

Project Olympus AMD® EPYC<sup>™</sup> Processor 3U Server Specification

Rev 1.1

Author: Mark D. Chubb

#### License

Contributions to this Specification are made under the terms and conditions set forth in Open Compute Project Contribution License Agreement ("OCP CLA") ("Contribution License") by: **ZT Group Int'l, Inc, dba ZT Systems** 

You can review the signed copies of the applicable Contributor License(s) for this Specification on the OCP website at http://www.opencompute.org/products/specsanddesign

Usage of this Specification is governed by the terms and conditions set forth in Open Compute Project Hardware License – Permissive ("OCPHL Permissive"), ("Specification License").

You can review the applicable Specification License(s) executed by the above referenced contributors to this Specification on the OCP website at http://www.opencompute.org/participate/legal-documents/

**Note**: The following clarifications, which distinguish technology licensed in the Contribution License and/or Specification License from those technologies merely referenced (but not licensed), were accepted by the Incubation Committee of the OCP:

NOTWITHSTANDING THE FOREGOING LICENSES, THIS SPECIFICATION IS PROVIDED BY OCP "AS IS" AND OCP EXPRESSLY DISCLAIMS ANY WARRANTIES (EXPRESS, IMPLIED, OR OTHERWISE), INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, FITNESS FOR A PARTICULAR PURPOSE, OR TITLE, RELATED TO THE SPECIFICATION. NOTICE IS HEREBY GIVEN, THAT OTHER RIGHTS NOT GRANTED AS SET FORTH ABOVE, INCLUDING WITHOUT LIMITATION. RIGHTS OF THIRD PARTIES WHO DID NOT EXECUTE THE ABOVE LICENSES, MAY BE IMPLICATED BY THE IMPLEMENTATION OF OR COMPLIANCE WITH THIS SPECIFICATION. OCP IS NOT RESPONSIBLE FOR IDENTIFYING RIGHTS FOR WHICH A LICENSE MAY BE REQUIRED IN ORDER TO IMPLEMENT THIS SPECIFICATION. THE ENTIRE RISK AS TO IMPLEMENTING OR OTHERWISE USING THE SPECIFICATION IS ASSUMED BY YOU. IN NO EVENT WILL OCP BE LIABLE TO YOU FOR ANY MONETARY DAMAGES WITH RESPECT TO ANY CLAIMS RELATED TO, OR ARISING OUT OF YOUR USE OF THIS SPECIFICATION, INCLUDING BUT NOT LIMITED TO ANY LIABILITY FOR LOST PROFITS OR ANY CONSEQUENTIAL, INCIDENTAL, INDIRECT, SPECIAL OR PUNITIVE DAMAGES OF ANY CHARACTER FROM ANY CAUSES OF ACTION OF ANY KIND WITH RESPECT TO THIS SPECIFICATION, WHETHER BASED ON BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), OR OTHERWISE, AND EVEN IF OCP HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## Table of Contents

License			2
Table of Co	ontents		3
Table of Fig	gures		6
Table of Ta	ables		8
Scope			9
Overview.			10
1. Rack	Compatibility		11
2. Syste	em Requirements		11
2.1 Sei	rver Configurations		11
2.2 Mo	otherboard		12
2.3 PC	le Configuration		14
2.3.1	PCle Port Mapping		14
2.3.2	PCIe Card Physical Numbering		15
2.4 2-5	Slot Active x16 PCIe Riser Card		16
2.4.1	Major Component Placement		16
2.4.2	Block Diagram		17
2.4.3	PCIe Riser Configuration		17
2.4.4	PCIe Riser Power Requirements		19
2.4.5	PCIe Riser I2C Requirements		20
2.4.6	PCIe Riser Clock Requirements		21
2.5 2-9	Slot Passive x32 PCIe Riser Card		21
2.5.1	Major Component Placement		22
Date: Augu	ist 2019	Page	3

2	.5.2	Block Diagram23
2	.5.3	PCle Riser Configuration2
2	.5.4	PCIe Riser Power Requirements2
2	.5.5	PCIe Riser I2C Requirements20
2	.5.6	PCIe Riser Clock Requirements2
2.6	PCI	e Card Requirements22
2	.6.1	Power Brake (PWRBRK_N)27
2	.6.2	IPMI Capable I2C27
3.	Serve	er Management Requirements 28
3.1	BM	C Requirements2
3.2	12C	Block Diagram2
4.	Therr	nal Design Requirements
4.1	The	rmal Design Requirements30
4.2	PCI	e Riser Temp Sensor Locations34
5.	Mech	anical Requirements
5.1	Cha	ssis Requirements:
5.2	Cha	ssis Views3!
6.	Syste	m Power Requirements
6.1	Syst	tem Power Budget39
6.2	Pov	ver Distribution Board40
6.3	Maj	jor Component Placement40
6.4	PDE	3 Power Block Diagram4
6.5	PDE	3 Configuration42

6.6	PDB Power Delivery Requirements	42
6.7	PDB Miscellaneous IO Requirements	43
6.8	Split Cable Power Harness for PCIe Cards	44
6.9	Front Panel PSU LED Cable	45
6.10	System Power Capping ("Quick Response")	46

# Table of Figures

Figure 1: Project Olympus AMD® EPYC™ Processor 3U Server System	12
Figure 2: Motherboard Block Diagram	13
Figure 3: Motherboard PCIe Mapping Layout	14
Figure 4: XPO200 3UA PCIe Bus Connectivity for 2-Slot Active & 2-Slot Passive Riser Config	15
Figure 5: Physical Card Numbering for PCIe Cards	15
Figure 6: 2-Slot Active Riser Layout (Top and Backside)	16
Figure 7: 2-Slot Active Riser Block Diagram	17
Figure 8: 2-Slot Active Riser PCIe Port Mapping and Routing Layers (View from Backside)	18
Figure 9: 2-Slot Active Riser Power Delivery Diagram	20
Figure 10: I2C Block Diagram for 2-Slot Active PCIe Riser	20
Figure 11: Clock Block Diagram for 2-Slot Active PCIe Riser	21
Figure 12: 2-Slot Passive x32 Riser Layout (Topside)	22
Figure 13: 2-Slot Passive x32 Riser Block Diagram	23
Figure 14: 2-Slot Passive x32 Riser PCIe Port Mapping and Routing Layers	24
Figure 15: 2-Slot Passive Riser Power Delivery Diagram	26
Figure 16: I2C Block Diagram for 2-Slot Passive PCIe Riser	27
Figure 17: I2C Block Diagram for XPO200 3UA	29
Figure 18: 3U Chassis System Fan Position Diagram with Closed Fan Flappers	31
Figure 19: 3U Chassis with Open Fan Flappers	31
Figure 20: System Fan Connectivity Diagram	32
Figure 21: 3U Chassis Cross Section Showing Middle Riser Baffle (to block airflow)	33
Figure 22: 2-Slot Active Riser Temp Sensor Locations	34
Figure 23: 2-Slot Passive x32 Riser Temp Sensor Locations	34

