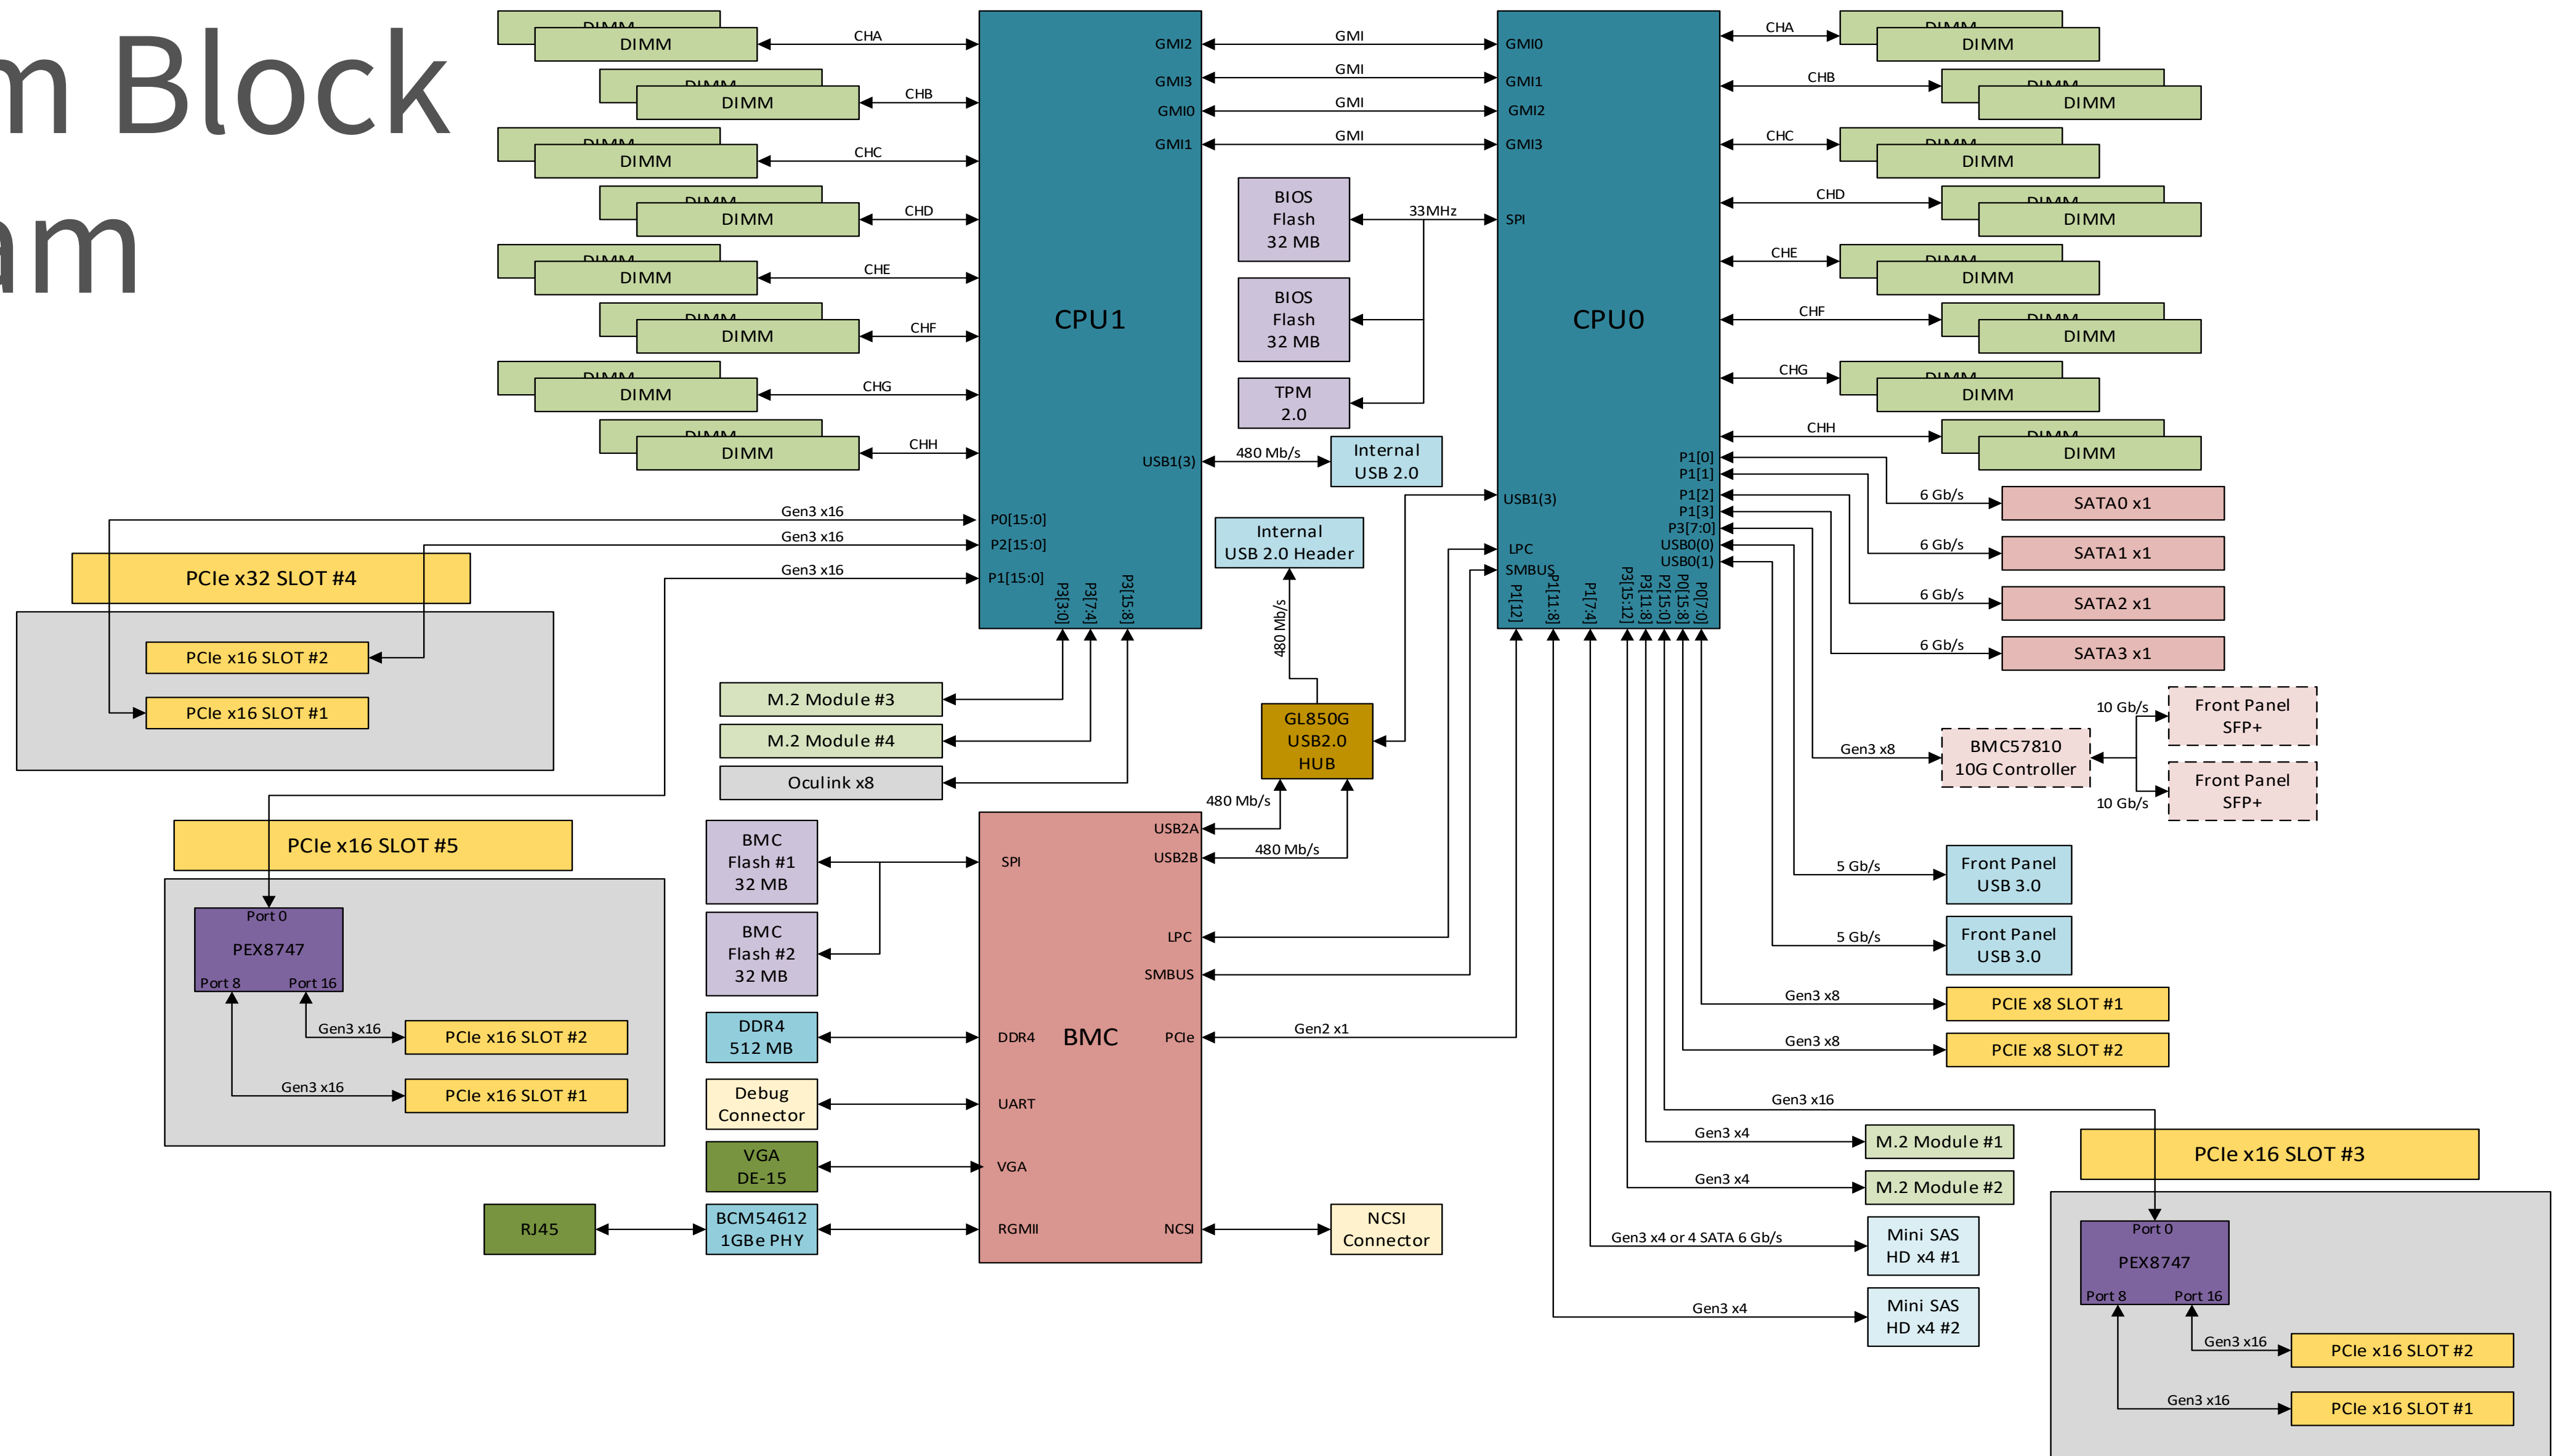
ZT Systems: ZT-XPO200-3UN

ZT Systems in process of submitting these 3U Server Design Specifications to OCP Community!



