



OPEN
Compute Project

Grand Canyon Storage System Specification

Revision 1.0
9/14/2022

Authors:

Abe Garcia, Hardware Engineer, Meta

Chenyu Xu / Dmitriy Shapiro, Mechanical Engineer, Meta

Dan Zhang, Software Engineer, OpenBMC, Meta

Jeff Puglis, Hardware Engineer, Meta

Madhavan Ravi, Hardware Engineer, Meta

Rajesh Kasukurthy, Thermal Engineer, Meta

Table of Contents

1. LICENSE	9
1.1. OCP CLA	9
2. OCP TENETS COMPLIANCE	10
2.1 Openness	10
2.2. Efficiency	10
2.3. Impact	10
2.4. Scale	11
3. REVISION TABLE	12
4. SCOPE	13
5. OVERVIEW	13
5.1 System Architecture	13
5.1.1 System Design and Flexibility	13
5.1.2 System Improvements from Bryce Canyon	15
5.2 System Components	16
5.2.1 Barton Springs (Codename for the Integrated 1S Server Card)	16
5.2.2 Storage Controller Card	17
5.2.4 IOC Module	19
5.2.5 E1.S SSD	21
5.2.6 OCP3.0 NIC	22
5.2.7 Drive Plan Board (DPB)	22
5.2.8 Power Transfer Board (PTB)	25
5.2.9 Power Cable-Track and Power Delivery Board	25
5.3 System Configurations	26
5.3.1 Dual Storage Server	28
5.3.2 Single Storage Server	28
5.3.3 JBOD	29
5.4 Systems Management Overview	30
6. RACK COMPATIBILITY	33
7. RACK IMPLEMENTATION	33
7.1 Type 5	176
7.2 Type 7	34
7.3 Compatibility with Open Rack Standard	35

8. SYSTEM INTERACTION AND SERVICEABILITY	35
8.1 Labeling of Logical Domains	35
8.2 LEDs and Buttons	38
8.2.1 Front Panel Buttons and LEDs	38
8.2.1.1 Power Button	39
8.2.1.2 Power LED	39
8.2.1.3 Reset Button	39
8.2.1.5 USB Debug Port	40
8.2.1.6 SSD LEDs	40
8.2.1.7 NIC LEDs	41
8.2.1.8 IOC Module LEDs	41
8.2.3 HDD LED behaviors	41
8.2.4 Fan Module LEDs	42
8.3 Hard Drive Serial Number Window	42
8.4 Service Operations	43
8.4.1 HDD Replacement	43
8.4.2 SSD Replacement	44
8.4.3 SCC Replacement	45
8.4.4 UIC Replacement	46
8.4.5 UIC TPM & BSM Replacement	47
8.4.6 NIC Replacement	48
8.4.7 Barton Springs Replacement	49
8.4.8 IOC Module Replacement	50
8.4.9 Fan Module Replacement	51
8.4.10 Fan Louver & Honeycomb Replacement	52
8.4.11 Drive Plan Board Replacement	54
8.4.12 Cable Track and Power Delivery Board Replacement	56
9. SYSTEM PERFORMANCE	58
9.1 Designing the system interfaces without bottlenecks	58
9.2 HDD performance degradation due to vibrations	59
9.2.1 Normal Operating Environment (sustained operating conditions)	60
9.2.2 Unsustained Operating Conditions	60
9.2.3 Design Criteria to Enable Chassis Longevity	60
10. ELECTRICAL DESIGN	60
10.1 PCIe Topology	60
10.1.1 PCIe Mapping	60
10.1.2 PCIe Clocks and Resets	62

10.2 System Management	64
10.2.1 Device Management signaling	64
10.2.2 I2C Topology	65
10.2.3 USB, SPI, JTAG	67
10.2.4 Network Controller Sideband Interface (NC-SI) requirements	68
10.2.5 UART	68
10.2.6 GPIO (Insert) Topology	69
10.2.8 Fan control	70
10.2.9 System Configurations	72
10.2.9.1 Integrated (Dual) Configuration	72
10.2.9.2 Integrated (Single) Configuration	72
10.2.9.3 JBOD Configuration	73
10.3 Signal Integrity requirements	73
10.3.1 SAS/SATA	73
10.3.2 PCIe	73
10.4 ESD Compliance	73
11. POWER	73
11.1 Power Topology	73
11.2 System Grounding Topology	74
11.2.1 Mechanical Interface Isolation Distance	75
11.2.2 Connector Housing Ground Topology	76
11.3 12V Busbar Cable Input	76
11.4 Drive Power Requirements	77
11.4.1 Staggered spin up requirements	77
11.4.2 Hot-Swap Requirements	77
11.4.3 5V Conversion Failure Domains	77
11.4.4 5V Conversion Loading Requirements	78
11.5 Hot-Swap Requirements	78
11.5.1 Capacitive Load	78
11.6 IOCM Power Topology	78
11.7 Storage Controller Card Power Topology	80
11.8 User Interface Card Power Topology	81
11.9 Power Budget	81
11.9.1 System Power Budget	81
11.9.2 SCC Connector Power Budget	82
11.9.3 IOCM Connector Power Budget	82
11.9.4 UIC Connector Power Budget	82

11.10 Power Sequencing	82
11.10.1 SCC Power Sequencing	82
11.10.2 Barton Springs Power Sequencing	83
11.10.3 IOCM Power Sequencing	83
11.10.4 UIC Power Sequencing	84
12 INDIVIDUAL PCBA SPECIFICATIONS	85
12.1 Drive Plane Boards (RDPB & FDPB)	85
12.1.1 Connectors	85
12.1.1.1 Compute Slots	176
12.1.1.2 SCC Connectors	176
12.1.1.3 E1.S/IOCM Connectors	85
12.1.1.4 NIC Connectors	85
12.1.1.5 UIC Connectors	85
12.1.1.6 HDD Connectors	176
12.1.1.7 12V System Power Input	86
12.1.1.8 DPB to DPB Power Connectors	86
12.1.1.9 DPB to DPB Signal Connectors	86
12.1.1.10 Fan Connectors	86
12.1.2 LEDs	86
12.1.2.1 Drive Activity	86
12.1.2.2 Drive Fault/Identification	86
12.1.3 Thermal Sensor Placement	86
12.1.4 I2C Component Placement	87
12.1.5 Drawer Detection Sensor	87
12.1.6 Drive Control Signals	87
12.1.7 PCB Details	88
12.1.7.1 Front DPB	88
12.1.7.1.1 Important Components and Features	88
12.1.7.1.2 PCB Key Dimensions	89
12.1.7.2 Rear DPB	91
12.1.7.2.1 Important Component and Features	91
12.1.7.2.2 PCB Key Dimensions	92
12.1.7.2.3 PCB Stackup	93
12.2 Storage Controller Card (SCC)	93
12.2.1 Storage Controller Card Configuration Diagram	94
12.2.2 Storage Controller Card Configurations	176
12.2.2.1 Connector Configurations	176

12.2.5 Storage Controller Card PCB	95
12.2.5.1 Important Components and interfaces	95
12.2.5.2 PCB Key Dimensions	97
12.2.5.3 PCB Stack up	97
12.3 IOCM	98
12.3.1 IOCM Configuration Diagram	98
12.3.2 IOCM Connectors	176
12.3.3 IOCM Card PCB	176
12.3.3.1 Important Components and Interfaces	100
12.3.3.2 PCB Key Dimensions	101
12.3.3.3 PCB Stackup	102
12.4 User Interface Card (UIC)	102
12.4.1 UIC Variants	102
12.4.2 BMC	176
12.4.3 CPLD	176
12.4.4 Connectors	104
12.4.5 USB Debug	105
12.4.6 BSM and TPM 2.0	105
12.4.7 User Interface Card PCB	106
12.4.7.1 PCB Key Dimensions	106
12.4.7.2 PCB Stack up	106
12.5 Power Delivery Board (PDB)	107
12.5.1 Configuration Diagram	107
12.5.2 Connectors	176
12.5.3 PDB PCB	176
12.5.3.1 Important Components and Interfaces	108
12.5.3.2 PCB Key Dimensions	109
12.5.3.3 PCB Stack up	109
12.6 Power Transfer Board (PTB)	110
12.6.1 Configuration Diagram	110
12.6.2 Connectors	176
12.6.3 PTB PCB	111
12.6.3.1 Important Components and Interfaces	111
12.6.3.2 PCB Key Dimensions	113
12.6.3.3 PCB Stack up	113
13. INTERNAL CABLES	114
13.1 JBOD Connection Bracket	114

13.2 Single Server SAS Cable	115
14. EXTERNAL CABLES	115
14.1 JBOD configuration	115
14.2 Cable Details	116
15 BMC OVERVIEW	117
15.1 Introduction	117
15.2 Front Panel	117
15.2.1 OCP Debug Card	117
15.3 Power Control	117
15.4 Remote Console	120
15.5 IPMI FRUID	120
15.6 Sensor Information	121
15.7 Firmware Information/Update	121
15.8 Log Information	122
15.9 Fan Speed Control	122
15.10 Enclosure Management	122
15.10.1 HDDs Status	122
15.10.2 Error Code	122
15.11 Event List Along With The Error Code	128
15.12 IPMI Command Support List	132
15.13 Expander IPMB Command Support List	145
16. EXPANDER FIRMWARE OVERVIEW	151
16.1 Drive LEDs	151
16.1.1 Drive Link LED	151
16.1.2 Drive Fault / Identification LED	152
16.1.3 Drawer Open Thermal Warning Behavior	152
16.2 Power Readings	153
16.3 Runtime System Configurations	153
16.4 Report Status of GPIOs	154
16.5 Even Report and Log	154
16.6 Security Features	154
17. THERMAL DESIGN SPECIFICATION/REQUIREMENTS	154
17.1. Data Center Environmental Conditions	155
17.1.1 Location of Data Center/Altitude	155
17.1.2 Cold-Aisle Temperature	155
17.1.3 Cold-Aisle Pressurization	155

17.1.4 Relative Humidity	155
17.2 System Thermal Operational Conditions	155
17.2.1 Inlet Temperature	155
17.2.2 Pressurization	155
17.2.3 Fan Redundancy	156
17.2.4 Service Mode	156
17.2.5 Delta T	156
17.2.6 System Airflow or Volumetric Flow	156
17.2.7 Thermal Margin	156
17.2.8 Thermal Sensor	156
17.3 Grand Canyon Thermal Design Requirements	176
17.3.1 Fans	157
17.3.2 Cooling Zones	157
17.3.3 92mm Fan Module	157
17.3.4 Air Baffles	159
18. MECHANICAL	162
18.1. Overview	162
18.2. Overall Dimensions	176
18.3. Rotational Vibration Requirements	176
18.4 Drive Plane Board	176
18.5 Drawer Structural Walls	164
18.6 Slide Rail	176
18.7 Bus Bar Power Module & Cable Track	176
18.8. Drawer Release Handles	176
18.9. Rack Release Latching	176
18.10 Chassis Lift Handles	168
18.11 Drive Retention	169
18.12 UIC and IOC Module	171
18.13 Storage Controller Card	171
19. ENVIRONMENTAL AND REGULATIONS	172
19.1 Regulatory Compliance	172
19.2 Product Safety Compliance	172
20. ENVIRONMENTAL REQUIREMENTS	173
20.1. Vibration & Shock	176
20.2 Regulations	176
20.3 Chassis Labels and Markings	176

20.4 Data Center Design Requirements	174
20.5 Mean Time Between Failure (MTBF) Requirements	175
15.5.1 MTBF Prediction	175
20.5.2 MTBF Demonstration Test	175
20.5 Certifications	175
21. PRESCRIBED MATERIALS	175
21.1 Disallowed Components	176
21.2 Capacitors & Inductors	176
21.3 Component Derating	176
21.4 Sustainable Materials	176
22. SYSTEM FIRMWARE	177
23. HARDWARE MANAGEMENT	177
23.1 Redfish Compliance	177
23.2 Source Availability	177
24. SECURITY	178
25. REFERENCES (OPTIONAL)	178
APPENDIX A - REQUIREMENTS FOR IC APPROVAL (to be completed Contributor of Baseline Spec)	178

1. LICENSE

1.1. OCP CLA

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: **Meta, Inc.**

You can review the signed copies of the applicable Contributor License(s) for this Specification on the OCP website at <https://www.opencompute.org/legal-documents>.

Usage of this Specification is governed by the terms and conditions set forth in **Open Compute Project Hardware License – Permissive (“OCPHL Permissive”)** also known as a **“Specification License”**.

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: None.

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.

2. OCP TENETS COMPLIANCE

2.1 Openness

Grand Canyon system is open to enable large scale deployments in hyperscale type of environments leveraging OpenBMC, OCP NIC, industry standard interfaces and a modular form factor design concept. It provides flexible and open interfaces for future expandability.

2.2. Efficiency

Grand Canyon offers efficiency gains across the thermal management system due to mechanical improvements, performance efficiency gains due to increases in performance/watt as well as performance/TB, improved vibration mitigation to allow interoperability with future HDD capacities.

2.3. Impact

The Grand Canyon platform demonstrates flexibility of configurations and modularity of design. It

improves performance/watt, performance/TB, provides improved vibration for reliability while being serviceable at scale. It should be noted that the HDD slot vibration mitigation techniques enable the system to be leveraged in the future for multiple HDD generations.

2.4. Scale

Grand Canyon enables serviceability gains as all major system components are serviceable to be maintained at scale. It is designed such that CPU Board (Barton Springs) can be easily upgraded to the next generation of CPUs.

3. REVISION TABLE

Date	Revision	Author	Description
9/14/2022	1.0	Jeff Puglis	Ready for OCP release

4. SCOPE

This document defines the technical specifications for the Grand Canyon Storage System, hereinafter referred to as “Grand Canyon,” used in the Open Compute Project.

Grand Canyon is Meta’s latest HDD storage server chassis as a follow-on to the previous generation storage platform, Bryce Canyon. The Grand Canyon design retains the high level system architecture from Bryce Canyon - 4OU tall, single drawer dense storage chassis that contains up to 72 3.5” hard disk drives (HDDs), 2 compute modules, 1 NIC and 2 SSDs per compute module. This design will have an upgraded 1S Server card, UIC, Storage controller card and is Open Rack v2 (ORv2) compliant.

5. OVERVIEW

Grand Canyon is a refresh to the previous generation, Bryce Canyon. There were 4 major reasons that gave motivation to the need for the Grand Canyon project.

1. The need to address the shrinking runway for component and commodity end-of-life on Bryce Canyon and to enable continued HDD efficiency at Meta.
2. Improve the System Performance headroom (compute, memory bandwidth, interface bandwidths, etc) to ensure the HDD storage server SKU can push drives to their limits and allow introduction of software features that consume system resources but enable efficient use of drives.
3. Improve the HDD vibration damping features (both structural and acoustic) in the chassis via add-on options to allow for dense HDD capacity adoption in the chassis (driven by recent projections from HDD vendors using energy-assisted recording technology).
4. Introducing the 9.5mm heat sink E1.S form factor SSD

5.1 System Architecture

5.1.1 System Design and Flexibility

The architectural intent behind the Grand Canyon design is a flexible HDD-Storage hardware platform that enables multiple configurations for Meta’s needs. From a high-level architecture perspective, Grand Canyon will be an upgrade to Bryce Canyon while retaining a lot of similarities.

The Grand Canyon storage system is a 4OU system, with a majority of the components installed and serviced from the top. To enable this, there is an inner drawer design that fully extends the unit out of the rack, while it is still operational. A mechanical overview is shown in Figure 5.1

Multiple configurations of the chassis exist to satisfy various applications. The architecture is designed to maximize common components between configurations and modularize the components that differ. The configurations supported in Grand Canyon are dual storage server, single storage server and

JBOD. In its dual storage server configuration, the chassis supports up to 72 drives and two compute modules. Each compute module connects to a Storage Controller Card (SCC) that controls 36 drives and is logically separate from the other compute module and the 36 drives it controls.

The only common component shared by these two compute-storage groups is the Drive Plane Board assembly (like a backplane) and incoming power (12V) delivered by the bus bar clip. Grand Canyon will also design in an OCP3.0NIC (SFF) and 2x E1.S SSDs per compute node in the chassis. The chassis is designed such that the highest failure rate components (HDDs, SSDs, NICs, server card, other PCBAs) are hot-swappable without affecting the other host that resides in the same enclosure.




Figure 5.1: System Overview

5.1.2 System Improvements from Bryce Canyon

Table 5.1: System Improvements

	CONSIDERATION	DETAILS
Storage Solution	Implement front-serviceable data SSD with provision to hotswap i.e., without system power off	Pivot to front loadable E1.S form factor Hot-swappable data SSD improves system availability
	Designing with HDDs and SSDs as the server's IO bottlenecks for all use cases	Grand Canyon will start off with 50Gbps NIC and provide an upgradability option to use 100Gbps NIC, if needed
	Designing for Dual actuator adoption	Power delivery, thermal and vibration control will be designed into the chassis to support Dual actuator adoption. CPU server card and NIC upgrades to support dual-actuator drives will be drop-in replacements
DC Operations	Power cable track redesign to improve serviceability and reliability	Addition of a new power delivery board in the mid-section to keep the power cable constrained in the mid-section entirely
	Better serviceability of the Drive Plane Board	Designed for slide-out and slide-in replacement of the Drive Plane Board (DPB)
	Design for both ORv2 and ORv3 rack configurations	Design flexibility to ensure Grand Canyon can MP with Orv2 rack configuration but have the flexibility to be upgraded to Orv3 in future. <ul style="list-style-type: none"> Power delivery flexibility to move from 12V to 48V busbar Mechanical flexibility to move to Orv3
Components	CPU	Intel Cooperlake 26 core - TDP =91W (including PCH)
	DIMM configuration	Dual Storage server (T5) – 4x16GB DDR4 Single Storage server (T7) – 4x32GB DDR4
	Boot SSD	m.2 2280 NVMe SSD
	Data SSD configuration	2x 2TB E1.S SSD per compute node inside Grand Canyon i.e., total of 4x 2TB E1.S SSDs per Grand Canyon chassis
	IOC	Broadcom Avenger - SAS4016W
	Expander	Broadcom Margay - SAS4x48

	BMC	Aspeed AST2620
	NIC	OCP3.0 single host multi-sourced NICs
	Fans	Separate cages per Grand Canyon chassis Rack Type 5 – 1 dual-rotor fan per cage Rack Type 7 – 1 single-rotor fan per cage Integrate FRU eeprom inside fans for traceability and I2C system interconnect to Fans
Vibration	Improving vibration dampening from structural and acoustic vibration sources in the chassis.	The chassis shall design in the following options: <ol style="list-style-type: none"> 1. Fan cage isolation from chassis via elastomer padding, fan grommets 2. Enable acoustic foam installation at strategic locations 3. DPB air gap seal to reduce air leakage and bring max fan PWM lower than 100%. 4. Lower fan blade count and better geometries 5. New material on HDD latches, if needed 6. Options to mount fan louvers to reduce air leakage on fan-failure and lower fan speeds to be less than 100% 7. Retain wireframe fingerguard from Bryce Canyon

5.2 System Components

This section outlines the major PCBAs and FRU-able commodities that make up the system, other than the HDDs.

5.2.1 Barton Springs (Codename for the Integrated 1S Server Card)

Barton Springs is the integrated 1S server card. It is a top-accessible PCBA once the drawer is pulled out of the chassis tub. As Compute workload and needs for DRAM bandwidth, connected SSD, special instructions on the CPU, etc varies quite a bit from Storage, advantages are achievable only when the Compute card designs are flexible and agile enough to adapt to evolving needs. In order to reduce the churn on a Storage chassis design tied to a server card form-factor that cannot be predicted to be constant, Grand Canyon breaks away from reusing an existing design and will standardize on a new form-factor defined by the maximum remaining volume in the chassis while accommodating 72 HDDs.

Barton Springs shall be designed to this form-factor specifically for Grand Canyon chassis to ensure we can achieve dense HDD packing within the chassis for greater advantages, make the server card form-factor modular and long-lasting to allow for CPU upgrades on the same chassis. Shown below in Figure

5.2 is a rendering of Barton Springs card with the CPU w/ heatsink, PCH w/ heatsink, a client m.2 SSD for boot purposes, and 4 DIMM sockets (1 DIMM per channel).




Figure 5.2: Barton Springs PCBA

The design details of the server card will be covered under the Barton Springs 1S server card design specification document.

5.2.2 Storage Controller Card

The Storage Controller Card (SCC) is a top-accessible module (with the drawer extended) that contains the ICs required to translate PCIe connectivity from the Server card to SAS/SATA for the HDDs. The SCC supports:

- SAS IOC, a.k.a., HBA
- SAS Expander
- SAS connections to the front bulkhead (to enable JBOD [Just a Bunch Of Disks] functionality)

The SCC connects to Barton Springs via a PCIe Gen3 x16 link. There are 2 configurations of the SCC, normal (IOC+Expander populated) and JBOD (only Expander is populated).




Figure 5.3: Storage Controller Card (SCC)

The User Interface Card is a low speed management card that houses the BMC SoC, Trusted platform Module (TPM), and a removable BMC Storage module or BSM in an m.2 2230 form factor. Since the E.1S SSD and OCP3.0 NICs are front pluggable form factors, we envision these to plug into the Front drive plane board directly and the need for an IO module card (used in Bryce Canyon) now reduces to that of the UIC.

In addition to these components, the UIC also brings out on its front panel, a debug USB Type-A connector, a bicolor status LED and a mechanical power button for one half or one node of the chassis.

- The debug USB-A connector should multiplex the UART from the BMC, Expander and host and toggle between them using the UART-SEL button on the OCP debug card.
- The bicolor status LED should indicate good health or fault status from BMC, host, and expander.
- The mechanical power button controls the enable of the Hot swap controller (HSC) that enables power for each logical half of the chassis.




Figure 5.4: User Interface Card (UIC)

The UIC is a front accessible PCBA which plugs into the FDPB using SFF-TA-1002 card edge connector. Since the BMC sub-system is fully contained in the UIC, Grand Canyon now has the provision to upgrade the BMC solution without affecting the rest of the chassis design.

The dual and the single storage server configurations will deploy the full blown UIC.

There needs to be a second version called 'UIC w/o BMC' without the BMC, BSM, TPM circuitry. This version brings out the front panel elements only for use on the single storage server and JBOD configurations.

5.2.4 IOC Module

The single storage server configuration used in Type 7 rack utilizes the IOC module to allow for SAS expansion from the single storage server to 2x JBOD chassis. The IOC module is a front accessible PCBA that plugs into the front drive plane board (FDPB) in-place of the 4x E1.S SSDs.




Figure 5.5: IOC Module

5.2.5 E1.S SSD

As explained in the System Improvements section above, Grand Canyon will be designed with 4x E1.S SSDs per chassis, 2x per Barton Springs. E1.S has separate thermal case options. To make the thermal margins acceptable and to ensure we have a low fan RPM to ensure less vibration impact on the HDDs, it is imperative that we choose an option that's better than the bare PCB E1.S option.

To stay within the 4OU chassis form factor, the only front pluggable form-factors that work in Grand Canyon are the heatspreader option and the 9.5mm symmetric case. The heat-spreader option requires Meta to design its own bracket to mount the SSD and make it a front-loadable FRU. In lieu of this, Grand Canyon will opt to design in the 9.5mm E1.S SSD case option.




Figure 5.6: E.1S SSD

5.2.6 OCP3.0 NIC

Grand Canyon hosts 2x single host OCP3.0 NICs in the small form factor (SFF), one NIC per Barton Springs.




Figure 5.7: OCP 3.0 NIC

5.2.7 Drive Plan Board (DPB)

The Drive Plane Boards (DPBs) sit at the base of the chassis. There are physically two Printed Circuit Boards (PCBs) that make up the drive plane board assembly, and they are connected via high speed connectors and an air baffle to create a single Field Replaceable Unit (FRU). The DPBs connect all of the boards in the system together, along with the drives, fans, etc. Its function is to support connectivity for the following:

- Single port SAS/SATA connections to HDDs (up to 72)
- PCIe Gen3 x4 connectivity to E1.S SSDs (up to 4) or PCIe Gen3 x8 connectivity to an IOC module (for JBOD expansion)
- PCIe Gen3 x16 to NICs (up to 2)
- Dual-rotor counter-rotating Fans (up to four)
- Drive power control (individual 5V/12V power control to each HDD and SSD slot)
- Sensors (ambient temperature, drawer open/intrusion)

- Incoming power (12V) distribution






Figure 5.8: Drive Plane Board (DPB) Assembly

The population of these components is dependent on the overall system configuration (see Section). Shown below is a snapshot of the population options for the front accessible devices on the DPB. To enable this, there are 2 PCBA versions of the DPB assembly – one for Type 5 rack and one for Type 7 rack.



Type 5 (Dual Storage Server) Front Access Device Configuration



Type 7 Headnode (Single Storage Server) Front Access Device Configuration



Type 7 JBOD Front Access Device Configuration

Figure 5.9: DPB Front Access Device Configurations (Bottom View)

5.2.8 Power Transfer Board (PTB)

The DPB shall be designed to distribute 12V as the main power plane to all loads in the system. In ORv2 rack, this DPB 12V can tie directly into the busbar 12V supply.

Grand Canyon PCBAs should be designed in such a way that it is flexible enough to be reused for the ORv3 design with 48V busbar. The Power Transfer Board serves the purpose of abstracting the busbar voltage change away from the rest of the system and delivering only a constant 12V to the rest of the system. It is recommended to physically fit in a 48V to 12V power conversion solution (either a Brick or Discrete) in a thermally optimal location behind the DPB so as to enable the same DPB and system PCBAs to be reused for an ORv3 configuration in the future. The expectation is to reuse the chassis, PCBAs and layout as-is without change, so Grand Canyon designed today should allocate room for the power conversion solution and enable a bypass option for use in the ORv2 case.

5.2.9 Power Cable-Track and Power Delivery Board





Figure 5.10: Power Delivery Board (PDB) & Power Transfer Board (PTB)

5.3 System Configurations


This section outlines the different possible configurations of the overall system. There is a base Grand Canyon chassis configuration, but a few Field Replaceable Units (FRUs) can be installed to create simple JBODs, or single and dual integrated storage servers. Table 2 shows the current configurations considered for production.




Configuration Name	Rack Type	1S Server Card				Storage Controller Card				IOC module		OCP 3.0 single		E1.S SSD		Fans		Drives	
		# Server Cards	# CPU Core	Mem. Capacity	Mem. Config	# SCC	SCC Type	Expander	IOC	# IOC module	SAS IOC	# NIC	Data rate	# SSD	SSD capacity	# Fans	Fan size	# Drives	# Drives Spun up
Dual Storage Server	Type V	2	TBD	64GB	4x 16GB	2	Normal	Yes	Yes	0		2	50Gbps	4	2TB	4	92mm	72	72
Single Storage server	Type VII Headnode	1	TBD	128GB	4x 32GB	2	1x Normal, 1x JBOD	Yes	*Yes	1	1	1	50Gbps	0		4	TBD	72	~5
JBOD	Type VII JBOD	0				2	JBOD	Yes	No	0		0		0		4	TBD	72	~5

Figure 5.11: Grand Canyon System Types and Configurations


The 3 configurations are differentiable from their faceplate as well as can be seen on the following pages.



Type 5 (Dual Storage Server) Configuration Faceplate



Type 7 Headnode (Single Storage Server) Configuration Faceplate



Type 7 JBOD Configuration Faceplate

Figure 5.12: Faceplate Configurations

5.3.1 Dual Storage Server

The nominal configuration described thus far is the Dual Storage Server configuration. This configuration is used for Type 5 racks. The block diagram for this configuration is shown in Figure 5.13.

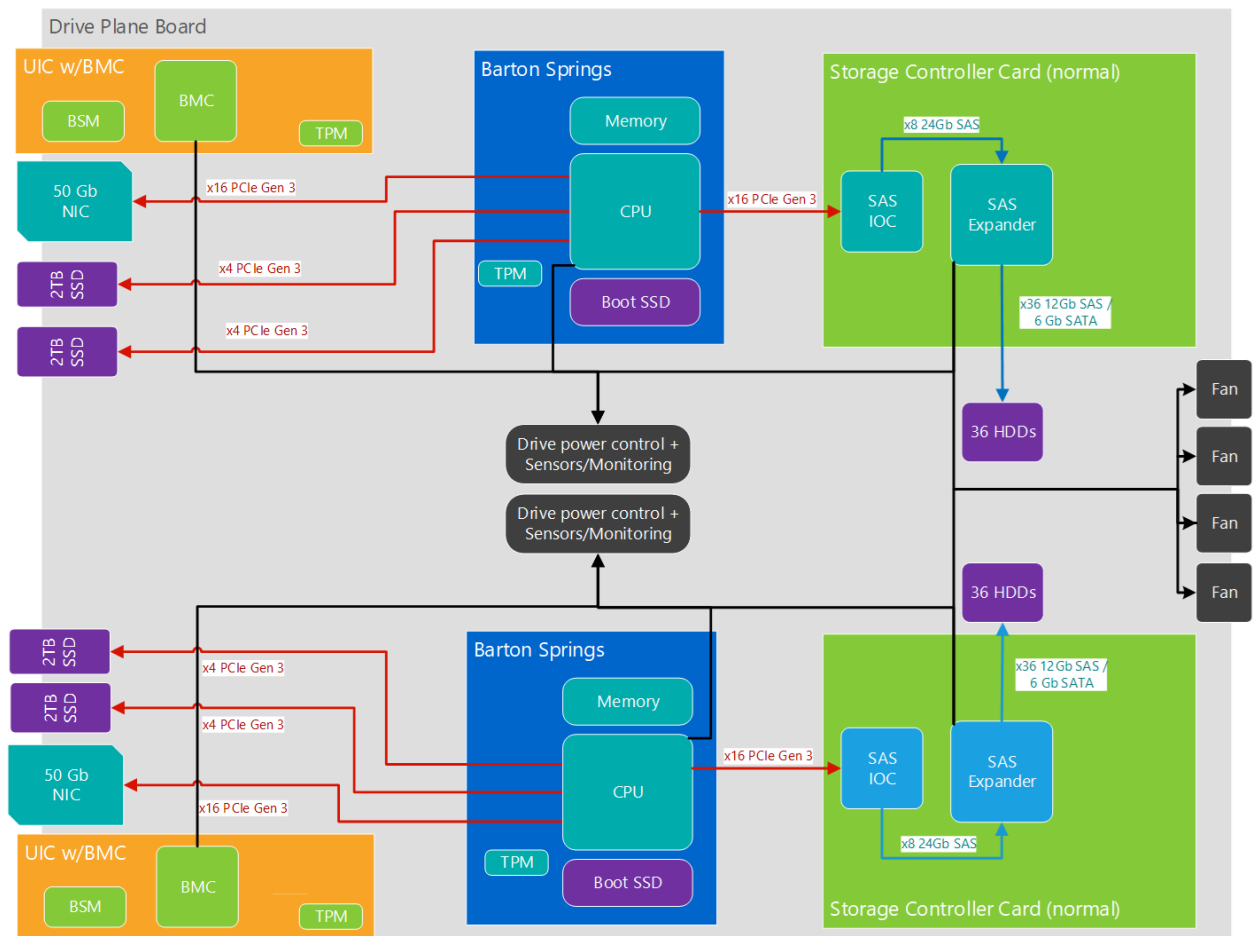


Figure 5.13: Dual Storage Server System Configuration Block Diagram

5.3.2 Single Storage Server

The Single Storage Server configuration is intended for applications that use a lower compute-to-storage ratio i.e., more disks per compute card. The block diagram is shown in Figure 5.14. This configuration serves as the head node for a Type 7 configuration, with several JBODs connected downstream to increase the number of drives behind a single server node. To connect all 72 drive slots within the chassis, the internal MiniSAS HD connector on both SCCs can be connected via an x4 cable.