Figure 24:	3U Chassis/System Layout (Iso View)35
Figure 25:	3U Chassis/System Layout (Top View)
Figure 26:	3U Chassis/System and PSU Layout (Rear View)
Figure 27:	3U Chassis/System with Split Top Cover and King Slide Rails
Figure 28:	3U Chassis/System in 19" Project Olympus Rack
Figure 29:	3U Chassis/System in 19" Project Olympus Rack (Extended 67%)
Figure 30:	Power Distribution Board Layout40
Figure 31:	Power Distribution Board Power Block Diagram41
Figure 32:	XPO200 3UA 3U Server Power Supply Load Sharing Diagram (Max Supported Power)42
Figure 33:	Power Distribution Board Miscellaneous IO Block Diagram44
Figure 34:	Split Cable Power Harness for 300W High Power PCIe Card (MI25)45
Figure 35:	Front Panel PSU LED Cable

## Table of Tables

Table 1: Project Olympus Specifications	9
Table 2: Supported System Configuration Options	11
Table 3: 2-Slot Active Riser Stack-up	18
Table 4: 2-Slot Active Riser Power Budget	19
Table 5: 2-Slot Passive Riser Stack-up	25
Table 6: 2-Slot Passive Riser Power Budget	25
Table 7: I2C Bus Mapping for XPO200 3UA	29
Table 8: XPO200 3UA Server Power Budget	40
Table 9: PSU Status LED Descriptions	45

#### Scope

This document defines the technical specifications for the Project Olympus AMD® EPYC<sup>™</sup> 3U PCIe Expansion Server. The ZT model number for this system is the XPO200 3UA, which will be used to identify the 3U Server throughout this Specification.

This XPO200 3UA 3U Server implementation is compatible with existing Project Olympus building blocks. Other Project Olympus elements which are utilized in this system are outlined in their respective specifications. These elements include, but are not limited to, the Power Supply Unit (PSU), Power Management Distribution Unit (PMDU), Universal Motherboard, Server Rack, and Rack Manager (RM). Specifications for the overall Project Olympus Rack are posted here...

http://www.opencompute.org/wiki/Server/ProjectOlympus

Specification title	Description
Project Olympus Server Rack Specification	Describes the mechanical rack hardware used in
	the system
Project Olympus Server Mechanical	Describes the mechanical structure for the
Specification	server used in the system.
Project Olympus Universal Motherboard	Describes the server motherboard general
Specification	requirements.
Project Olympus PSU Specification	Describes the Power Supply Unit (PSU) used in
	the server
Project Olympus Power Management	Describes the Power Management Distribution
Distribution Unit Specification	Unit (PMDU).
Project Olympus Rack Manager Specification	Describes the Rack Manager PCBA used in the
	PMDU.

#### Table 1: Project Olympus Specifications

#### Overview

This specification defines a Project Olympus based 2 Socket, 3U Server (XPO200 3UA) that is intended to support multiple double-wide PCIe Cards. This server design is based on the Project Olympus 3U Server Base Specifications and fits within the Project Olympus OCP framework. The OCP Approved Project Olympus Motherboard based on the AMD® EPYC<sup>™</sup> platform can be used in this system along with various PCIe Riser Options to provide an IO rich platform for up to 5 x Double-wide Full-Length Cards plus 1 x Single-wide Half-Length Card. Due to the increased power requirements in this XPO200 3UA Server, three of the Project Olympus 3-phase 1kW Power Supply Silver Boxes will also be used in this system. Each Silver Box includes 3 internal Power Supply Units (PSUs) providing a total of 9 PSUs in an 8 + 1 redundancy scheme. Each Project Olympus vertical Power Managed Distribution Unit (PMDU).

#### 1. Rack Compatibility

This XPO200 3UA PCIe Expansion Server is compatible with the Project Olympus 19" EIA310-D Rack. Due to the Server's weight, it requires custom slide rails, which allow for smoother travel during Serviceability when sliding in and out of the Rack.

#### 2. System Requirements

The system requires a Mount Olympus AMD® EPYC<sup>™</sup> generation Motherboard, which can be combined with several different PCIe Riser Options for system configuration flexibility. It includes a Power Distribution Board (PDB) to allow load sharing of the 3 Project Olympus 1kW PSUs. These components reside within a 3U chassis and are compatible with the Project Olympus ecosystem.

#### 2.1 Server Configurations

Below are the targeted XPO200 3UA Compute Node Server SKUs known at the time of this System Architecture Specification release. All XPO200 3UA Server Configurations are subject to change based on Customer requirements. Additional Configurations may be added at a later time based on Customer requirements.

Qty	Description
1	3U, 19" EIA310-D Compliant supporting Project Olympus PMDU
	connections
1	Mount Olympus 2-Socket AMD® EPYC <sup>™</sup> Motherboard
2	AMD® EPYC™ 7551 Processor
16	32GB DDR4, DR, 2667 R-DIMMs; Total System Memory: 512GB
2	2-Slot Active PCIe x16 Riser Card with 48-lane PCIe Switch
1	2-Slot Passive PCIe x32 Riser Card
4	AMD® Radeon Instinct™MI25 GPU, 300W PCIe x16 Card
1	10G Single port SFP+ PCIe 2.0 x8 5GT/s
4	M.2 960GB NVMe SSD, PCIe x4 110mm (sourced from CPU)
1	TPM2.0 SPI Module
6	60mmx56mm Dual Rotor Fans
3	Project Olympus 1020W 3-Phase, non-LES PSU
1	PDB and Cable Harnesses to support 12V Power to MB, Risers, PCIe
	Qty 1 2 16 2 1 4 1 4 1 6 3 1

#### Table 2: Supported System Configuration Options



Figure 1: Project Olympus AMD® EPYC™ Processor 3U Server System

#### 2.2 Motherboard

The motherboard used in the XPO200 3UA System is the Dual Socket Project Olympus AMD® EPYC<sup>™</sup> Motherboard (MB), which is the computational element in the Project Olympus Server. This is a Dual Socket EPYC<sup>™</sup> generation server board leveraged from the 2S AMD® EPYC<sup>™</sup> OCP Olympus MB Specification.