Open. Together.

System Block Diagram



2-Slot Active PCIe x16 Riser Card Layout

2 x PCIe x16 slots supporting (Slots #1-2, Bottom to Top):

- 2 x FHFL Double-Wide PCIe cards
- Requires PWRBRK_N (pin B30) signal for High Power PCIe Card throttling

Broadcom® 48-Lane PCIe Gen3 Switch (PEX8747)


- Switch Bifurcated into 2 x16 Ports (1 x16 Port to each x16 Slot)
- Switch is located on Backside as required by Chassis Thermals (Heatsink is required)
- Use of all 48-lanes and requires 15.5W

Differential Clock Buffer on board for additional Clock Fanout

I2C MUX used to avoid Address contention (see I2C block diagram)

8 Layer Stackup, Mid-Loss Material (0.062")

Enclosed in a Riser Cage for easy Serviceability and Retention



2-Slot Passive PCIe x32 Riser Card Layout

For Project Olympus AMD® EPYC™ MB Only (Middle Riser)

Uses 2 x HSEC8 connectors in-line for x32 Lanes

1 x PCIe x16 slot supporting (Slot #1, Bottom):


- 1 x FHHL Single-Wide PCIe card
- Designated for the Network Card

1 x PCIe x16 slot supporting (Slot #2, Top):

- 1 x FHFL Single-Wide PCIe Card OR
- 1 x FHFL Double-Wide PCIe Card
- All x16 PCIe lanes sourced from PCIe Riser Slot #4 (No Switch Required)
- Requires PWRBRK_N (pin B30) signal for High Power PCIe Card throttling

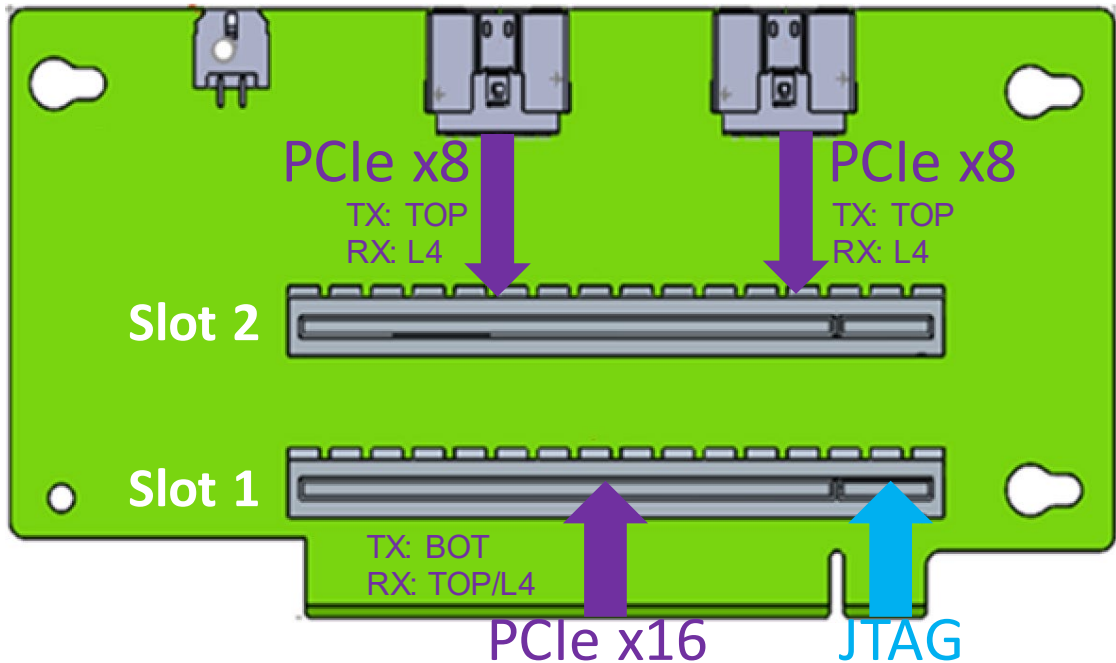
I2C MUX used to avoid Address contention (see I2C block diagram)