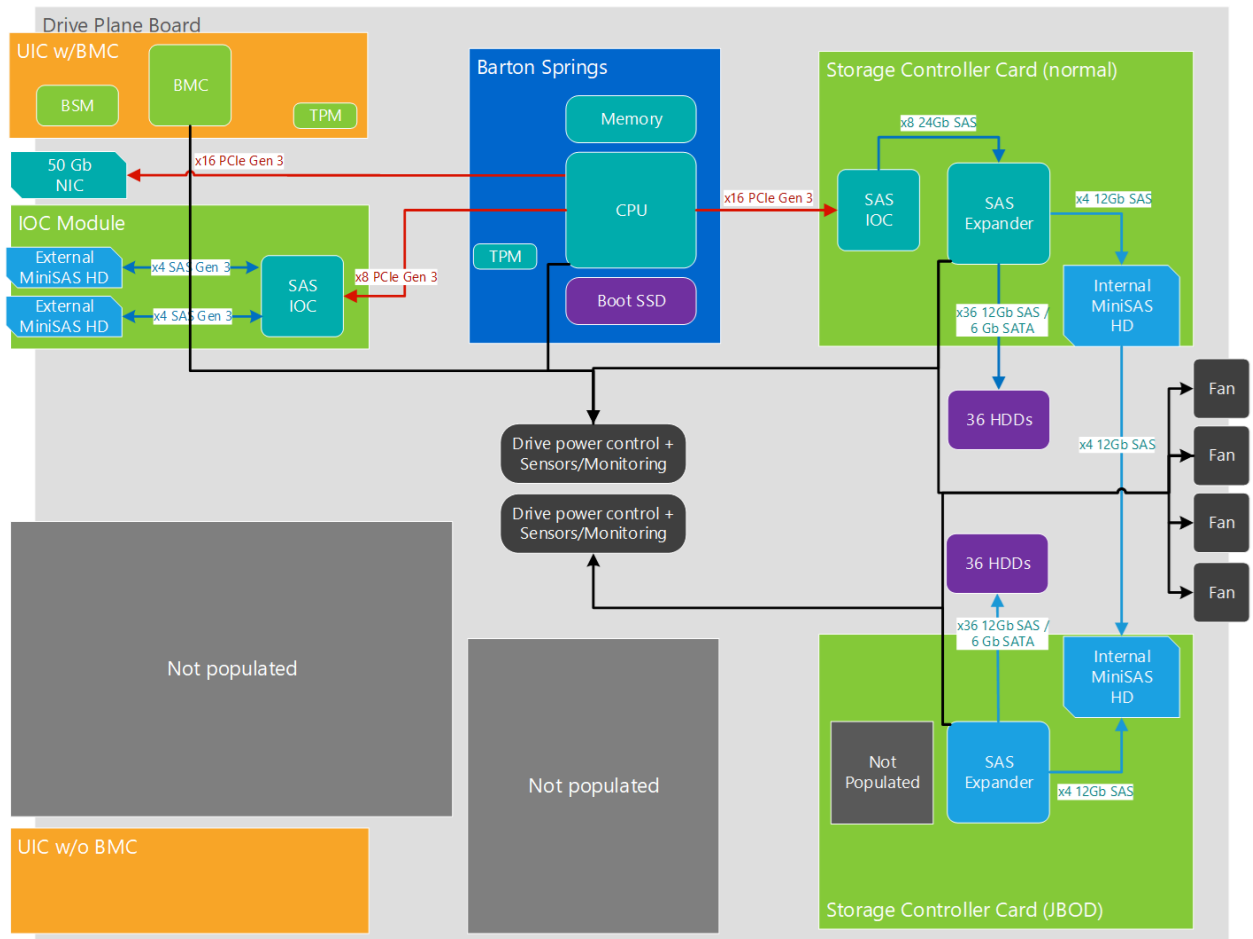


Figure 5.14: Single Storage Server System Configuration Block Diagram

5.3.3 JBOD

The JBOD configuration allows for an external compute module to connect to the Grand Canyon system. This compute module is cabled to the system via a SAS connector that is exposed on the front panel of the chassis.

In this configuration, the SCC does not have a SAS IOC populated. The server cards, NICs, SSDs are also not installed, and the fan control is coordinated by the SAS expander in lieu of the UIC w/o BMC.

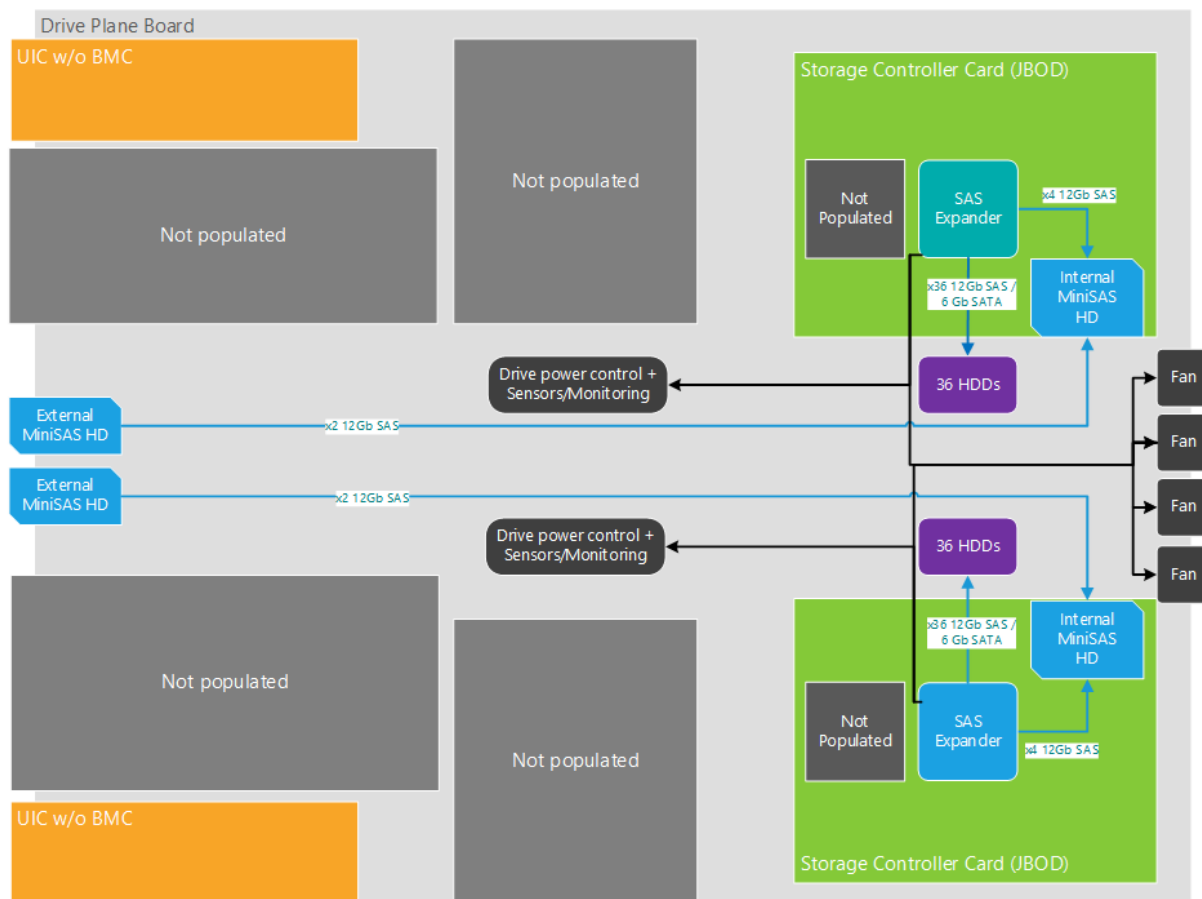


Figure 5.15: JBOD Server System Configuration Block Diagram

5.4 Systems Management Overview

The system is managed by two to four controllers (two BMCs and two expanders). Each of these controllers (referred to as peers) have ownership of part of the chassis. Interactions among these peers are limited to:

1. Insertion detection of all peers and Barton Springs
2. Heartbeat of all peers
3. Access to master hot swap controller (through multimaster I2C)
4. Access to TACH sensor (through multimaster I2C)
5. Fan speed control through PWM comparators

At a high level, the BMCs own the sensors on the UIC, Barton Springs, E1.S and the IOC chip on the SCC. The expanders own other sensors in the chassis, including HDD S.M.A.R.T. temperature.

The BMCs and SCCs all generate Pulse Width Modulation (PWM) for fan speed control using sensor inputs to determine fan speed. These PWM signals are compared using a PWM comparator and the highest requested fan speed is selected and used to control the setting for all 4 fan's PWM speed.

Each peer in the system monitors the insertion and heartbeat of other peers. If an expected peer is not presented or does not generate proper PWM, the firmware should increase fan PWM to a predefined value (e.g., 100%).

I2C device ownership information (per-configuration) can be found in the following Tables.

BMC owned I2C devices:

BMC I2C Bus #	I2C Devices (Slave Address)	Comments
2	BIC (0x40)	IPMB connection between BMC and BIC (Bridge IC) on Barton Spring Server Card
3	BS FPGA (0x1E)	Barton Springs FPGA Slave Address
	Barton Springs FPGA Version (0x80)	Barton Springs FPGA Version Register
4	TMP75 (0x94)	UIC Temperature Sensor
	M24C64 (0xA0)	UIC FRU EEPROM
5	UIC FPGA (0x1E)	UIC FPGA Slave Address
	UIC FPGA Version (0x80)	UIC FPGA Version Register
6	EEPROM (0xA8)	BSM FRU
7	PCA9555 (0x4E)	I/O Expander for Decoding for Debug Card

8	PLDM Sensor (0x3E)	NIC Temperature and PLDM sensors
	NIC FRU (0xA0)	NIC FRU EEPROM
9	ADC128 (0x3A)	E1.S/IOCM and NIC Voltage Sensor (Main Source)
	LTC2991 (0x90)	E1.S/IOCM and NIC Voltage Sensor (2nd Source)
	TCA9555 (0x42)	I2C to GPIO E1.S & MISC
	TCA9555 (0x48)	I2C to GPIO Insert & PWRGD
10	SCC Expander (0xE2)	SCC Expander Slave Address
11	SCC IOC (0x68)	SCC IOC Slave Address
12 (Type 5)	NVMe SMBUS Address (0xD4)	E1.S 0 Temperature Sensor
12 (Type 7)	IOCM IOC (0x68)	IOCM IOC Slave Address
13 (Type 5)	NVMe SMBUS Address (0xD4)	E1.S 1 Temperature Sensor
13 (Type 7)	ADS1015 (0x92)	IOCM Voltage Sensor (Main Source)
	LTC2990 (0x98)	IOCM Voltage Sensor (2nd Source)
	TMP75 (0x94)	IOCM Temperature Sensor

	M24C64 (0xA0)	IOCM FRU EEPROM
--	---------------	-----------------

6. RACK COMPATIBILITY

NOTE: To achieve interoperability, new specifications and products seeking recognition shall be compatible with an OCP adopted architecture: ORv2, ORv3, or OpenEDGE architectures. This is required for storage enclosures and systems, rack frames, and any enclosure which holds a server sled or edge device sled. Compatibility is also desired of server sleds and edge sleds. All other device types not listed are not required to support this.

The Grand Canyon platform will be used in a few configurations to meet the needs of several rack types, as outlined in future sections. Grand Canyon is ORv2 compliant.

7. RACK IMPLEMENTATION

The Grand Canyon platform will be used in a few configurations to meet the needs of several rack types, as outlined in the following sections.

7.1 Type 5

In this configuration, there are two storage servers with 36 drives each inside the Grand Canyon chassis. The compute portion is an integrated 1S server.



Figure 7.1: Type 5 Rack Configuration

7.2 Type 7

There will be two different Grand Canyon configurations populated within a rack of this type. There will be a main Grand Canyon unit with 1 compute card, and two Grand Canyon JBODs (without a compute card) connected behind this, for a total of 216 drives in each functional system. In the base unit, there will be a single 1S server that connects to all drives in the base chassis.

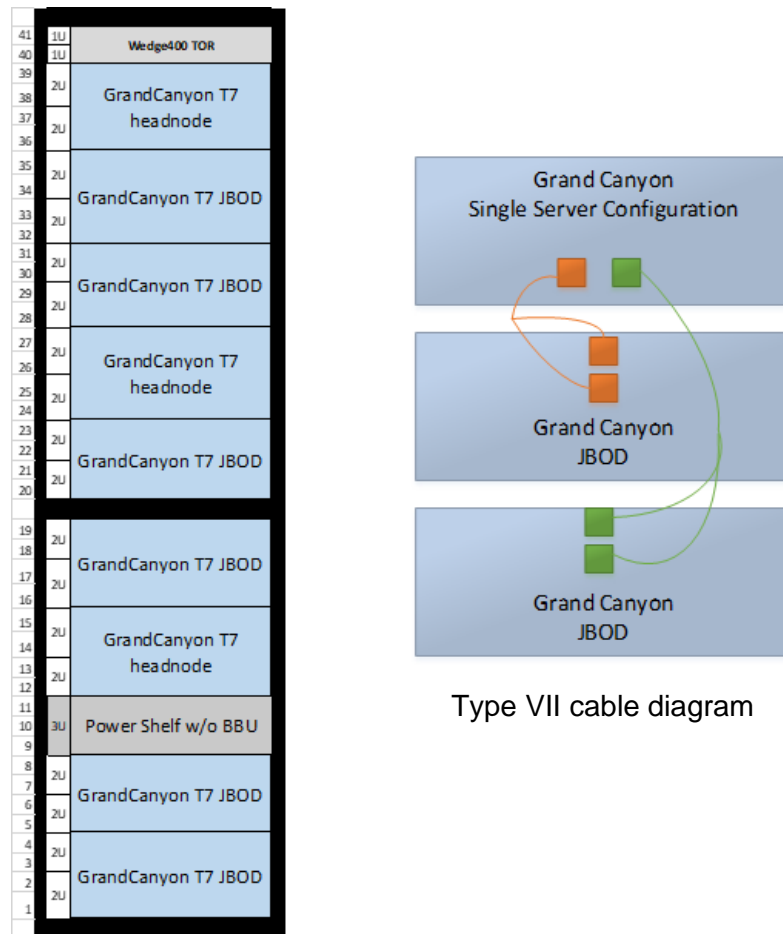


Figure 7.2: Type 7 Rack Configuration


7.3 Compatibility with Open Rack Standard

All specifications and products seeking product recognition for use in Open Rack shall comply with Open Rack specification revision 2.0.


8. SYSTEM INTERACTION AND SERVICEABILITY

8.1 Labeling of Logical Domains

The drive domains are physically divided between the front and rear of the chassis via the two DPBs. The compute modules are separated in a left/right fashion. On the following page, Figure 8.1 shows the 2 logically separate systems using yellow and blue highlights. Consequently, the correct logical division of the two domains must be made clear through highly visible labeling.



Type 5 (Dual Storage Server)



Type 7

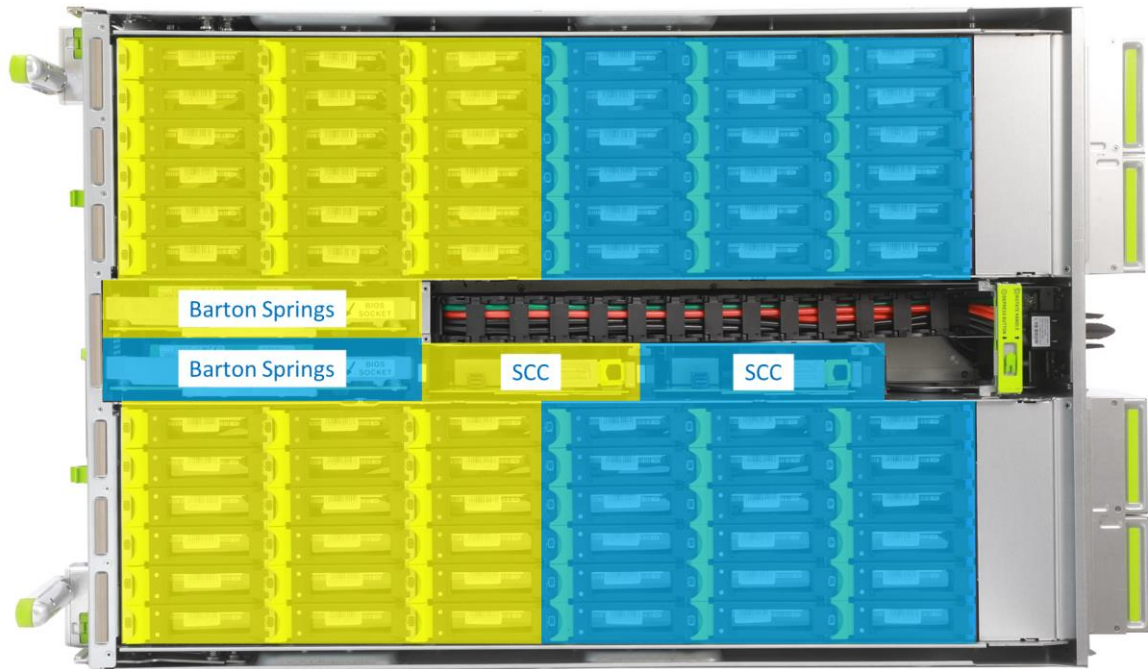


Figure 8.1: System Drive Domains

One important design point that makes drive swaps intuitive is to implement light pipes near the drives to bring the light from the activity/fault LEDs to a visible location. If the labels are metal-backed or are otherwise normally opaque, but clear over the light pipe area, this will help contain the light. Also, the drive numbers can also be opaque, essentially creating a backlit number for the drive slots.

The shape pointing to the drive should be a trapezoid instead of a triangle to maximize space for drive numbers for clear identification of the drive slot, while at the same time making the direction of the indicator obvious. A close-up of the drives showing these indicators is shown in Figure 8.2.




Figure 8.2: Drive Number Indicators

8.2 LEDs and Buttons

This section gives an overview of the various LEDs and buttons within the Grand Canyon system.

8.2.1 Front Panel Buttons and LEDs

In all chassis configurations, there will be a power button with integrated system power LED, reset button, a bicolor system status/fault LED, and a USB debug port per host/bmc/expander. There will be two instantiations of the front panel, one for each of the two logical systems in the physical chassis. These features will all be located on a User Interface Card (UIC).

There are two NIC slots per logical system. As seen in Figure 8.3 below each NIC, if populated, has an activity LED and a Link LED. There are four E1.S drives per system, 2 per logical domain. Each SSD has an activity LED and a fault LED.




Figure 8.3: LED Designations

8.2.1.1 Power Button

The functionality of the power button is to physically turn on/off the hot swap controller that gates the power to one server subsystem of the Grand Canyon system. This is used for service events, (such as server replacement or SCC replacement) to allow the technician to power off the server subsystem to perform this action. The button shall be a push-on/push-off type switch, so that the state is persistent until an operator pushes the button again. The system should ship with this button in the “ON” position.

8.2.1.2 Power LED

In order to minimize the space needed and to make servicing the system intuitive, the system power LED should be integrated into the power button. The LED shall be blue. The LED will have the ability to be controlled by both the SCC and the BMC to allow blink patterns for system identification. A blink pattern of 1s on, 1s off will indicate unit identification.

8.2.1.3 Reset Button

The reset button shall be a momentary on push button switch with a black plunger. This reset button is routed to the SCC card to serve as a hard reset signal for the expander. The reset button on the USB debug card can be used along with UART select function key to reset the BMC or the server card.

The system status/fault LED informs the user of the system health and whether there is a problem with the system. This will be driven by the expander or the BMC to indicate any faults that happen with the

sensors monitored by them. The LED shall be solid yellow during a fault and solid blue indicating everything's good.

BMC	Barton Springs	Event	Bicolor LED
off	off	system off	Off
Booting	off	BMC booting	Solid Yellow
on	off	BMC ready, Barton Springs power off	Blinking Yellow
on	on	Barton Springs ready, no faults	Solid blue
on	on	BMC Fault critical (BMC health, sensors health, NIC health)	Solid Yellow
on	on	Barton Springs Health/Fault Status Critical	Solid Yellow
on	off	BMC Fault critical (BMC health, sensors health, NIC health)	Solid Yellow
on	on	Expander Fault critical (Expander, sensors health)	Solid Yellow
on	on	System identification	NA

8.2.1.5 USB Debug Port

Grand Canyon Platform has one OCP debug USB connector per compute card/BMC/Expander located at the front panel. The connector supports the OCP USB3.0 debug card. Information in regards to the host and their POSTcodes/SEL/config are displayed thru the LCD of the OCP USB3 debug card. The UART select button on the card allows the users to walk through the UART console of the host, BMC or the Expander for their information.

The power and reset push buttons on the USB debug port should apply to the host, BMC or the expander based on the UART selector option. The USB port on the debug card is purely meant for BMC. Users may insert a thumb drive here to allow data transfer to/from USB etc.

8.2.1.6 SSD LEDs

The E1.S SSD has two LEDs – Power and Fault/Locate. The Power LED is a mono-colored blue LED which indicates if the system is currently powered. The Fault/Locate LED is a mono-colored amber LED which indicates SSD fault locations along with drive identification.

System Status	Drive Status	Power LED	Fault/Locate LED
Off	Service Allowed	Off	Off
On	Drive Okay	On	Off
Off	Fault	Off	On
On	Locate	On	Blink
Off	Locate	Off	Blink
On	Fault	On	On

8.2.1.7 NIC LEDs

The NIC LEDs will follow the OCP NIC 3.0 specification and will have two LEDs – Link and Activity. The Link LED is a bi-colored green/amber LED which denotes link status. The Activity LED is a mono-colored green LED which shows whether the link is in use.

Link Status	Link Activity	Link LED	Activity LED
No Link	No	Off	Off
Link at Max Speed	No	Green	Off
Link at Max Speed	Yes	Green	Blink ($\frac{1}{2}$ Hz - 5 Hz)
Linked below Max Speed	No	Amber	Off
Linked below Max Speed	Yes	Amber	Blink ($\frac{1}{2}$ Hz - 5 Hz)

8.2.1.8 IOC Module LEDs

There are 2 mini-SAS HD connectors on the IOC modules to connect the T7 head node to 2x JBODs.

There is a single blue/yellow LED per MiniSAS HD port. The behavior of the LED is described below on the following page.

- **Solid blue** – All four lanes up and linked at 12Gb SAS
- **Solid yellow** – Fault/warning
- **Off** – No cable present/link is down

8.2.3 HDD LED behaviors

Each drive slot has a single light pipe with an LED indication as to the state of the drives. The behaviors are described below, and summarized in a table to identify the various states:

- **Drive powered off/ not installed** – No LED illuminated
- **Drive powered on and PHY linked to Storage Controller Card** – Blue LED illuminated
- **Drive fault** – Yellow LED illuminated
- **Drive identify** – Blinking yellow LED. If a drive is powered on, this behavior will show the yellow LED illuminated for one second, blue LED illuminated for one second, repeating. If the drive is powered off, this would appear as a yellow LED on for one second, off for one second, repeating.

The following table identifies different conditions within the box, and the respective behaviors of the LEDs.

Status	LED Indication	Blue (ACT)	Yellow (Fault)	Mechanical	Expander code
No Drive	No	Off	Off	Empty Slot	No
Drive Not Seated (or backwards)	No	Off	Off	Latch Won't Close	No
Drive Powered On, no Link	Yes	Blink (1s on, 1s off)	Off		Yes
Drive Powered On, SAS/SATA Link established	Yes	On	Off		Yes
Drive Faulted	Yes	Off	On		Yes
Drive Identify (Drive Off / Not Present)	Yes	Off	Blink (1s on, 1s off)		Yes
Drive Identify (Drive On)	Yes	Blink (1s on, 1s off)	Blink (1s on, 1s off)		Yes

8.2.4 Fan Module LEDs

The fan modules have LED indications to indicate the health of the fans:

- **Blue** – When illuminated, indicates fan module health is good (actively controlled via the SCC)
- **Yellow** – When illuminated, indicates a fan fault

8.3 Hard Drive Serial Number Window

Grand Canyon features drive latches that have openings to allow the operator to scan the drive serial number if it confirms to the open window. The details of this window are shown in Figure 8.4.

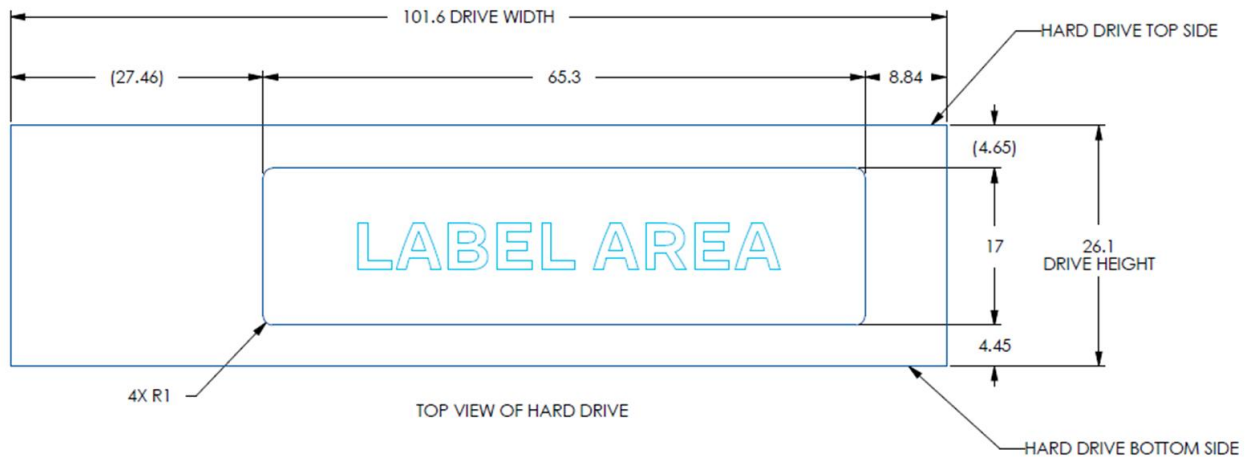


Figure 8.4: Hard Drive Serial Number Window

8.4 Service Operations

The Grand Canyon system provides hot swap controllers on the power rails, but the system will not support hot swap of the SCC, NIC or compute module due to PCIe surprise removal/insertion events. This functionality can be added later if needed, as the electrical subsystem supports this. The drive slots, however, all support hot plug.

8.4.1 HDD Replacement

The drives are all hot pluggable, and it is up to the service flow if the drives are spun down or not before removal. To replace a drive, the latch should be opened and pulled up to remove the drive from the system. By pulling up on the latch once it is opened to a vertical position, the drive is partially removed from the system to allow easy access to grab the drive. A new drive can then be inserted, and then the latch can be closed to lock the drive into place.




Figure 8.5: HDD Replacement

8.4.2 SSD Replacement

SSDs are hot-swappable so there is no need to power off the system before an SSD swap. To replace the E1.S SSD, release the ejector latch and pull the SSD out from the front panel. Replace a new SSD in its place and lock in the ejector latch.




Figure 8.6: SSD Replacement

8.4.3 SCC Replacement

To replace the SCC, the A/B domain should be shut down via the power button on the front panel module. The power button is illuminated blue while the power is on, and turns off when the power is off to that portion of the system. The SCC can then be removed, a new module installed, and then the power button can be pressed to turn the system back on.

To remove the SCC, first remove the air baffle above it. Pull the latch tab on the module and rotate the ejector until it is fully disengaged. Pull on the ejector handle to remove the SCC. To install, guide the SCC into the slot and rotate the ejector until it fully bottoms out. Place the air baffle back on.




Figure 8.7: SCC Replacement

8.4.4 UIC Replacement

To replace the UIC, the A/B domain should be shut down via the power button on the module. The power button is illuminated blue while the power is on, and turns off when the power is off for that portion of the system. The UIC can then be removed, and a new UIC module can be installed. The power button is pressed to turn the system back on. To replace the UIC, release the ejector latch and pull the UIC out from the front panel. Replace a new UIC in its place and lock in the ejector latch.




Figure 8.8: UIC Replacement

8.4.5 UIC TPM & BSM Replacement

The UIC has a removable TPM (Trusted Platform Module) and BSM (BMC Storage Module).

To replace the TPM, unfasten the screw holding it down and pull it away from the UIC. Replace it with a new TPM by aligning it with the connector and pushing it down. Secure with the screw.

To replace the BSM, pull back on the spring loaded latch to release the module. Pull the BSM out of the connector and away from the UIC. To replace the BSM, install it into the connector and push down on it near the spring loaded latch until it clicks into place.




Figure 8.9: UIC TPM & BSM Replacement

8.4.6 NIC Replacement

To replace the NIC, release the ejector latch and pull it to unseat the NIC from the connector and remove it out. Replace with a new NIC and follow the same steps in reverse order to latch the NIC onto the system. It is recommended for the NIC to be swapped only when the system is powered off.




Figure 8.10: NIC Replacement

8.4.7 Barton Springs Replacement

To replace the compute module, the A/B domain should be shut down via the power button on the front panel module. The power button is illuminated blue while the power is on, and turns off when the power is off to that portion of the system. The compute module can then be removed, a new module installed, and then the power button can be pressed to turn the system back on.

The SSD and DIMMS are FRUs, and can be replaced separately if they are identified as failed instead of the compute module itself. To remove the compute module, first remove the air baffle above it. Then, pull the latch tabs on both ejectors and rotate the ejectors until they are fully disengaged. Pull on the ejector handles to remove the compute module. To install, guide the module into the slot and rotate ejectors until they fully bottom out. Place the air baffle back on.




Figure 8.11: Barton Springs Replacement

8.4.8 IOC Module Replacement

On the T7 head node configuration, the IOC module can be replaced on the front panel after the system A/B power is turned off. The power button is illuminated blue while the power is on, and turns off when the power is off. Once the power is off, the IOC module can be ejected with the help of the ejector latch. Once a new card is loaded in, power can be turned back on for the system.




Figure 8.12: IOC Module Replacement

8.4.9 Fan Module Replacement

The fan modules are hot-swappable from the rear of the chassis. Each fan module consists of one fan integrated into an easy to service module. Each fan is a dual rotor setup and the fan module as is will be a FRU. In the Type 7 configuration there may only be a single rotor fan contained in the module.

If the fan is removed for too long, the remaining fan module may increase its fan speeds in an effort to keep the system within its thermal operating limits. In order to remove the fan, squeeze the two latches on the bottom of the module inward. Pull on the module until it comes out. To install a new fan, just slide it into the chassis until the latches click into place.



Figure 8.13: Fan Module Replacement

8.4.10 Fan Louver & Honeycomb Replacement

Behind the fans are a set of louvers and honeycomb. The purpose of the honeycomb is to reduce acoustic vibration from the fans. The purpose of the louvers is to reduce backflow in the case of fan failure or in a service scenario.

First, remove the fan module from the chassis. In order to service the louver module, press on the two tabs on top of the module and pull it out. Once the louver module is out, you can remove the honeycomb. To remove the honeycomb, pull on the tab at the top of the honeycomb module and pull it out. In order to replace both modules, repeat the steps in reverse.



Figure 8.14: Fan Louver Module Removal



Figure 8.15: Fan Honeycomb Module Removal

8.4.11 Drive Plan Board Replacement

There are two DPBs in the system, but these are always replaced as a single unit. For all intents and purposes, these should be considered as a single FRU – it is only divided due to manufacturing limitations. Replacing the DPB should be a rare event, but if it's required, perform the following:

1. Power down both sides of the system via the power buttons on the front panel.
2. Unplug the cables on the front of the system.
3. Pull the drawer out to its fully extended position.
4. Remove the SCCs, NICs, SSDs, UICs and compute modules.
5. Unlatch and lift up on all of the drives, which will hold them up and away from the DPB.
6. Pull out and rotate the retaining plunger on each side of the chassis until it is locked in the open position.
7. Unscrew four thumb screws on the bottom of the DPB assembly, and slide the DPB towards the front of the chassis to remove the board.
8. Install the new DPB assembly by aligning along the DPB card guide and sliding back in until the DPB engages with the power delivery board.
9. Follow steps 1-6 in the reverse order to put all of the components back into the system.




Figure 8.16: Drive Plane Board Replacement

8.4.12 Cable Track and Power Delivery Board Replacement

There is one bus bar power module consisting of medusa cable track and power delivery board. Perform the following to replace the bus bar power module.

1. Power down both sides of the system via the power buttons on the front panel.
2. Unplug the cables on the front of the system.
3. Remove the entire chassis from the rack.
4. Pull the drawer out to its fully extended position.
5. Remove the rear airflow baffle by lifting it at the finger touchpoints.
6. Disconnect the bus bar bracket and connector from the chassis. First, remove the two screws attaching the connector to the chassis. Then, release the two captive screws holding the bracket and lift it up, ensuring the two tabs on the bottom are disengaged.
7. Release the mechanism to remove the bus bar power module. Lift up the locking tab on the power module, press on the button under the handle, and push the handle forward.
8. Pull the bus bar power module out of the drawer.
9. Follow step 1-6 in the reverse order to put all of the components back into the system.