**NOTE:** The information in this section is for reference only. For the latest detailed Spec information please refer to the Project Olympus AMD® EPYC<sup>™</sup> MB Specification posted here...

http://www.opencompute.org/wiki/Server/ProjectOlympus

- Processor:
  - 2 Socket Spread Core Design using AMD® EPYC<sup>™</sup> Processors
  - Supports up to 180W TDP
  - o Includes 1U Remote Heatsink
  - Memory: 32 DDR4 DIMMs, 2 DIMMs per Channel
- PCIe Slots/Connectors:
  - PCIe x8 Connector (Slot #1) CPU0
  - PCIe x8 Connector (Slot #2) CPU0
  - PCIe x16 Riser Connector (Slot #3) CPU0

- PCIe x32 Riser Connector (Slot #4) CPU1
- PCIe x16 Riser Connector (Slot #5) CPU1
- PCIe x4 M.2 Connector (M.2 #1) CPU0
- PCIe x4 M.2 Connector (M.2 #2) CPU0
- PCIe x4 M.2 Connector (M.2 #3) CPU1
- PCIe x4 M.2 Connector (M.2 #4) CPU1
- PCIe x8 OCuLink Connector CPU1
- SATA Connectors:
  - 4 x SATA 7-pin Connectors (P1[3:0]) CPU0
  - 1 x SATA MiniSAS HD Connector (SATA P1[7:4] or PCIe[7:4]
  - 1 x SATA MiniSAS HD Connector (SATA P1[11:8] or PCIe[11:8]
- BMC: ASPEED AST2520
  - o All PCIe Slots are connected to BMC I2C Buses for PCIe Card telemetry
  - I2C MUXes are used to avoid I2C Address contention
- Security: SPI TPM2.0 Module
- Front IO Ports:
  - 1 x BMC Dedicated Management NIC Port
  - o 2 x USB3.0 Ports
  - Power Button (Pre-Production Only)
  - Reset Button (Pre-Production Only)
  - 1 x Video Port (EMPTY)
  - 1 x 10GbE SFP+ Connector (EMPTY)



Figure 2: Motherboard Block Diagram

## 2.3 PCIe Configuration

## 2.3.1 PCIe Port Mapping

The PCIe ports from each processor are mapped as shown in the Figure and Table below:



Figure 3: Motherboard PCIe Mapping Layout



Figure 4: XPO200 3UA PCIe Bus Connectivity for 2-Slot Active & 2-Slot Passive Riser Config

#### 2.3.2 PCIe Card Physical Numbering

Below is the PCIe Card Chassis Level Numbering scheme for the XPO200 3UA System configurations as viewed from the front of the Chassis.



Figure 5: Physical Card Numbering for PCIe Cards

#### 2.4 2-Slot Active x16 PCIe Riser Card

This is a PCIe x16 Active Riser Card with Broadcom/PLX 48-lane Gen3 PCIe Switch to 2 x16 PCIe Slots. There can be two 2-Slot Active Risers per system, each plugged into MB Riser Slots #3 & #5.

## 2.4.1 Major Component Placement



Figure 6: 2-Slot Active Riser Layout (Top and Backside)

## 2.4.2 Block Diagram



Figure 7: 2-Slot Active Riser Block Diagram

## 2.4.3 PCIe Riser Configuration

- 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 from MB CPLD routed to each Slot for High Power PCIe Card throttling
    - This feature can be used for system power capping
    - This signal is typically only supported on High Power PCIe Cards (i.e. Cards greater than 75W)
- PCIe Port Mapping and Routing Layers are shown in the figure below. PCIe Connectors (not shown) are on Topside:



Figure 8: 2-Slot Active Riser PCIe Port Mapping and Routing Layers (View from Backside)

Dielectric and Copper Properties											
	Copper		Dielectric Material					Layer			
										Design	Actual
								Actual	Actual	Thickness	Thickness
Name	Description	Type	Weight	Type	Type	Er	Df	Dk	Df/1MH	(mil)	(mil)
Top_SM	Top Soldermask	Soldermask			Soldermask	3.5	0.030	3.70		0.50	
L1top	Signal_L1top	Signal	0.5 oz. plated	STD						1.90	
PP_L1top_L2	Prepreg	PP			1080*1	3.6	0.009	3.50	0.007	2.70	
L2	GND_L2	GND	1.0 oz.	RTF						1.30	
Core_L2_L3	Core	Core			2116	3.6	0.009	3.50	0.007	4.00	
L3	Signal_L3	Signal	1.0 oz.	RTF						1.30	
PP_L3_L4	Prepreg	PP			7628*2	3.6	0.009	3.50	0.007	15.00	
L4	PWR_L4	PWR	2.0 oz.	RTF						2.60	
Core_L4_L5	Core	Core			2116X1	3.8	0.009	3.60	0.007	4.00	
L5	PWR_L5	PWR	2.0 oz.	RTF						2.60	
PP_L5_L6	Prepreg	РР			7628*2	3.6	0.009	3.50	0.007	15.00	
L6	Signal_L6	Signal	1.0 oz.	RTF						1.30	
Core_L6_L7	Core	Core			2116	3.6	0.009	3.50	0.007	4.00	
L7	GND_L7	GND	1.0 oz.	RTF						1.30	
PP_L7_L8bot	Prepreg	PP			1080*1	3.6	0.009	3.50	0.007	2.70	
L8bot	Signal_L8bot	Signal	0.5 oz. plated	STD						1.90	
Bot_SM	Bot Soldermask	Soldermask			Soldermask	3.5	0.030	3.70		0.50	
with soldermask, total thickness: 62.60										0.00	
without soldermask, total thickness: 61.60										0.00	

#### **PCIe Switch Configuration** 2.4.3.1.

- Broadcom/PLX 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 0
  - PCIe Switch Thermal Design Power Requirements: 0
    - Use of all 48-lanes requires 13.2W •

- Heatsink is required
- Please refer to the PEX8747 Datasheet and User Manual for additional design requirements