8 Layer Stackup, Mid-Loss Material (0.062")

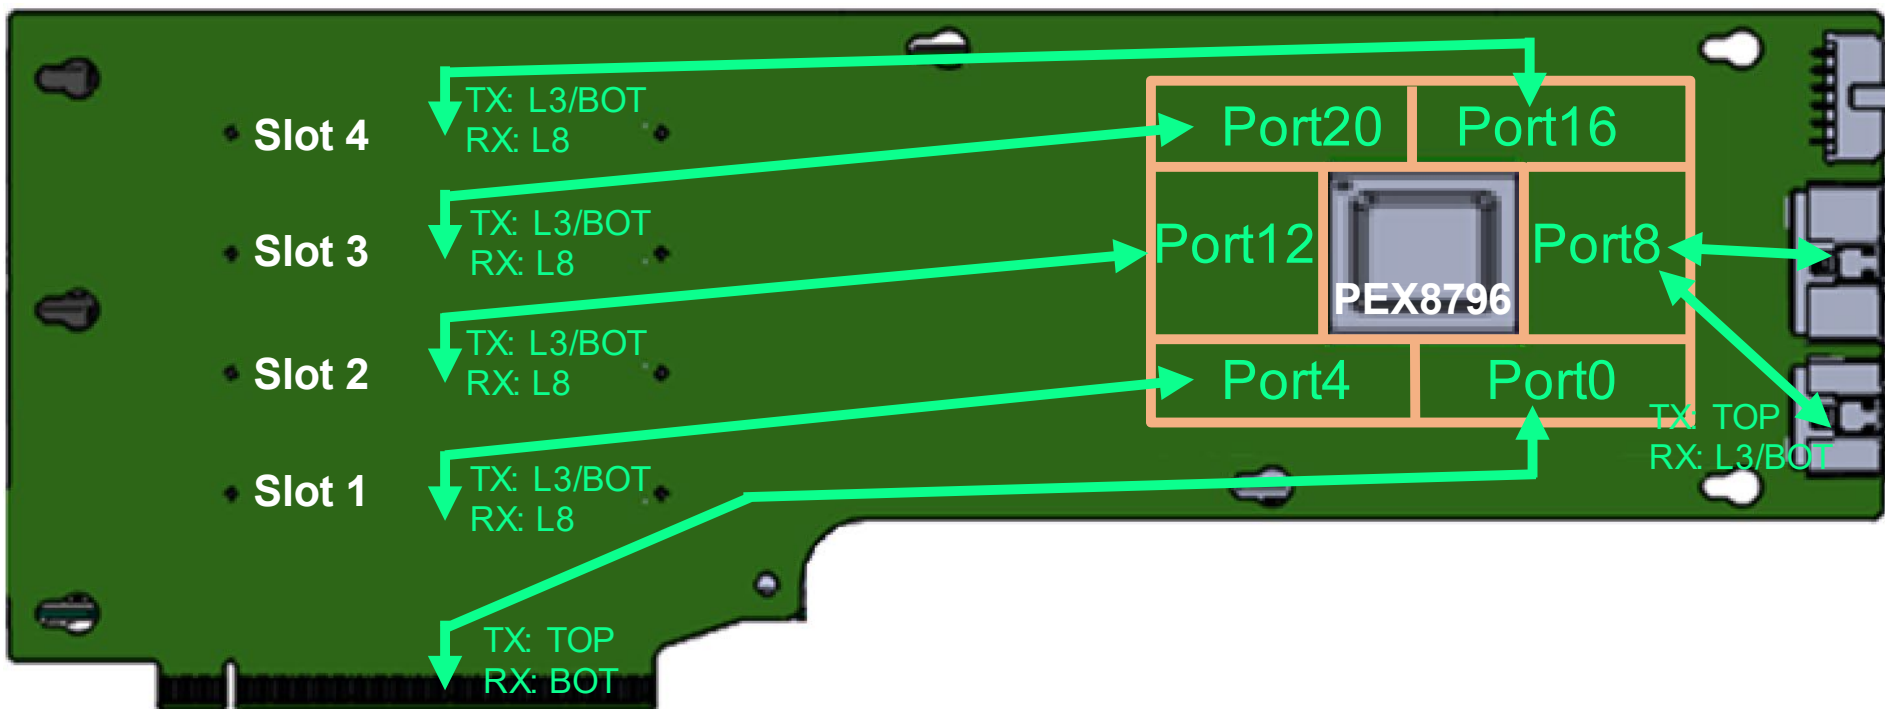


Other 3U Server PCIe Riser Card Options


2-Slot Passive PCIe x16 Riser Card




4-Slot Active PCIe x16 Riser Card



3-Slot Active PCIe x16 Riser Card




5-Slot Active PCIe x16 Riser Card



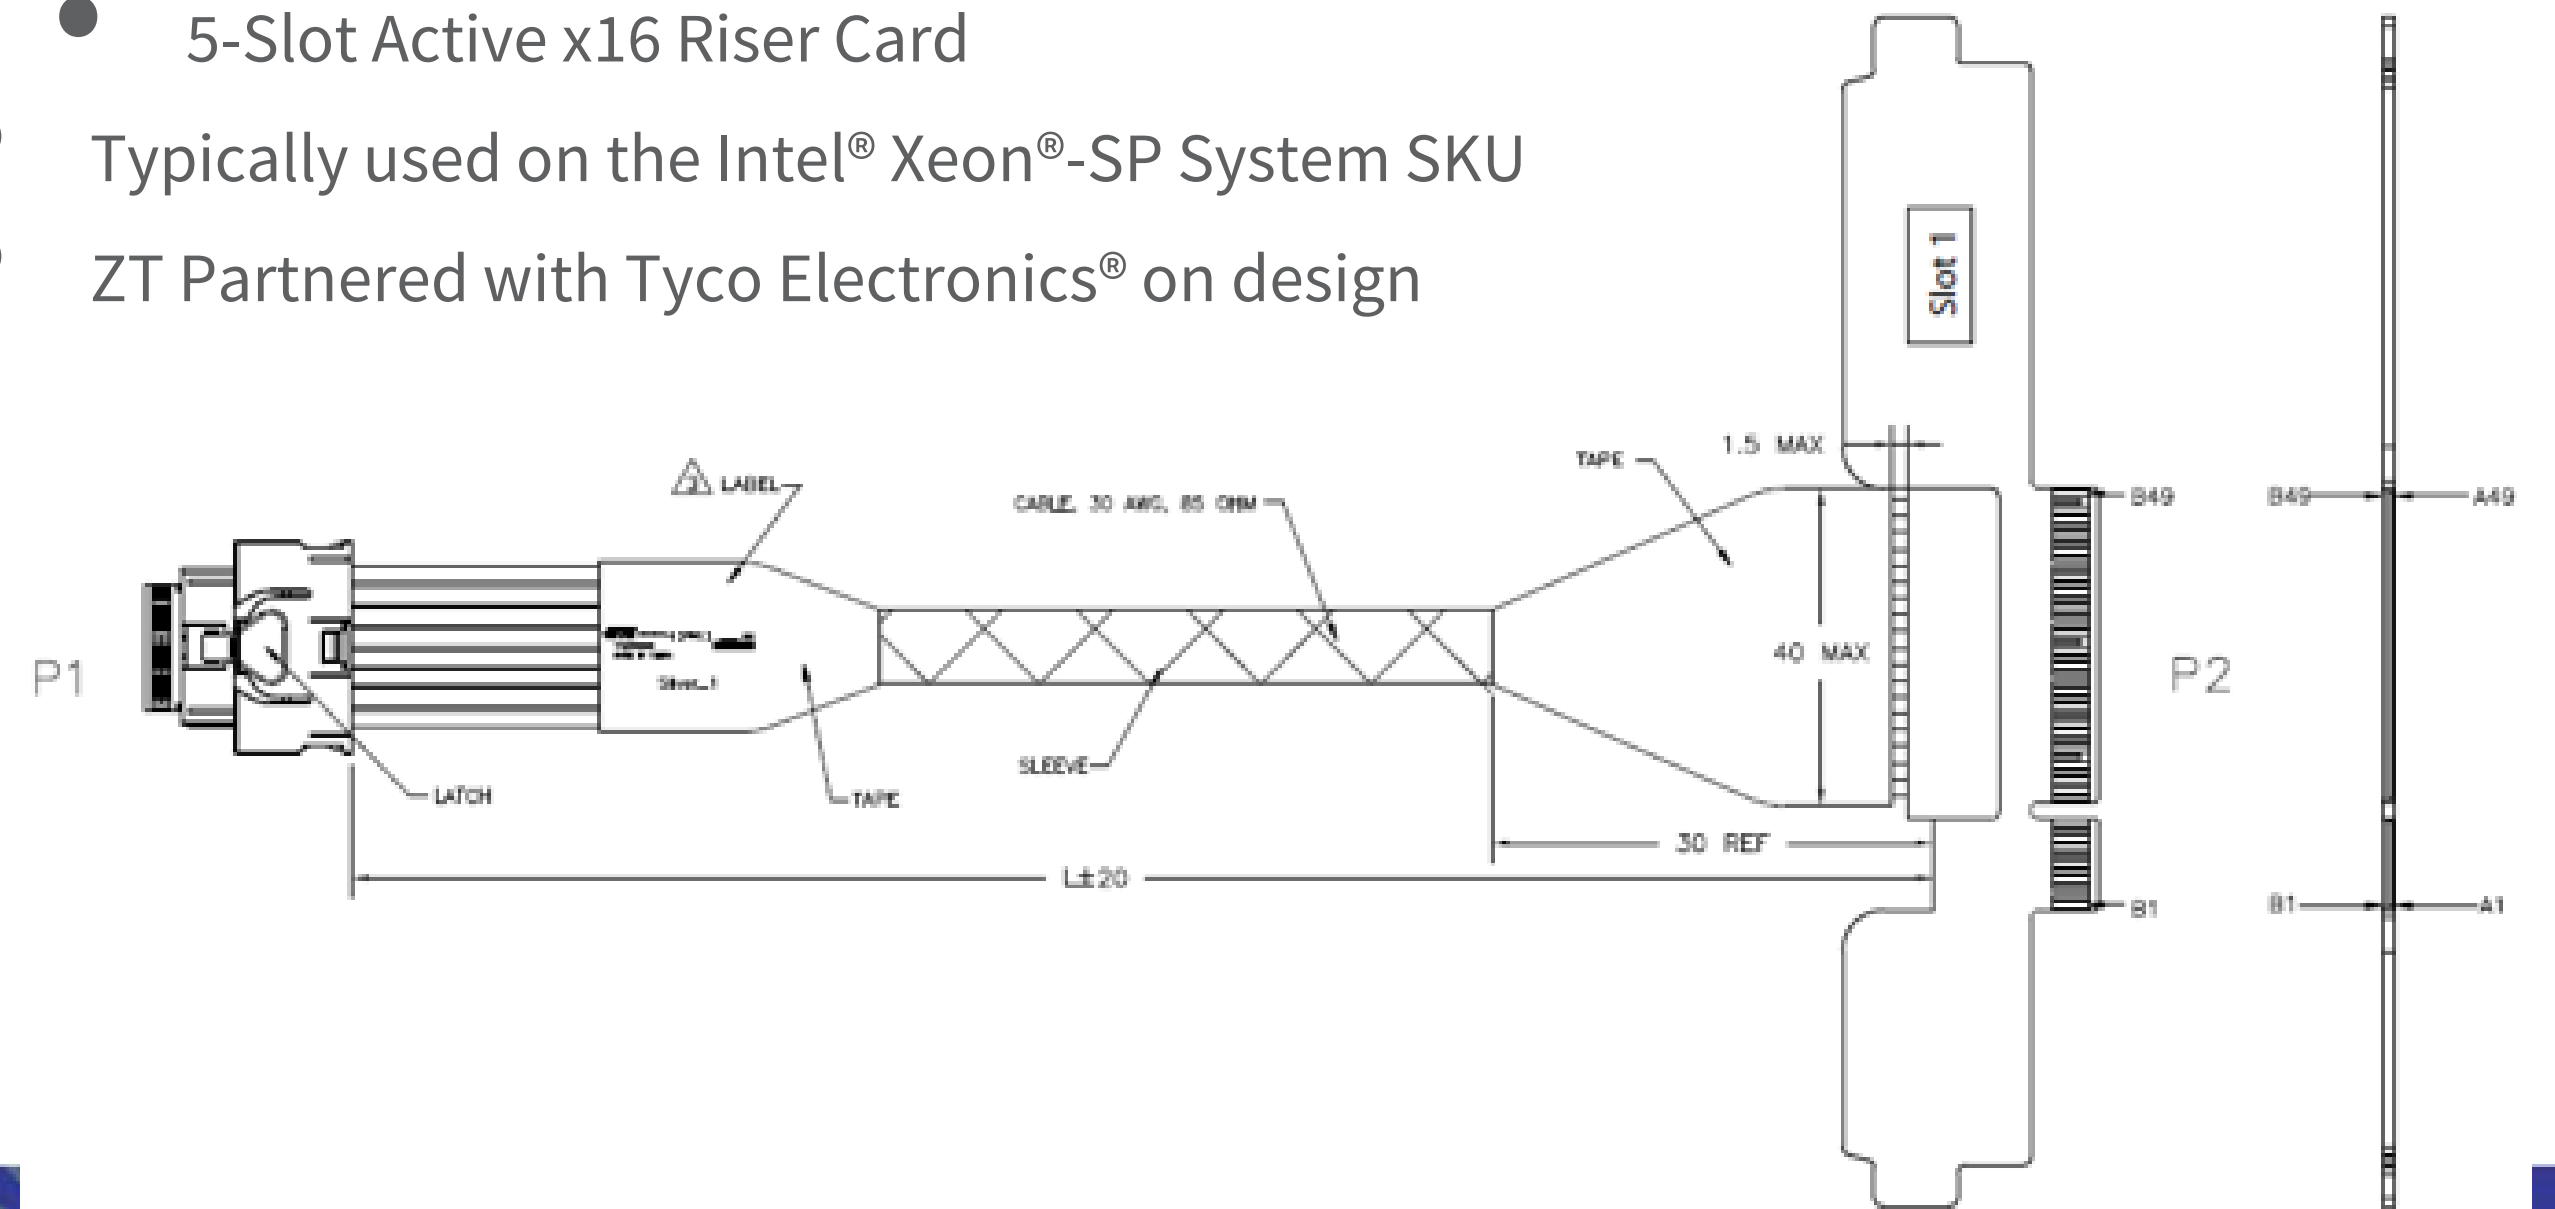
Provides PCIe Expansion System Flexibility!