Figure 8.17: Cable Track and Power Delivery Board Replacement

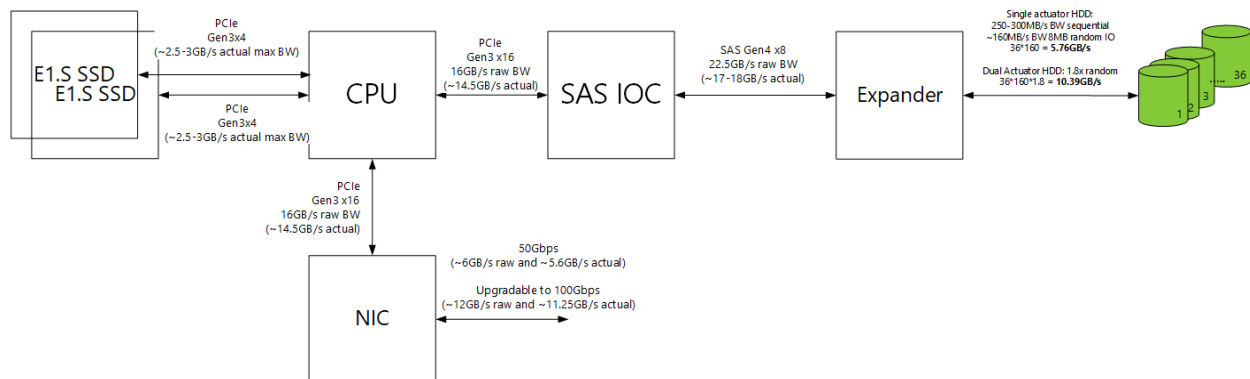
9. SYSTEM PERFORMANCE

Performance degradation at a system level can be due to interface bottlenecks like PCIe, SAS interfaces, or due to component bottlenecks like IOC IOPS limit or CPU performance or drive performance degradation due to vibrations. The system must be designed so that there is minimal degradation due to the system design from any of these aspects.

9.1 Designing the system interfaces without bottlenecks

Grand Canyon is a HDD storage server/JBOD and the data stored on HDDs are the IO movers on the system. Considering the typical and worst case workloads, the critical components were selected to ensure the Grand Canyon system has no interface bottlenecks for any of the configurations discussed so far.

Grand Canyon Type 5 configuration will be designed to accommodate Dual-actuator or Split-actuator HDDs and will provide optionality to allow for NIC upgrade to remove system interface bottlenecks.



Figure

9.1: Type 5 Data Path Diagram

The Grand Canyon Type 7 configuration conserves power by actively powering off HDDs. At any given point in time, there will be only 4 HDDs per Expander powered on in the Type 7 configuration. Out of the 4 HDDs, only 3 HDDs will be actively serving IO and the 4th HDD is reserved to take over when an already active HDD is deemed bad and needs replacement.

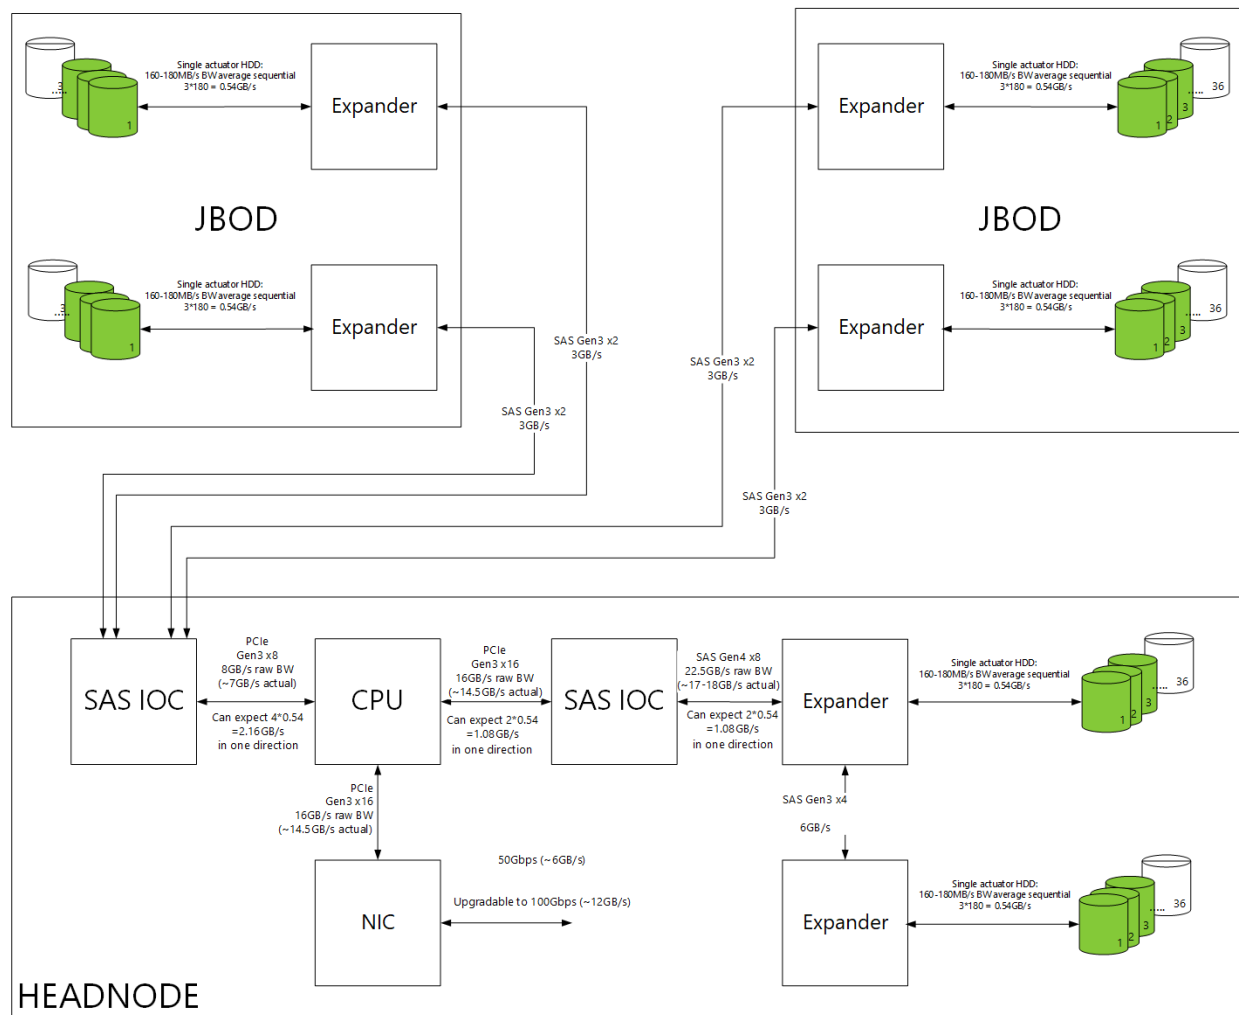


Figure 9.2: Type 7 Data Path Diagram

9.2 HDD performance degradation due to vibrations

Several areas contribute to negative drive performance impact, such as surrounding disk drives and the fans. The system must be designed so that over the operating conditions, all drives are able to perform within the specified limits as outlined below.

The only exception to the requirements described above are during drawer removal and insertion operations for service. During these events, it is acceptable for the drive performance to be affected, but there must be no damage to the drives themselves (for example, head touchdown, etc.).

9.2.1 Normal Operating Environment (sustained operating conditions)

The system should have performance degradation <5% average across the system (up to 35C, 1800m/6,000ft altitude), where every slot conforms to this <5% performance degradation requirement.

9.2.2 Unsustained Operating Conditions

The system should have performance degradation <10% max for any drive slot, and <10% average across the entire system. Examples of these operating conditions include single and multiple fan rotor failures, chassis intrusion/drawer pull-out during hot servicing of drives. This operating point includes the maximum temperature and altitude operating points (up to 35C, 1800m/6,000ft altitude). The maximum service time by design should accommodate 10 mins of drawer pullout with the system remaining operational and no components reaching critical thermal conditions.

9.2.3 Design Criteria to Enable Chassis Longevity

Grand Canyon will be designed to support dense HDDs (advent of energy assisted recording) and shall be designed to adopt the projected HDD capacities for a 5 year lifetime from its release date. To that point, adding support to the chassis for a drive capacity 3 or 4 or 5 years out in time is a waste today. Grand Canyon shall be designed with multiple options to dampen/attenuate structural and acoustic vibrations seen by the HDDs and enable phasing in these options with HDD capacity bump-up as needed, to ensure that the above performance degradation limits are met.

10. ELECTRICAL DESIGN

This section details the electrical design of the major system boards, the systems management topology, and other system details.

10.1 PCIe Topology

10.1.1 PCIe Mapping

The PCIe topology of the Grand Canyon system is shown in the diagram on the following page.

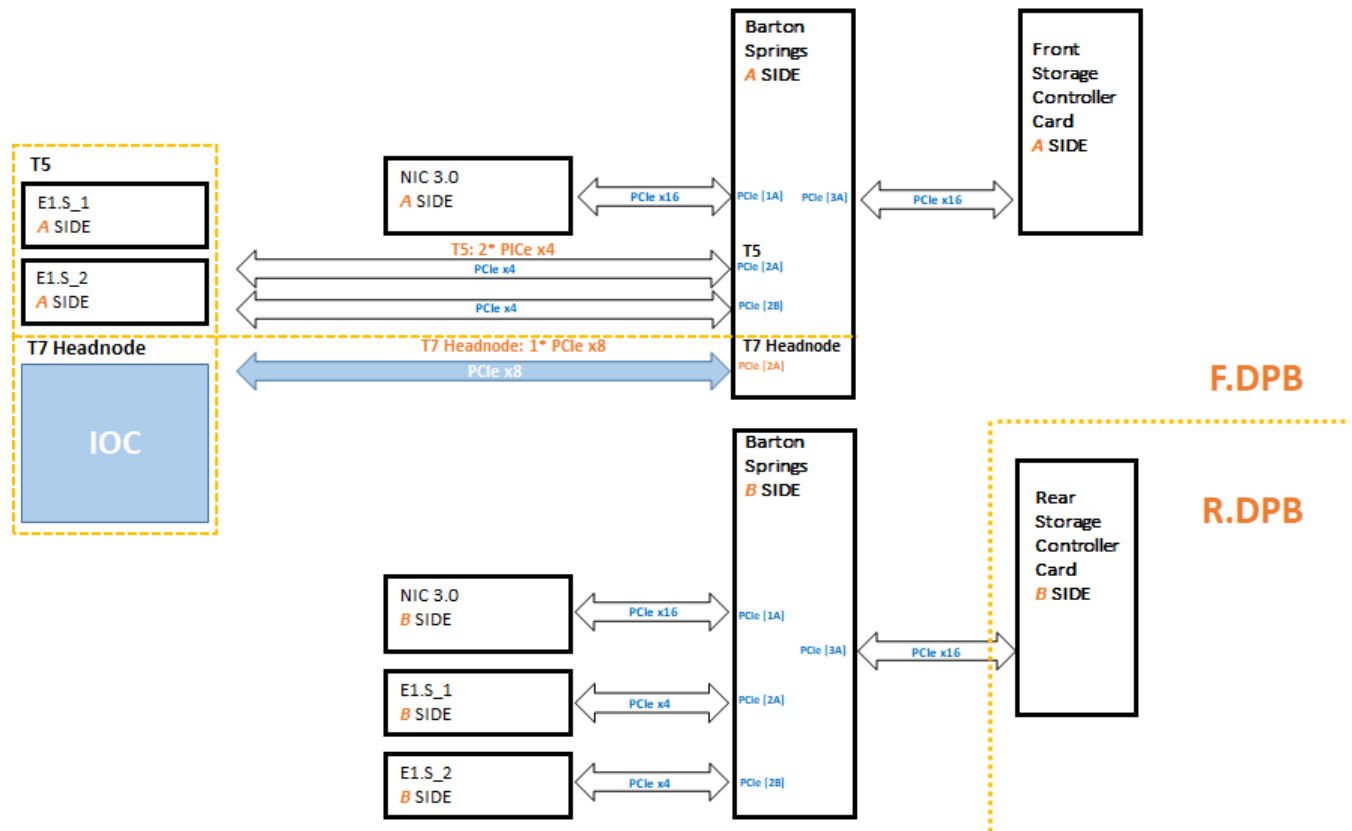


Figure 10.1 System PCIe Topology

The PCIe key commodity mapping is shown on the next page.

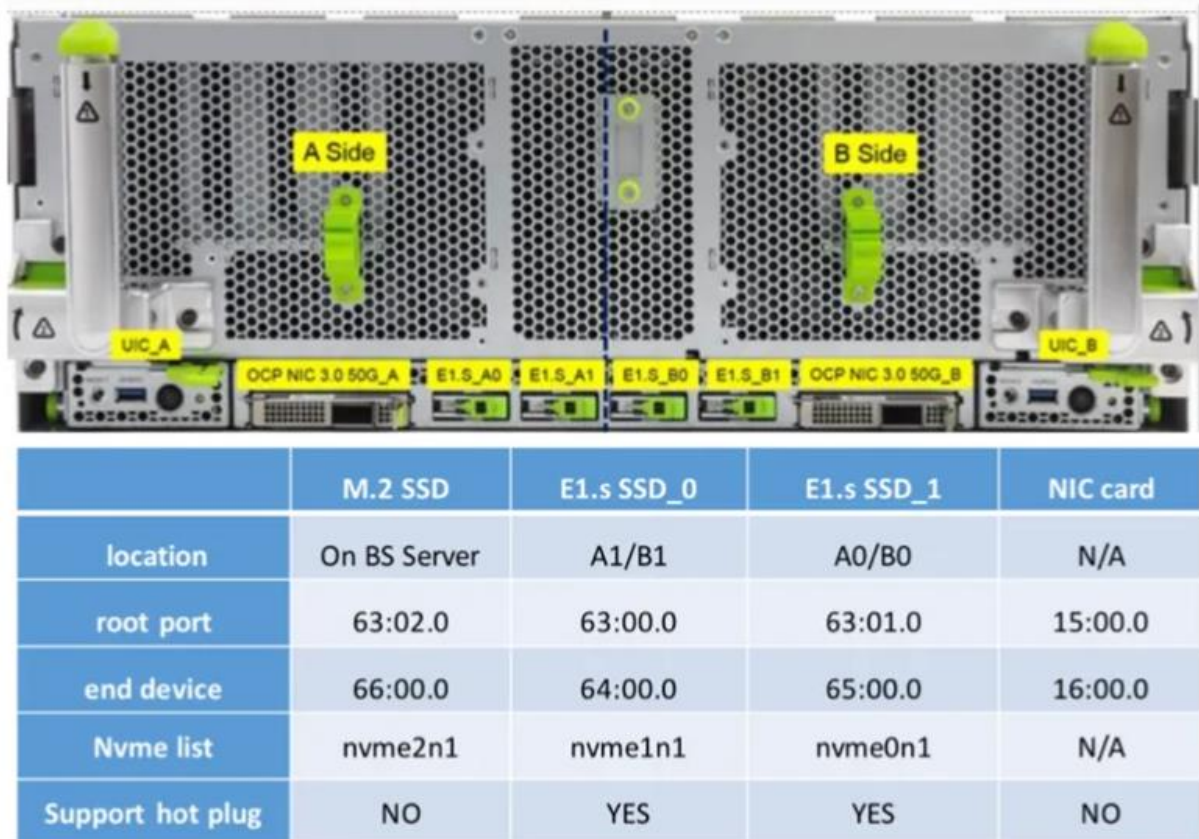


Figure 10.2: Key Commodity Mapping - PCIe Devices

10.1.2 PCIe Clocks and Resets

The PCIe clocks and resets originate from the compute server, and connect to the various PCIe devices in the system. The routing of the clocks and resets is shown below in Table 10.1 and Figure 10.1. Due to the nature of PCIe reset signals, buffers with Schmitt Triggers should be used if the traces are deemed too long to have a sufficient edge rate.

Table 10.1: PCIe Clock Info

PCH clock output	Destination
CLKOUT_SRC_3	SCC board (IOC)
CLKOUT_SRC_4	OCP NIC 3.0

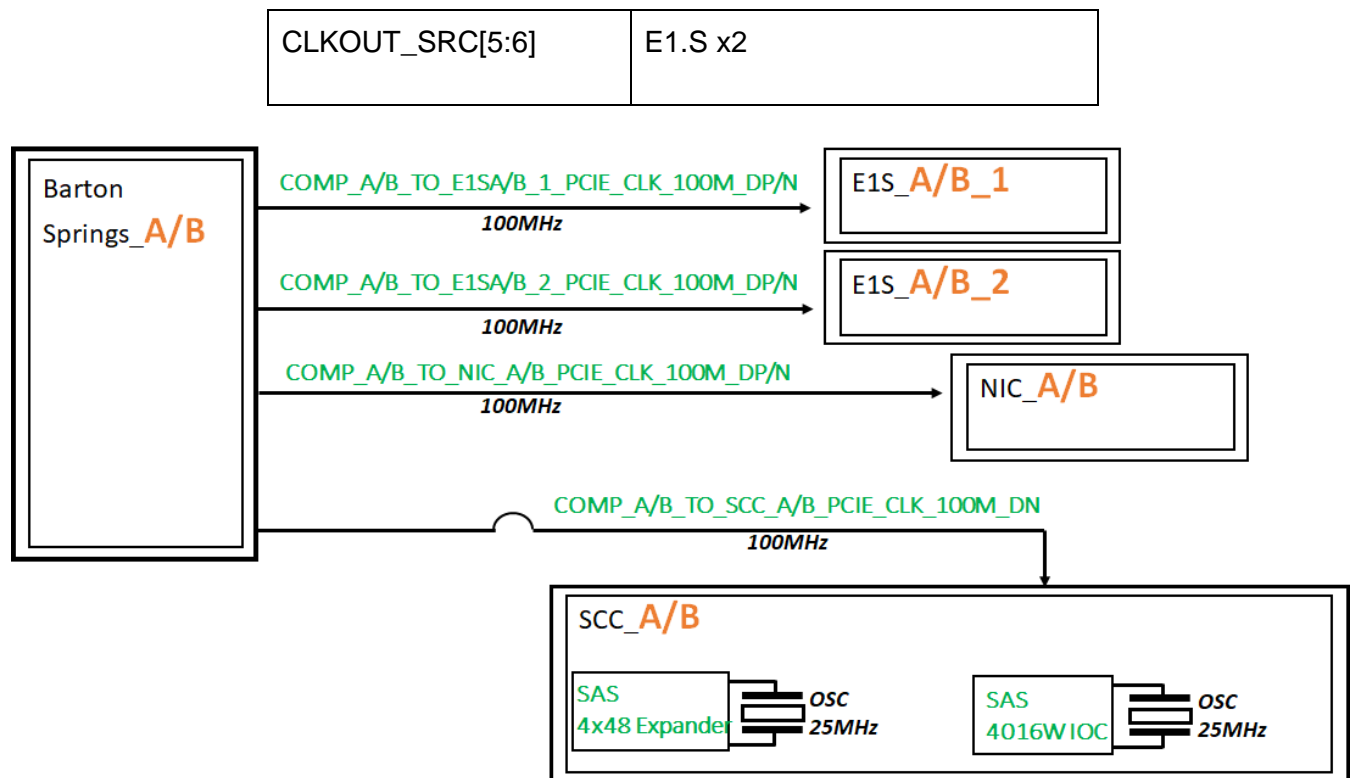


Figure 10.3: PCIe Clock Topology

Figure 10.4 describes all critical resets in the Grand Canyon system. Figure 10.4 also has the PCIe resets highlighted in green color. The compute node triggers the resets when appropriate. The SCC and E1.S drives receive the PERST signal directly from the compute node. The BMC is also sent a copy for informational purposes. The FPGA on the UIC card is responsible for the NIC power sequencing so it will ensure the NIC is powered on properly, receive the PERST signal from the CPU and relay the PERST to the NIC.

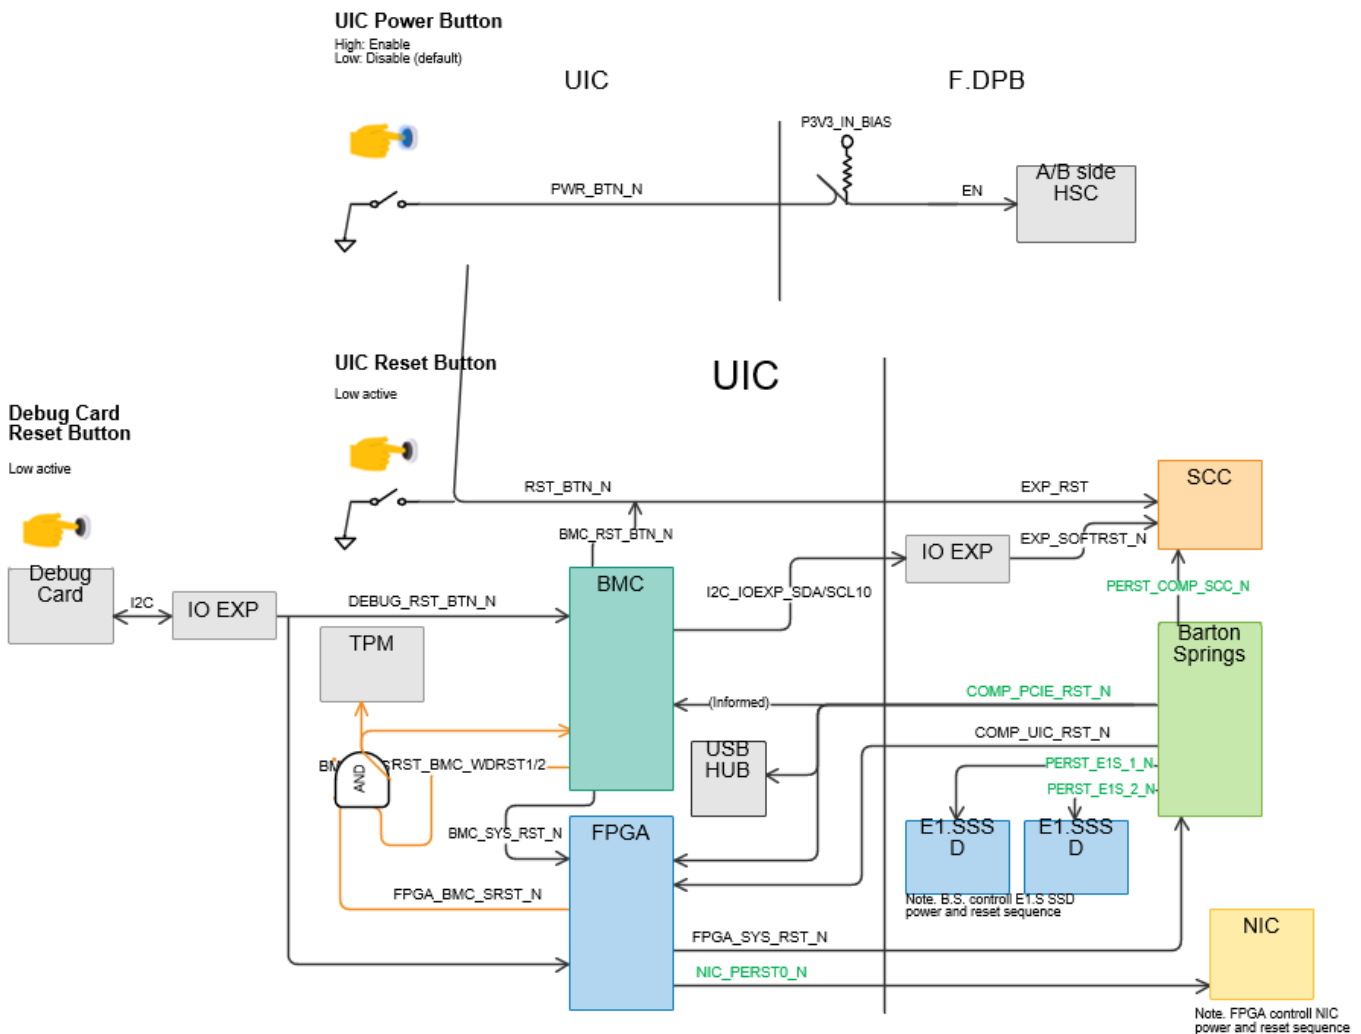


Figure 10.4: System Resets

10.2 System Management

The systems management on Grand Canyon must take into account its modularity. As covered earlier, there are several configurations that are possible: two BMCs and two SAS expanders (Dual Storage Server), a single BMC and two SAS expanders (Single Storage Server), and with SAS expanders alone (JBOD).

10.2.1 Device Management signaling

There are various connectivity protocols at work connecting the various devices and chips within the system. The Grand Canyon design revolves around the following three configurations:

- **Dual Storage Server:** Two BMCs and two SAS expanders reside in the box.
 - The BMCs perform the bulk of the systems management functions, while requesting some information from the expanders over dedicated I2C lines via IPMI protocol. The expanders manage the functions related to the drives, such as drive insert, power control, and LEDs. The expander also manages the IOC.
 - Between the BMC and expander, there are some shared resources – the elements on the SCC (FRUID, voltage monitor, etc.), as well as on the DPB (e.g., temperature sensors). Both the BMC and expander can initiate requests to poll the devices; however, this “multimaster” implementation can be augmented with GPIO lines between the BMC and expander. These lines can be used to request and grant access to these shared resources. Spare pins are provided for this reason.
 - Each BMC manages its connected Barton Springs server, SCC, and drives independently from the other BMC. The only shared resource should be the fans and the busbar HSC (P12V_CLIP).
- **Single Storage Server:** One BMC (e.g. BMC-A) and two expanders reside in the box. The singular BMC performs the bulk of the systems management functions. However, the BMC must be able to communicate to the second expander in order to request information that is not directly accessible to the BMC.
- **JBOD:** Two SAS expanders reside in the box. They take over the entire systems management.

To accommodate for the various configurations, the BMCs need to be aware if another healthy BMC resides in the box and also the configurations of the SCCs (to distinguish between Dual Storage Server and Single Storage Server configurations). The expanders also need to be aware if healthy BMCs are present (to distinguish between JBOD from the other configurations). Several measures are taken to accomplish this:

- **To establish presence,** the presence (or “insert”) pins are routed between all UICs and SCCs.
- **To self-identify location,** slot identification bits are used.
- **To identify hardware type,** type identification bits are used.
- **To establish health,** a heartbeat signal is run from each BMC and expander to each of the other devices within the box.

10.2.2 I2C Topology

There are several I2C devices in the system. The tables below break down the I2C connectivity from the 2 system management masters – BMC and Expander to the various devices in the system.

Table 10.2: I2C busses originating from BMC

Net Name	Destination	Description
I2C_UIC_SCL/SDA	BMC <-> UIC sensors	UIC temp sensor and FRU
I2C_DEBUG_SCL/SDA	BMC <-> Debug card	To debug card through USB connector.
I2C_CPLD_SCL/SDA	BMC <-> CPLD	To CPLD/FPGA on UIC
I2C_BSM_SCL/SDA	BMC <-> BSM	To BSM EEPROM
I2C_COMP_BIC_SCL/SDA	BMC <-> BS BIC	To Barton Springs BIC
I2C_COMP_CPLD_SCL/SDA	BMC <-> BS CPLD	To Barton Springs CPLD
I2C_SCC_SCL/SDA	BMC <-> SCC sensors (reserved)	SCC FRU, temp sensor, P12V_A/B HSC, and voltage sensors
I2C_EXP_SCL/SDA	BMC <-> Expander	To SAS expander
I2C_IOC_SCL/SDA	BMC <-> IOC	To IOC
I2C_DPB_SCL/SDA	BMC <-> DPB sensors (reserved)	DPB FRU, temp sensor, and voltage sensors
I2C_IOEXP_SCL/SDA	BMC <-> IO expanders	IO expanders on F.DPB
I2C_FAN_SCL/SDA	BMC <-> FAN / HSC	To Fan controller(TACH), I2C MUX(for FRUFAN), and P12V_IN HSC
I2C_NIC_SCL/SDA	BMC <-> NIC/IOC_moudule	To NIC/IOC module
I2C_E1S_1_SCL/SDA	BMC <-> E1.S SSD_1	To E1.S SSD_1
I2C_E1S_2_SCL/SDA	BMC <-> E1.S SSD_2	To E1.S SSD_2

Table 10.3: I2C busses originating from Expander

Net Name	Destination	Description
I2C_EXP_SCL/SDA	SCC EXP <-> BMC	To UIC BMC
I2C_IOC_SCL/SDA	SCC EXP <-> SCC IOC	To SCC IOC, UIC BMC (RSVD)
I2C_SCC_SCL/SDA	SCC EXP <-> SCC sensors	To SCC FRU, temp sensor, EEPROM, voltage sensors, HSC_P12V_SCC, UIC BMC(RSVD)
I2C_RMT_EXP_SCL/SDA	SCC EXP <-> RMT SCC EXP	To remote SCC EXP
I2C_DPB_SCL/SDA	SCC EXP <-> DPB sensors	To DPB FRU, temp sensor, and voltage sensors, HSC_P12V_A/B
I2C_HDD_CTRL_SCL/SDA	SCC EXP <-> IO expanders	To HDD power and insert, Fan insert, Drawer closed, RMT SCC INS
I2C_HDD_LED_SCL/SDA	SCC EXP <-> IO expanders	To HDD ACT/FLT LEDs, Fan LEDs
I2C_FAN_SCL/SDA	SCC EXP <-> Fan, HSC_CLIP	To fan control, HSC_P12V_CLIP

- Shown below is a comprehensive system level block diagram of the I2C buses.



Figure 10.6 gives a detailed description of the USB, SPI and JTAG connections in the Grand Canyon system.

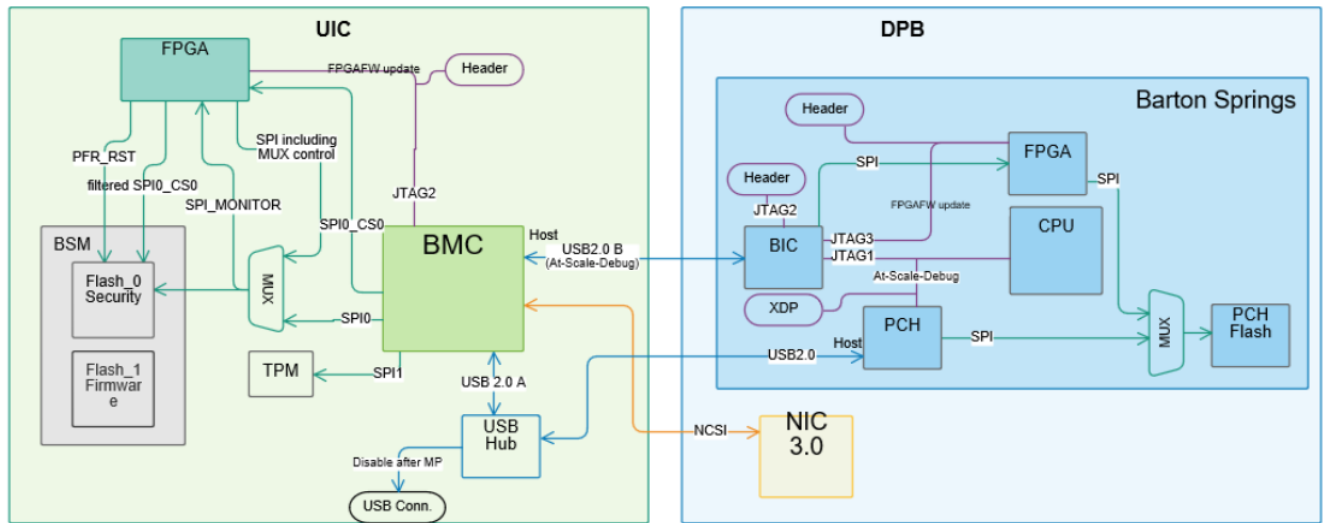


Figure 10.6 System USB, SPI and JTAG Architecture

10.2.4 Network Controller Sideband Interface (NC-SI) requirements

NC-SI is used to provide a high speed sideband between the BMC and the NIC

10.2.5 UART

Grand Canyon will support the OCP debug card that plugs into the modified USB3.0 connector (routes the UART onto the SuperSpeed pins, and provides for a UART channel select).

All UART ports on Grand Canyon are fed into the FPGA which has a multiplexer or MUX in it to route appropriately to enable various channels for monitoring, observability and ad hoc debug reasons.

The UART channel select has the following behavior: The microcontroller on the debug card controls a GPIO expander on Grand Canyon via I2C. The GPIO expander then drives an output that mirrors the channel select button on the debug card. That is, when the button is pressed, the GPIO expander will pull the output low, and when the button is released, the GPIO expander will drive the output high. This UART channel select pin is routed to the BMC on the UIC which then drives the MUX logic in the FPGA to switch UART connections from various devices in the system to the front panel debug port. BMC console UART should be the default UART exposed to the debug port on the UIC. The following is a list of all UART ports available in the system:

- The IOC and Expander provide two UARTs – SMART UART or firmware UART, and a Serial debug or SDB UART. The SDB UART for IOC is purely for debugging, while the SDB UART for Expander is also for out-of-band firmware recovery.