## 2.4.4 PCIe Riser Power Requirements

- Each PCIe slot supports up to 75W as defined for a standard PCIe x16 card.
- A 12V Auxiliary connector (2x6 pin cabled from PDB) is required on the Riser to support PCIe Slot power, PEX8747 VRs, and 3.3V VR:
  - The required 12V Current for this Riser is 14.1A (assume 18A for conn derating).
  - $\circ$  12V AUX Connector uses 5 x 12V pins: 18 Amps/5 pins = 3.6 Amps per pin.
- Additional 12V power required by high power PCIe cards will cable directly from the Power Distribution board (PDB) to the cards themselves.
- Different Wire Harnesses can be used depending on the PCIe Card's 12V Auxiliary Power connector
- 3.3V Slot power is required for all PCIe Cards:
  - Per the PCIe Spec, 3.3V requires 3.0A per slot (2 x 3A = 6A), which cannot be pulled from only three Riser Slot 3.3V Pins.
  - 3.3V must be sourced from an on-board VR
- 3.3V\_Stby may be required by some PCIe Cards:
  - Per the PCIe Spec, 3.3V\_Stby requires 375mA per slot (2 x 375mA = 750mA), which can be pulled from the single Riser Slot 3.3V\_Stby pin.
  - No need for an on-board 3.3V\_Stby VR

2-Slot Active Riser 12V Budget	Voltage (V)	Current (A)	Qty	Total Current (A)	Power (W)	VR eff	12V Power Req'd (W)	
PCIe 12V	12	5.5	2	11	132	100%	132	
PCIe 3.3V	3.3	3	2	6	19.8	90%	22	
DB600Z 3.3V	3.3	0.081	1	0.081	0.2673	90%	0.30	
Misc Logic/Components	3.3	0.1	1	0.1	0.33	90%	0.37	
PEX8747 VDD09	0.9		1		8.48	90%	9.42	
PEX8747 VDD09A	0.9		1		3.746	90%	4.16	1011
PEX8747 VDD18A	1.8		1		0.981	85%	1.15	12V Current
PEX8747 VDD18	1.8		1		0.022	85%	0.03	(A)
						Totals	169	14.1



Figure 9: 2-Slot Active Riser Power Delivery Diagram

## 2.4.5 PCIe Riser I2C Requirements

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.



Figure 10: I2C Block Diagram for 2-Slot Active PCIe Riser

## 2.4.6 PCIe Riser Clock Requirements

A Differential Clock Buffer (DB600Z) is used to fanout additional PCIe Clocks. The Clock Buffer is set to PLL Bypass Mode by default. Recommend including resistor stuffing Options to change the PLL Bandwidth Mode if needed to improve Clock Jitter performance:

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

		CLK0	SLOT1, PCle x16
_	CLK	CLK1 — NA	
		CLK2 —••• <sup>27.4</sup>	PEX8747
		CLK3 —•••• <sup>27.4</sup>	SLOT2, PCle x16
	CLK	CLK4 <mark>—</mark> NA	
	Buffer	CLK5 — NA	



#### 2.5 2-Slot Passive x32 PCIe Riser Card

This is a PCIe x32 Passive Riser Card with 2 x16 PCIe Slots supported in Riser Slot #4. The x32 card edge requires two Samtec HSEC8 connectors (200 pin & 60 pin) placed in-line on the MB. The Bottom x16 PCIe Slot is sourced from the lower order x16 lanes on the PCIe Riser Slot, and the Top x16 PCIe Slot is sourced from the higher order x16 lanes.

## 2.5.1 Major Component Placement



Figure 12: 2-Slot Passive x32 Riser Layout (Topside)

#### 2.5.2 Block Diagram



Figure 13: 2-Slot Passive x32 Riser Block Diagram

## 2.5.3 PCIe Riser Configuration

- 1 x PCIe x16 slot supporting (Slot #1, Bottom):
  - 1 x FHHL Single-Wide PCIe card
  - Designated for a Network Card
  - JTAG signals routed through PCIe JTAG pins of Slot #1 for debug purposes
  - All x16 PCIe lanes are sourced from MB PCIe Riser Slot #4
- 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 are sourced from MB PCIe Riser Slot #4
- Requires PWRBRK\_N (pin B30) signal from MB CPLD to these Slots for High Power PCIe Card throttling
  - This feature can be used for system power capping

#### OCP • Project Olympus AMD® EPYC™ Processor 3U Server Specification

- This signal is typically only supported on High Power PCIe Cards (i.e. Cards greater than 75W)
- PCIe Port Mapping and Routing Layers pictured below:



Figure 14: 2-Slot Passive x32 Riser PCIe Port Mapping and Routing Layers

Dielectric and Copper Properties											
Layer			Copper		Dielectric Material				Layer		
										Design	Actual
								Actual	Actual	Thickness	Thickness
Name	Description	Type	Weight	Type	Type	Er	Df	Dk	Df/1MH	(mil)	(mil)
Top_SM	Top Soldermask	Soldermask			Soldermask	3.5	0.030	3.70		0.50	
L1top	Signal_L1top	Signal	0.5 oz. plated	STD						1.90	
PP_L1top_L2	Prepreg	PP			1080*1	3.6	0.009	3.50	0.007	2.70	
L2	GND_L2	GND	1.0 oz.	RTF						1.30	
Core_L2_L3	Core	Core			2116	3.6	0.009	3.50	0.007	4.00	
L3	Signal_L3	Signal	1.0 oz.	RTF						1.30	
PP_L3_L4	Prepreg	PP			7628*2	3.6	0.009	3.50	0.007	15.00	
L4	PWR_L4	PWR	2.0 oz.	RTF						2.60	
Core_L4_L5	Core	Core			2116X1	3.8	0.009	3.60	0.007	4.00	
L5	PWR_L5	PWR	2.0 oz.	RTF						2.60	
PP_L5_L6	Prepreg	PP			7628*2	3.6	0.009	3.50	0.007	15.00	
L6	Signal_L6	Signal	1.0 oz.	RTF						1.30	
Core_L6_L7	Core	Core			2116	3.6	0.009	3.50	0.007	4.00	
L7	GND_L7	GND	1.0 oz.	RTF						1.30	
PP_L7_L8bot	Prepreg	PP			1080*1	3.6	0.009	3.50	0.007	2.70	
L8bot	Signal_L8bot	Signal	0.5 oz. plated	STD						1.90	
Bot_SM	Bot Soldermask	Soldermask			Soldermask	3.5	0.030	3.70		0.50	
with soldermask, total thickness: 62.60 0.00											