Middle PCIe Riser Card connections



Sliver Connector to PCIe x8 Edge Connector Cable

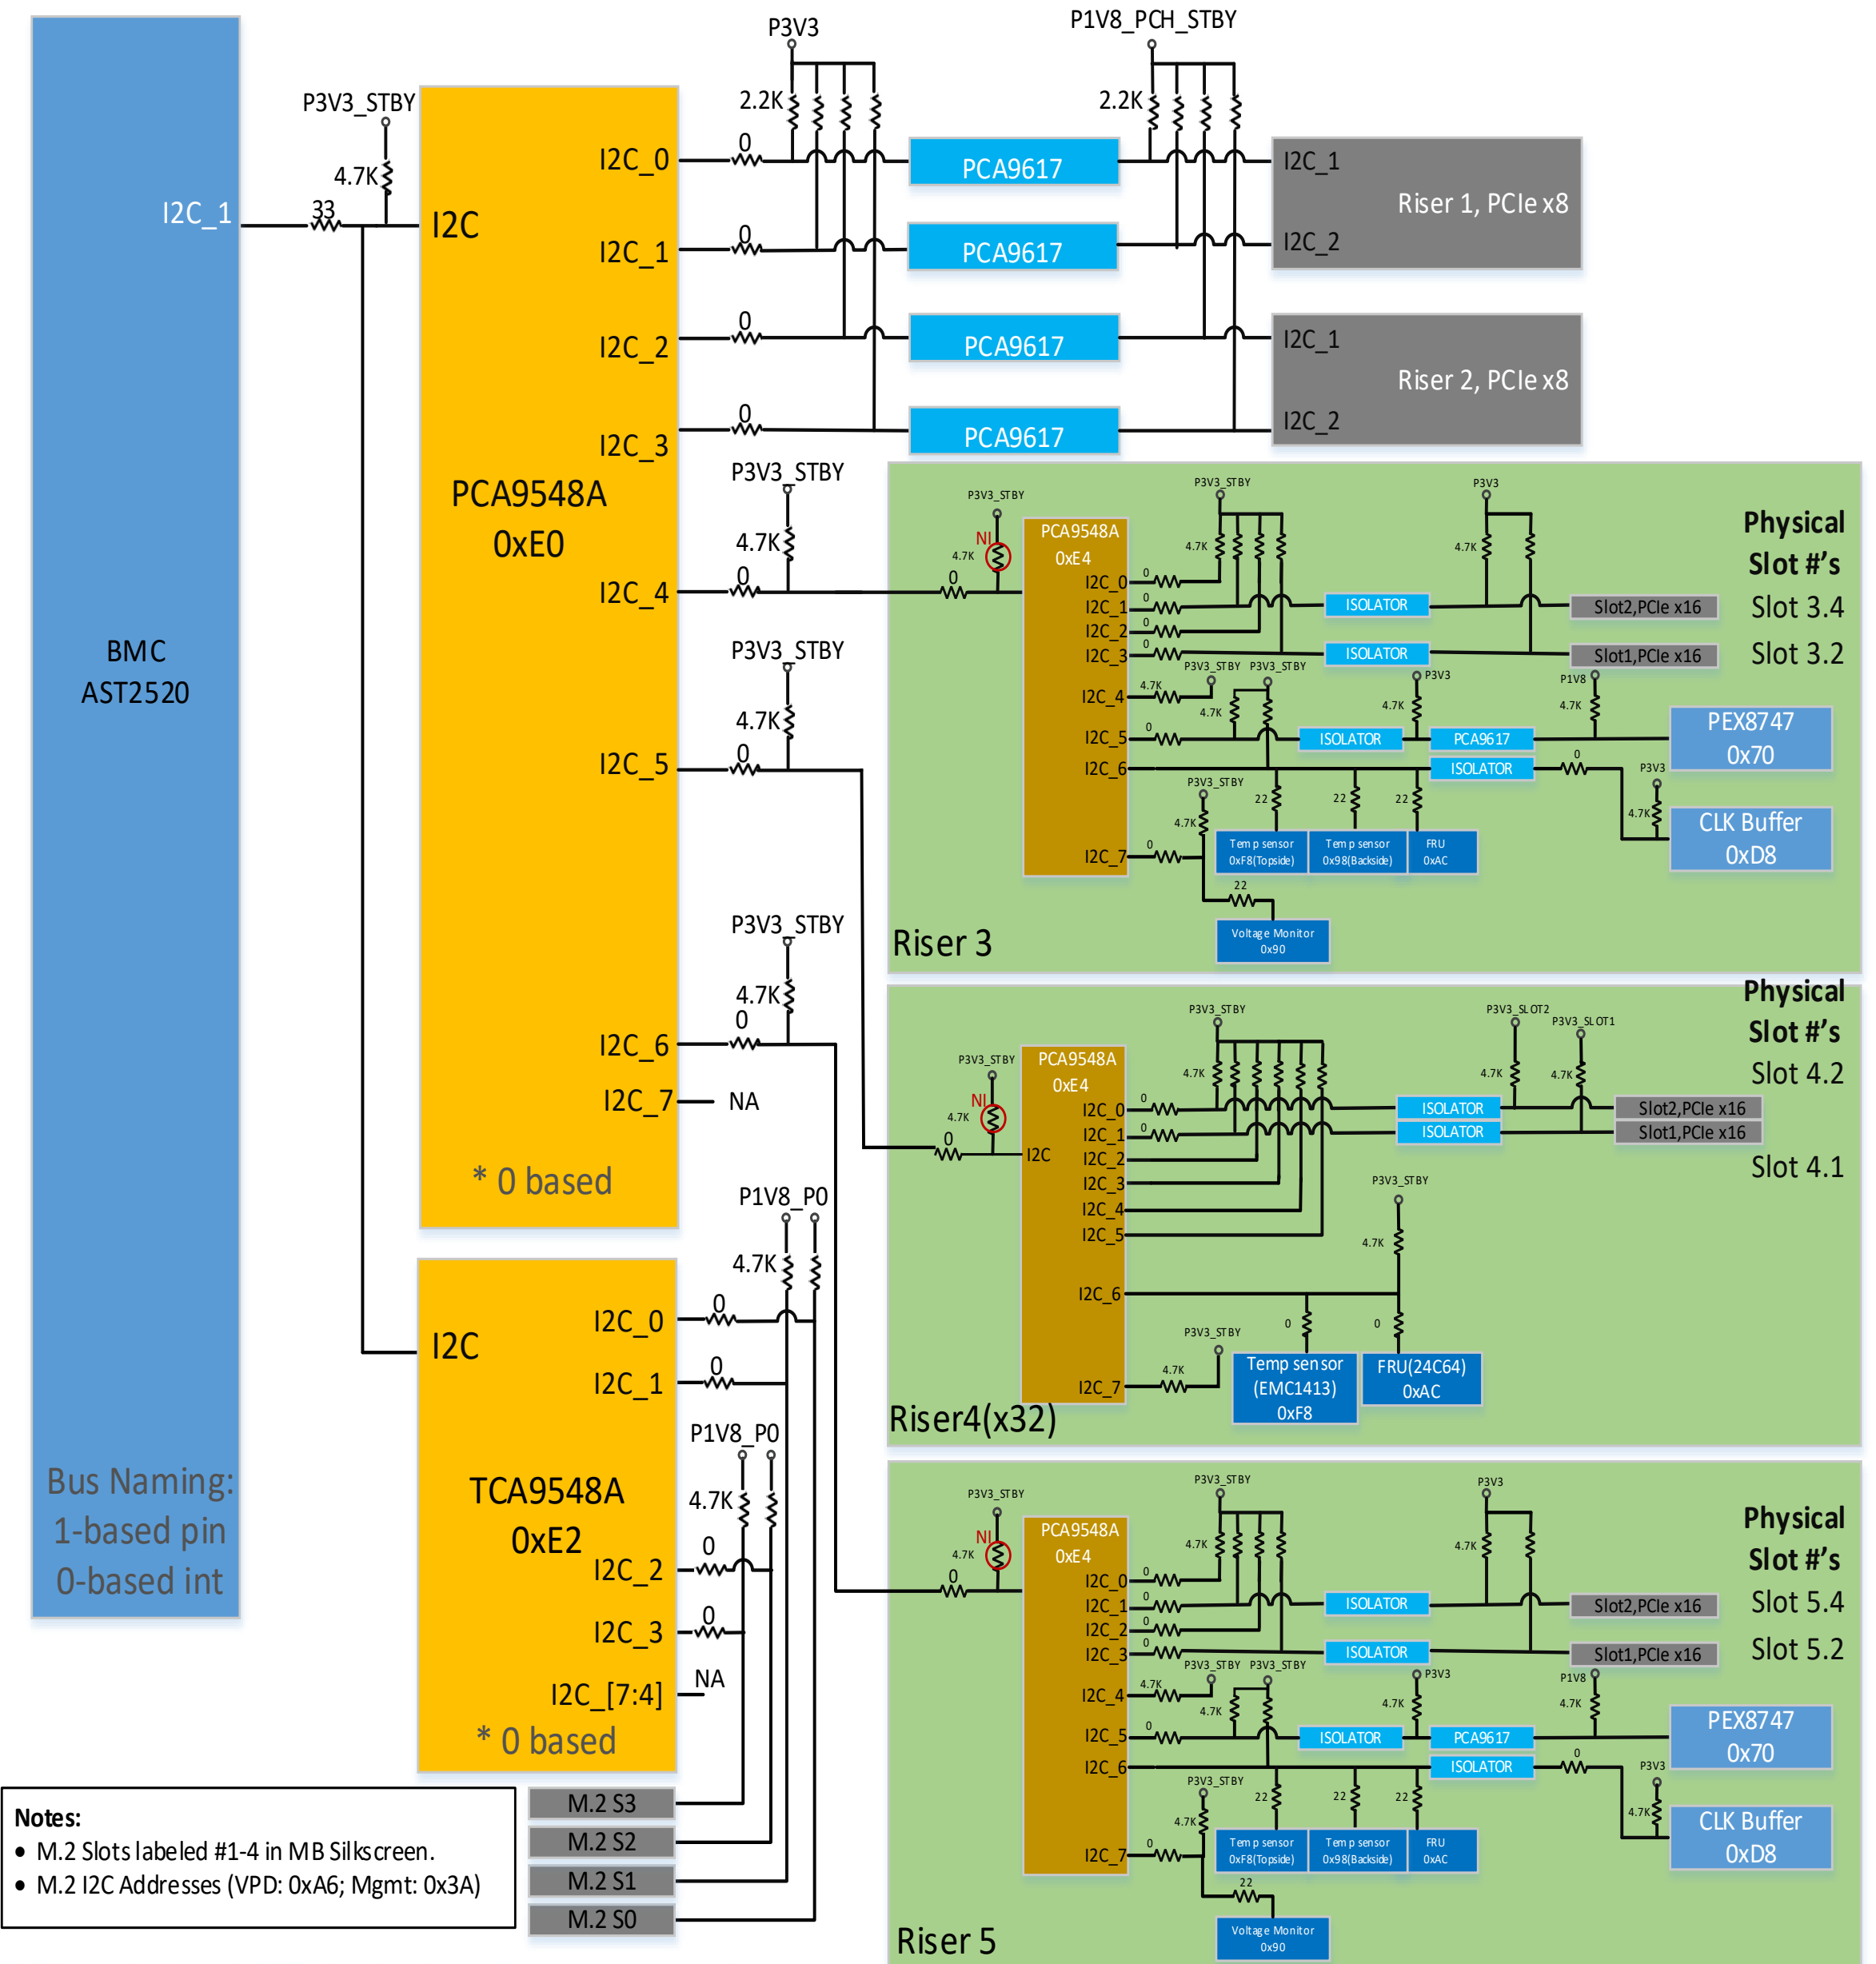
- Provides additional PCIe Lane connections from unused PCIe Slots
- 2 Cables provide PCIe x16 Lanes to Riser PCIe Slot or Riser PCIe Switch
- Supports Middle Riser Options:
 - 2-Slot Passive x16 Riser Card
 - 3-Slot Active x16 Riser Card
 - 5-Slot Active x16 Riser Card
- Typically used on the Intel® Xeon®-SP System SKU
- ZT Partnered with Tyco Electronics® on design



PCIe I2C Block Diagram

All PCIe Connectors connect to the BMC I2C for PCIe Card telemetry.

- Use the PCIe designated pins for I2C (B5 & B6).
- To avoid I2C Address contention between PCIe Cards, an I2C MUX is needed on each riser
- The design must add appropriate voltage translation/isolation between different I2C voltage domains.