- CPU console or host shell UART0 port. CPU's UART1 port is used for out of band (OOB) firmware recovery of the Expander ASIC.
- Bridge IC or BIC UART port is also routed to the FPGA. For details on the BIC, please refer to Barton Springs specification document.

The UART topology inside the FPGA is complex but is necessary to enable initial bringup and validation. We may decide to disable certain mux paths and make it simpler for later stages of the design.

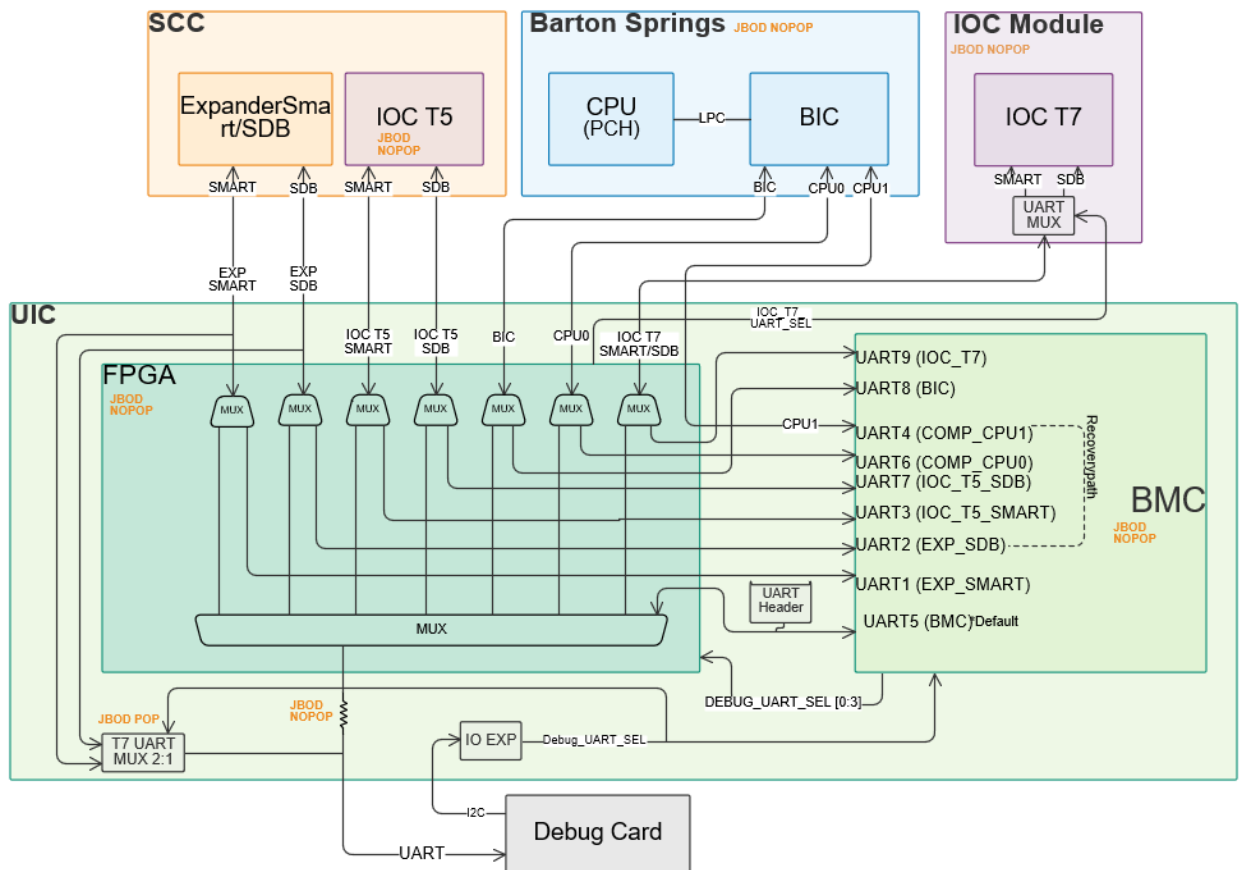
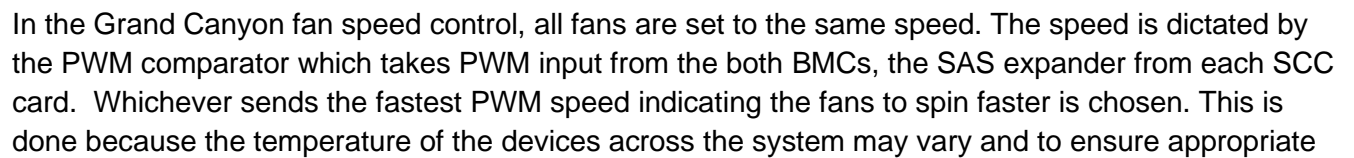


Figure 10.7: System UART Architecture

10.2.6 GPIO (Insert) Topology

The block diagram shown below in Figure 10.6 represents the Grand Canyon presence connections. Each removable PCBA, commodity and fan has a presence pin which the system can detect and alert if a component is not installed for the appropriate T5 or T7 configuration.



cooling for the hottest components, the highest PWM needs to be chosen. The TACH speed from each fan is sent to a TACH monitor chip and this information can be read out from the system.

The SAS expanders and BMC also detect fan presence. If a fan is not present or there is some sort of fan failure, the fans will ramp to near maximum speed to ensure system thermal safety.

The system is designed to remain in a safe operating temperature condition in the event of a single fan failure. The remaining 3 fans will ramp their speed up to compensate for the fan that failed.

There are I2C buses to each fan so in the future if fans are upgraded to ones that have FRU information it can be utilized without making any PCB changes. This feature allows the implementation of the fan specific FSC. System should identify the fans and adapt accordingly. FSC tables should have a default and generic fan table that applies for any 92mm fan from the selected vendors, this default fan table can be used in the scenario where the system can not identify the fans. Fans from different selected vendors should have similar performances.

Figure 10.9 below shows the overall fan system architecture.

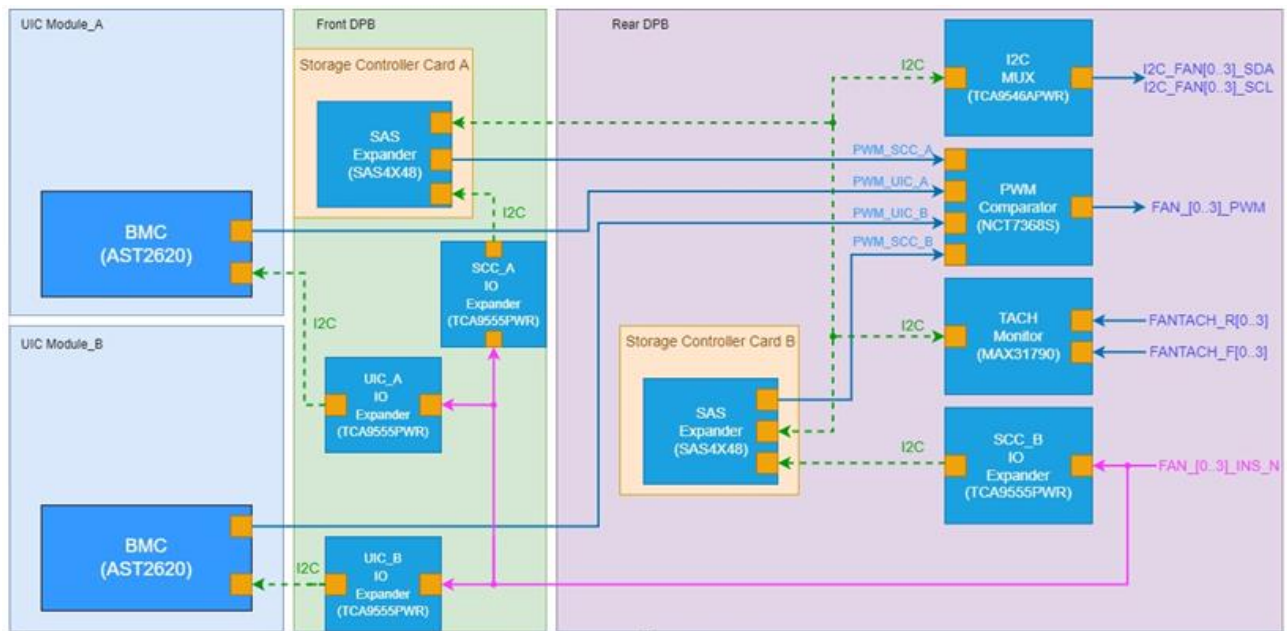


Figure 10.9 System Fan Control Architecture

10.2.9 System Configurations

10.2.9.1 Integrated (Dual) Configuration

In the Dual Storage Server configuration, the two BMCs handle the majority of systems management for the system. Each BMC manages a connected compute server, SCC, drives and power/system signaling independently of the other side of the system.

The SAS expanders handle communication with individual drives, monitor certain sensors and indicators through the system. Fan control and thermal regulation is controlled mutually from each chip. There are some shared resources as well and the BMC and expander can initiate requests to these devices over I2C or IPMB to the SAS expander on the SCC.

10.2.9.2 Integrated (Single) Configuration

In the Single Storage Server configuration, a single BMC controls the majority of systems management on its own side of the system, while the opposite-side SCC manages the opposite side of the system (its side).

10.2.9.3 JBOD Configuration

In this configuration, there are no BMCs. The two SCCs installed only have SAS expanders present.

The SCC has resistor strappings to enable it to turn on automatically once power is applied. Upon successful sequencing of these DC-DC converters, SAS expander B is brought out of power on reset.

10.3 Signal Integrity requirements

10.3.1 SAS/SATA

All drive slots must be capable of 12Gbps and 6Gbps SAS, and tested to a BER of 1 in 10^{-15} . When testing SATA 6Gbps, the BER compliance should be tested to a BER of 1 in 10^{-15} .

10.3.2 PCIe

All drive slots must be capable of PCIe Gen3 and tested to a BER of 1 in 10^{-15} .

10.4 ESD Compliance

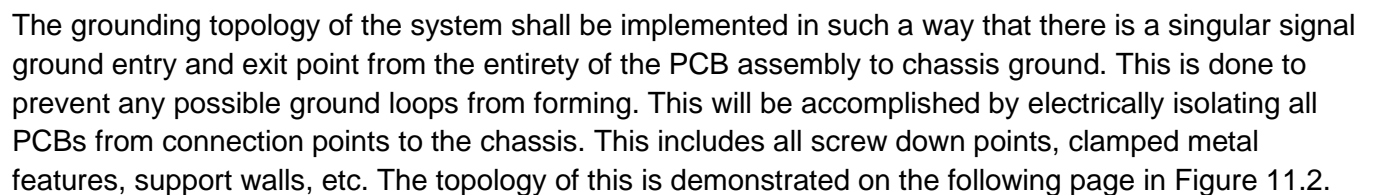
The Grand Canyon Chassis shall be designed such that all ESD events that may occur while interacting with the buttons, LEDs and connectors accessible from the front of the chassis, will be handled gracefully. This will be implemented with appropriate ESD protection devices. This includes, but is not limited to, the nets interfacing with the power and reset buttons, the USB and UART signals on the debug headers, as well as all externally interfaced LEDs that do not go through a long light pipe that would isolate them from ESD events. These ESD protection devices must be placed on the PCB as close to the component as feasible, with the ground path to dissipate the ESD event routed as short as possible to the ground plane

11. POWER

11.1 Power Topology

Figure 11.1 shows the overall power topology of the system. The specific implementation in the other boards will be covered in later sections. The power topology breaks the two logical domains into separate power domains, each of which are protected by a hot-swap controller. Isolating the side power from the side power reduces the chance of failure on one side propagating to the other.

Because the fans are supplied by a voltage rail common to both A/B sides, care must be taken such that hot-swap of the fans will not bring down the voltage to either A/B side.



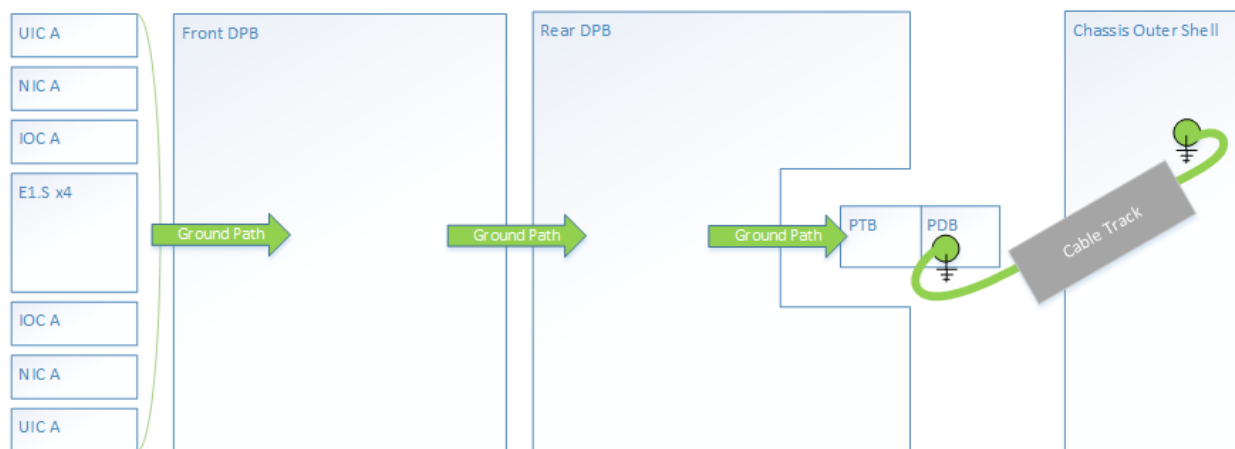


Figure 11.2: System Grounding Topology

In order to have a good ground reference connection for the Grand Canyon sliding drawer assembly, the grounding connection point will be implemented as a green 8 AWG grounding wire to the connector on the PTB. This ground wire will run through the cable track and be connected to the outer shell of the Grand Canyon chassis through a solidly connected bolt. The ground path will then travel through the chassis shell, into the support knives and into the rack.

This additional grounding wire is necessary to provide a direct ground connection from the Grand Canyon drawer to the outer chassis shell, and then to the earth grounded rack. This is necessary. There are metal rails attaching the drawer to the shell of the chassis, the rails are an electrically and mechanically dynamic system that will not provide a consistent or continuous ground path due to the rolling ball bearings, nylon bearing ball carriers and insulating grease coating all contact points.

11.2.1 Mechanical Interface Isolation Distance

Robust grounding isolation must be maintained between the PCB assembly signal ground, and all mechanical interfaces between the PCB assembly and the chassis. This restriction applies to all non-electrically connected screw and bolt holes, key-hole features, PEM® nuts and mounting points. In order to achieve this, a standard gap of 2mm, and an absolute minimum gap of 1 mm is necessary for all nets and planes on all layers of the PCB. This gap extends radially from anywhere there will be an interface between mounting hardware and the PCBs.

For anything contacting the top and/or bottom of the PCB, such as the maximum contact area of bolt head, brackets, nuts, etc., there shall be a 2mm isolation gap extending out of this maximum contact area. For any mechanical holes penetrating the PCB assemblies, the 2mm isolation gap extends outward radially from the hole on all layers of the PCB. These restrictions are illustrated on the following page in Figure 11.3.




Figure 11.3: Contact Area Restrictions

11.2.2 Connector Housing Ground Topology

In order to further ensure that there is a singular signal ground connection point to the system/chassis ground, the outer shells of all connectors on the front of the Grand Canyon chassis must be isolated from signal ground on all PCBs through the connection of an RC filter network. This network serves two purposes first, to electrically isolate the shells from chassis ground, and secondly, to provide a tuning mechanism in order to optimize and reduce EM emissions from the connectors. This will be implemented similar to Figure XX below.




Figure 11.4: Connector Housing Ground Topology

The values of the resistor used should be 1 MΩ, and the parallel capacitor should have a value of 1000pF with a voltage rating on the order of 2KV.

11.3 12V Busbar Cable Input

To enable the drawer to remain powered while the drawer is extended, a power cable is routed between a busbar clip and the drive plane board. One end of this cable has an ORv2 busbar clip, as detailed in the “Cubby Three bay shelf for Open Rack V2” specification¹.

The cables leading from the clip are to be terminated into a squeeze-to-release receptacle (P/N: TE 1600798-1, or equivalent), which mates with a vertical plug on the Rear Drive Plane Board. The cable construction should contain six (7) 8AWG high-flex wires, three used for 12V, and 4 for GND.

¹ <http://files.opencompute.org/oc/public.php?service=files&t=db28490bdd3afbff048dbcceae8d9f44>




Figure 11.5: 12V Busbar Cable Connector

11.4 Drive Power Requirements

This section outlines the requirements of power supplies to the hard drives.

11.4.1 Staggered spin up requirements

The system must implement a staggered spin up of the hard drives. The spin up algorithm should spin up the drives as fast as possible, while keeping the total chassis power under 1200W. To do this, the drives will need to be spun up in groups. The first group can have many drives, due to plenty of power being available, but care needs to be taken not to draw too much from the onboard regulators.

11.4.2 Hot-Swap Requirements

The HDDs, SSDs must be capable of being electrically hot-swapped.

11.4.3 5V Conversion Failure Domains

5V conversion on the DPB for HDDs should be sized to provide enough current capacity to support dual-actuator HDDs in the chassis. To maintain failure domains at an acceptable size, no more than nine HDDs shall be present behind a single 5V conversion, unless there is evidence of better efficiency by having fewer HDDs per converter.

11.4.4 5V Conversion Loading Requirements

The 5V conversion should be able to maintain regulation at both full-load, as well as no-load (when all drives are powered off, or not inserted; this is distinctly separate from idle).

11.5 Hot-Swap Requirements

In addition to drives, all compute server cards, and SCCs should be capable of electrical hot-swap. That is, each of these components should be able to be removed without bringing down any voltage rails in the chassis. Furthermore, the overall chassis should be able to be plugged into a live busbar without dropping the main 12V rail of the rack. To achieve this, each hot-swappable component is protected by a hot-swap controller.

All hot-swap controllers must support the following features:

- Overcurrent protection (OCP), and short circuit protection
- Over voltage protection (OVP)
- Under voltage protection (UVP)
- In-rush current control when the system/board is inserted and powered up.
- MOSFETs must be kept within a safe operating area during all operational conditions such as power on/off and fault conditions.
- Default HSC response for fault conditions shall be latch off with retry as stuff option
- Implements a fast (<20us) overcurrent monitoring scheme that generates an alert based on a remotely programmable threshold.

11.5.1 Capacitive Load

To minimize the inrush current applied to the Open Rack V2 Power Shelf during initial rack power on and assertion of Yosemite V3 sleds into a live bus bar, special design consideration around input capacitance must be considered:

1. Between the bus bar and input of any load switch or HSC, the placement of input capacitors is disallowed (if truly necessary, this must be thoroughly investigated and validated).
2. The placement of input capacitors at the output of any load switch or HSC is allowed, but the total sum of capacitance from a single Yosemite V3 sled shall not exceed TBD mF (maximum of 226mF capacitive load per power zone).

11.6 IOCM Power Topology

The IOCM board power topology is shown in the diagram below.

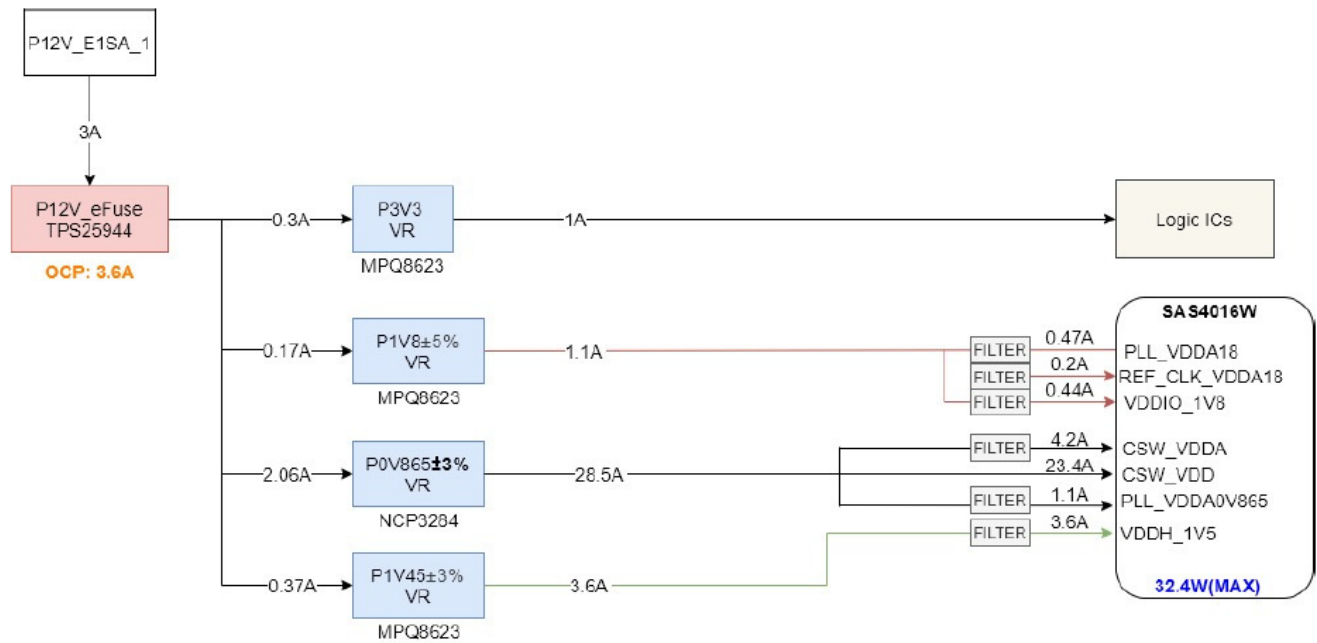


Figure 11.6: IOCM Power Topology

11.7 Storage Controller Card Power Topology

The SCC power topology is shown in the figure below.

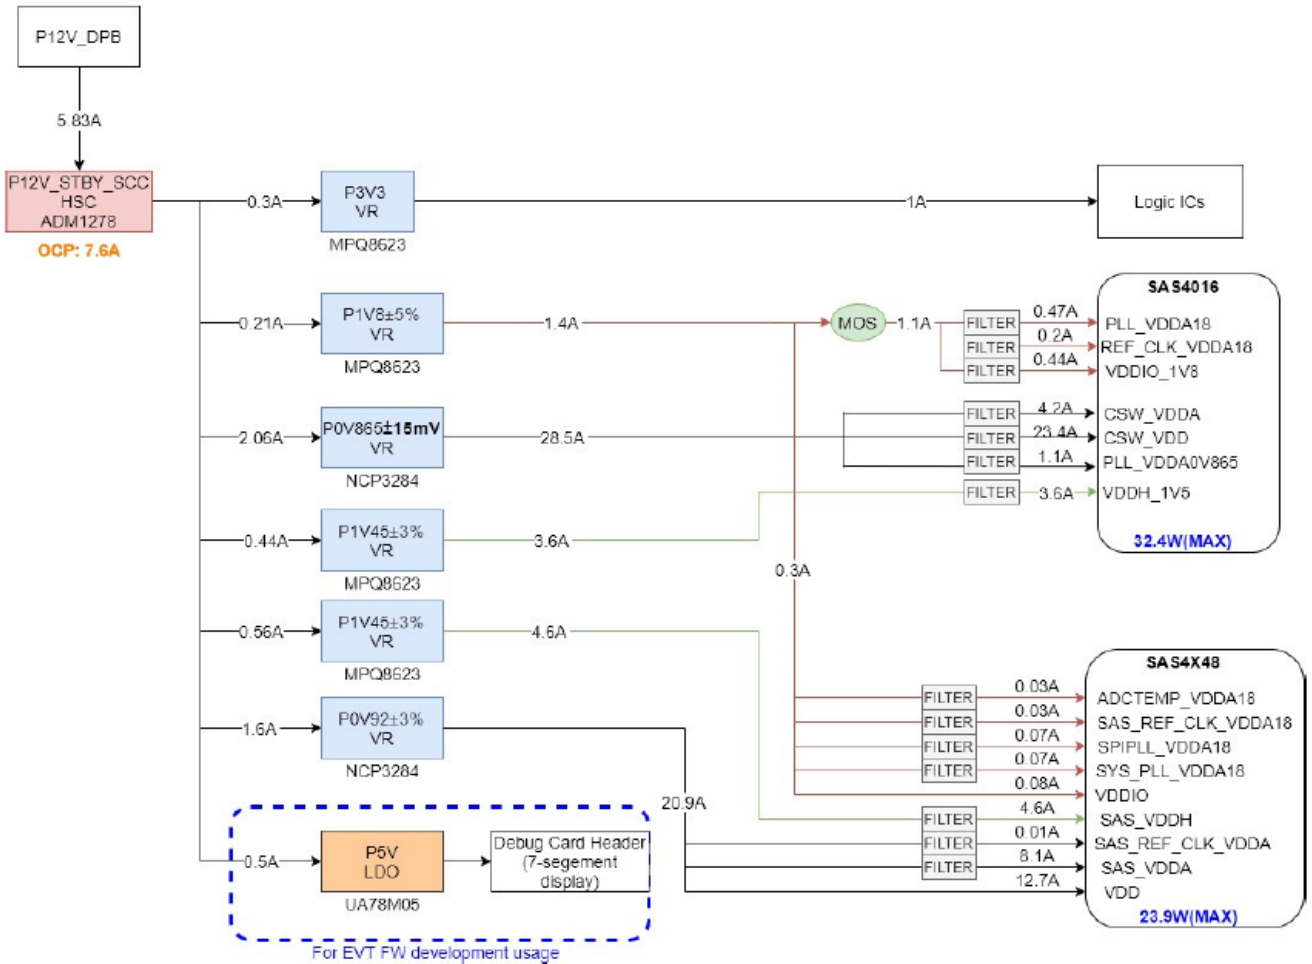


Figure 11.7: SCC Power Topology

11.8 User Interface Card Power Topology

The UIC power topology can be seen in the figure below.

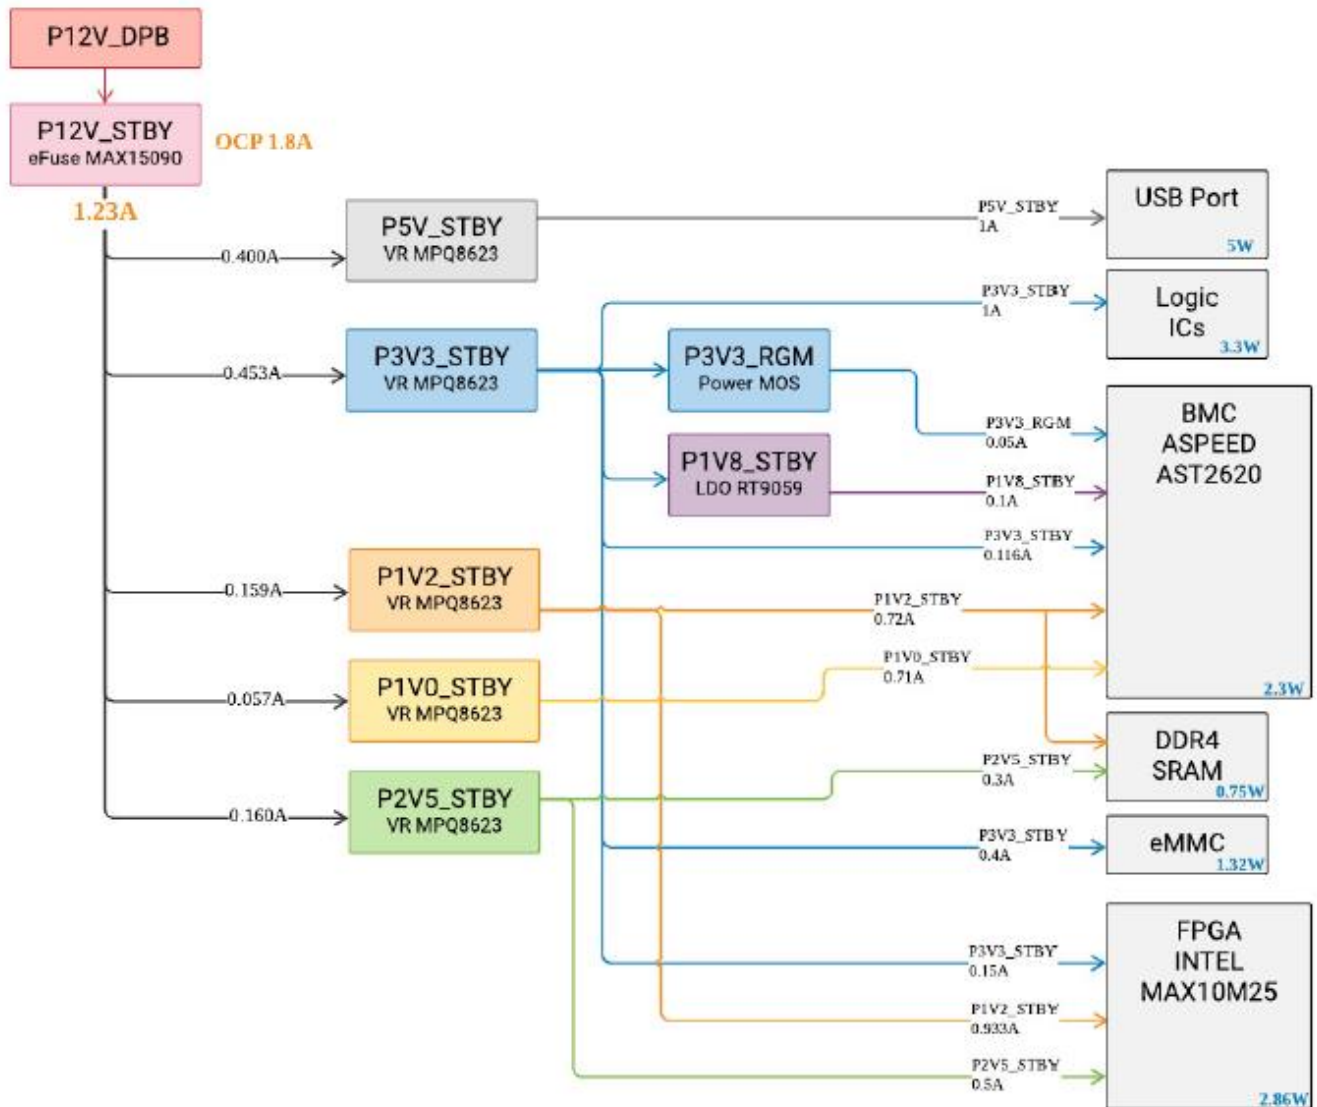


Figure 11.8: UIC Power Topology

11.9 Power Budget

11.9.1 System Power Budget

The system power budget is limited by the busbar clip ampacity rating, which is 150A.

11.9.2 SCC Connector Power Budget

The power budget of the SCC connector is driven by the worst-case draw of the IOC and SAS expanders, and the connector should be sized to support at least 5.9A @ 12V.

11.9.3 IOCM Connector Power Budget

The power budget of the IOCM connector is driven by the worst-case draw of the SAS expander and the connector should be sized to support at least 3A @12V.

11.9.4 UIC Connector Power Budget

The power budget of the UIC connector is driven by the worst-case draw of the FPGA, BMC, USB hub, as well as other miscellaneous circuits and the connector should be sized to support at least 1.25A @ 12V.

11.10 Power Sequencing

11.10.1 SCC Power Sequencing

The SCC power sequencing must be compliant with the IOC and Expander sequencing requirements. The figures below are the power sequencing requirements for each chip utilized on the SCC board.