#### Table 5: 2-Slot Passive Riser Stack-up

without soldermask, total thickness:

#### 2.5.4 PCIe Riser Power Requirements

- Each PCIe slot supports up to 75W as defined for a standard PCIe x16 card.
- 12V and 3.3V to PCIe Slot #1 is supported from the MB Riser Slot
- Standard PCIe Slot current rating is 1.1A per pin (5 x 12V pins = 5.5A per Slot)
- A 12V/3.3V Auxiliary connector (2x2 pin cabled from PDB) is required on the Riser to support PCIe Slot #2 power
  - The required 12V Current for this Riser is ~5.5A (assume 8A for conn derating).
  - 12V AUX Connector uses 1 x 12V pin: 8A/1 pin = 8A per pin.
  - The required 3.3V Current for this Riser is ~3.14A.
  - $\circ$  12V AUX Connector uses 1 x 3.3V pin: 3.14A/1 pin = 3.14A per pin.
- Additional 12V power required by high power PCIe cards will cable directly from the PDB to the cards themselves.
- Different Wire Harnesses can be used depending on the PCIe Card's 12V Auxiliary Power connector
- Only Slot #2 can support a High Power, double-wide PCIe Card
- 3.3V\_Stby may be required by some PCIe Cards:
  - Per the PCIe Spec, 3.3V\_Stby requires 375mA per slot (2 x 375mA = 750mA), which can be pulled from the single Riser Slot 3.3V\_Stby pin.
  - No need for an on-board 3.3V\_Stby VR.

2-Slot x32 Riser 12V Budget	Voltage (V)	Current (A)	Qty	12V Current Req'd (A)
PCIe 12V	12	5.5	1	5.5

Table 6: 2-Slot	Passive Riser	Power	Budget
-----------------	---------------	-------	--------

61.60

0.00

#### OCP • Project Olympus AMD® EPYC™ Processor 3U Server Specification

2-Slot x32 Riser 3.3V Budget	Voltage (V)	Current (A)	Qty	3.3V Current Req'd (A)
PCIe 3.3V	3.3	3	1	3
Misc Logic/Components	3.3	0.1	1	0.1
			Total	3.1



Figure 15: 2-Slot Passive Riser Power Delivery Diagram

## 2.5.5 PCIe Riser I2C Requirements

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





## 2.5.6 PCIe Riser Clock Requirements

Up to six PCIe Clocks are routed through the MB PCIe Riser pins. Therefore, no need for a Differential Clock Buffer on the Passive x32 PCIe Riser. Route two of the six PCIe Clock pairs to PCIe Slots #1 and #2 and leave the other four Clock pairs floating.

**NOTE:** In order to minimize radiated emissions (EMI), it is recommended that BIOS disable the unused PCIe Clocks routed to the MB Riser Slot #4.

## 2.6 PCIe Card Requirements

This section provides any special requirements for the PCIe Cards in order to support the System features of the Project Olympus server.

## 2.6.1 Power Brake (PWRBRK\_N)

Power Brake (PWRBRK\_N; Pin B30) is an optional feature in the PCIe Spec that allows the system to throttle the PCIe Card. This feature is useful for System Power Capping when the 12V current in the system exceeds a predefined Over Current (OC) threshold in the PSU. With the Power Brake feature, the System has the ability to throttle PCIe Cards in order to avoid shutting down and ride through the OC event. It is required that all high-power PCIe Cards used in the XPO200 3UA support the Power Brake function in order to support the Project Olympus Universal MB Power Capping feature.

The response time of the PCIe Card throttling relative to a PWRBRK\_N signal assertion must be approximately 10ms or less. The PCIe Card should throttle down to at least 50% of its max thermal design power (TDP). The PCIe Card must return to full speed after PWRBRK\_N has been de-asserted.

**NOTE:** This feature is not required for lower power PCIe Cards (< 150W) but is strongly recommended.

#### 2.6.2 IPMI Capable I2C

Any PCIe Cards used in the XPO200 3UA should support IPMI based I2C commands for BMC access to I2C devices and sensors.

Date: August 2019

#### 3. Server Management Requirements

The following section provides a high-level overview of the XPO200 3UA Server Management Requirements. Additional Server Management Requirements can be found in the Project Olympus BMC FW Specification.

#### 3.1 BMC Requirements

- Must generate the SHA-256 (Secure Hash Algorithm-256bit) Hash of the Signed BMC FW
- Rack Manager supported in the PMDU on the Project Olympus Rack will assign a "U" number to the Server BMC through PSU1 NODE\_ID/SLOT\_ID Bits only. PSU2 and PSU3 NODE\_ID/SLOT\_ID signals are not connected.
- System Fan Speed Control (FSC) mainly based on CPU and PCIe Card Temperature Sensors
- PSU Fan Speed Relationship with System Fan Speed developed to boost PSU fan speed relative to different system fan speeds:
  - If System Fans are less than 80% PWM, PSU can use its own FSC
  - If System Fans are greater than or equal to 80% PWM, PSU Fans must be boosted to 100% to avoid air recirculation through PSUs
- 1U Project Olympus Power Supply FW Updates from the BMC must be performed one PSU (Silver Box) at a time. The PSU will ensure that each of its internal 340W Modules are updated in sequence. This flow will ensure that only one internal 340W module out of nine is powered down/reset across the three Silver Boxes.
- Due to Battery capacity limitations, no Battery related features will be supported on this system.
- When a PSU ALERT# signal is asserted, this is an indication that a critical event has occurred in one or more of the PSUs. The BMC FW must read all 3 of the PSU status registers and log a SEL Event.
- Must support the Remote Debug-At-Scale feature to access CPU debug information through the BMC dedicated management port.

## 3.2 I2C Block Diagram

The MB BMC must be able to access all PCIe Cards through I2C. Below are the I2C Device Mappings for the XPO200 3UA Server Configuration.