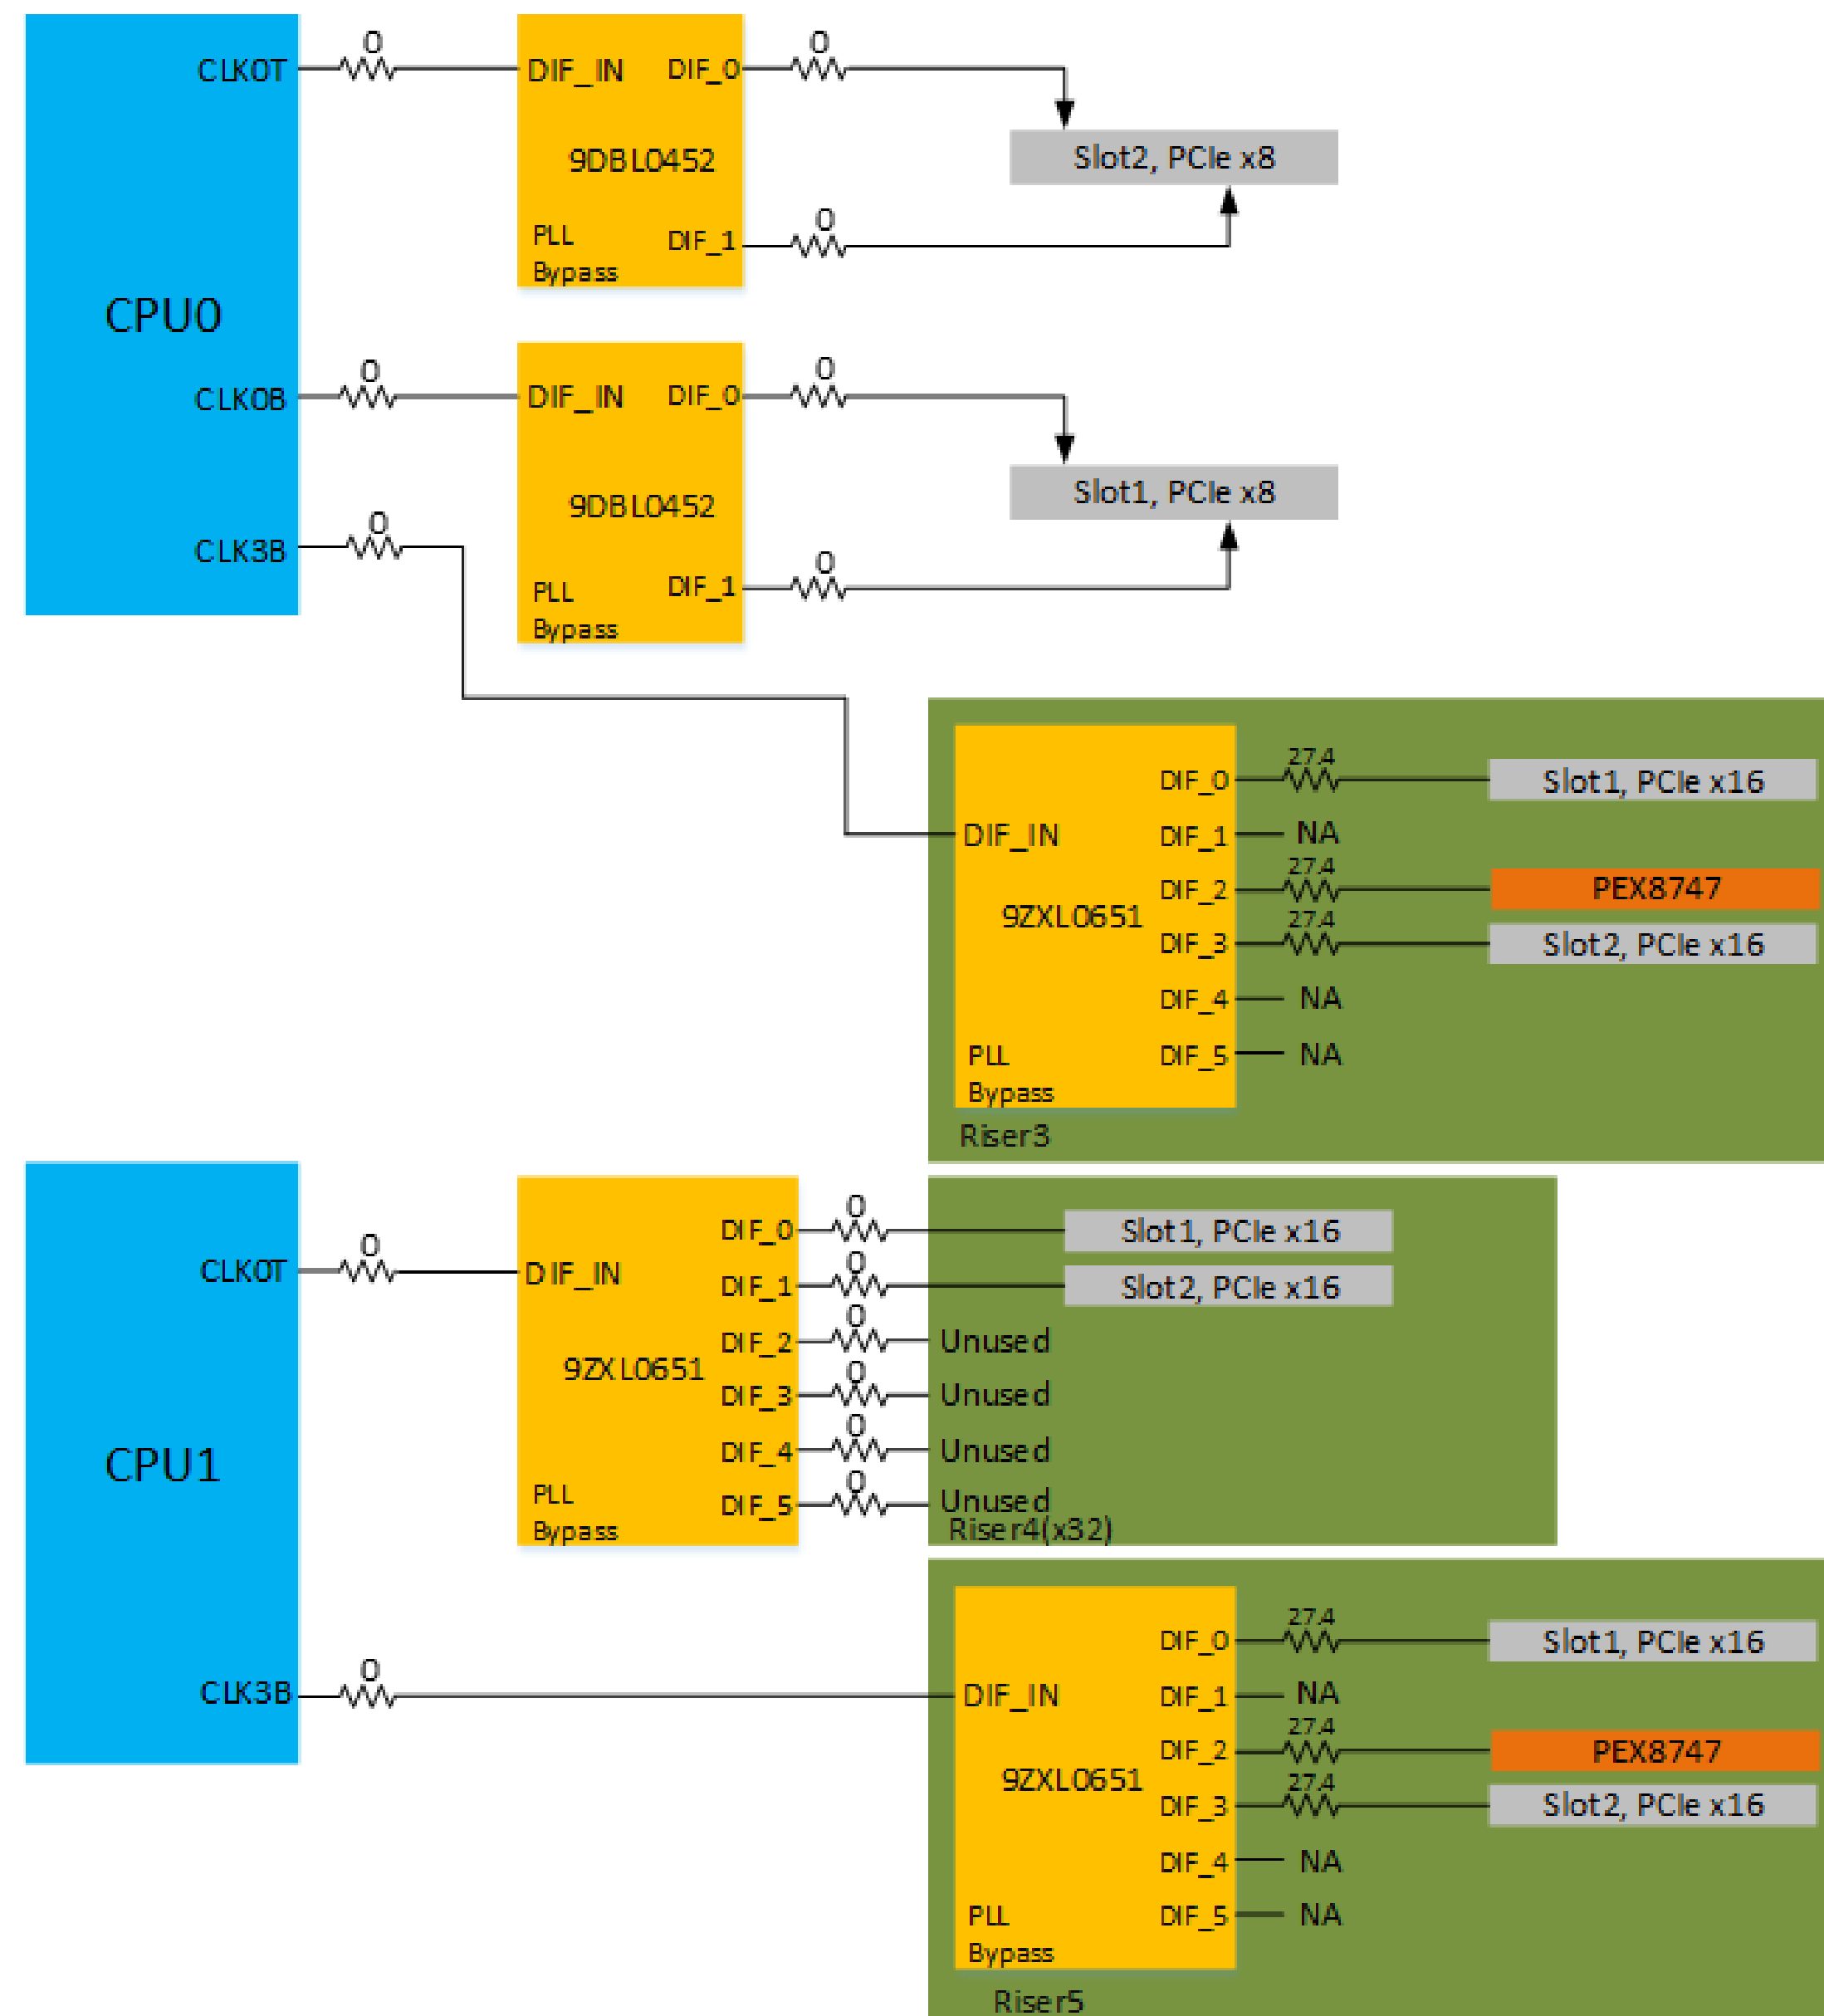
PCIe Clock Block Diagram

Differential Clock Buffers used on the Risers to fanout additional PCIe Clocks.

The Clock Buffers are set to PLL Bypass Mode by default (shown to provide best margins)

Included resistor stuffing Options to change the PLL Bandwidth Mode if needed to adjust Clock Jitter performance:

- PLL High BW Mode
- PLL Low BW Mode
- PLL Bypass Mode (default)