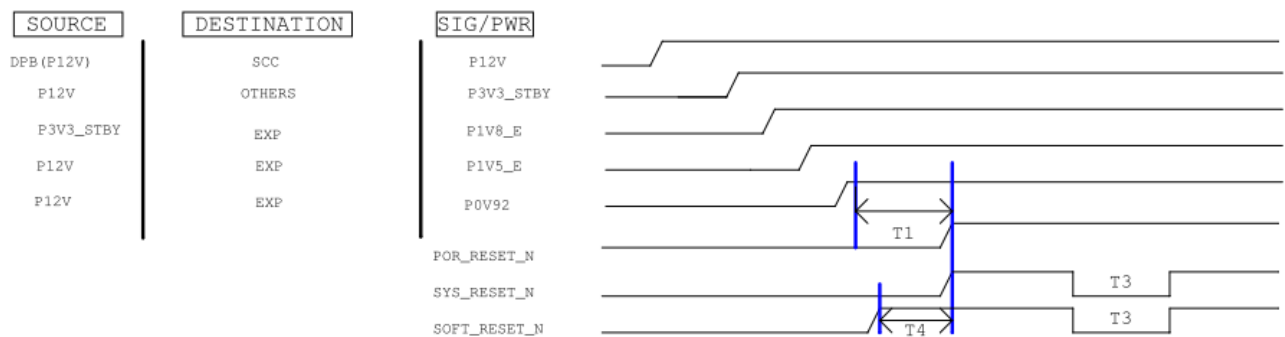


Figure 11.9: Expander Power Sequencing

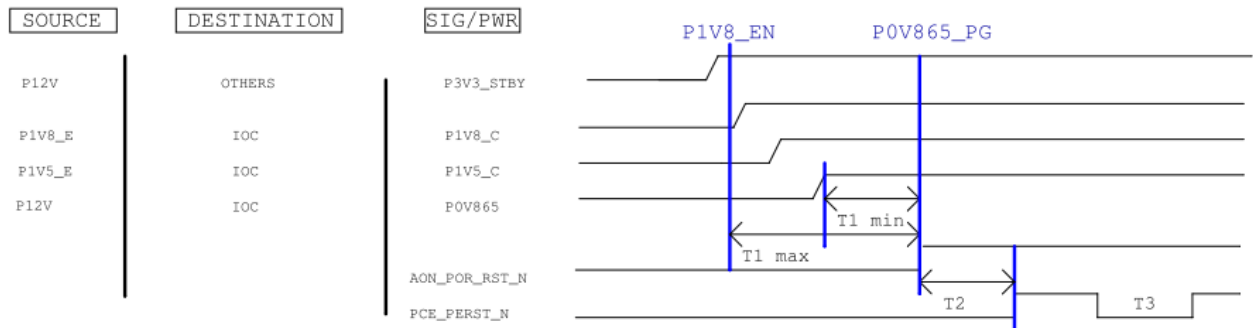


Figure 11.10: IOC Power Sequencing

Symbol	Parameter	Min.	Max.
t1	Delay between VDD reaching 0.9V before the board supply supervisor circuit deasserts POR_RESET_N.	1 ms	
t2	Delay between board deasserting SYS_PWR_ON_RST_N and board deasserting SYS_RST_N.	0	—
t3	RESET_N and SOFT_RESET_N pulse width.	5 μ s	—
t4	Delay between SOFT_RESET_N deasserting and RESET_N deasserting.	5 μ s	—

Figure 11.11: Power Sequencing Timing Values

11.10.2 Barton Springs Power Sequencing

The Barton Springs power sequencing will be covered in a separate Barton Springs specification document.

11.10.3 IOCM Power Sequencing

Expander power sequencing can be seen below in Figure 11.11 and Figure 11.12.

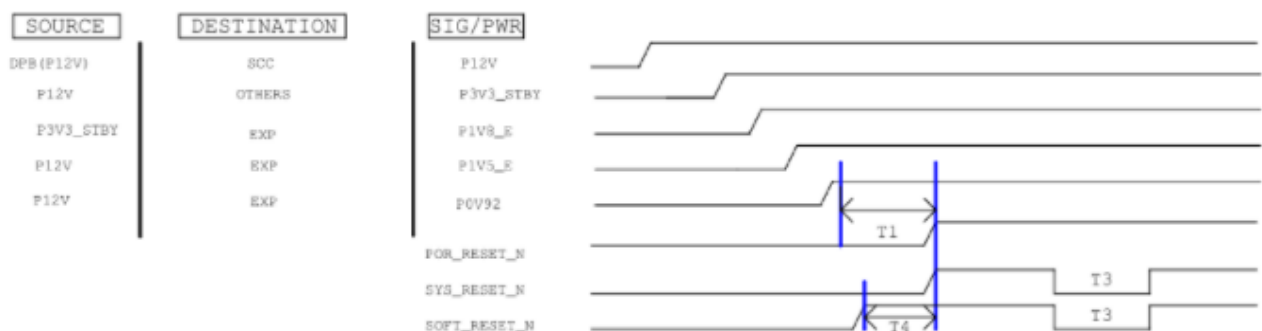


Figure 11.12: Power Sequencing Timing


Symbol	Parameter	Min.	Max.
t1	Delay between VDD reaching 0.9V before the board supply supervisor circuit deasserts POR_RESET_N.	1 ms	
t2	Delay between board deasserting SYS_PWR_ON_RST_N and board deasserting SYS_RST_N.	0	—
t3	RESET_N and SOFT_RESET_N pulse width.	5 μ s	—
t4	Delay between SOFT_RESET_N deasserting and RESET_N deasserting.	5 μ s	—

Figure 11.13: Power Sequencing Timing Values

11.10.4 UIC Power Sequencing

The only power sequencing requirements on the UIC board are from the BMC chip. The sequence order and timing can be found in Figure 11.13.

ASPEED AST2620



Power-up Sequence Parameters

PARAMETER	SYMBOL	MIN	TYP	MAX	UNIT
PV33D to PV18D	T_{A1}	0			ms
PV18D to PV33D_RGM	T_{A2}	20			us
PV33D_RGM to PV12D	T_{A3}	0			ms
PV12D to PV10D	T_{A4}	0			ms
CLKIN stable to SRST# inactive	T_B	500			us
Power stable to SRST# inactive	T_C	1			ms

Figure 11.14: Power Sequencing Timing and values

12 INDIVIDUAL PCBA SPECIFICATIONS

12.1 Drive Plane Boards (RDPB & FDPB)

The Drive Plane Boards serve to connect all the PCBs in the system to create the platform. The boards are divided into a Rear Drive Plane Board and a Front Drive Plane Board, joined with a coplanar board connector. This is to allow the DPB PCBs to be small enough to easily manufacture. The division of components between the boards is shown in Figure 12.1.

12.1.1 Connectors

12.1.1.1 Compute Slots

There are two compute server slots in the Grand Canyon chassis. Each microserver is connected via two connectors, one for power and one for signals. The signal connector must support PCIe Gen 3 interfaces and also must be press fit or SMT type connectors.

12.1.1.2 SCC Connectors

The SCC connectors are comprised of several modules. These include an 85 ohm impedance module for PCIe, a 100 ohm module for the SAS and GPIO signals, a power module, and a guide pin.

12.1.1.3 E1.S/IOCM Connectors

The IOCM module is used only in the T7 configuration and utilizes two of the 4 E1.S connectors (EDSFF/SFF-TA-1006 spec) on the FDPB to make power connections and a signal path back to each of the compute server cards.

12.1.1.4 NIC Connectors

The NIC connectors on the FDPB utilize the standard OCP3.0 form factor connector (SFFTA-1002).

12.1.1.5 UIC Connectors

The UIC connectors utilize the industry standard 56 pin SFF-TA-1002 (P/N: Amphenol ME1005610202081, or equivalent).

12.1.1.6 HDD Connectors

There are 36 SAS drive connectors used on each DPB. These connectors shall be surface mount, 29-pin 12Gb SAS connectors. The contacts must be plated with a minimum of 30m There shall be at least two sources for this connector that are fully footprint-compatible.

12.1.1.7 12V System Power Input

The 12V cable from the busbar clip terminates into the RDPB with a squeeze-to-release connector that plugs into a receptacle (P/N: Amphenol JX412-50194, or equivalent).

12.1.1.8 DPB to DPB Power Connectors

There are two 12V power rails that connect through the DPB to DPB power connectors. To minimize the power loss, and ensure the best PCB routing practices should be followed. The PCB power planes should be routed on 2oz power planes.

12.1.1.9 DPB to DPB Signal Connectors

There is a signal connector for high speed and management signals to go between the FDPB and RDPB.

12.1.1.10 Fan Connectors

There are four connectors to support dual counter-rotating fans in the back of the chassis. These connectors are only on the Rear DPB. These should be Molex Mini-fit 0438100279 or 0438100059, or equivalent.

12.1.2 LEDs

Each drive slot has two LEDs that will illuminate a light pipe to allow the user easy identification of the drive status.

12.1.2.1 Drive Activity

The drive activity LED is a blue LED, and is controlled by the SAS expander on the SCC. The LED will be illuminated when the PHY link is established with the HDD. The expander should turn off the LED if the drive is not linked, spun down, etc. In the future, if there is a requirement to blink the LED to indicate activity, the expander can emulate this behavior.

In order to save power, the activity LED should only be illuminated when the drawer is pulled out for service. This is possible due to the drawer insertion signal that alerts the SCC that the drawer is open.

12.1.2.2 Drive Fault/Identification

The drive fault indicator is a yellow LED. This is controlled by the SAS expander on the SCC via I2C GPIO devices located on the DPB. The DPB must implement a circuit to ensure that when the fault LED is illuminated, the drive activity LED is not illuminated.

12.1.3 Thermal Sensor Placement

There will be thermal sensors placed on the FDPB and UIC card to monitor the air inlet temperature. There will be 4 thermal sensors on the FDPB and one on each UIC card as shown in the figures below.

Note that the front access devices (UIC, NIC, E1.S) don't sit in front of the FDPB but rather plug in to connectors on the bottom side. The diagram is slightly misleading.




Figure 12.1: DPBs and Connector Interfaces

12.1.4 I2C Component Placement

Due to the large number of I2C devices on the DPB, care should be taken when placing the components to reduce the overall bus capacitance. For the drive related I2C components, such as the drive insert, drive fault, and drive power control GPIO expanders, it would be preferable to group these close to the SCC connector to reduce etch length and capacitance. Care should also be taken to reduce branches in the routing that add to overall capacitance and reflections.

12.1.5 Drawer Detection Sensor

The Rear DPB must implement a drawer detection sensor. One state from this sensor should indicate that the drawer is completely closed, and the other state should indicate that it is between fully open and almost closed. This signal will go to the BMCs and SAS expanders to allow the chassis to understand that the drawer is open. The system is unable to cool drives adequately if the drawer is left open, and there may need to be a thermal throttling or emergency shutdown if left open too long.

12.1.6 Drive Control Signals

There are four different HDD control signals routed through the DPB - Drive Insert, Fault LED, Drive Power, and Drive Activity. The first three of the signals are Insert, Fault and Power.

12.1.7 PCB Details

12.1.7.1 Front DPB

12.1.7.1.1 Important Components and Features

The following figure details the major functional blocks of the Front DPB, and their placement on the PCB. More detailed information on the specific functional blocks can be found in the schematics as well as elsewhere in this document.




Figure 12.2: DPB Functional Blocks and Placement

12.1.7.1.2 PCB Key Dimensions

The figure below shows the PCBA outline for the Front DPB.




Figure 12.3: FDPB Dimensions

12.1.7.1.3 PCB Stackup

Board Number:		20K055-1A		Version		01	
Project Name:		Grand Canyon					
Model Name:		Grand Canyon DVT FDPB					
Layer Count:		14					
Date:		2021/3/10					
Low Loss Advanced Material:		TU-S83 +VLP/EM-528 +HVL/EM-528 +RTF2 (Tg>170, Td>365)					
Gold Finger (Y/N):		No					
Thickness (mm):		2.36					
Product Type:		Halogen Free(HF)					
EE / SI Engineer:							
Back Drill Via / Panel Rotation 10° :		Yes / No					

Layer		Cu Weight Wiwynn/Vendor	Thickness		Glass/Copper Type Wiwynn/Vendor	Er (1GHz)		DF @4GHz 80C	
			Wiwynn	Vendor		Wiwynn	Vendor	Wiwynn	Vendor
Top	Mask		0.50						
	Signal	0.5 oz +plating	2.10		HTE				
	Prepreg		4.00		3313 * 1	3.2		0.00637	
L2	Ground	1.0 oz	1.30		VLP				
	Core		5.00		2116 * 1	3.2		0.00925	
L3	Signal	1.0 oz	1.30		VLP				
	Prepreg		9.00		2116 * 2	3.2		0.00925	
L4	Ground	1.0 oz	1.30		VLP				
	Core		4.00		3313 * 1	3.2		0.00925	
L5	Signal	1.0 oz	1.30		VLP				
	Prepreg		7.00			3.2		0.00925	
L6	Ground	2.0 oz	1.30		RTF				
	Core		4.00		3313 * 1	3.2		0.00925	
L7	Power	2.0 oz	2.60		RTF				
	Prepreg		4.00		3313 * 1	3.2		0.00925	
L8	Power	2.0 oz	2.60	0.00	RTF				
	Core		4.00	0.00	3313 * 1	3.2	0	0.00925	0
L9	Ground	2.0 oz	1.30	0.00	RTF				
	Prepreg		7.00	0.00	0	3.2	0	0.00925	0
L10	Signal	1.0 oz	1.30	0.00	VLP				
	Core		4.00	0.00	3313 * 1	3.2	0	0.00925	0
L11	Ground	1.0 oz	1.30	0.00	VLP				
	Prepreg		9.00	0.00	2116 * 2	3.2	0	0.00925	0
L12	Signal	1.0 oz	1.30	0.00	VLP				
	Core		5.00	0.00	2116 * 1	3.2	0	0.00925	0
L13	Ground	1.0 oz	1.30	0.00	VLP				
	Prepreg		4.00	0.00	3313 * 1	3.2	0	0.00637	0
Bottom	Signal	0.5 oz +plating	2.10	0.00	HTE				
	Mask		0.50	0.00	0	3.8	0	0.0398	0
Thickness requirement : 2.36 ± 10% (mm)		mil (with solder mask)	93.40	0.00					
		mm (with solder mask)	2.37	0.00					

Figure 12.4: FDPB Stackup

Single Ended Impedance (ohm)					Differential Impedance (ohm)															
Imp Variation	50				Imp Variation	85								100						
	Inner/Outer+/-10%					Inner/Outer+/-10%								Inner/Outer+/-10%						
Bus	Single Ended Signals				Bus	PCIe								SATA, SAS						
Type	SE				Type	DP								DP						
Layers	1,3,5,10,12,14				Layers	1,3,5,10,12,14								1,3,5,10,12,14						
	Wiwynn		Vendor			Wiwynn				Vendor				Wiwynn			Vendor			
	Width	Imp	Width	Imp		Width	Space	Imp	Width	Space	Imp	Width	Space	Imp	Width	Space	Imp			
S1	6.3(L2)	50	?	?	S1	6.5(L2)	5	85	?	?	?	4.9(L2)	6.5	100	?	?	?			
P2	reference layer				P2	reference layer								reference layer						
S3	6.1(L2/L4)	50	?	?	S3	6.4(L2/L4)	5.1	85	?	?	?	5.1(L2/L4)	7.8	100	?	?	?			
P4	reference layer				P4	reference layer								reference layer						
S5	4.9(L4/L8)	50	?	?	S5	6.9(L4/L8)	5.9	85	?	?	?	4.9(L4/L8)	8.5	100	?	?	?			
P6	reference layer				P6	reference layer								reference layer						
P7	reference layer				P7	reference layer								reference layer						
P8	reference layer				P8	reference layer								reference layer						
P9	reference layer				P9	reference layer								reference layer						
S10	4.9(L9/L11)	50	?	?	S10	7.4(L9/L11)	6.4	85	?	?	?	4.9(L9/L11)	8.9	100	?	?	?			
P11	reference layer				P11	reference layer								reference layer						
S12	6.1(L11/L13)	50	?	?	S12	6.4(L11/L13)	5.1	85	?	?	?	5.1(L11/L13)	7.8	100	?	?	?			
P13	reference layer				P13	reference layer								reference layer						
S14	6.9(L13)	50	?	?	S14	6.5(L13)	5	85	?	?	?	4.9(L13)	6.6	100	?	?	?			

Figure 12.5: FDPB Stackup

12.1.7.2 Rear DPB

12.1.7.2.1 Important Component and Features

The following figure details the major functional blocks of the Front DPB, and their placement on the PCB. More detailed information on the specific functional blocks can be found in the schematics as well as elsewhere in this document.




Figure 12.6: RDPB Functional Blocks and Placement

12.1.7.2.2 PCB Key Dimensions

The figure below shows the PCBA outline for the Rear DPB.




Figure 12.7: RDPB Dimensions

12.1.7.2.3 PCB Stackup

Board Number:

20X054 1A

Project Name:

Grand Canyon

Model Name:

Grand Canyon DVT RDPB

Layer Count:

10

Date:

2021/4/9

Low Loss Material:

EM 526 +RTF2/TU 863 plus +VLP (Tg>170, Td>365)

Solid Finger (Y/N):

No

Thickness (mm):

2.36

Product Type:

Halogen Free(HF)

EE / SI Engineer:

Back Drill Via / Panel Rotation 10° :

Yes / No

Version
01

Layer	Cu Weight	Thickness	Glass/Copper Type	Er (1GHz)	DF @4GHz 80C
	Wh/wh/wh/Vendor	Wh/wh/wh/ Vendor	Wh/wh/wh/Vendor	Wh/wh/wh/ Vendor	Wh/wh/wh/ Vendor
Top	Mask	0.50		3.8	0.032674
	Signal	0.5 oz + plating	HTE		
	Prepreg	4.00	2116 * 1	3.3	0.0078418
L2	Ground	1.0 oz	VLP		
	Core	5.00	2116 * 1	3.3	0.0095
	Signal	1.0 oz	VLP		
L3	Prepreg	15.00	2116 * 3	3.3	0.0095
	Ground	1.0 oz	VLP		
	Core	4.00	Designed by vendor	3.3	0.0095
L5	Prepreg	19.00	Designed by vendor	3.3	0.0095
	Prepreg	19.00	Designed by vendor	3.3	0.0095
	Prepreg	19.00	Designed by vendor	3.3	0.0095
L6	Prepreg	19.00	Designed by vendor	3.3	0.0095
	Prepreg	19.00	Designed by vendor	3.3	0.0095
	Prepreg	19.00	Designed by vendor	3.3	0.0095
L7	Prepreg	19.00	Designed by vendor	3.3	0.0095
	Prepreg	19.00	Designed by vendor	3.3	0.0095
	Prepreg	19.00	Designed by vendor	3.3	0.0095
L8	Prepreg	19.00	Designed by vendor	3.3	0.0095
	Prepreg	19.00	Designed by vendor	3.3	0.0095
	Prepreg	19.00	Designed by vendor	3.3	0.0095
L9	Prepreg	19.00	Designed by vendor	3.3	0.0095
	Prepreg	19.00	Designed by vendor	3.3	0.0095
	Prepreg	19.00	Designed by vendor	3.3	0.0095
Bottom	Prepreg	19.00	Designed by vendor	3.3	0.0095
	Prepreg	19.00	Designed by vendor	3.3	0.0095
	Prepreg	19.00	Designed by vendor	3.3	0.0095
	Mask	0.50		3.8	0.032674
Thickness requirement : 2.36 ± 10% (mm)	mil (with solder mask)	93.20	0.00		
	mm (with solder mask)	2.37	0.00		

Figure 12.8: RDPB Stackup

Single Ended Impedance (ohm)					Differential Impedance (ohm)												
Imp Variation	S0 Inner/Outer+/- 10%				Imp Variation	85 Inner/Outer+/- 10%					100 Inner/Outer+/- 10%						
Bus	Single ended Signals				Bus	PCIe					SATA, SAS						
Type	SE				Type	DP					DP						
Layers	1,3,8,10				Layers	1,3,8,10					1,3,8,10						
	Wharbyn		Vendor			Wharbyn		Vendor			Wharbyn		Vendor				
	Width	Imp	Width	Imp		Width	Space	Imp	Width	Space	Imp	Width	Space	Imp			
S1	76.21	90	?	?	S1	6.56(2)	5	85	?	?	?	4.96(2)	6.6	100	?	?	?
P2	reference layer				P2	reference layer					reference layer						
S3	76.21(4)	90	?	?	S3	6.66(2)(4)	4.9	85	?	?	?	5.46(2)(4)	7.8	100	?	?	?
P4	reference layer				P4	reference layer					reference layer						
P5	reference layer				P5	reference layer					reference layer						
P6	reference layer				P6	reference layer					reference layer						
P7	reference layer				P7	reference layer					reference layer						
S8	76.21(8)	90	?	?	S8	6.66(2)(8)	4.9	85	?	?	?	5.46(2)(8)	7.8	100	?	?	?
P9	reference layer				P9	reference layer					reference layer						
S10	76.21	90	?	?	S10	6.56(2)	5	85	?	?	?	4.96(2)	6.6	100	?	?	?

Figure 12.9: RDPB Stackup

12.2 Storage Controller Card (SCC)

The following diagrams detail the major functional blocks of the SCC and their placement on the PCBs. This is detailed for both versions of the SCC, the Type 5 version with the expander and IOC, as well as the JBOD version with just the expander where the IOC is unpopulated on the PCB.



The SCC has two configurations. The first is fully populated and is used in the T5 and T7 headnode. In the T7 headnode case, only one SCC has the IOC and the other does not. In the T7 JBOD configuration, both SCCs do not have an IOC chip and rely on the headnode for the IOC functionality.

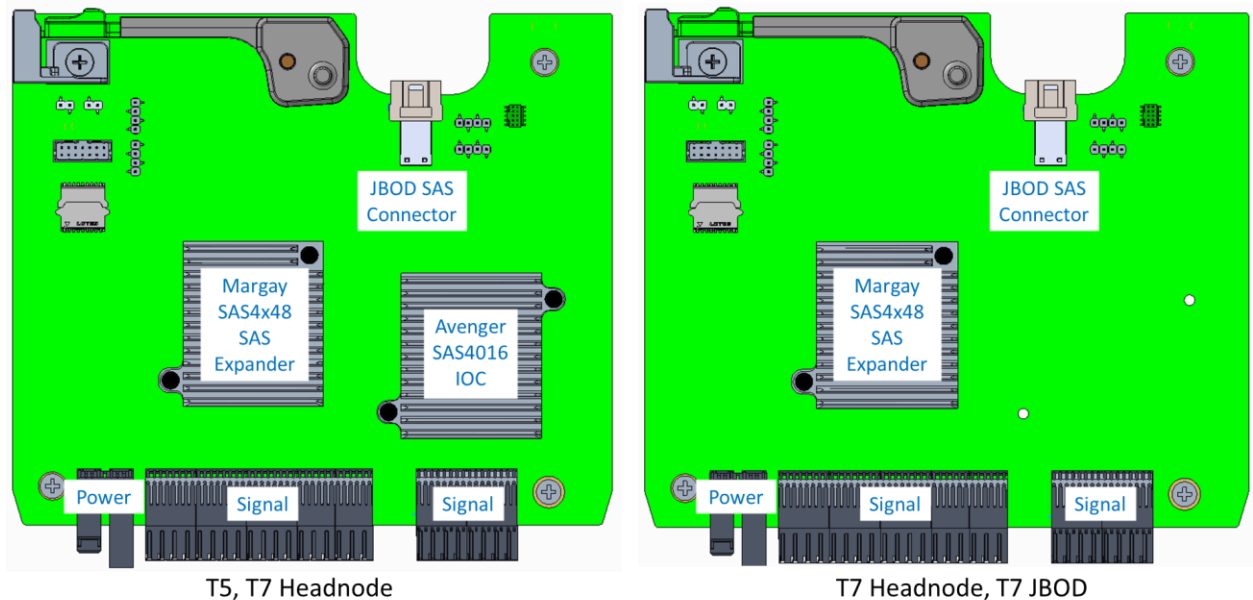


Figure 12.11: SCC T5 and T7 Configurations

12.2.2.1 Connector Configurations

Amphenol - (power connector) MPN: JX310-50387

Amphenol - (signal connector) MPN: JX310-50558

Amphenol - (signal connector) MPN: JX310-50704

Amphenol - (MiniSAS HD connector) MPN: G40H11331HR-W

12.2.5 Storage Controller Card PCB

12.2.5.1 Important Components and interfaces

Important SCC components can be seen in the PCB rendering on the following page in Figure 12.12.




Figure 12.12: SCC Important Components

12.2.5.2 PCB Key Dimensions




Figure 12.13: SCC Dimensions

12.2.5.3 PCB Stack up

Board Number: 20X049-1A		Version 01	
Project Name: Grand Canyon			
Model Name: Storage Controller Card			
Layer Count: 8			
Date: 2021/3/11			
Low Loss Material: EM-526 +RTF2/NPG-170D +V2/TU-863 plus +RTF/EM-526 +RTF (Tg>170, Td>365)			
Gold Finger (Y/N): No			
Thickness (mm): 1.57			
Product Type: Halogen Free(HF)			
EE / SI Engineer:			
Back Drill Via / Panel Rotation 10° :		No / No	

Layer	Cu Weight	Thickness		Glass/Copper Type	Er (1GHz)		DF @8GHz 80C	
		Wiwynn	Vendor		Wiwynn	Vendor	Wiwynn	Vendor
Top	Mask		0.50			3.8		0.034931
	Signal	0.5 oz+ plating	1.90	HTE				
	Prepreg		3.00	1080 *1	3.3		0.008384	
L2	Ground	1.0 oz	1.30	VLP				
	Core		4.00	3313 *1	3.3		0.00825	
L3	Signal	1.0 oz	1.30	VLP				
	Prepreg		16.20	7628 *2 +1080 *1	3.3		0.00825	
L4	Ground	1.0 oz	1.30	VLP				
	Core		4.00	3313 *1	3.3		0.00825	
L5	Power	1.0 oz	1.30	VLP				
	Prepreg		16.20	7628 *2 +1080 *1	3.3	0	0.00825	0
L6	Signal	1.0 oz	1.30	VLP				
	Core		4.00	3313 *1	3.3	0	0.00825	0
L7	Ground	1.0 oz	1.30	VLP				
	Prepreg		3.00	1080 *1	3.3	0	0.008384	0
Bottom	Signal	0.5 oz+ plating	1.90	HTE				
	Mask		0.50	0	3.8	0	0.034931	0
	Thickness requirement : 1.57 ± 10% (mm)							
		mil (with solder mask)	63.00	0.00				
		mm (with solder mask)	1.60	0.00				

Figure 12.14: SCC Stackup

Single Ended Impedance (ohm)					Differential Impedance (ohm)																
Imp Variation	SE				Imp Variation	85														100	
Bus	Inner/Outer +/- 10%				Bus	Inner/Outer +/- 10%														Inner/Outer +/- 10%	
Type	N/A				Type	PCIe														SATA, SAS	
Layers	SE				Type	DP														DP	
	1.3,6,8				Layers	1.3,6,8														1.3,6,8	
	Width		Vendor			Width		Vendor			Width		Vendor			Width		Vendor			
	Width	Space	Imp	Width	Space	Imp	Width	Space	Imp	Width	Space	Imp	Width	Space	Imp	Width	Space	Imp	Width	Space	Imp
S1	5.20(2)	6.8	85	?	?	?	5.20(2)	6.8	92.5	?	?	?	6.70(2)	11.8	100	?	?	?	?	?	?
P2	reference layer				P2	reference layer				P2	reference layer				P2	reference layer					
S3	5.40(3.6,6)	6.8	85	?	?	?	5.30(3.6,6)	6.8	92.5	?	?	?	4.90(3.6,6)	7.5	100	?	?	?	?	?	?
P4	reference layer				P4	reference layer				P4	reference layer				P4	reference layer					
P5	reference layer				P5	reference layer				P5	reference layer				P5	reference layer					
S6	5.40(5.6,7)	6.8	85	?	?	?	5.30(5.6,7)	6.8	92.5	?	?	?	4.90(5.6,7)	7.5	100	?	?	?	?	?	?
P7	reference layer				P7	reference layer				P7	reference layer				P7	reference layer					
S8	5.20(2)	6.8	85	?	?	?	5.20(2)	6.8	92.5	?	?	?	6.70(2)	11.8	100	?	?	?	?	?	?

Figure 12.15: SCC Stackup

12.3 IOCM

The following diagrams detail the major functional blocks of the IOCM and their placement on the PCBs.

12.3.1 IOCM Configuration Diagram

Figure 12.16 shows the IOCM high level block diagram.

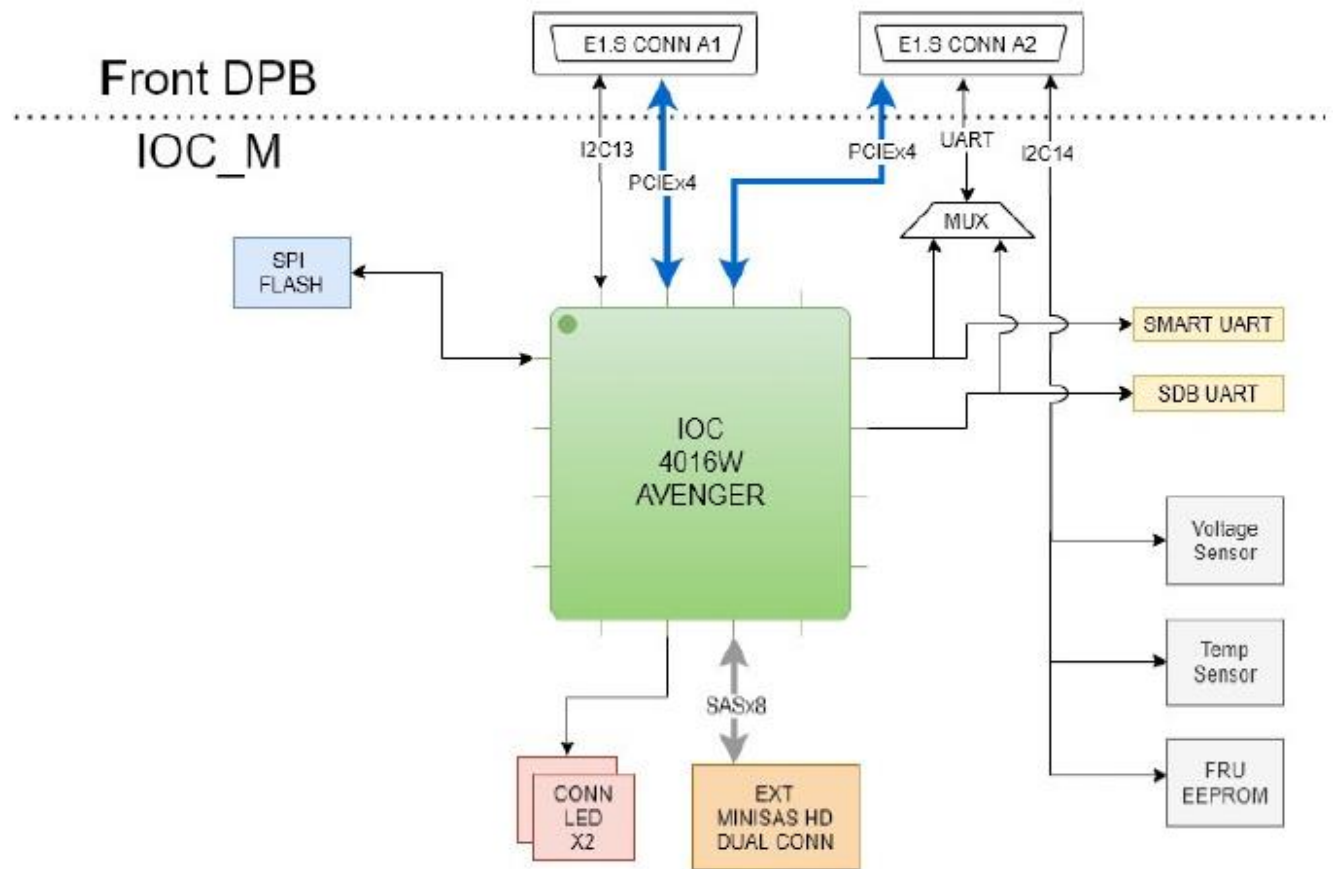


Figure 12.16: IOCM Block Diagram

12.3.2 IOCM Connectors

MiniHD SAS - Amphenol: MPN U92-A210-1001-70

12.3.3 IOCM Card PCB

12.3.3.1 Important Components and Interfaces




Figure 12.17: IOCM Important Components

12.3.3.2 PCB Key Dimensions




Figure 12.18: IOCM Dimensions

12.3.3.PCB Stackup

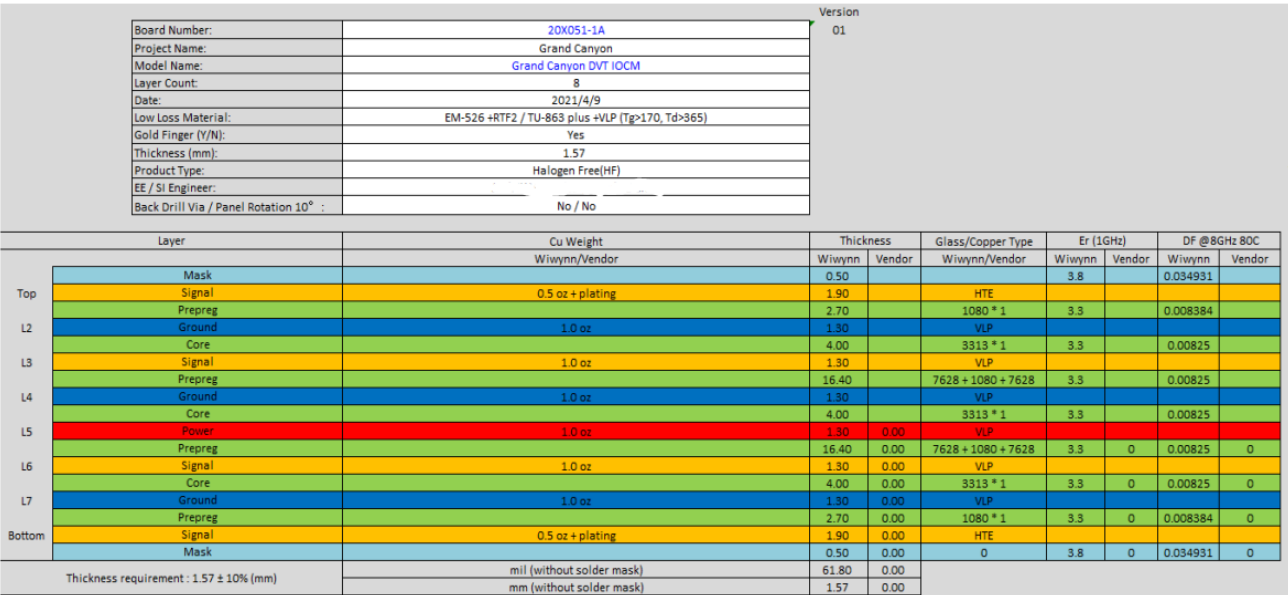


Figure 12.19: IOCM Stackup

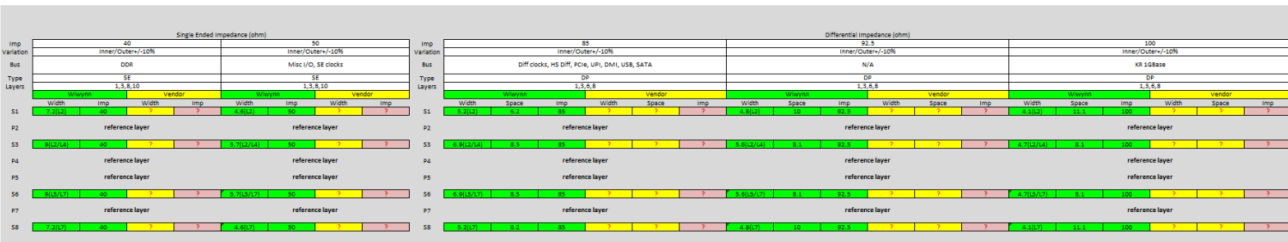


Figure 12.20: IOCM Stackup

12.4 User Interface Card (UIC)

12.4.1 UIC Variants

There are two variants of the UIC, one with a BMC and one without. The BMC will be removed from the Type 7 system configuration.