#### OCP • Project Olympus AMD® EPYC™ Processor 3U Server Specification



Figure 17: I2C Block Diagram for XPO200 3UA

Valid/Invalid	PCle Riser Slot # (Silkscreen	PCle Physical Slot #	I2C Bus (From BMC side, 0-	I2C Channel on First I2C MUX	I2C Channel on Second I2C MUX
	on Riser)	(1 to 5 bottom to top)	based)	(0-based)	(0-based)
Invalid	N/A	3.5	N/A	N/A	N/A

Table 7: I2C Bus Mapping for XPO200 3UA

	Valid/Invalid	PCle Riser Slot # (Silkscreen on Riser)	PCle Physical Slot # (1 to 5 bottom to top)	I2C Bus (From BMC side, 0- based)	I2C Channel on First I2C MUX (0-based)	I2C Channel on Second I2C MUX (0-based)
PCIo	Valid	2	3.4	0	4	1
Riser in	Invalid	N/A	3.3	N/A	N/A	N/A
Slot 3	Valid	1	3.2	0	4	3
01010	Invalid	N/A	3.1	N/A	N/A	N/A
PCle	Invalid	N/A	4.5	N/A	N/A	N/A
Riser in	Invalid	N/A	4.4	N/A	N/A	N/A
Slot 4	Invalid	N/A	4.3	N/A	N/A	N/A
	Valid	2	4.2	0	5	0
	Valid	1	4.1	0	5	1
PCle	Invalid	N/A	5.5	N/A	N/A	N/A
Riser in	Valid	2	5.4	0	6	1
Slot 5	Invalid	N/A	5.3	N/A	N/A	N/A
	Valid	1	5.2	0	6	3
	Invalid	N/A	5.1	N/A	N/A	N/A

## 4. Thermal Design Requirements

The following section provides a high-level overview of the XPO200 3UA Server Thermal Design Requirements.

## 4.1 Thermal Design Requirements

- Mechanically Supports up to 8 Dual Rotor 60mm x 56mm System Fans (6 Supported Electrically)
- Supports 158 CFM/kW @ 25C, 30C and 35C Ambient
- Supports Single-Fan Zone for System Components (not including PSUs)
- Supports two rotor Fan-Fail scenario: N+2 condition (CFM/kW Spec not required during Fan-Fail)
- The system must operate at full functionality in an "N" Fan condition at the highest workload and highest ambient defined.
- For dual rotor Fan numbering: "a" Fans on inlet side and "b" Fans on exhaust side
- Rear of chassis supports plastic Fan Flappers for Fan-Fail condition to prevent air recirculation.

#### OCP • Project Olympus AMD® EPYC™ Processor 3U Server Specification



Figure 18: 3U Chassis System Fan Position Diagram with Closed Fan Flappers



Figure 19: 3U Chassis with Open Fan Flappers



Figure 20: System Fan Connectivity Diagram

- Supports isolated PSU Fan Zone with sheet metal and plastic ducting
- For lightly loaded configs where a limited number of PCIe cards are plugged into PCIe Riser 4 (Middle Riser) a plastic baffle can be used to block airflow through the middle of the Chassis. The baffle forces more airflow across the high powered PCIe Cards on the outer Riser Slots.

#### OCP • Project Olympus AMD® EPYC™ Processor 3U Server Specification



Figure 21: 3U Chassis Cross Section Showing Middle Riser Baffle (to block airflow)

## 4.2 PCIe Riser Temp Sensor Locations



Figure 22: 2-Slot Active Riser Temp Sensor Locations



Figure 23: 2-Slot Passive x32 Riser Temp Sensor Locations

#### 5. Mechanical Requirements

#### 5.1 Chassis Requirements:

- 3U 19" EIA310-D Compliant Chassis
- Supports Project Olympus 42U or 48U Racks
- Supports Custom King Slide Rails (Does NOT support Project Olympus T-Pin Tracks)
- Supports Project Olympus PMDU power connections
- Supports OCP Approved AMD® EPYC<sup>™</sup> Olympus Motherboard
- Supports quick release front latches for Blade Hot Swap from Rack
- Supports Split Top Cover for convenient Serviceability of components in front MB Bay
- Supports clean cable management with strategically placed cable tie mounts
- Supports Power Distribution Board on elevated shelf in Rear Bay
- Supports 2.5" or 3.5" HDD/SSD on elevated shelf in Rear Bay
- Supports 3 x Project Olympus 1kW PSUs
- Supports up to 8 x Dual Rotor 60mm Fans (6 fans supported electrically)
- Supports 3 x Riser Cages and retention guides/brackets for PCIe Risers and single-wide/double-wide cards respectively

#### 5.2 Chassis Views



Figure 24: 3U Chassis/System Layout (Iso View)



Figure 25: 3U Chassis/System Layout (Top View)



Figure 26: 3U Chassis/System and PSU Layout (Rear View)

Due to the weight of the XPO200 3UA Server (90lbs), the standard Project Olympus T-Pin Slide Rails will not allow the Chassis to pull in and out smoothly. As a result, a custom King Slide Rail Kit is used to support easier travel during Servicing. However, since the width of the chassis is 17.36" (44.10 cm) thinner Slide rails were required to fit into the 19" Project Olympus Rack, which limits how far the XPO200 3UA Server can be pulled out of the Rack. The XPO200 3UA Server can only be pulled out about 67%,

with 650mm-travel ball bearing sliding rails. For this reason, the Top cover is split into two sections (Front and Rear). The 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 access.



Figure 27: 3U Chassis/System with Split Top Cover and King Slide Rails



Figure 28: 3U Chassis/System in 19" Project Olympus Rack



Figure 29: 3U Chassis/System in 19" Project Olympus Rack (Extended 67%)

#### 6. System Power Requirements

The XPO200 3UA Server will support the 1U Project Olympus 1kW Power Supply. Due to System Power Requirements, multiple Power Supplies in parallel will be required. This section defines how the power delivery sub-system will be designed, controlled and monitored.

#### 6.1 System Power Budget

The following section provides System Power Budgets for the 3U PCIe Expansion Server Configurations. Specific System SKUs will require separate Power Budget analysis to ensure they remain within the Max TDP and Max Peak loading envelope.