Power Distribution Board Layout

PDB supports single 12V Power Domain from 3 x Project Olympus PSUs

- Uses Passive Droop Sharing
- No V_Sense or I_Share signals used


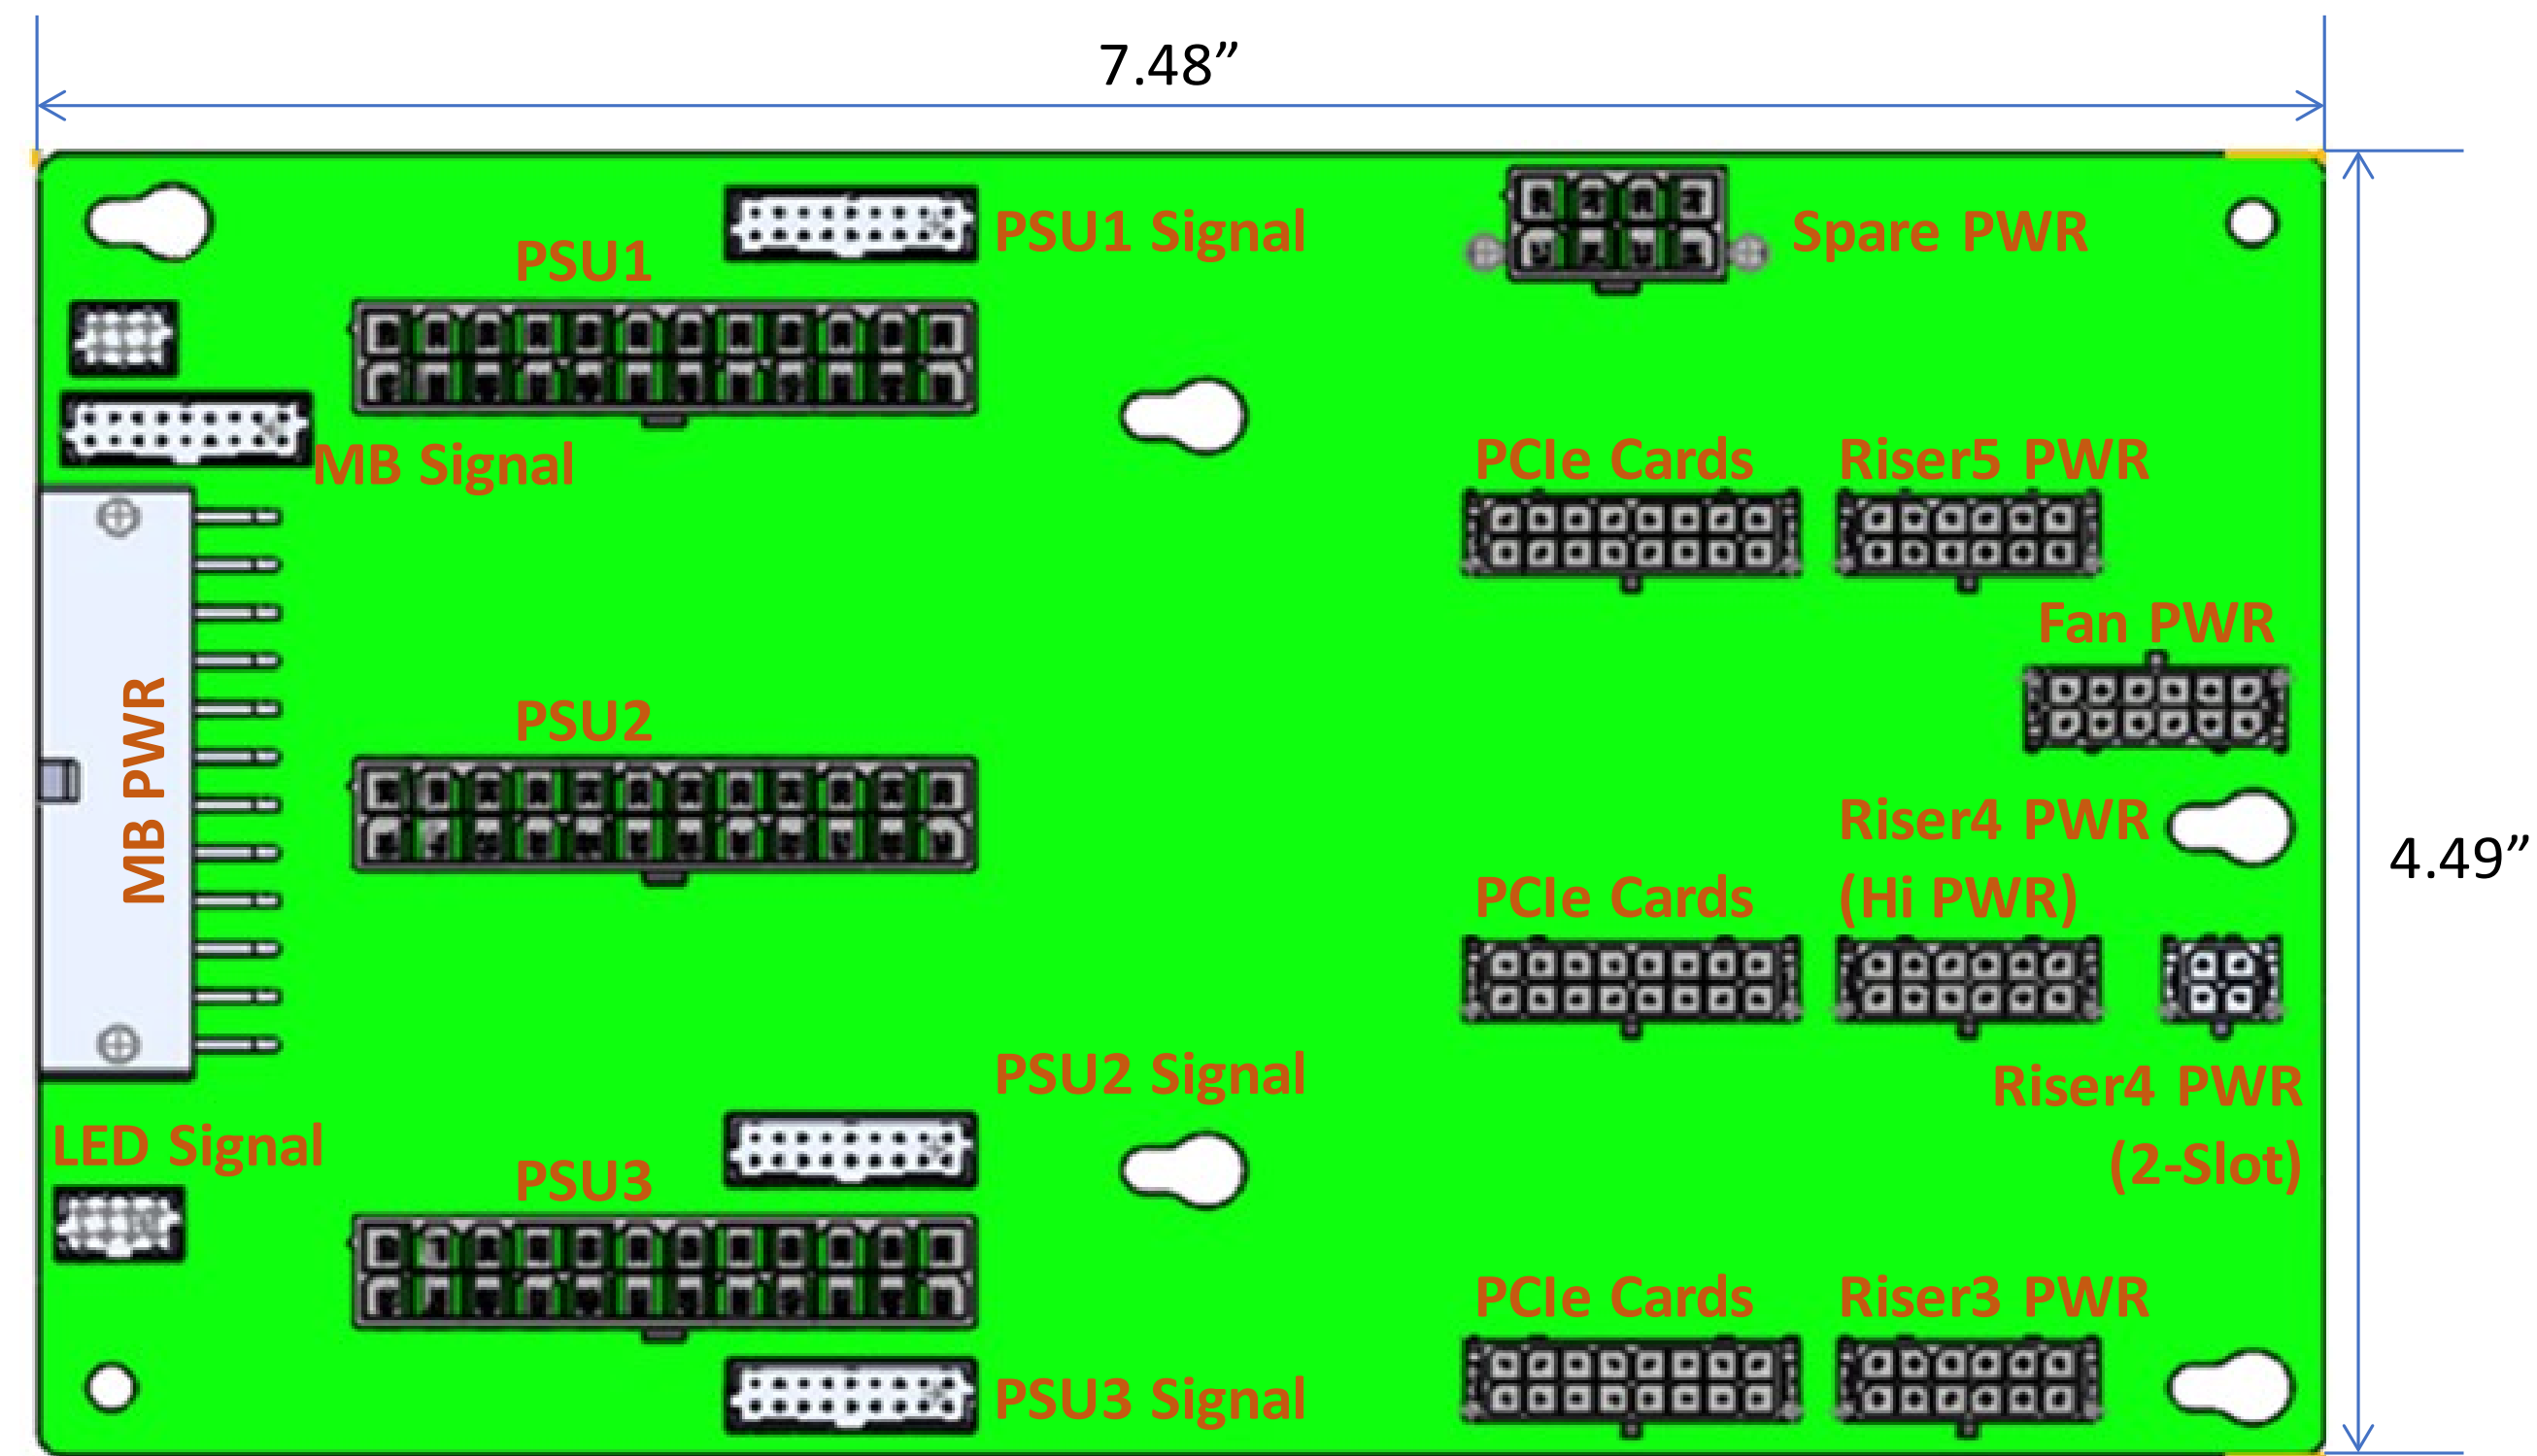
Power Delivery Requirements

- Hot Swap Controller for Power Monitoring and Isolation
- 12V Current Sense Ckts used for each branch
- Supports Power to entire System (MB, Risers, PCIe Cards)
- Fan Power also comes from PDB (none from MB)