- Type 5 - 2 UIC with BMC
- Type 7 headnode - one UIC with BMC and one UIC without BMC
- Type 7 JBOD - two UIC without BMC

The block diagram for the UIC card can be seen in Figure 12.21 below.

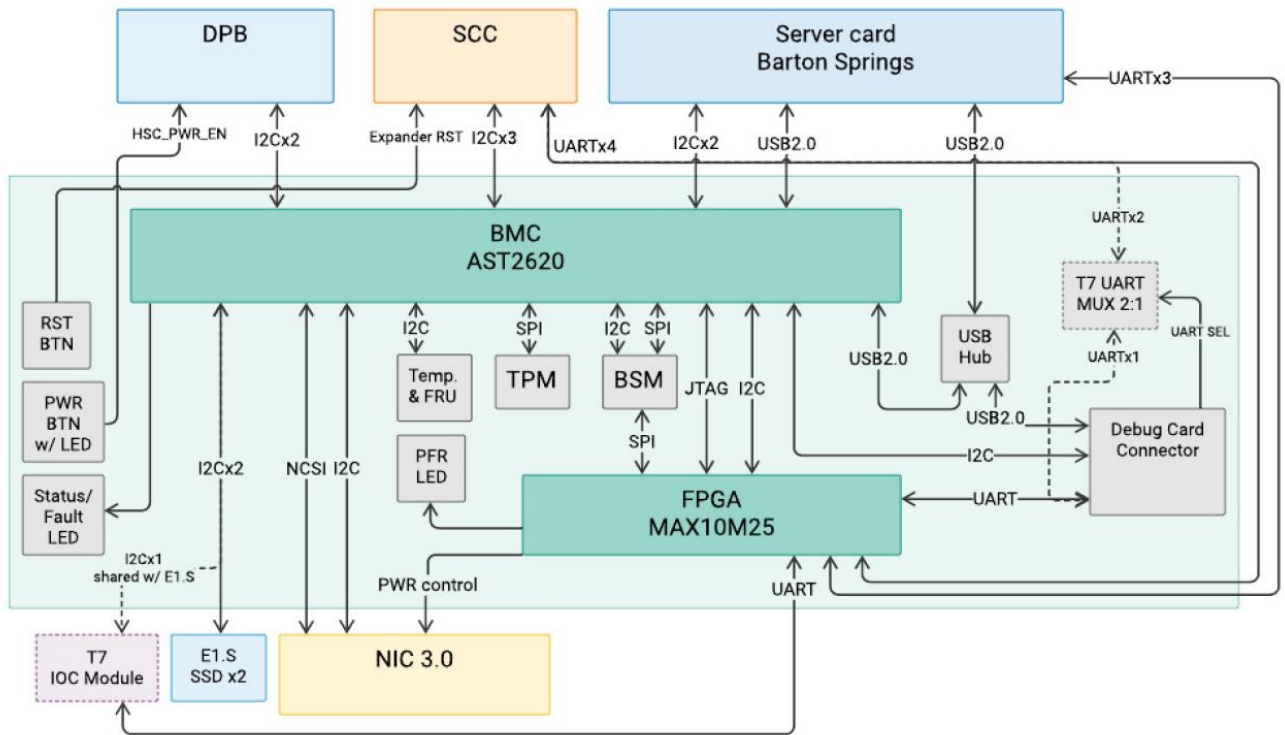


Figure 12.21: UIC Block Diagram

The major components can be seen below in figure 12.22 below.




Figure 12.22: UIC Key Components

12.4.2 BMC

The UIC card will utilize the Aspeed AST2600 series BMC. See BMC section for additional info about features and functionality.

12.4.3 CPLD

The CPLD will utilize the Intel Max10. For specific design details and implementations see the schematic and layout.

12.4.4 Connectors

The UIC is a front accessible PCBA which plunges into the FDPB using the SFF-TA-1002 card edge connector. Since the BMC subsystem is fully contained in the UIC, Grand Canyon now has the ability to upgrade the BMC solution without impacting the rest of the design.

There is one connector, a USB type A connector that is used for debug purposes.
MFG: Singatron - P/N: 2UB4001-170101F - (Or equivalent).

12.4.5 USB Debug

The UIC card has a USB3.0 connector available for debug purposes. This connector is intended to work in conjunction with the OCP debug card to gather diagnostic information. The USB2.0 pins are routed to a USB hub located on the UIC and the USB 2.0 are repurposed for debug signaling.

12.4.6 BSM and TPM 2.0

The UIC utilizes the BMC Storage Module (BSM). It is mounted to provide security features as well as hold the BMC flash images. The BSM is in an m.2 2230 form factor. Figure 12.13 below contains the BSM block diagram for reference.

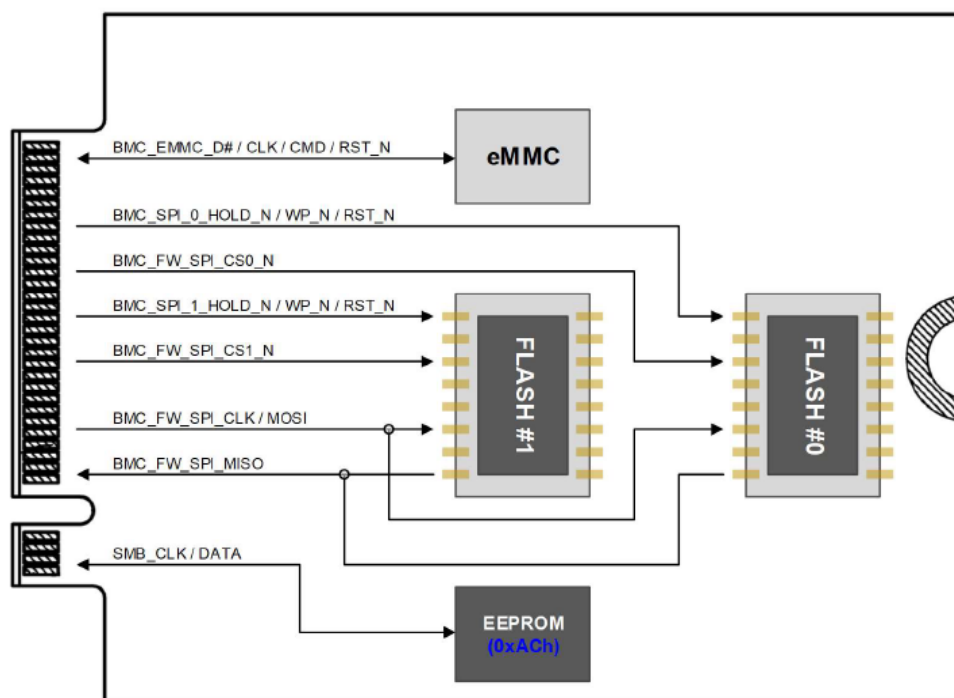


Figure 12.23: BSM Key Interfaces

The UIC also has a separate connector where a TPM2.0 module can be plugged in for system security purposes. The UIC utilizes the Infineon TPM 2.0 SLB9670 series parts.

12.4.7 User Interface Card PCB

12.4.7.1 PCB Key Dimensions




Figure 12.24: UIC Dimensions

12.4.7.2 PCB Stack up

Board Number: 20X052-SB		Version 01	
Project Name: Grand Canyon			
Model Name: Grand Canyon EVT UIC			
Layer Count: 10			
Date: 2020/9/24			
Middle Loss Material: IT-170GRA1/NPG-171/EM-828G +RTF (Tg>170, Td>365)			
Gold Finger (Y/N): Yes (see Note 10)			
Thickness (mm): 1.57			
Product Type: Halogen Free(HF)			
EE / SI Engineer:			
Back Drill Via / Panel Rotation 10° :		No / No	

Layer	Cu Weight Wiyynn/Vendor	Thickness		Glass/Copper Type Wiyynn/Vendor	Er (1GHz)		DF @4GHz 25C	
		Wiyynn	Vendor		Wiyynn	Vendor	Wiyynn	Vendor
Top	Mask		0.50					
	Signal	0.5 oz + plating	1.90	HTE	3.8		0.02599	
	Prepreg		2.70	1080	3.5		0.01356	
L2	Ground	1.0 oz	1.30	RTF				
	Core		4.00	2113x1	3.5		0.0145	
	Signal	1.0 oz	1.30	RTF				
L3	Prepreg		5.00	1080*2	3.5		0.0145	
	Ground	2.0 oz	2.60	RTF				
	Core		4.00	2113x1	3.5		0.0145	
L5	Signal	2.0 oz	2.60	RTF				
	Prepreg		12.00	1506*2	3.5		0.0145	
	Power	2.0 oz	2.60	RTF				
L6	Core		4.00	0.00	2113x1	3.5	0	0.0145
	Ground	2.0 oz	2.60	RTF				
	Prepreg		5.00	0.00	1080*2	3.5	0	0.0145
L8	Signal	1.0 oz	1.30	0.00	RTF			
	Core		4.00	0.00	2113x1	3.5	0	0.0145
	Ground	1.0 oz	1.30	0.00	RTF			
L9	Prepreg		2.70	0.00	1080	3.5	0	0.01356
	Signal	0.5 oz + plating	1.90	0.00	HTE			
	Mask		0.50	0.00	0	3.8	0	0.02599
Thickness requirement : 1.57 ± 10% (mm)		mil (without solder mask)	62.80	0.00				
		mm (without solder mask)	1.60	0.00				

Figure 12.25 UIC Stackup

Single ended impedance (ohm)												
Imp	40				50				55			
Variation	Inner/Outer +/- 10%				Inner/Outer +/- 10%				Inner/Outer +/- 10%			
Bus	N/A				Single ended signals				DDR4			
Type	S8				S8				S8			
Layers	1,3,8,10				1,3,8,10				1,3,8,10			
Layer	Microvia			v.door	Microvia			v.door	Microvia			v.door
	Width	Imp	Space		Width	Imp	Space		Width	Imp	Space	
S1	0.400	40	1	7	4.400	50	1	7	4.800	55	1	7
P2	reference layer				reference layer				reference layer			
S3	0.300	40	1	7	0.300	50	1	7	0.300	55	1	7
P4	reference layer				reference layer				reference layer			
S5	0.400	40	1	7	0.400	50	1	7	0.400	55	1	7
P6	reference layer				reference layer				reference layer			
P7	reference layer				reference layer				reference layer			
S8	0.300	40	1	7	0.300	50	1	7	0.300	55	1	7
P9	reference layer				reference layer				reference layer			
S10	0.400	40	1	7	0.400	50	1	7	0.400	55	1	7

Differential impedance (ohm)												
Imp	85				100				110			
Variation	Inner/Outer +/- 10%				Inner/Outer +/- 10%				Inner/Outer +/- 10%			
Bus	PCIe, USB				SATA, SAS, DDR4				SATA, SAS, DDR4			
Type	S8				S8				S8			
Layers	1,3,8,10				1,3,8,10				1,3,8,10			
Layer	Microvia			v.door	Microvia			v.door	Microvia			v.door
	Width	Space	Imp		Width	Space	Imp		Width	Space	Imp	
S1	0.400	7.5	85	7	0.400	7.5	100	7	0.400	7.5	110	7
P2	reference layer				reference layer				reference layer			
S3	0.300	7.5	85	7	0.300	7.5	100	7	0.300	7.5	110	7
P4	reference layer				reference layer				reference layer			
S5	0.400	7.5	85	7	0.400	7.5	100	7	0.400	7.5	110	7
P6	reference layer				reference layer				reference layer			
P7	reference layer				reference layer				reference layer			
S8	0.300	7.5	85	7	0.300	7.5	100	7	0.300	7.5	110	7
P9	reference layer				reference layer				reference layer			
S10	0.400	7.5	85	7	0.400	7.5	100	7	0.400	7.5	110	7

Figure 12.26: UIC Stackup

12.5 Power Delivery Board (PDB)

12.5.1 Configuration Diagram




Figure 12.27: PDB Design

12.5.2 Connectors

PDB to PTB:

Amphenol - MPN: 10106121-7000A01LF

TE - MPN: 2322124-9

PDB to Power Cable:

Amphenol - MPN: 10127402-23S1000LF

Alltop - MPN: C29404-10643-Y

12.5.3 PDB PCB

12.5.3.1 Important Components and Interfaces

The PDB is a completely passive, pass-through PCB. The intent is that if the Grand Canyon system needs to migrate to the ORV3 rack, a 48V to 12V converter could be implemented on this board to support the ORV3 architecture.

Below is the overall block diagram for the PDB shown in Figure 12.28.

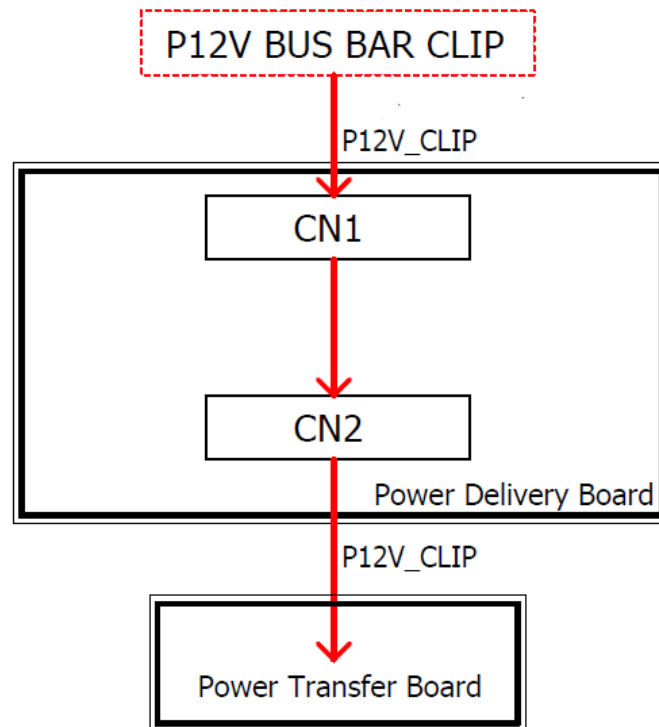


Figure 12.28: PDB Block Diagram

12.5.3.2 PCB Key Dimensions




Figure 12.29: PDB Dimensions

12.5.3.3 PCB Stack up

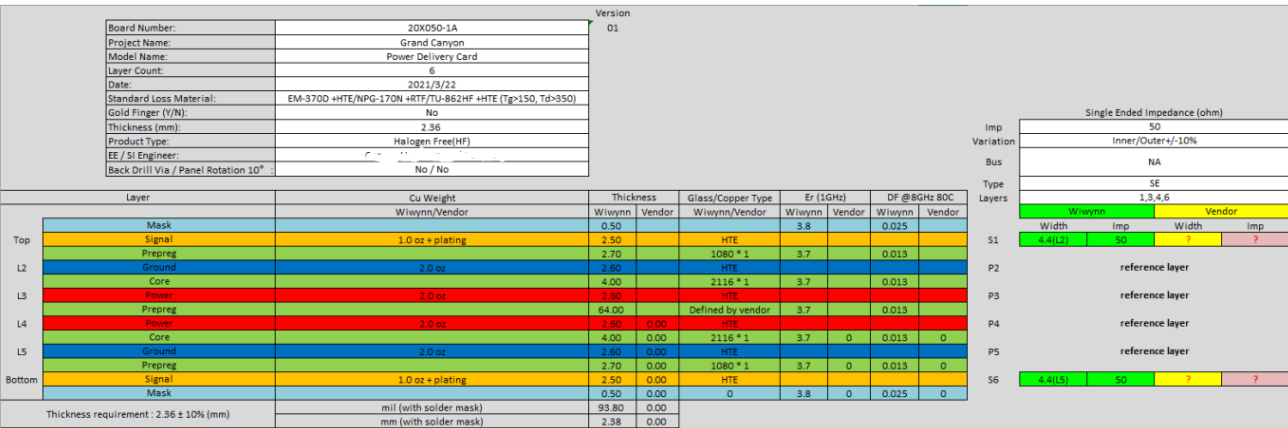


Figure 12.30: PDB Stackup

12.6 Power Transfer Board (PTB)

12.6.1 Configuration Diagram

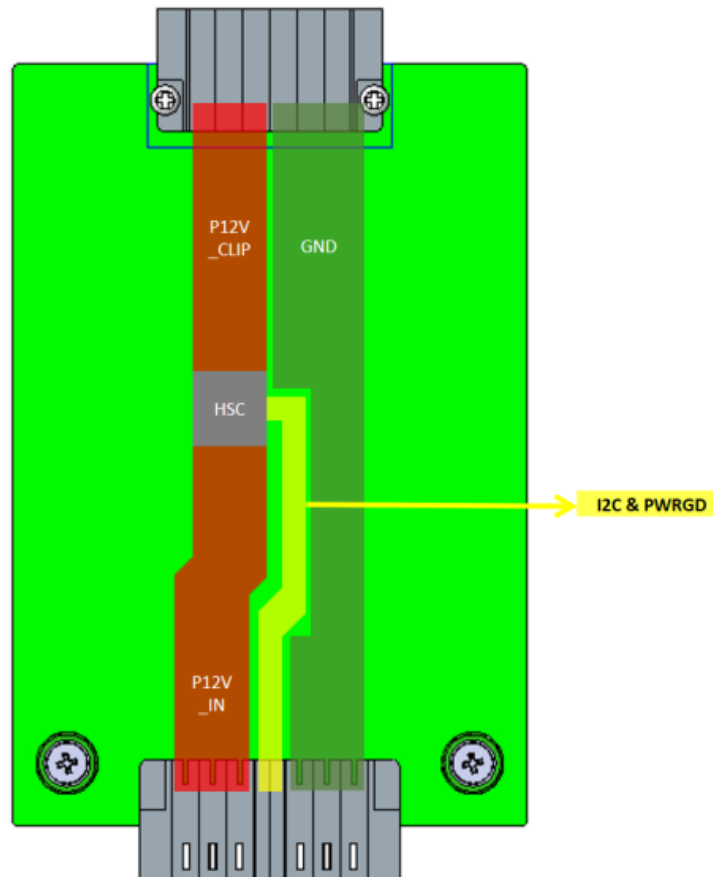


Figure 12.31: PTB Configuration

12.6.2 Connectors

PTB to RDPB::

Amphenol - MPN: 10127396-33H1410LF

Alltop - MPN: C29507-10643-Y

PTB to PDB:

Amphenol - MPN: 10127396-26S1000LF

Alltop - MPN: C29506-10643-Y

12.6.3 PTB PCB

12.6.3.1 Important Components and Interfaces

On the PTB board there is only one major component, the HSC. The HSC is responsible for ensuring when the Grand Canyon system is slid into the bus bar, power is applied properly not damaging any components or interfaces. The HSC has a power good signal that is used to start the power up sequence of the rest of the Grand Canyon System.

The HSC is also needed to protect the system in the event of a short or failure. Telemetry like current, voltage, power, etc. can be read from the HSC through an I2C interface.

A block diagram of the critical interfaces can be seen below in Figure 12.32.

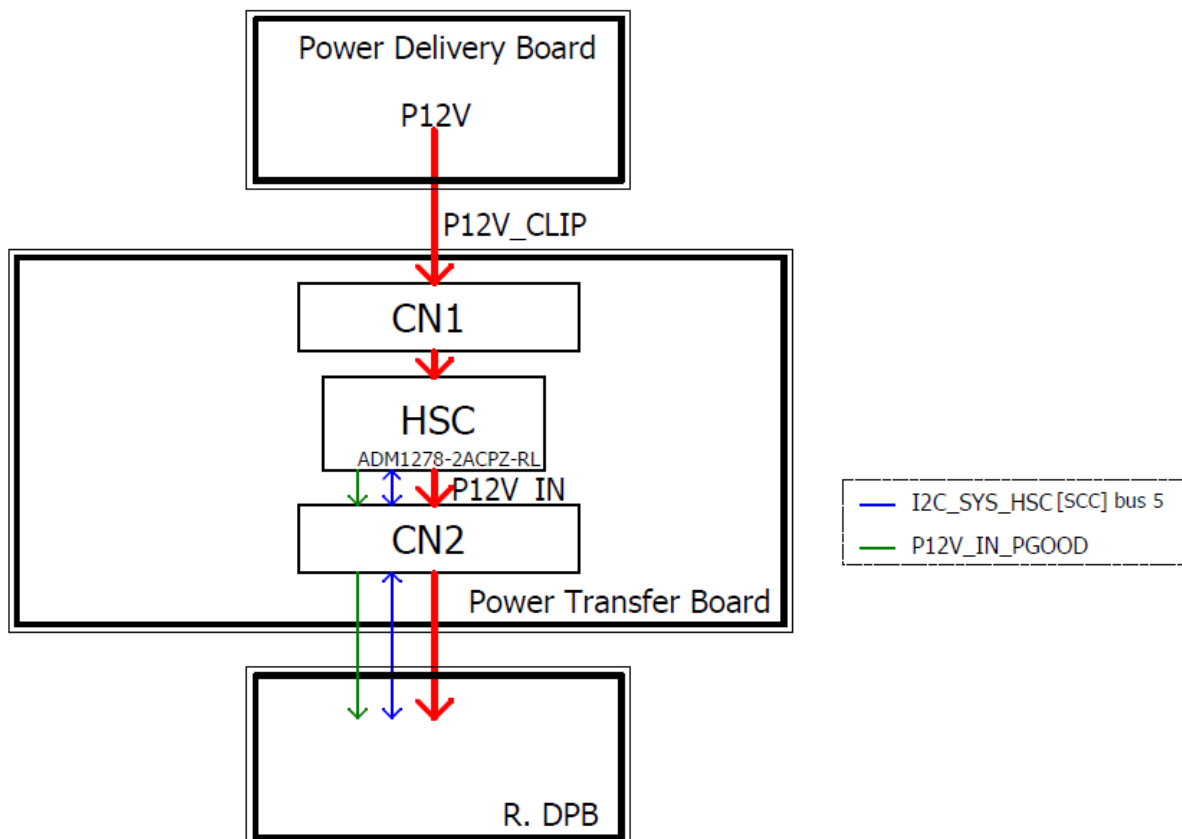


Figure 12.32: PTB Block Diagram

The power circuit diagram can be seen below in Figure 12.33.

Power

Standard voltage: 12.5V
OVP: 13.73V
UVP: 11V
OCP: 166A

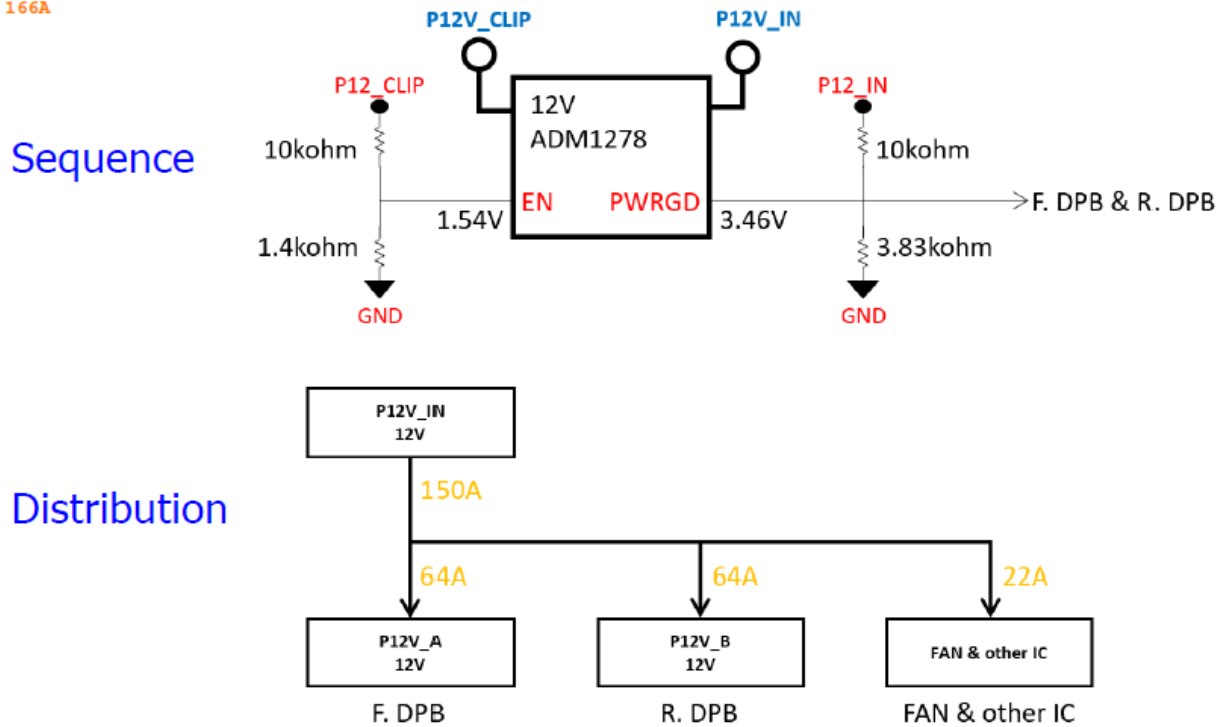


Figure 12.33: PTB Sequence and Distribution Diagram

12.6.3.2 PCB Key Dimensions




Figure 12.34: PTB Dimensions

12.6.3.3 PCB Stack up

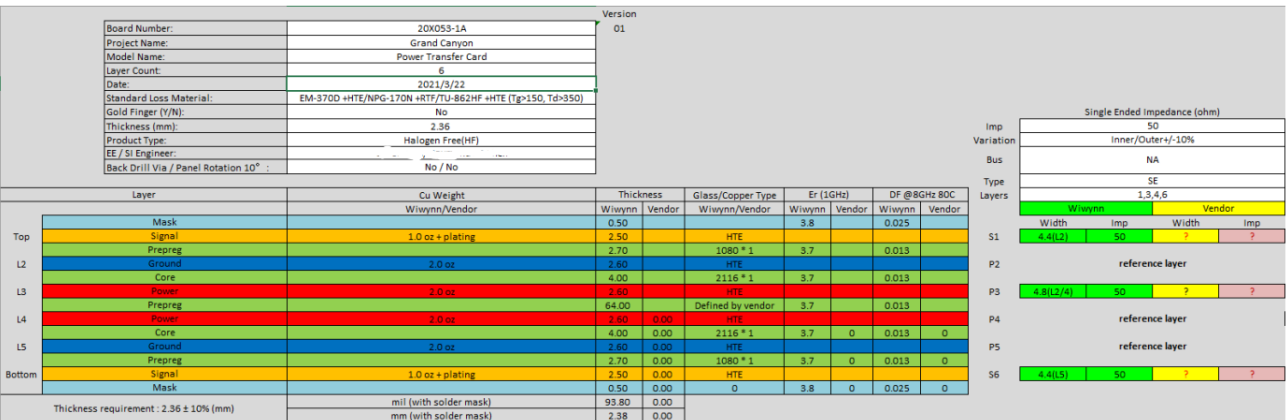


Figure 12.35: PTB Stackup

13. INTERNAL CABLES

13.1 JBOD Connection Bracket

When the Grand Canyon system is used as a JBOD, a bracket is installed with external MiniSAS HD connectors that route into the SCCs. An image of this connector is shown below.




Figure 13.1: JBOD Mini-SAS Front Panel Connector

This connector will terminate to an internal MiniSAS HD that plugs directly into the SCC board. There will be two cable lengths, one that routes to the front SCC, and one that routes to the rear SCC.

The two different lengths of cable are due to the physical positioning of the SCC boards as shown below in Figure 13.2.



Figure 13.2: JBOD Mini-SAS Cable

13.2 Single Server SAS Cable

When the Grand Canyon is used as a single storage server, the single compute modules need to access all 72 drive slots in the headnode. To do this, there will be an internal x4 SAS cable connecting the two SCCs. In order to make the cabling easier, this cable should have a 90-degree angle on both ends of the cable assembly.

14. EXTERNAL CABLES

14.1 JBOD configuration

The cabling between a Grand Canyon with 1x Barton Springs and its two corresponding Grand Canyon JBODs is shown in FIGURE 14.1. A custom Y-cable is used to connect the single server configuration to the JBODs. Not shown is the internal cabling - one cable in the server connects the two SCCs together (i.e., daisy-chained expander ASICs), and two cables in each JBOD are used to connect the individual SCC cards inside the JBOD to the front panel MiniSAS HD ports.




Figure 14.1: JBOD External Cabling

14.2 Cable Details

Two different lengths of the Y-cable will be used to accommodate the different distances of the two JBOD systems. Details of the cable drawing can be found in the design collateral packages.

15 BMC OVERVIEW

15.1 Introduction

There are a maximum of two BMCs in one Grand Canyon system (depending on the SKU). Each of the BMCs are physically connected to a server and an expander. The firmware running on the BMC is OpenBMC; the source code is available at <https://github.com/Meta/openbmc>.

Refer to the “Grand Canyon OpenBMC Specification” for details about the OpenBMC firmware interface, utilities and guidelines. At a high level, the BMC should support the following functions:

- Front Panel
 - OCP Debug Card
 - LEDs
- Power Control
- Remote Console
- IPMI FRUID
- Sensor Information
- Firmware Information/Update
 - Server card Firmware Update
 - Expander Firmware Recovery
- Log Information
- Fan Speed Control
- Enclosure Management
 - Enclosure Information
 - HDD status
 - Enclosure Error Code

15.2 Front Panel

The front panel has LEDs, buttons and interfaces that can be used to interact with the system. The functions of these LEDs and buttons has been explained earlier in this specification.