Function	Device		Watts (W)	21	VREff	Subtotal (W)
CB1	AND EDVC 18/0W 32C 2/CH	TDP	180	2	90%	400.0
020	7MD EP 10 10011,320, 2012	PMAX	240	2	90%	533.3
System Memory	32GB DDR4, 2400MHz, 2Rx4 RDIMM		δ	16	90%	9.88
PCIe Riser Card	PCIe 2 x16 S bt Riser Card with PLX PEX8747 PCIe S witch		15.5	2	90%	34.4
GPU Card Double Mide	AMD M(25 (Veca 10) GRU Card	TDP	300	4	100%	1200.0
0.0.00000000000000000000000000000000000	rived in zo (vega iv) or o caro	PMAX	552	4	100%	2208.0
Network	PCIe SFP+Nic		7.22	1	100%	7.2
BMC	AST2520 with DDR4 Memory		1.7	1	90%	1.9
USB	US83.0		2.5	2	90%	5.6
SSD	M2 Samsung P M963 960 GB PCI e SSD		7.8	8	90%	66.7
S∳stem Fan s	Delta 6066 GFC0612SS-03ZGF		32.4	6	100%	194.4
Other	Drivers, logic, pulkups, etc		10	1	90%	11.1
					Total (TDP)	2010
	With Turbo Mode E	inabled		т	obl(PMAX)	3151
				PS	U Efficiency	92%

#### Table 8: XPO200 3UA Server Power Budget

#### 6.2 Power Distribution Board

This is a Power Distribution Board (PDB) designed for the XPO200 3UA Server, which supports load sharing of 12V Power from the 3 x Project Olympus Power Supplies (PSUs) to the entire system. The PDB also supports the pass through and combining of various management signals required for Project Olympus Rack Manager & BMC identification, control and monitoring.

#### 6.3 Major Component Placement



Figure 30: Power Distribution Board Layout

#### 6.4 PDB Power Block Diagram



Figure 31: Power Distribution Board Power Block Diagram

#### 6.5 PDB Configuration

The XPO200 3UA Server will use the PDB to Load Share across 3 x Project Olympus PSUs. The PDB uses Passive Droop Sharing between PSUs, meaning no V\_Sense or I\_Share signals are required. The PSU Setpoint should be 12.35V. The Project Olympus PSUs support 3 x internal, phase-balanced 340W PSUs that also support Passive Droop Sharing.

Since the XPO200 3UA Server supports 3 PSUs with 3 internal PSU Modules, this equals 9 total PSU Modules per Server. As a result, the Server supports an 8+1 redundancy model, where the system will ride through any single PSU module failure. Depending on the loading conditions a System may be able to ride through losing 2 PSU Modules. To maintain 8+1 PSU Module redundancy with margin, the max supported System Power Budget is approximately 2400W (DC).

Project Olympus PSUs typically support an integrated battery pack (680W max). However, the combined battery power of 3 x 680W packs does not meet the XPO200 3UA Server max System Power requirements. Consequently, the Non-Battery PSU will be used in this system and therefore, no Battery related features will be supported.



Figure 32: XPO200 3UA 3U Server Power Supply Load Sharing Diagram (Max Supported Power)

#### 6.6 PDB Power Delivery Requirements

The PDB supports a Hot Swap Controller for Power Telemetry and isolation if the current sense circuits are tripped. A Hot Swap Controller that can be used for this application is the ADM1278. All P12V System components are behind this Hot Swap Controller except for the MB, since the MB already has its own Hot Swap Controller. The PDB Hot Swap Controller sits on the same BMC I2C bus as the MB Hot Swap Controller, but with a different I2C address.

A 12V current sense chip that can be used for each power delivery branch, except for the MB power delivery branch, is the INA301:

- Branch P12V\_A: Riser #3 and 2 x High Powered PCIe Cards
- Branch P12V\_B: Riser #4 and 2 x High Powered PCIe Cards
- Branch P12V\_C: Riser #5 and 2 x High Powered PCIe Cards
- Branch P12V\_D: Fan Power

**NOTE:** All Fan Power comes directly from the PDB, so the Fan power connectors from the MB are not used.

Below is a list of additional components and features supported on the PDB:

- IO Expander: Device that supports exposing additional GPIOs from the PDB to the BMC through I2C.
- FRU: Contains identification information about the PDB.
- Temp Sensor: Supports Temperature Sensor readings that can be used by the BMC to ensure Thermal Levels remain within an acceptable range.
- 12V Vsense Pins (Reserved for Future Use): Used to support 12V remote sense for more optimal load sharing between Power Supplies. Not used on current version of the 3U PCIe Expansion Server or Project Olympus PSUs.

#### 6.7 PDB Miscellaneous IO Requirements

This section describes the miscellaneous power related IO signals that traverse the PDB from the Project Olympus Rack Manager to the Project Olympus MB and corresponding BMC.

- PMBUS: Supports communication channels from BMC to PSUs and PDB I2C devices
  - PSU1\_PMBUS has a private connection to BMC and PCH
  - PSU2\_PMBUS & PSU3\_PMBUS share a connection through an I2C MUX to avoid address contention
- PSU\_ALERT\_N: PSU output that represents a change in status in any one of the three PSUs. Although each PSU signal is Open Drain, they include 1kOhm internal pull-ups. These pull-ups in combination with the 4.7K pull-up on the MB create a strong equivalent pull-up and make it difficult for the PSU drivers to drive the PSU\_ALERT\_N signal low enough to meet the Vil\_max of the receivers on the MB. As a result, each PSU\_ALERT\_N signal has been isolated by its own Open Drain Buffer. If any PSU asserts its PSU\_ALERT\_N signal, the BMC should read the status registers of all three PSUs (through PMBUS) to find out what happened. The PDB PSU\_ALERT\_N signal feeds both PSU1\_ALERT\_N and PSU2\_ALERT\_N pins to the BMC and on-board Throttle logic:
  - PSU1\_ALERT\_N (thru 1U Conn): Feeds the CPU PROCHOT\_N pins for immediate CPU throttling.
  - PSU2\_ALERT\_N (thru 2U Conn): Feeds the PCIe Slot PWRBRK\_N pins for immediate PCIe Card throttling.
- PS\_ON\_N: PSU input used to Power-on/Power-off the PSUs. Each signal is Wire-OR'd together so that all three PSUs can be Powered-on/Powered-off at the same time.
- BLADE\_EN\_N: Pass-through signal from Rack Manager (RM) that supports initiating a Power-On/Power-Off event through assertion/de-assertion of PS\_ON\_N signal. Only BLADE\_EN\_N from PSU1 is tied to PS\_ON\_N. BLADE\_EN\_N from PSU2 and PSU3 can be monitored by the BMC (through I2C IO Expander) as a provision in case separate System actions are required; but are not currently supported in the 3U PCIe Expansion Server BMC FW.