Provides passthrough of PMDU signals between MB and RM

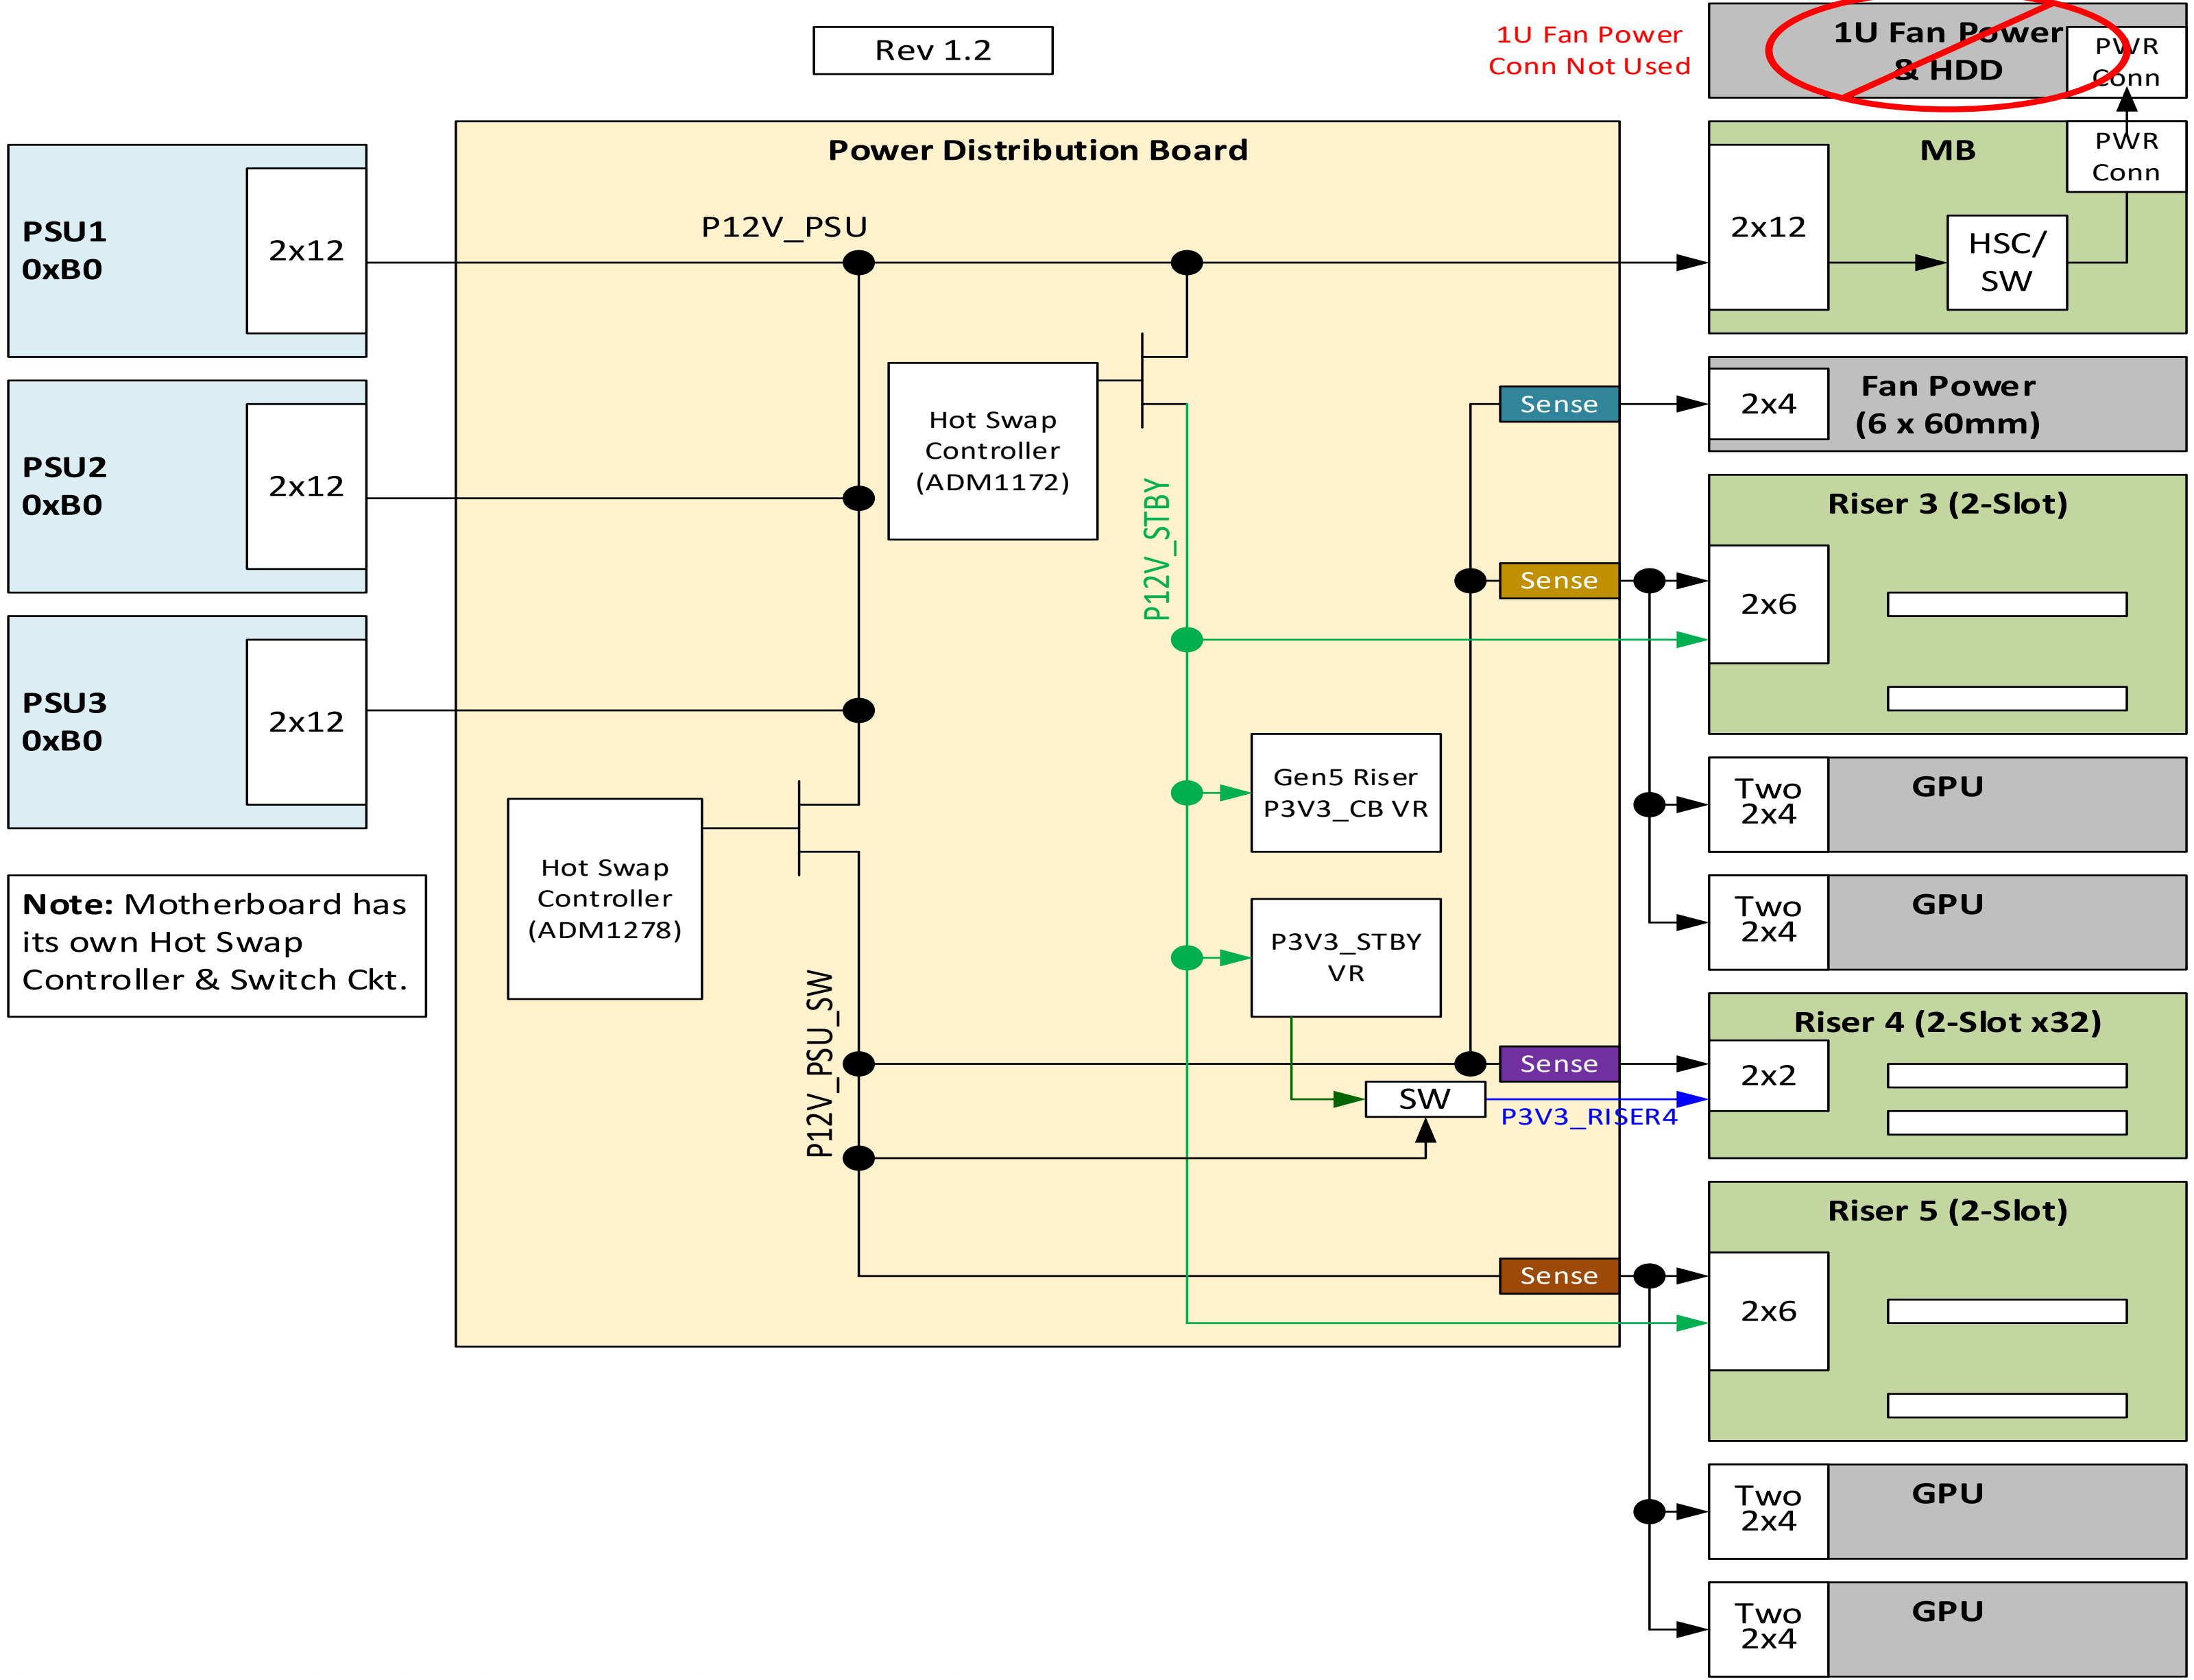
6 Layer Stackup (0.093”), 2oz copper planes