15.2.1 OCP Debug Card

Grand Canyon front panel USB ports support the use of the OCP LCD debug card and all of its features. Please refer to the OCP LCD debug card specification for more details.

15.3 Power Control

The BMC is responsible to control power to all the components of the system. In addition to controlling power to the various components, the BMC also provides a mechanism via console for a user to control

power to the server module or to the entire subsystem, including itself. Note that turning off power to the server also turns off power to the PCIe endpoints present in the system that are connected to that server, to adhere to PCIe Card Electro-Mechanical (CEM) specifications. The table below shows the details on the power control options for the server:

		BMC	Host(S5)	BS (Host (S0)+ BIC)	SCC	NIC	E1.S IOCM	HDD	SLED (UIC+BS+SCC+ HDD+NIC+E1S)
	Index	1	2	3	4	5	6	7	8
Power-on remote	1	Not support	[Ctrl 2-1]	[Ctrl 3-1]	[Ctrl4-1]	[Ctrl5-1]	Not support	[Ctrl7-1]	Not support
Power-on local	2	Not support	[Ctrl2-2]	Not support	Not support	Not support	Not support	Not support	[Ctrl 8-2]
Power-off remote	3	Not support	[Ctrl2-3]	[Ctrl 3-3]	[Ctrl4-3]	[Ctrl5-3]	Not support	[Ctrl7-3]	Not support
Power-off local	4	Not support	[Ctrl2-4]	Not support	Not support	Not support	Not support	Not support	[Ctrl 8-4]
Reset remote	5	[Ctrl 1-5]	[Ctrl2-5]	[Ctrl 3-5]	[Ctrl4-5]	Not support	Not support	Not support	[Ctrl 8-5]
Reset local	6	Not support	[Ctrl2-6]	Not support	Not support	Not support	Not support	Not support	Not support

For the supported portions:

(1) BMC

[Ctrl 1-5] reset remote
reboot

(2) Host (S5)

[Ctrl 2-1] power-on remote
power-util server on
[Ctrl 2-2] power-on local
Press the power button on debug card
[Ctrl 2-3] power-off remote
power-util server off
[Ctrl 2-4] power-off local
Press the power button on debug card
[Ctrl 2-5] reset remote
power-util server reset
[Ctrl 2-6] reset local
Press the reset button on debug card

(3) BS (Host (S0) + BIC)

- [Ctrl 3-1] power-on remote
 - # power-util server 12V-on*
 - (power on stand-by only, but if the power policy is 'on' or 'lps' with the last power state is 'on', will continue to power on Host OS)
- [Ctrl 3-3] power-off remote
 - # power-util server 12V-off*
 - (power off both standby and normal)
- [Ctrl 3-5] reset remote
 - # power-util server 12V-cycle*
- (4) SCC
 - [Ctrl 4-1] power-on remote
 - # gpiocli -c aspeed-gpio -s SCC_STBY_PWR_EN set-value 1*
 - [Ctrl 4-3] power-off remote
 - # gpiocli -c aspeed-gpio -s SCC_STBY_PWR_EN set-value 0*
 - [Ctrl 4-5] reset remote
 - # gpiocli -c aspeed-gpio -s BMC_RST_BTN_N_R set-value 0*
 - # sleep 1*
 - # gpiocli -c aspeed-gpio -s BMC_RST_BTN_N_R set-value 1*
- (5) NIC
 - [Ctrl 5-1] power-on remote
 - # gpiocli -c aspeed-gpio -s BMC_NIC_FULL_PWR_EN_R 1*
 - [Ctrl 5-3] power-off remote
 - # gpiocli -c aspeed-gpio -s BMC_NIC_FULL_PWR_EN_R 0*
- (7) HDD
 - [Ctrl 7-1] power-on remote
 - From host:
 - # fbjbod hdd --hdd-on \$hdd_slot /dev/sgX (\$hdd_slot: 0~35)*
 - [Ctrl 7-3] power-off remote
 - From host:
 - # fbjbod hdd --hdd-off \$hdd_slot /dev/sgX (\$hdd_slot: 0~35)*
- (8) SLED (UIC+BS+SCC+HDD+NIC+E1S)
 - [Ctrl 8-2] power-on local
 - Press the power button on UIC
 - [Ctrl 8-4] power-off local
 - Press the power button on UIC
 - [Ctrl 8-5] reset remote
 - # power-util sled-cycle*

15.4 Remote Console

The BMC can provide access to the Server remote console which is commonly known as Serial Over LAN (SOL). There can be multiple sessions connected to the Server console. It also buffers the console history of the server.

15.5 IPMI FRUID

The BMC allows the user to read and write the FRUID information to all the components of the system. FRUID information can vary from manufacturer name, manufacturing date, serial number, part number and sometimes to custom fields to differentiate hardware also. Each PCBA in the system is a FRU and has its own FRUID and associated identification data that BMC can access and write also. Here is the list of the FRUs whose information can be fetched by the BMC.

The utility '*fruid-util*' can be used to read/write/modify the FRUID information of each board. Support to modify individual fields in FRUID using "fruid-util <FRU> --modify --<field> <path to binary>".

-- dump dumps the FRU information raw data of a particular board into a *file* mentioned in the command.

--write flash new FRU image *file* into the eeprom of the board.

--modify modifies the FRU contents and stores it in another FRU image *file*. It only modifies the FRU image *file*. The user should execute '**--write**' command to write the new FRU image to eeprom if necessary.

FRU ID mapping table:

FRU ID	Board Name
0	All
1	Server
2	UIC
3	DPB
4	SCC
5	NIC
6	E1S_IOCM

15.6 Sensor Information

There are several temperature, voltage, power sensors throughout the system that are read and monitored by the BMC. The BMC can read the sensor values from the server, DPB, IOCM, SCC and NIC. In addition to this, the BMC constantly monitors the sensors on the IOM, NIC and Server every two seconds to look for any anomalies in the system.

15.7 Firmware Information/Update

The BMC has the mechanism to read the current firmware version and push firmware updates when available to the critical components in the system. The BMC is able to read the version and update the firmware of the following components:

- BIOS
- ME
- CPLD
- Bridge-IC
- Bridge-IC Boot loader

- BMC itself
- Expander firmware

15.8 Log Information

The BMC provides logs which refer to the various errors and events that happened on the system. In concept, this is similar to IPMI's System Event Log (SEL). These logs are stored in persistent memory.

15.9 Fan Speed Control

The Fan Speed Control (FSC) is designed based on the system design, ambient temperature and altitude (environment variations), and desired margin for each component in the system from its allowable maximum temperature. Based on the various temperature sensors in the system, the BMC sets the PWM for the fans. The initial default fan speed (PWM) is set to 50% duty cycle.

15.10 Enclosure Management

The utility '**enclosure-util**' provides the enclosure information, HDD status, flash status, and enclosure error codes.

15.10.1 HDDs Status

Users will get the 4 bytes response data when sending the command with "**--hdd-status**" option. The array device status element is below:

15.10.2 Error Code

The **--error** option will inform users whether there is an error in the system. The error code depends on whether the error is monitored from Expander or BMC, thus the error code is defined by both Expander and BMC.

Currently, we inherit the definition of Expander Error Code from Bryce Canyon.

Expander Error Code Definition (GrandCanyon_ErrorCode_EventLogList_0501):

EXPANDER			
Error Code	Description	Error Code	Description
0	No error	50	HDD20 fault request
1	I2C bus 1 crash	51	HDD21 fault request
2	I2C bus 2 crash	52	HDD22 fault request
3	I2C bus 3 crash	53	HDD23 fault request
4	I2C bus 4 crash	54	HDD24 fault request
5	I2C bus 5 crash	55	HDD25 fault request
6	I2C bus 6 crash	56	HDD26 fault request
7	I2C bus 7 crash	57	HDD27 fault request
8	I2C bus 0 crash	58	HDD28 fault request
9	IOC HeartBeat loss	59	HDD29 fault request
10	Reserved	60	HDD30 fault request
11	SCC voltage critical	61	HDD31 fault request
12	SCC_HSC voltage critical	62	HDD32 fault request

13	DPB voltage critical	63	HDD33 fault request
14	DPB_HSC voltage critical	64	HDD34 fault request
15	PTB_HSC_CLIP voltage critical	65	HDD35 fault request
16	HDD X voltage critical	66	HDD X fault sensed
17	SCC current critical	67	Reserved
18	DPB current critical	68	Reserved
19	PTB_HSC_CLIP current critical	69	Reserved
20	Reserved	70	Internal Mini-SAS Link Loss
21	SCC_Expander_Temp critical	71	Internal SAS Link Loss
22	Reserved	72	SCC I2C device loss
23	SCC_Temp1 critical	73	DPB I2C device loss
24	SCC_Temp2 critical	74	PTB I2C device loss
25	DPB_INLET_Temp1 critical	75	UIC I2C device loss
26	DPB_INLET_Temp2 critical	76	Reserved
27	DPB_OUTLET_Temp critical	77	Fan 0 front fault
28	HDD X SMART temp. critical	78	Fan 0 rear fault

29	UIC_Temp critical	79	Fan 1 front fault
30	HDD0 fault request	80	Fan 1 rear fault
31	HDD1 fault request	81	Fan 2 front fault
32	HDD2 fault request	82	Fan 2 rear fault
33	HDD3 fault request	83	Fan 3 front fault
34	HDD4 fault request	84	Fan 3 rear fault
35	HDD5 fault request	85	Reserved
36	HDD6 fault request	86	Reserved
37	HDD7 fault request	87	Reserved
38	HDD8 fault request	88	Drawer Pull out
39	HDD9 fault request	89	Peer SCC Plug out
40	HDD10 fault request	90	UICA Plug out
41	HDD11 fault request	91	UICB Plug out
42	HDD12 fault request	92	Fan 0 Plug out
43	HDD13 fault request	93	Fan 1 Plug out
44	HDD14 fault request	94	Fan 2 Plug out

45	HDD15 fault request	95	Fan 3 Plug out
46	HDD16 fault request	96	Reserved
47	HDD17 fault request	97	Reserved
48	HDD18 fault request	98	Reserved
49	HDD19 fault request	99	H/W Configuration/Type Not Match

BMC			
Error Code	Description	Error Code	Description
0xCE	BMC CPU utilization exceeded	0xEF	I2C bus 8 hang
0xCF	BMC Memory utilization exceeded	0xF0	I2C bus 9 hang
0xE0	ECC Recoverable Error	0xF1	I2C bus 10 hang
0xE1	ECC Unrecoverable Error	0xF2	I2C bus 11 hang
0xE2	Barton Springs Missing	0xF3	I2C bus 12 hang
0xE3	SCC Missing	0xF4	I2C bus 13 hang
0xE4	NIC Missing	0xF5	I2C bus 14 hang

0xE5	E1.S Missing	0xF6	I2C bus 15 hang
0xE6	IOC Module Missing	0xF7	Server health is bad
0xE7	I2C bus 0 hang	0xF8	UIC health is bad
0xE8	I2C bus 1 hang	0xF9	DPB health is bad
0xE9	I2C bus 2 hang	0xFA	SCC health is bad
0xEA	I2C bus 3 hang	0xFB	NIC health is bad
0xEB	I2C bus 4 hang	0xFC	BMC remote heartbeat health is abnormal
0xEC	I2C bus 5 hang	0xFD	SCC local heartbeat health is abnormal
0xED	I2C bus 6 hang	0xFE	SCC remote heartbeat health is abnormal
0xEE	I2C bus 7 hang	0xFF	Reserved

BMC error code definition:

For the error codes from 0xE7 to 0xF6, GC won't enable I2C bus monitoring until the unified bus status reporting solution is defined.

15.11 Event List Along With The Error Code

DESCRIPTION	ERROR CODE	CRITERIA	LOG STRING
FRU EEPROM is empty	N/A	If all of the data header is 0xff in EEPROM	Critical-FRU header <FRU_BIN_FILE> is empty
New FRU file's checksum is invalid	N/A	New FRU file's checksum is invalid	Critical-New FRU data checksum is invalid
BMC CPU utilization exceeded	0xCE	If BMC CPU utilization exceeds the threshold.	Warning-ASSERT: BMC CPU utilization (<CURRENT_CPU>%) exceeds the threshold (<CPU_THRESHOLD>%). Warning-DEASSERT: BMC CPU utilization (<CURRENT_CPU>%) is under the threshold (<CPU_THRESHOLD>%).
BMC Memory utilization exceeded	0xCF	If BMC Memory utilization exceeds the threshold.	Critical-ASSERT: BMC Memory utilization (<CURRENT_MEM>%) exceeds the threshold (<MEM_THRESHOLD>%). Critical-DEASSERT: BMC Memory utilization (<CURRENT_MEM>%) is under the threshold (<MEM_THRESHOLD>%).
ECC Recoverable Error	0xE0	If ECC Recoverable Error Counter over thresholds (0%, 50%, and 90%).	ECC Recoverable Error occurred (<ECC_THRESHOLD>%) Counter = <ECC_COUNTER> Address of last recoverable error = <MCR5C>

ECC Unrecoverable Error	0xE1	If ECC Unrecoverable Error Counter over thresholds (0%, 50%, and 90%).	Critical-ECC Unrecoverable Error occurred (over <ECC_THRESHOLD>%) Counter = <MCR50> Address of first unrecoverable ECC error = <MCR58>
Barton Springs Missing	0xE2	Check Barton Springs is present or not via IO expander P00.	Critical-ASSERT: server missing Critical-DEASSERT: server missing
SCC Missing	0xE3	Check if SCC is present or not via IO expander. (should confirm the GPIO number) Type 5: UICA check P06, UICB check P07 Type 7: check P06 and P07	Critical-ASSERT: scc missing Critical-DEASSERT: scc missing
NIC Missing	0xE4	Check Mezz card is present or not via BMC GPIOAA1. (should confirm the GPIO number)	Critical-ASSERT: nic missing Critical-DEASSERT: nic missing
E1.S Missing	0xE5	Check E1.S is present or not via IO expander. E1.S 0: P16 E1.S 1: P17	Critical-ASSERT: e1.s <SIDE><NUM> missing Critical-DEASSERT: e1.s <SIDE><NUM> missing
IOC Module Missing	0xE6	Check E1.S is present or not via IO expander P16 and P17.	Critical-ASSERT: iocm missing Critical-DEASSERT: iocm missing
I2C bus 0 hang	0xE7	I2C bus crash/I2C driver hang	For the error codes from 0xE7 to 0xF6, GC won't enable I2C bus monitoring until the unified bus status reporting solution is defined. Critical-ASSERT: ASSERT: I2C(<BUS_ID>) bus is locked
I2C bus 1 hang	0xE8	I2C bus crash/I2C driver hang	
I2C bus 2 hang	0xE9	I2C bus crash/I2C driver hang	

I2C bus 3 hang	0xEA	I2C bus crash/I2C driver hang	<p>(Master Lock or Slave Clock Stretch). Recovery error. (I2C bus index base 0)</p> <p>Critical-ASSERT: I2C(<BUS_ID>) bus is locked (Master Lock or Slave Clock Stretch). Recovery timed out. (I2C bus index base 0)</p> <p>Critical-ASSERT: I2C(<BUS_ID>) Slave is dead (SDA keeps low). Bus recovery error. (I2C bus index base 0)</p> <p>Critical-ASSERT: I2C(<BUS_ID>) Slave is dead (SDAs keep low). Bus recovery timed out. (I2C bus index base 0)</p> <p>Critical-ASSERT: I2C(<BUS_ID>) Undefined case. (I2C bus index base 0)</p> <p>Critical-DEASSERT: I2C(<BUS_ID>) Bus recovered. (I2C bus index base 0)</p> <p>Critical-I2C(<BUS_ID>) bus had been locked (Master Lock or Slave Clock Stretch) and has been recovered successfully. (I2C bus index base 0)</p> <p>Critical-I2C(<BUS_ID>) Slave was dead and bus has been recovered successfully. (I2C bus index base 0)</p>
I2C bus 4 hang	0xEB	I2C bus crash/I2C driver hang	
I2C bus 5 hang	0xEC	I2C bus crash/I2C driver hang	
I2C bus 6 hang	0xED	I2C bus crash/I2C driver hang	
I2C bus 7 hang	0xEE	I2C bus crash/I2C driver hang	
I2C bus 8 hang	0xEF	I2C bus crash/I2C driver hang	
I2C bus 9 hang	0xF0	I2C bus crash/I2C driver hang	
I2C bus 10 hang	0xF1	I2C bus crash/I2C driver hang	
I2C bus 11 hang	0xF2	I2C bus crash/I2C driver hang	
I2C bus 12 hang	0xF3	I2C bus crash/I2C driver hang	
I2C bus 13 hang	0xF4	I2C bus crash/I2C driver hang	
I2C bus 14 hang	0xF5	I2C bus crash/I2C driver hang	
I2C bus 15 hang	0xF6	I2C bus crash/I2C driver hang	
Server health	0XF7	Open Server health file failed or Server status is abnormal	

UIC health	0xF8	Open UIC health file failed or UIC status is abnormal	
DPB health	0xF9	Open DPB health file failed or DPB status is abnormal	
SCC health	0xFA	Open SCC health file failed or SCC status is abnormal	
NIC health	0xFB	Open NIC health file failed or NIC status is abnormal	
BMC remote heartbeat health	0xFC	Type 5 only. BMC remote heartbeat is lost continuous 3 minutes	Critical-BMC remote heartbeat is abnormal
SCC local heartbeat health	0xFD	SCC local heartbeat is lost continuous 3 minutes	Critical-SCC local heartbeat is abnormal
SCC remote heartbeat health	0xFE	Type 7 only. SCC remote heartbeat is lost continuous 3 minutes	Critical-SCC remote heartbeat is abnormal

15.12 IPMI Command Support List

COMMAND NAME	NETFN	CODE	DATA
Chassis (0x00)			
CHASSIS_GET_STATUS	0x00	0x01	Ref IPMI spec 28.2 (only support get power policy and current status)
CHASSIS_SET_POWER_RESTORE_POLICY	0x00	0x06	Ref IPMI spec 28.8
Sensor/Event (0x04)			
SENSOR_PLAT_EVENT_MSG	0x04	0x02	Ref IPMI spec 29.3
App (0x06)			
APP_GET_DEVICE_ID	0x06	0x01	Ref IPMI spec 20.1
APP_COLD_RESET	0x06	0x02	Ref IPMI spec 20.2
APP_GET_SELFTEST_RESULTS	0x06	0x04	Ref IPMI spec 20.4
APP_MANUFACTURING_TEST_ON	0x06	0x05	Ref IPMI spec 20.5 (send “sled-cycle”ASCII string)
APP_GET_DEVICE_GUID	0x06	0x08	Ref IPMI spec 20.8
APP_RESET_WDT	0x06	0x22	Ref IPMI spec 27.5
APP_SET_WDT	0x06	0x24	Ref IPMI spec 27.6
APP_GET_WDT	0x06	0x25	Ref IPMI spec 27.7

APP_SET_GLOBAL_ENABLES	0x06	0x2E	Ref IPMI spec 22.1
APP_GET_GLOBAL_ENABLES	0x06	0x2F	Ref IPMI spec 22.2
APP_CLEAR_MESSAGE_FLAGS	0x06	0x30	Ref IPMI spec 22.3
APP_GET_SYSTEM_GUID	0x06	0x37	Ref IPMI spec 22.14
APP_SET_SYS_INFO_PARAMS	0x06	0x58	Ref IPMI spec 22.14a
APP_GET_SYS_INFO_PARAMS	0x06	0x59	Ref IPMI spec 22.14b
Storage (0x0A)			
STORAGE_GET_FRUID_INFO	0x0A	0x10	Ref IPMI spec 34.1
STORAGE_READ_FRUID_DATA	0x0A	0x11	Ref IPMI spec 34.2
STORAGE_GET_SDR_INFO	0x0A	0x20	Ref IPMI spec 33.9
STORAGE_RSV_SDR	0x0A	0x22	Ref IPMI spec 33.11
STORAGE_GET_SDR	0x0A	0x23	Ref IPMI spec 33.12
STORAGE_GET_SEL_INFO	0x0A	0x40	Ref IPMI spec 31.2
STORAGE_RSV_SEL	0x0A	0x42	Ref IPMI spec 31.4
STORAGE_GET_SEL	0x0A	0x43	Ref IPMI spec 31.5
STORAGE_ADD_SEL	0x0A	0x44	Ref IPMI spec 31.6
STORAGE_CLR_SEL	0x0A	0x47	Ref IPMI spec 31.9
STORAGE_GET_SEL_TIME	0x0A	0x48	Ref IPMI spec 31.10

STORAGE_GET_SEL_UTC	0x0A	0x5C	Ref IPMI spec 31.11a
Transport (0x0C)			
TRANSPORT_SET_LAN_CONFIG	0x0C	0x01	Ref IPMI spec 23.1
TRANSPORT_GET_LAN_CONFIG	0x0C	0x02	Ref IPMI spec 23.2; IPv6 parameter: 197
TRANSPORT_GET_SOL_CONFIG	0x0C	0x22	Ref IPMI spec 26.3

OEM (0x30)			
CMD_OEM_GET_SENSOR_REAL_READING	0x30	0x20	Request: Byte 1 - FRU Byte 2 - sensor number Response: Byte 1 - Completion Code Byte 2~3 - sensor reading in decimal integer Byte 4 - sensor reading in decimal point
OEM_BYPASS_CMD	0x30	0x34	Request: Byte 1 - Bypass Command Target 00h: Bridge-IC 01h: ME 02h: reserved and not supported 03h: NCSI

			<p>04h: Network</p> <p>Byte 2 ~ Byte N -</p> <p>If Byte 1 = 00h or 01h: Request IPMI Data</p> <p>If Byte 1 = 03h: Byte 2: Network Interface Byte 3: Channel Byte 4: Command Code Byte 5 ~ Byte N: Request NCSI Data</p> <p>If Byte 1 = 04h: Byte 2: Network Interface Byte 3: Action 00h: Enable 01h: Disable</p> <p>Response: Byte 1 - Completion Code Byte 2 - Standard cmd: ref IPMI spec OEM cmd: ref corresponding IPMI OEM command NCSI cmd: ref NCSI spec</p>
OEM_SET_BOOT_ORDER	0x30	0x52	<p>Request:</p> <p>Byte 1 - Boot Mode Bit 0 - 0 : Legacy, 1 : UEFI Bit 1 - CMOS Clear (Optional, BIOS implementation dependent) Bit 2 - Force Boot into BIOS Setup (Optional, BIOS implementation dependent) Bit 6:3 - reserved Bit 7 - Boot Flags Valid</p>

			<p>Byte 2 ~ Byte 6 - Boot Sequence</p> <p>Bit 2:0 - Boot Device ID</p> <p>000b: USB Device</p> <p>001b: Network</p> <p>010b: SATA HDD</p> <p>011b: SATA-CDROM</p> <p>100b: Other Removable Device</p> <p>Bit 7:3 - reserve for boot device special request</p> <p>If Bit 2:0 is 001b (Network), Bit3 is IPv4/IPv6 order</p> <p>Bit3 = 0b: IPv4 first</p> <p>Bit3 = 1b: IPv6 first</p> <p>Response:</p> <p>Byte 1 - Completion Code</p>
OEM_GET_BOOT_ORDER	0x30	0x53	<p>Request:</p> <p>None</p> <p>Response:</p> <p>Byte 1 - Completion Code</p> <p>Byte 2 - Boot Mode</p> <p>Bit 0 - 0 : Legacy, 1 : UEFI</p> <p>Bit 1 - CMOS Clear</p> <p>(Optional, BIOS implementation dependent)</p> <p>Bit 2 - Force Boot into BIOS Setup</p> <p>(Optional, BIOS implementation dependent)</p> <p>Bit 6:3 - reserved</p> <p>Bit 7 - Boot Flags Valid</p> <p>Byte 3 ~ Byte 7 - Boot Sequence</p> <p>Bit 2:0 - Boot Device ID</p>

			000b: USB Device 001b: Network 010b: SATA HDD 011b: SATA-CDROM 100b: Other Removable Device Bit 7:3 - reserve for boot device special request If Bit 2:0 is 001b (Network), Bit3 is IPv4/IPv6 order Bit3 = 0b: IPv4 first Bit3 = 1b: IPv6 first
OEM_GET_PLAT_INFO	0x30	0x7E	Request: None Response: Byte 1 - Completion Code Byte 2 - Platform Configuration Number 91h: Grand Canyon Type 5A 99h: Grand Canyon Type 5B A1h: Grand Canyon Type 7
OEM_SET_PPR	0x30	0x90	DDR4 PPR/sPPR support;Need to repair DDR4 memory by using extra row as part of repair flow Request: Byte 1 - PPR Command Selector Byte 2..N - Configuration parameter data per table below Response: Byte 1 - Completion Code Command: PPR Action Selector: 0x01

			<p>Byte 1 - Enable or disable PPR function in next host system reboot</p> <p>Bit[7]: 0 for disable;1 for enable</p> <p>Bit[6..0] - 01h for Soft PPR; 02h for Hard PPR</p> <p>If there is no candidate row for repair: When get, response data is zero for PPR function When Set, response completion code: D5h for parameter not support in current state</p> <p>Command: PPR Candidate Row Count</p> <p>Selector: 0x2</p> <p>Byte1 - PPR Candidate row count, range from 0 to 100</p> <p>0h - No candidate row</p> <p>Command: PPR Candidate Row Address</p> <p>Selector: 0x3</p> <p>Byte1 - Set Selector: PPR candidate row index, from 0 to ("PPR Candidate Row Count" - 1).</p> <p>Byte 2..7 - PPR candidate row address</p> <p>Byte 2:</p> <p>Bit[7:2] - Reserved</p> <p>Bit[1:0] - Logical Rank</p> <p>Byte 3:</p> <p>Bit[7:5] - Socket Number</p> <p>Bit[4:2] - Channel Number</p> <p>Bit[0:1] - DIMM Number</p>
--	--	--	---

			<p>Byte 4: Device ID</p> <p>Byte 5:</p> <p> Bit[7:4] - Bank</p> <p> Bit[3:0] - Bank Group</p> <p>Byte 6: Row (LSB)</p> <p>Byte 7: Row (MSB)</p> <p>Command: PPR History Data</p> <p>Selector: 0x04</p> <p> Byte1: Set Selector: PPR history index, from 0 to ("PPR History Count" - 1)</p> <p> Byte 2..17: PPR history data</p> <p> Byte 2..5 - PPR Table</p> <p>Signature 0xbfeeeefb</p> <p> Byte 6 - Checksum</p> <p> Byte 7..10 - Seconds since Epoch for First PPR Entry</p> <p> Byte 11..14 - Seconds since Epoch for the Most Recent Entry</p> <p> Byte 15 - Number of Successful PPR (LSB)</p> <p> Byte 16 - Number of Successful PPR (MSB)</p> <p> Byte 17 - DIMM ID</p>
OEM_GET_PPR	0x30	0x91	<p>DDR4 PPR/sPPR support;Need to repair DDR4 memory by using extra row as part of repair flow</p> <p>Request:</p> <p> Byte 1 - PPR Command Selector</p> <p> Byte 2 - Set Selector. Selects a given set of parameters under a parameter selector value. 00h if parameter does not require a set selector</p>

			<p>Response:</p> <p>Byte 1 - Completion Code</p> <p>Byte 2..N - Configuration parameter data per table below</p> <p>Command: PPR Action</p> <p>Selector: 0x01</p> <p>Byte 1 - Enable or disable PPR function in next host system reboot</p> <p>Bit[7]: 0 for disable;1 for enable</p> <p>Bit[6..0] - 01h for Soft PPR; 02h for Hard PPR</p> <p>If there is no candidate row for repair:</p> <p>When get, response data is zero for PPR function</p> <p>When Set, response completion code: D5h for parameter not support in current state</p> <p>Command: PPR Candidate Row Count</p> <p>Selector: 0x2</p> <p>Byte1 - PPR Candidate row count, range from 0 to 100</p> <p>0h - No candidate row</p> <p>Command: PPR Candidate Row Address</p> <p>Selector: 0x3</p> <p>Byte1 - Set Selector: PPR candidate row index, from 0 to</p>
--	--	--	--

		<p>("PPR Candidate Row Count" - 1).</p> <p>Byte 2..7 - PPR candidate row address</p> <p>Byte 2:</p> <p>Bit[7:2] - Reserved</p> <p>Bit[1:0] - Logical Rank</p> <p>Byte 3:</p> <p>Bit[7:5] - Socket Number</p> <p>Bit[4:2] - Channel Number</p> <p>Bit[0:1] - DIMM Number</p> <p>Byte 4: Device ID</p> <p>Byte 5:</p> <p>Bit[7:4] - Bank</p> <p>Bit[3:0] - Bank Group</p> <p>Byte 6: Row (LSB)</p> <p>Byte 7: Row (MSB)</p> <p>Command: PPR History Data</p> <p>Selector: 0x04</p> <p>Byte1: Set Selector: PPR history index, from 0 to ("PPR History Count" - 1)</p> <p>Byte 2..17: PPR history data</p> <p>Byte 2..5 - PPR Table</p> <p>Signature 0xbfeeeefb</p> <p>Byte 6 - Checksum</p> <p>Byte 7..10 - Seconds since Epoch for First PPR Entry</p> <p>Byte 11..14 - Seconds since Epoch for the Most Recent Entry</p> <p>Byte 15 - Number of Successful PPR (LSB)</p> <p>Byte 16 - Number of Successful PPR (MSB)</p> <p>Byte 17 - DIMM ID"</p>
--	--	--

CMD_OEM_GET_USB_CDC_STATUS	0x30	0xB0	Request: None Response: Byte 1 - Completion Code Byte 2 - Status 00h: Disable 01h: Enable
CMD_OEM_CTRL_USB_CDC	0x30	0xB1	Request: Byte 1 - Status Control 00h: Disable 01h: Enable Response: Byte 1 - Completion Code
OEM_GET_PCIE_CONFIG	0x30	0xF4	Request: None Response: Byte 1 - Completion Code Byte 2 - PCIe Configuration Number 06h: Type 5 08h: Type 7
OEM (0x32)			
OEM_ADD_STRING_SEL	0x32	0x30	Request: Byte 1 - For BMC Reserved Byte 2 - String Length Byte 3 ~ Byte N - Event Log String (ASCII) Response: Byte 1 - Completion Code
OEM_SETUP_EXP_UART_BRIDGING	0x32	0x40	Request: None Route Host UART0 to Expander Debug Port Response:

			Byte 1 - Completion Code
OEM_TEARDOWN_EXP_UART_BRIDGING	0x32	0x41	Request: None Disable UART route of Host UART0 to Expander Debug Port Response: Byte 1 - Completion Code
OEM_SET_IOC_FW_RECOVERY	0x32	0x42	Request: Byte1 - Component 00h: SCC 01h: IOC Module Byte 2 - Set the IOC Recovery State 00h: Disable 01h: Enable Response: Byte 1 - Completion Code
OEM_GET_IOC_FW_RECOVERY	0x32	0x43	Request: Byte1 - Component 00h: SCC 01h: IOC Module Response: Byte 1 - Completion Code Byte 2 - Current Status of IOC Recovery 00h: Disable 01h: Enable
OEM_SET_IOC_WWID	0x32	0x50	Request: Byte1 - Component 01h: IOC Module Byte2 - Byte9 - WWID Response:

			Byte 1 - Completion Code
OEM_GET_IOC_WWID	0x32	0x51	Request: Byte1 - Component 00h: SCC 01h: IOC Module Response: Byte 1 - Completion Code Byte 2 - Byte9 - WWID
OEM (0x38)			
Reset BMC from BIC	0x38	0x16	Request: Byte 1 ~ Byte 3 - Manufacturer ID 009C9Ch, LS byte first Response: Byte 1 - Completion Code

[AC1]update

15.13 Expander IPMB Command Support List

EXPANDER SUPPORT COMMAND						
Command Name	NetFn	Code	oBMC Request Data	Expander Response Data		Description
OEM						
Get Expander Error Code	0x30	0x11	None	Byte 1	Expander error codes 0-7	Return total of 13 bytes with all Expander error code. Please refer to Expander FW Specification for detailed error code description.
				Byte 2	Expander error codes 8-15	
				Byte 3	Expander error codes 16-23	
				Byte 4	Expander error codes 24-31	
				Byte 5	Expander error codes 32-39	
				Byte 6	Expander error codes 40-47	
				Byte 7	Expander error codes 48-55	
				Byte 8	Expander error codes 56-63	
				Byte 9	Expander error codes 64-71	
				Byte 10	Expander error codes 72-79	
				Byte 11	Expander error codes 80-87	