- BLADE\_THROTTLE: Pass-through signal from RM that supports immediate throttling of the System. Only PSU1 BLADE\_THROTTLE is supported/required on the PDB since it ties directly into the PROCHOT\_N and PWRBRK\_N inputs of the CPU and GPU cards respectively.
- NODE\_ID[5:0]/SLOT\_ID[5:0]: Pass-through signal from RM that assigns a unique ID number to each Server (stored in its BMC). Each ID corresponds to a "U" location in the Rack. Only PSU1 passes the NODE\_ID bits to a given 3U Server/BMC. The other PSU NODE\_ID bits are not needed and subsequently left floating on the PDB.
- PSU\_LED[1:0]: PSU output that provides status indications for its corresponding PSU. Since the MB only supports 2 sets of PSU\_LEDs, the PDB supports a separate PSU\_LED header/cable, which houses 3 x Bi-Color LEDs; one for each PSU. See following section for more details on the PSU\_LED Cable.
- BLADE\_PRESENT\_N: Passthrough signal from BMC to RM that ensures that each Server is present. In the case of this system, only BLADE\_PRESENT\_N from PSU1 is used. PSU2 & PSU3 BLADE\_PRESENT\_N signals are left floating.



Figure 33: Power Distribution Board Miscellaneous IO Block Diagram

## 6.8 Split Cable Power Harness for PCIe Cards

The section below describes the cable power harnesses used for the high power PCIe cards (150W, 225W, 300W). Depending on the pinout of the power connector on the PCIe Add-In card, a different power harness may be required. 12V Cable Power Harnesses are not required for standard 75W PCIe Cards.

Different 12V Split Cable Power Harnesses will be designed to support delivering power from the PDB to two high power PCIe Cards per PCIe Riser. Each high power PCIe Card will include one or more of 2x3 or 2x4 12V Auxiliary connectors to support additional power above the base 75W PCIe Slot power.

#### OCP • Project Olympus AMD® EPYC™ Processor 3U Server Specification

Below is the 12V Split Cable Power Harness for the high AMD® MI25 power PCIe GPU Card that supports two 2x4 12V Auxiliary connectors.





#### 6.9 Front Panel PSU LED Cable

A small Front Panel LED Cable will support 3 x Bi-Color LEDs to provide status on the 3 different Project Olympus PSUs. The Front Panel PSU LED Cable will be attached to the PDB and snap into apertures in the front right side of the 3U Chassis. All System diagnostic LEDs must be visible from the Front of the Chassis (cold aisle).

LED Signal Name	Color	Description
PSUx_GREEN_LED	Green	Solid On = AC and DC Power Good Blinking = Battery Power Good
PSUx_YELLOW_LED	Yellow	Solid On = Failure of 1 PSU Phase Blinking = Failure of 2 PSU Phases

Table 9: PSU	Status LED	Descriptions
--------------	------------	--------------

NOTE: Since this system only supports non-LES PSUs, the Battery LED states are not applicable.

#### OCP • Project Olympus AMD® EPYC™ Processor 3U Server Specification



Figure 35: Front Panel PSU LED Cable

## 6.10 System Power Capping ("Quick Response")

The Project Olympus 3U Server supports Power Capping as an extension of the Project Olympus Universal MB feature. For MB Power Capping triggers and flows, please refer to the Project Olympus Universal MB Specifications. This section only covers PSU and PDB level triggers of the 3U Server, which are intended to support the "Quick Response" Power Capping feature (i.e. less than 10ms).

The main trigger of a "Quick Response" System Power Capping event in the 3U Server is when a 12V Over Current (OC) event (Fault or Warning) is indicated in the STATUS IOUT command of any one of the three Project Olympus PSUs. In most cases, the Power Capping feature will not be needed, but it is intended as a precaution to avoid shutting down the System, until the event is over (if temporary), or the System can be properly serviced (if permanent). If an OC Event should happen, one of the three PSUs will assert its PSU ALERT N pin as a notification to the System to take action. As described in the PDB Miscellaneous IO Section of this Specification, the PDB combines the PSU ALERT N pins of all three Project Olympus PSUs and cables both the PSU1 ALERT N (drives CPU PROCHOT N pins) and PSU2 ALERT N (drives PCIe PWRBRK N pins) pins on the MB for instant CPU and PCIe Card throttling. The CPU and PCIe Cards are two of the highest power consumers in the system. So, by throttling these devices, the overall 12V current should be reduced enough for the System to ride through the OC Event without shutting down, thus maintaining high availability. While the BMC monitors this System Power Capping Event, it is not involved in the flow. For System Power Capping to help the 3U Server ride through an OC Event, the response time of the CPU and PCIe Card throttling must be approximately 10ms or less. The Power Capping flow through the BMC would not fall within this 10ms window, therefore the use of the immediate hardware PSU\_ALERT\_N signal is needed, or the System would shut down.

There are various causes of an OC Event that occur when System Devices drive up the power based on peak workloads, failed components, or some combination of these. One such condition that can trigger an OC Event is when one or more PSU Modules fail, and the remaining PSU modules are expected to

carry the load of the System. The 3U Server supports three PSUs, each with three internal PSU Modules for a total of 9 x PSU Modules. The 3U Server supports an 8+1 redundancy model, but under lighter loading conditions could also support a 7+2 or even a 6+3 redundancy model. Whatever the case may be, any time a PSU module is lost, the current draw on the remaining PSUs will increase. Depending on the System Loading conditions, or how many PSU Modules fail, the current can increase enough to cross the PSU OC Limit, thereby asserting the PSU\_ALERT\_N signal and forcing the System Power to be reduced.

**NOTE:** There is no PSU\_ALERT\_N signal assertion from the Project Olympus PSU when one or more PSU Modules fail. The PSU\_ALERT\_N signal is only asserted for OC events. It is not asserted for any other PSU failure events (over temp, Vin fault, Fan failure, etc...). If one or more PSU Modules fail, a SEL Event is logged by the BMC FW (during 1 second polling intervals), but there is no System Power Capping triggered; provided the remaining PSU Modules can handle the load. System Power Capping is only triggered when the PSU OC limit is tripped, which could potentially be caused by losing one or more PSU Modules depending on the System loading conditions. This means Power Capping is only used when the System needs it to protect the PSUs from shutting down in an over load situation.