Different power harnesses can be used for different High-Power PCIe Cards



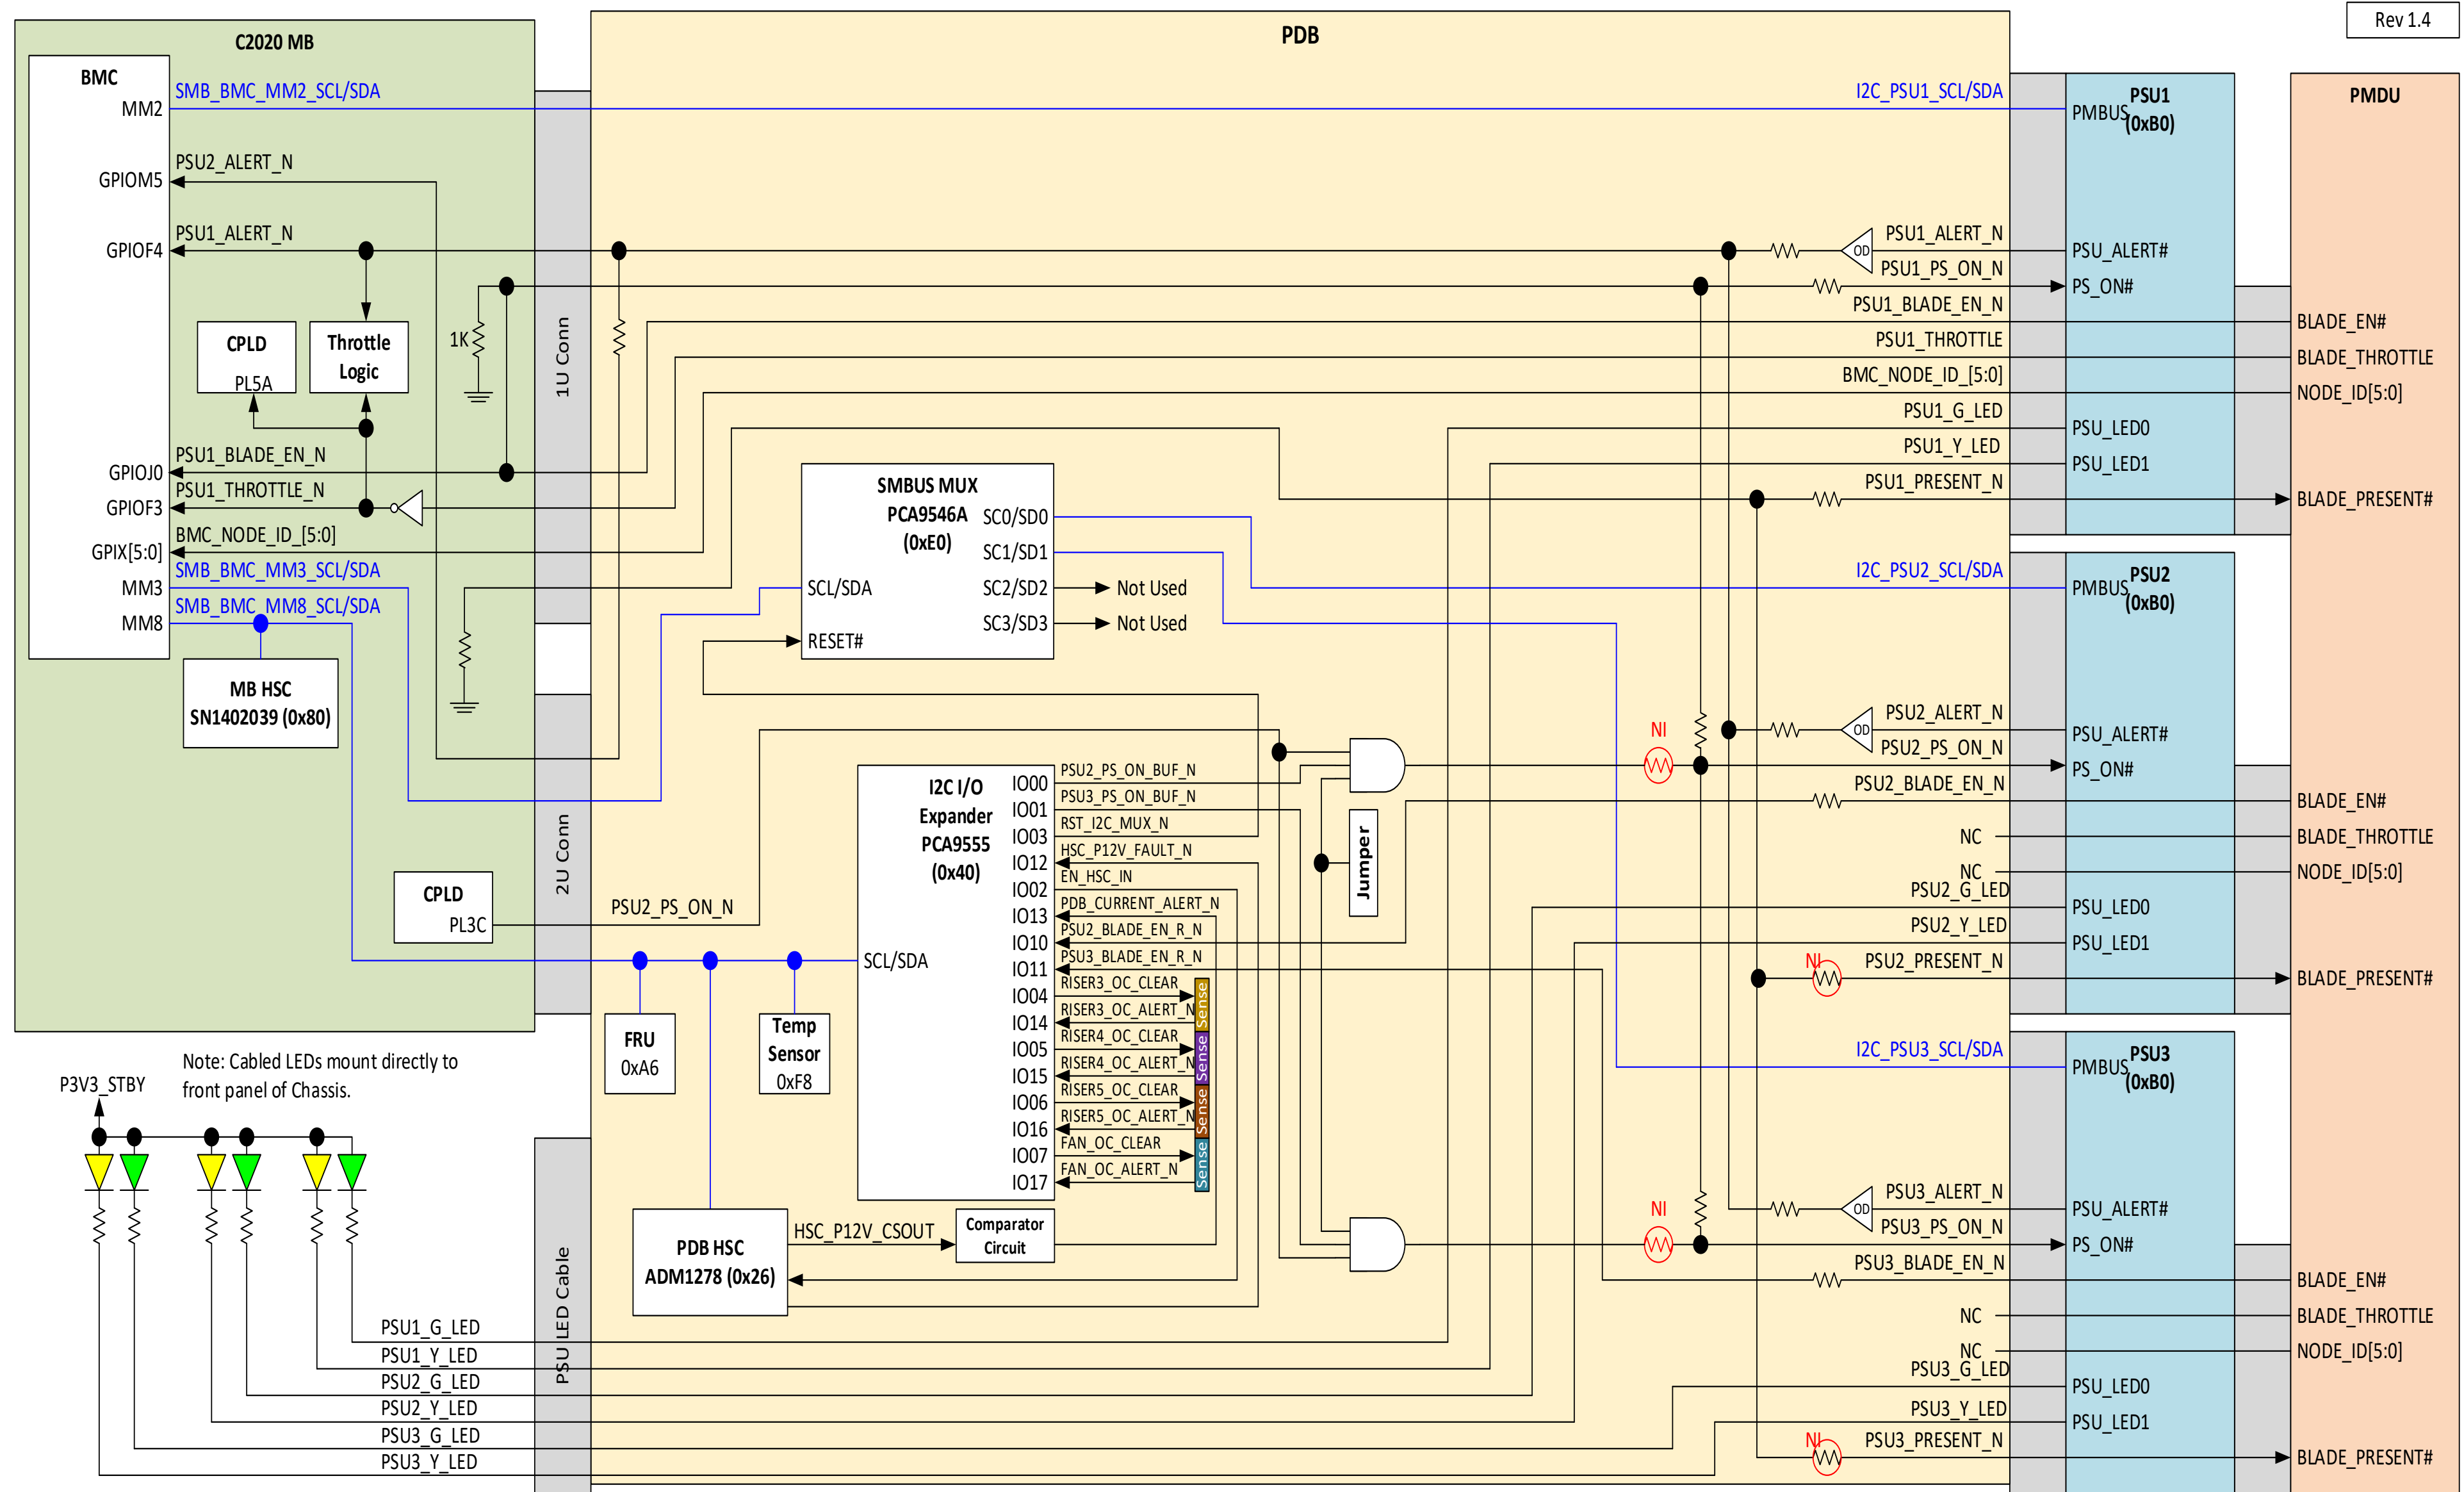
Power Harness for 2 x AMD® MI25 GPUs

PDB Power Distribution Block Diagram




PDB Misc Signal Block Diagram

Provides passthrough of PMDU signals between MB and RM



3U Server Slide Rails



Supports Slide Rails for smoother travel during servicing


Due to the weight of C2310 Server (90lbs), standard Project Olympus T-Pin Slide Rails will not allow the Chassis to pull in and out smoothly.

Width of Chassis prevents thicker rail slides (only 67% pull-out)

Top cover is split into two sections (Front and Rear).

Front Top Cover section allows access to most of the System hardware, leaving only the System Fans, PDB and Power Cables under the Rear Top Cover without in-Rack access.

Using Slide Rails from King Slide®




Plans for OCP Recognition of Project Olympus 3U Server

3U Server Base Spec submitted by Microsoft® (Siamak Tavallaei)
– Targeting OCP Acceptance by Mid 2019

3U AMD® EPYC™ Server Product Spec & Design Collateral
submitted by ZT – Targeting OCP Acceptance by Q3 2019

3U Intel® Xeon®-SP Server Product Spec & Design Collateral
submitted by ZT – Targeting OCP Acceptance by Q3 2019

Both Products to be Available on
Open Compute Marketplace 2H 2019



Visit us at Booth A17

Or www.ztsystems.com/ocp



Open. Together.



Open. Together.

OCP Global Summit | March 14–15, 2019