				Byte 12	Expander error codes 88-95	
				Byte 13	Expander error codes 96-103 (error codes 100 to 103 are reserved)	
Get Expander FW Version	0x30	0x12	None	Byte 1	Region 0 (Bootloader) Status	Totally return 20 bytes of this command.
				Byte 2	Region 0 (Bootloader) Major Version	
				Byte 3	Region 0 (Bootloader) Minor Version	Expander has 1 bootloader region, 2 firmware regions and 1 configuration region.
				Byte 4	Region 0 (Bootloader) Unit Version	
				Byte 5	Region 0 (Bootloader) Dev Version	
				Byte 6	Region 1 (Firmware) Status	Each region has 1 status byte and 4 region version bytes.
				Byte 7	Region 1 (Firmware) Major Version	
				Byte 8	Region 1 (Firmware) Minor Version	
				Byte 9	Region 1 (Firmware) Unit Version	
				Byte 10	Region 1 (Firmware) Dev Version	
				Byte 11	Region 2 (Firmware) Status	
				Byte 12	Region 2 (Firmware) Major Version	
				Byte 13	Region 2 (Firmware) Minor Version	
				Byte 14	Region 2 (Firmware) Unit Version	

					Byte 15	Region 2 (Firmware) Dev Version	
					Byte 16	Region 3 (Configuration) Status	
					Byte 17	Region 3 (Configuration) Major Version	
					Byte 18	Region 3 (Configuration) Minor Version	
					Byte 19	Region 3 (Configuration) Unit Version	
					Byte 20	Region 3 (Configuration) Dev Version	
OEM Get Sensor Readings	0x30	0x2D	Byte 1	NrSR: number of sensors request from oBMC	Byte 1	Completion Code	
			Byte 2	sensor number #1	Byte 2	NrSR: number of sensors response from Expander	
			Byte 3	sensor number #2	Byte 3	Data1 of Sensor #1	
			Byte 4	sensor number #3	Byte 4	Data2 of Sensor #1	
			Byte 5	sensor number #4	Byte 5	Data3 of Sensor #1	
			Byte 6	Data4 of Sensor #1	
			Byte N	sensor number #n	Byte 7	Data5 of Sensor #1	
					
					Byte 2+ ((n-1)*5)	Data1 of Sensor #n	

					Byte 2+ ((n-1)*5)+1	Data2 of Sensor #n	
					Byte 2+ ((n-1)*5) + 2	Data3 of Sensor #n	
					Byte 2+ ((n-1)*5) + 3	Data4 of Sensor #n	
					Byte 2+ ((n-1)*5) + 4	Data5 of Sensor #n	
Get Array Device Status	0x30	0xC0	None		Byte 0	Completion Code	
					Byte 1	HDD Slot 0	
					Byte 2	Array Device Element Status [0] of HDD 0	
					Byte 3	Array Device Element Status [1] of HDD 0	
					Byte 4	Array Device Element Status [2] of HDD 0	
					Byte 5	Array Device Element Status [3] of HDD 0	
					...		
					Byte N*4 + 1	HDD Slot N (N: Max HDD Slot Number - 1)	
					Byte N*4 + 2	Array Device Element Status [0] of HDD N	
					Byte N*4 + 3	Array Device Element Status [1] of HDD N	
					Byte N*4 + 4	Array Device Element Status [2] of HDD N	

				Byte N*4 + 5	Array Device Element Status [3] of HDD N		
Sled cycle	0x30	0x30	None	None		Control front or rear DPB HSC	
System cycle	0x30	0x31	None	None		Control the CLIP HSC of PTB	
Get UIC location	0x30	0x40	None	Byte 0	Completion Code		
				Byte 1	UIC location (01h: UIC A; 10h: UIC B)		
Storage							
Read FRU Data	0x0A	0x11	Byte 0	FRU ID	Byte 0	Completion Code	
			Byte 1	Offset: Low Byte	Byte 1	Data Length	
			Byte 2	Offset: High Byte	Byte 2	Data [0]	
			Byte 3	Data Length	...		
					Byte (DataLen + 2 -1)	Data [DataLen-1]	

oBMC SUPPORT COMMAND							
Command Name	NetF n	Cod e	Expander Request Data		oBMC Response Data		Description
Sensor							
SetDrawerPullOutPushedInEvent	0x04	0x02	Byte 0	0x04	Byte 0	Completion Code	
			Byte 1	0x05: BMC_SDR_CHASSIS (sensor type)			
			Byte 2	0x00			
			Byte 3	0x6F: Assert 0xEF: De-assert			
			Byte 4	0x00			
			Byte 5	0xFF			
			Byte 6	0xFF			
Storage							
	0x32	0x30	Byte 0	0x00	Byte 0	Completion Code	

OEM Set Platform Event			Byte 1	0x00			
			Byte 2	0x00			
			Byte 3	0x00			
			Byte 4	Event data length			
			Byte 5~N	Event Data			
			Byte 6				

16. EXPANDER FIRMWARE OVERVIEW

The expander has the following functions:

- Sensors readout
- Calculate and report full chassis power
- HDD power control
- HDD LED control
- Generate PWM for fans
- Read and clear PHY error counters
- Read and write FRU IDs
- Report status of GPIOs
- Trigger chassis identify LED as beacon
- Runtime system configuration
- Event report and log
- Staggered spin up of drives

16.1 Drive LEDs

16.1.1 Drive Link LED

The drive Link LED is a blue LED, and is controlled by the SAS expander on the SCC. The LED will be illuminated when the PHY link is established with the HDD. The expander should turn off the LED if the drive is not linked, spun down, etc. In the future, if there is a requirement to blink the LED to indicate activity, the expander can emulate this behavior.

16.1.2 Drive Fault / Identification LED

The drive fault indicator is a yellow LED. This is controlled by the SAS expander on the SCC via I2C GPIO devices located on the DPB. The DPB must implement a circuit to ensure that when the fault LED is illuminated, the drive link/activity LED is not illuminated. The fault LED implements the following patterns:

- Fault bit is set in SES page: on
- Ident bit is set in SES page: blink three seconds on, one second off

To save power, the activity LEDs should only be illuminated when the drawer is pulled out for service. This is possible with the drawer insertion signal that alerts the SCC that the drawer is open.

16.1.3 Drawer Open Thermal Warning Behavior

The Grand Canyon system is required to keep all of the components in the system within the nominal temperatures. Both the expander firmware and BMC need to monitor when the drawer is opened for service, and must provide a visual indication to the user when one or more components in the system is about to hit a thermal trip point. This indication must start a minimum of 30 seconds before the expected trip point is hit.

The behavior should be implemented as shown below in the expander. All the fault LEDs in the first row should illuminate, followed by the second row one second later, and then the third row one second after that continuing until the 6th row is reached. At this point, the behavior should start again from the first step.

If one or more drives has previously been marked faulted (the reason for the drawer was opened) then the fault LED for this drive should be blink at 2Hz continuously during this mode so it is still easily identified.




Figure 16.1: Thermal Race Horse Pattern

Note: The LEDs should not light up when the drawer is closed to save power. The LEDs light up only after the drawer is opened for servicing the system.

16.2 Power Readings

The expander should provide power consumption of the whole chassis (based on current and voltage reading from the appropriate Hot-Swap-Controller). The expander should present two power readings:

- Real Time Power: the expander should read current from HSC at least once every two seconds. The Real-Time Power should be calculated based on the last current reading.
- Average Power: The Average Power reading is based on average current from HSC for a configurable time window (from 1 to 255 seconds).

16.3 Runtime System Configurations

The following parameters of the system must be managed by the expander firmware at runtime:

- Current fan speed profile
- Time window of power calculation
- HDD temperature polling interval

16.4 Report Status of GPIOs

The expander firmware uses a few GPIOs to identify the status of the system. The firmware should also report them to the host through a SCSI buffer.

These GPIOs should include:

- SCC type (JBOD, server, etc.)
- SCC location (front, back)
- IOCM insertion
- BMC heartbeat
- Peer SCC insertion
- Peer expander heartbeat
- Chassis pulled out or not

16.5 Even Report and Log

The expander firmware should provide status report and log for the following events:

- Sensor warnings
- I2C bus errors
- HDD errors
- SAS link errors
- Change of GPIO status
- Insert/remove of HDDs

The expander should keep timestamps for the event log. The expander gets timestamp from either BMC (for server configuration only) or from host utility (SCSI buffer write).

16.6 Security Features

Grand Canyon has secure boot with hardware root of trust on the IOC, Expander, SSDs, HDDs, NICs and BMC. To upgrade firmware, signed firmware will need to be used on the IOC, Expander, SSDs, HDDs, NICs and BMC.

17. THERMAL DESIGN SPECIFICATION/REQUIREMENTS

This section outlines the thermal design requirements for the Grand Canyon platform.

To meet thermal reliability requirements, the cooling solution should dissipate heat from the components when the system is operating up to its maximum thermal power.

The final thermal solution of the system should be most optimized and energy efficient under data center environmental conditions with the lowest capital and operating costs. The thermal solution should not allow any overheating issue for any components in the system.

17.1. Data Center Environmental Conditions

This section outlines data center operational condition requirements.

17.1.1 Location of Data Center/Altitude

Maximum altitude is 6,000 feet above sea level. Any variation of air properties or environmental difference due to the high altitude needs to be deliberated into the thermal design.

17.1.2 Cold-Aisle Temperature

Data centers generally maintain cold aisle temperatures between 18°C and 30°C (65°F to 85°F). The mean temperature in the cold aisle is usually 25°C with 3°C standard deviation. The cold aisle temperature in a data center may fluctuate minutely depending on the outside air temperature. Every component must be cooled and must maintain a temperature below its maximum specification temperature in the cold aisle.

17.1.3 Cold-Aisle Pressurization

Data centers generally maintain cold aisle pressure between 0 inches H₂O and 0.005 inches H₂O. The thermal solution of the system should consider the worst operational pressurization possible, which generally is 0 inches H₂O and 0.005 inches H₂O with a single fan (or rotor) failure.

17.1.4 Relative Humidity

Data centers generally maintain a relative humidity between 20% and 90%.

17.2 System Thermal Operational Conditions

17.2.1 Inlet Temperature

The inlet air temperature will vary. The cooling system should cover inlet temperatures 20°C, 25°C, 30°C, and 35°C. Cooling above 30°C is beyond the Meta operational condition but used during validation to demonstrate the thermal reliability and design margin. Any degraded performance is not allowed over the validation range 0°C-35°C.

17.2.2 Pressurization

Except for the condition when one rotor in a server fan fails, the thermal solution should not consider extra airflow from data center cooling fans. If and only if one rotor in a server fan fails, the negative or positive DC pressurization can be considered in the thermal solution in the hot aisle or the cold aisle, respectively.

17.2.3 Fan Redundancy

The Grand Canyon system should support N+1 fan redundancy. Under single fan failure, system performance should not be impacted within normal operating temperature range.

17.2.4 Service Mode

Grand Canyon shall be designed to provide sufficient cooling to support all the required service modes including the replacement of fan, HDD, and other FRUs.

17.2.5 Delta T

The Delta T is the air temperature difference across the system, or the temperature difference between the outlet air temperature and the inlet air temperature. The Delta T for Type 5 Grand Canyon Rack must be greater than 12.2°C (22°F) when running within the data center operational condition. The Delta T for Type 7 Storage Grand Canyon Rack must be greater than 7.8°C (14°F) when running within the data center operational condition.

17.2.6 System Airflow or Volumetric Flow

The unit of airflow (or volumetric flow) used for this spec is cubic feet per minute (CFM). The CFM can be used to determine the thermal expenditure or to calculate the Delta T of a system in approximation. The thermal expenditure is quantified by the metric CFM/W, which is calculated by the following formula:

$$\text{Thermal Expenditure} = \frac{\text{System airflow}}{\text{Total system power consumption, including fans}} \quad [\text{CFM/W}]$$

At sea level, the maximum allowable airflow per watt in Type 5 Grand Canyon rack is 0.145 at 30°C inlet temperature under the normal load. For the Type 7 Grand Canyon rack, the maximum allowable airflow per watt is 0.228 at 30°C inlet temperature, sea level. At temperatures above 30°C, the CFM/W requirement, although preferred, will not be enforced.

17.2.7 Thermal Margin

All the critical components shall have minimum 7% thermal margin at 30°C inlet temperature or less and 4% thermal margin over 30°C inlet temperature.

17.2.8 Thermal Sensor

All the critical components shall have the thermal sensors directly monitoring their health. The maximum allowable tolerance of thermal sensors is ±2°C, and +/- 1°C is pre.

17.3 Grand Canyon Thermal Design Requirements

17.3.1 Fans

Grand Canyon shall use I²C fans which enables smart fan features. The T5 configuration of Grand Canyon shall use four 92mm dual-rotor counter-rotating fans mounted on the back of the system. The T7 configuration of Grand Canyon shall use four 92mm single rotor fans with less airflow and pressure performance to reduce the airflow and power of the system. Fan blade design should also be taken into consideration to reduce the acoustic vibrations that may impact drive performance. Structural vibration of the fans should be optimized to minimize the impact on drive performance. Under normal operating conditions, fan power consumption should not exceed five percent of the total system power.

17.3.2 Cooling Zones

The Grand Canyon system has two cooling zones to the fans. These are referred to as Zone C for the center compute section, and Zone D for the drive areas. There are two fans connected to each zone. The BMC on the IOM only outputs PWM to Zone C, and the SCC controls both Zones C & D. Fans will take the maximum PWM provided by the controllers.




Figure 17.1: Grand Canyon Cooling Zones

17.3.3 92mm Fan Module

The fan modules are hot-swappable from the rear of the chassis. Each fan module consists of one fan integrated into an easy to service module. Each fan is a dual rotor setup and the fan module as is will be a FRU. In the T7 configuration there may only be a single rotor fan contained in the module. In order to remove the fan, squeeze the two latches on the bottom of

the module inward. Fans shall overhang on rubber material to isolate fan rotational vibrations that affect hard drive performance.

To further isolate the fan module from the system, the interface surfaces with the surface shall have a layer of dampening material. The fan cage should be compliant with all the safety requirements as per UL standards. Fan modules should use louvers as shown in Figure 17.2 to eliminate any potential recirculation during fan servicing and fan failure scenarios. Airflow should control the louvers automatically without any external actuators.



Figure 17.2: Fan louvers



Figure 17.3: Fan Module Front and Back ISO Views




Figure 17.4: Fan Module Dimensions

17.3.4 Air Baffles

One of the most important aspects of the thermal design is to keep the drives cool. To assist this, ducting may be implemented to redirect fresh (cold) air from the front of the chassis to help cool the back three rows of drives. Channels would be cut out of the PCB and an air baffle would be installed in the base of the chassis, as shown in Figure 17.5. Alternatively, perforation on the front panel of zone C could be designed (low percentage opening) to increase the airflow through zone D for drives. However, it should still provide adequate air to keep all the components in zone C within their respective thermal specification.



Figure 17.5: Air Baffles covering Barton Springs and SCCs




Figure 17.6: Air Baffles at the bottom tray for fresh air

For Grand Canyon Type 7 Head Node or Type 7 JBOD, a dummy server may be required to limit the air bypass through the empty slot(s).




Figure 17.7: Dummy Server configurations for Type 7

To make hot servicing of HDDs possible when the drawer is pulled out all the way, there is Mylar added to the bottom of the IOC module tray and transparent Mylar added to the HDD latch to emulate the top and bottom air guides which are normally present when the drawer is fully pushed into the tub.

The transparent Mylar is recommended so that labels can be scanned without removing the drive from the slot. Shown below are renderings of these enhancements to improve the airflow and keep the HDD and the NIC on the IOM well below their max temperature specifications even when the drawer is pulled out.




Figure 17.8: Mylar on drive retention latch and on bottom of IOC Module

18. MECHANICAL

This section outlines the mechanical design requirements of the system.

18.1. Overview



Figure 18.1: System Overview

18.2. Overall Dimensions

The system shall fit into the OCP rack V2.0 without any modifications and fit into the OCP rack V3.0 with minimum mechanical modifications.

The system dimensions are:

- 537mm wide, 189.5mm tall (40U height), 940.91mm deep
- 1782mm deep with drawer fully extended



Figure 18.2: Overall Dimensions

18.3. Rotational Vibration Requirements

In order to reduce rotation vibration from the fan from affecting the HDD normal operation, fan cage and fan tray design shall have:

- **Structural vibration isolation:** Fan is mounted on damping grommets; damping rubber is installed between fan cage, fan tray and chassis (Features shall be included to prevent over-compression of the damping rubber).
- **Acoustic vibration absorption:** Acoustic foams are installed close to the last row HDD to absorb acoustic vibration from fans.
- **Fan vent:** Fan cage vent guard that minimizes turbulence of air flow.
- **Distance of fans from HDD:** Four inches from fan rotor to HDD.

- In order to reduce rotational vibration from adjacent HDDs during normal HDD operation, HDD latch and HDD partition design shall have:
 - **Structural vibration isolation:** HDD is mounted on damping rubber or double-shot damping materials on HDD latch and HDD partition.

18.4 Drive Plane Board

To enable sliding mechanism of DPB serviceability, DPB design shall satisfy the requirements below:

- The Spool shall be mounted on the DPB through SMT or SMT and screw mount, for the vertical constraint of DPB.
- The DPB guide shall be mounted on both sides of the chassis wall to control the DPB sliding path. Guiding features shall be added for both sliding in and out motion to prevent accidental damage of connectors and components on DPB.
- The DPB sag shall be controlled within 0.85mm by adding plastic structural support to improve its stiffness.
- The DPB shall have adequate support structures under high load connectors (Barton Springs and SCC) to minimize PCBA warpage.

18.5 Drawer Structural Walls

To minimize weight requirements, sheet metal in inner structural walls is cored out to reduce weight. Subcomponents that do not contribute to structural rigidity are replaced with thinner sheet metal or plastic materials.




Figure 18.3: Drawer Structural Walls

18.6 Slide Rail

Rail design for the drawer type storage shall satisfy the requirements below:

- **Rail type:** Heavy duty telescopic rail

- **Dimension:** 785mm in length, 1642.6mm in full extension
- **Support:** Rail can support drawer weight of 90kg
- **Overhang weight:** Rail can support additional overhang weight of 147kg at the end of the rail
- **Drawer lock:** Drawer has two independent locks activated by a pair pull handles
- **Vibration and shock:** See vibration and shock section
- Full lock-out and lock-in mechanism to ensure the drawer is properly secured at all times
- Shock absorber inclusion on the rail to ensure HDDs do not experience shock up to 20G during pull-out and push-in




Figure 18.4: Slide Rail Dimensions

18.7 Bus Bar Power Module & Cable Track

Bus bar power module mechanism shall satisfy the requirements below:

- **Safety:** No Medusa cable pinching points allowed. The mechanism shall not open during accidental hand touch, in case of losing power of the whole system.
- **Serviceability:** The bus bar power module shall be serviceable without removing DPB. Guiding features shall be included to ease its serviceability.
- **Airflow:** There shall be venting holes on both sides of the module to allow airflow to go through the middle channel.

The cable track is present to control the 12V input power cables. This shall be a plastic material with the appropriate length to allow the drawer to fully extend. Cable tracks shall rotate freely and be able to fold and unfold themselves without any resistance.




Figure 18.5: Bus Bar Power Module

18.8. Drawer Release Handles

The drawer handles serve the purpose of latching and unlatching the drawer from the system tub/chassis. The handles shall have no sharp edges, shall be light so as to not add to the chassis total weight, and shall be strong enough to withstand the drawer pull out and insertion forces.

These drawer release handles are not intended for lifting the chassis. Potential security buttons shall be added to mitigate the safety risk of drawer unlatching from the system tub/chassis by accidentally lifting the chassis with the handles.

The handles operate by pushing on the green tabs and rotating the handles inward.



Figure 18.6: Drawer Handles

18.9. Rack Release Latching

The rack release latch serves to hold the chassis affixed to the rack. The rack latch shall be designed so that sliding and pushing the chassis all the way into the ORv2 rack clicks it in place and locks the latch without having to press the levers that open the latch. The rack latch shall be capable of holding the entire chassis in place even when the drawer filled with HDDs is pulled out all the way.

The rack latch shall be compatible for both ORv2 and ORv3 rack to minimize mechanical modifications and risk. Although the latch itself shall remain the same between ORv2 and ORV3, its position within the shelf can change depending on the design requirements.

The latch design for Grand Canyon features two pieces: the main latch and an extra spring plate for added retention.




Figure 18.7: Rack Release Latch



Figure 18.8: Rack Release Latch Operation

18.10 Chassis Lift Handles

The chassis features a handle that is able to withstand a minimum of 50kg each. The handle shall be folded away all the time to ensure the system slides freely into the rack. The handle shall be rubberized and green in color (Pantone 375C) for easy handling.



Figure 18.9: Chassis Lift Handles

18.11 Drive Retention

Drive retention shall be designed around the form factor outlined in SFF-8301 specification, rather than a specific drive model. The latching mechanism must support all SFF compliant 3.5" HDDs adhering to the 26.1mm max height.




Figure 18.10: SFF-8301 Standard

The hard drive latch shall be designed in such a way as to allow for the barcode of an installed hard drive to be scanned without unseating the drive, thereby providing a means to verify the identity of a drive prior to removing it for service.

In addition, the latching feature must be able to lift the drive a minimum of 20mm above the adjacent drive slots and/or any other obstructions to make the removal and insertion of the drive easier. Gaps between two adjacent hard drive latches and gaps between hard drive latches and partitions shall be minimized to prevent air flow leakage.



Figure 18.11: Drive Retention Latch




Figure 18.12: System Cross-Section Showing HDDs & Latches

18.12 UIC and IOC Module

The UIC and IOC have a pinch release lever for the tray lock latch. It shall be made with green plastic for easy identification. Releasing the lever and opening the latch to full open position should almost eject the UIC and IOC PCBA from the chassis connector, so that it is easy to pull them out of the chassis.

While installing the UIC and IOC back in, the card guides should be used to slide the PCBA in and the force from the latch should be used to secure the PCBA to the DPB connector.

No extra force on the UIC and IOC sheet metal front panel is necessary for installation. See Section 8.4 for installation/removal instructions.




Figure 18.13: IOC Module and UIC

18.13 Storage Controller Card

The SCC PCBA is fitted with one release/install latch. The latch locks in when the card is fully installed and allows for some overtravel to ensure proper wipe. The latch is released using the green plastic pinch release levers. The SCC connectors mating into the DPB are wide and the insertion force is high.

The SCC latches should be able to support the insertion force without deforming or putting any load on the PCBA. There should also be adequate support under the DPB to minimize deformation during insertion.




Figure 18.14: SCC Latch Operation

19. ENVIRONMENTAL AND REGULATIONS

Responsible for successful completion and demonstration of compliance to the FCC Authorization Program Requirements, EU Directive requirements, and NRTL Product Safety requirements specified below for L10 product.

19.1 Regulatory Compliance

- FCC Compliance: The product shall comply with FCC subpart 15 b class A. Characterization documentation and test data shall be provided.
- CE Compliance: The Product shall comply with the EU Directives noted below. Meta shall be provided with a CE Declaration of Conformity and complete technical reports and other documentation files for:
 1. The Low Voltage Directive; 2014/35/EU
 2. The EMC Directive; 2014/30/EU
 3. RoHS 2 Directive; 2011/65/EU, unless there are legal exemptions allowed
 4. REACH Regulation; 1907/2006

19.2 Product Safety Compliance

NRTL Certificate and complete technical report(s) per EN 60950-1, UL 60950-1, CAN/CSA-C22.2 No. 60950-1; Safety of Information Technology Equipment, General Requirements per Edition 2 and Amendment 2.

20. ENVIRONMENTAL REQUIREMENTS

Grand Canyon must meet the following environmental requirements:

- Gaseous Contamination: Severity Level G1 per ANSI/ISA 71.04-1985
- Ambient operating temperature range: 0°C to +35°C
- Operating and storage relative humidity: 10% to 90% (non-condensing)
- Storage temperature range: -40°C to +70°C (long-term storage)
- Transportation temperature range: -55°C to +85°C (short-term storage)
- Operating altitude with no de-rating to 1000m (6000 feet)

20.1. Vibration & Shock

Grand Canyon must meet shock and vibration requirements according to the following IEC specifications: IEC78-2-(*) and IEC721-3-(*) Standard and Levels. The testing requirements are summarized in the table below.

	OPERATING	NON-OPERATING
Vibration	0.4g acceleration, 5 to 500 Hz, 10 sweeps at 1 octave / minute per each of the three axes (one sweep is 5 to 500 to 5 Hz)	1g acceleration, 5 to 500 Hz, 10 sweeps at 1 octave / minute per each of the three axes (one sweep is 5 to 500 to 5 Hz)
Shock	6g, half-sine 11mS, 5 shocks per each of the three axes	12g, half-sine 11mS, 6 shocks per each of the three axes

20.2 Regulations

20.3 Chassis Labels and Markings

The chassis shall carry the following adhesive bar coded labels in locations that can easily be scanned during integration:

- Vendor P/N, S/N REV (revision would increment for any approved changes)
- Meta P/N (or OCP customer P/N)
- Date Code (industry standard: WEEK/YEAR)
- The assembly shall be marked “THIS SIDE UP,” “TOP SIDE,” “UP,” or other approved marking in bright, large characters in a color to be defined by ODM and Meta (or OCP customer). This printing may be on the PCB itself, or on an installed component such as an air baffle. The label should be clear and easy to read in low light conditions, when viewed from above or below from two feet away, and at an angle of approximately 60 degrees horizontal.

Exact label locations will be agreed upon between the vendor and Meta (or OCP customer).

20.4 Data Center Design Requirements

- Seismic tie downs – GR63
- All casters will swivel
- Push/tilt/angle ramp testing – see OCP spec at www.opencompute.org.
- Rack shall have an entry and exit angle of 15 degrees minimum.
- Rack force (kgw) required to push the rack from a non-moving position along a smooth, flat cement floor shall be less than 5% the total combined weight (kgw) of the rack and IT Gear (note: this is a goal, not a hard requirement).
- The following tests will be performed with the rack loaded at the maximum rated IT load.
 - Roll over a 6mm vertical step with each caster independently at 0.5m/s
 - Transitioning a 1” wide gap in the floor while fully loaded at 0.5m/s
 - Transition a 5-degree ramp
 - Roll a minimum of 800m on a concrete floor at 0.8m/s
- Transportation shock/vibration testing per ASTM D4169-09, Assurance Level II, Mechanized handling, truck, no stacking, distribution cycle 2. Alternative testing may be acceptable, with Meta approval.
- Seismic testing: NEBS GR63 zone 2 with the rack leveling feet lowered. Racks may be ganged together prior to test. NEBS GR63 zone 4 with the rack leveling feet lowered. Racks may be ganged together prior to test.
- Leveling feet requirements: The rack will provide four leveling feet to remove the weight off of the casters once deployed in the data center. The feet shall:
 - Be delivered from the rack supplier in the raised condition and stay in that condition under ASTM D4169-09.
 - Be capable of raising and lowering the feet to their maximum height using only an electric driver from the front and/or rear (side access not allowed).

- Be capable of cycling three times up and down under maximum cabinet load without damage to the rack or the leveler.
- Be capable of supporting the rack under maximum weight under the GR63 zone 4.
- Have a swivel base to prevent the rack from “walking” when the feet are deployed.
- Be capable of raising the rack a minimum of 35mm off of the floor.
- Cabinet life: Rack should be designed for an expected life of 10 years under the following environmental conditions:
 - Humidity: 85% max, 42°F dew point minimum
- Cabling and grounding: Cables should be retained out of the path of the FRUs so service events will not damage cables during installation and removal.
- Cables can be added and removed from the cable retention without tools.
- No sharp edges, and burrs must be removed around cable routing areas to prevent damage.

20.5 Mean Time Between Failure (MTBF) Requirements

15.5.1 MTBF Prediction

The board design shall have a minimum calculated MTBF of 300K hours at 95% confidence level at typical usage conditions (25°C ambient temperature).

- Recommended software is Relex
- For non-critical parts, use standard / default number in Relex
- For high failure rate parts, get actual data from each vendor

20.5.2 MTBF Demonstration Test

The board design shall also demonstrate the MTBF requirement above by running test as below:

- At full load for 50% of time, and performing AC cycling test 50% of time
- At 45°C ambient temperature for acceleration
- Total test duration needs to be calculated, discussed with and approved by Meta

20.5 Certifications

The Vendor needs to provide CB reports of the Grand Canyon system at the component level.

21. PRESCRIBED MATERIALS

21.1 Disallowed Components

The following components are not to be used in the design of the system:

- Components disallowed by the European Union's Restriction of Hazardous Substances directive (RoHS 6)
- Trimmers and/or potentiometers
- Dip switches

21.2 Capacitors & Inductors

The following limitations apply to the use of capacitors:

- Only aluminum organic polymer capacitors made by high quality manufacturers are used; they must be rated 105°C.
- All capacitors have a predicted life of at least 50,000 hours at 45°C inlet air temperature, under the worst conditions.
- Tantalum capacitors using manganese dioxide cathodes are forbidden.
- SMT ceramic capacitors with case size >1206 are forbidden (size 1206 are still allowed when installed far from the PCB edge and with a correct orientation that^[1]_{SEP} minimizes risk of cracks).
- Ceramic material for SMT capacitors must be X7R or better material (COG or NP0 type are used in critical portions of the design).
- Conditional usage of X5R ceramic material must be based on evaluation of worst-case thermal conditions and upon approval from Meta.
- Only SMT inductors may be used. The use of through hole inductors is disallowed.

21.3 Component Derating

For inductors, capacitors, and FETs, de-rating analysis is based on at least 20% derating.

21.4 Sustainable Materials

Materials and finishes that reduce the life cycle impact of servers should be used where cost and performance are not compromised. This includes the use of non-hexavalent metal finishes, recycled and recyclable base materials and materials made from renewable resources, with associated material certifications.

Meta identified plastic alternatives including polypropylene plus natural fiber (PP+NF) compounds that meet functionality requirements while reducing cradle-to-gate environmental impact when compared to PC/ABS. GreenGranF023T is one acceptable alternate material. JPSECO also offers a PP+NF material that is acceptable; the model number will be available at a later date. It is strongly preferred that such alternatives are identified and used. If a vendor is unable to use this, or a similar alternate material, the vendor will provide a list of materials that were considered and why they were not successfully incorporated.

The supplier shall use Halogen Free (IEC 61249-2-21), arsenic-free (less than 1000 ppm, or 0.1% by weight) and phthalate-free (less than 1000 ppm, or 0.1% by weight) material by default and discuss with Meta and get written approval if there is difficulty with sourcing in L10.

22. SYSTEM FIRMWARE

All Server products must have a completed OSF Tab in the [2021 Supplier Requirements](#). At this time, this is not required for any other product types. Please check with the OSF leadership for updates.

The completed checklist shall be uploaded and available at:

[https://github.com/opencomputeproject/OpenSystemFirmware/\[vendor_name\]/\[product_name\]/](https://github.com/opencomputeproject/OpenSystemFirmware/[vendor_name]/[product_name]/)

23. HARDWARE MANAGEMENT

23.1 Redfish Compliance

All Products shall have a completed Hardware Management Tab in the [2021 Supplier Requirements](#).

23.2 Source Availability

All Products shall have a completed BMC Tab in the [2021 Supplier Requirements](#) if applicable.

The BMC management source code is available at:

<https://github.com/facebook/openbmc>

The quick build instruction as follows:

```
$ git clone -b helium https://github.com/facebook/openbmc.git
$ cd openbmc
$ ./sync_yocto.sh
```

```
$ source openbmc-init-build-env grandcanyon
```

```
$ bitbake grandcanyon-image
```

The build process automatically fetches all necessary packages and builds the complete image. The final OpenBMC image is in

`openbmc/build/tmp/deploy/images/grandcanyon/flash-grandcanyon.`

The root password will be `OpenBmc`

24. SECURITY

All products shall have a completed Security Tab in the [2021 Supplier Requirements](#).

25. REFERENCES (OPTIONAL)

[1] “Title”, publication year, publication journal/conference/standard, volume, pages, link to publication if available

APPENDIX A - REQUIREMENTS FOR IC APPROVAL (to be completed Contributor of Baseline Spec)

List all the requirements in one summary table with links from the sections.

Requirements	Details	Link to which Section in Spec
Contribution License Agreement	OCP CLA	Grand Canyon Storage System Specification - v1.0
Tenets	Openness Efficiency Impact Scale	Grand Canyon Storage System Specification - v1.0
Supplier available within 120 days	Wiwynn	Grand Canyon Storage System Specification - v1.0
Will they apply for OCP product recognition?	OCP Accepted™	Grand Canyon Storage System Specification - v1.0