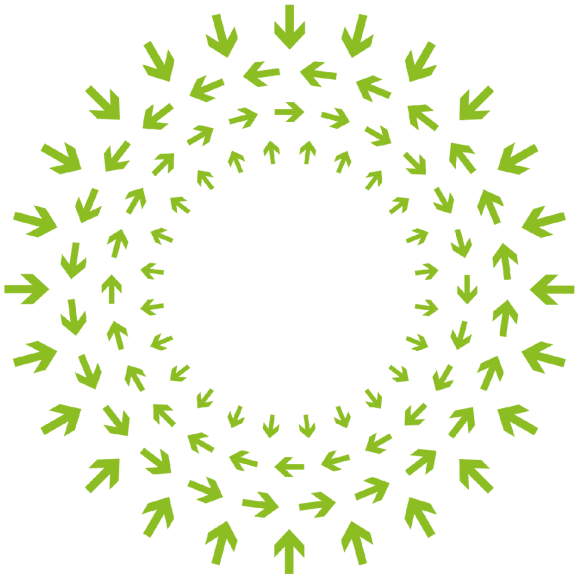


OPEN

Compute Project

Clemente Compute Tray Specification

Version 1.01

Effective September 4th, 2025

Authors:
Matt Bowman, Hardware Engineer, Meta
Hao Shen, Hardware Engineer, Meta
Anthony Chan, Power Engineer, Meta
Devika Vishwanath, Power Engineer, Meta
Enrique Beuses, Mechanical Engineer, Meta
Jaret Wyatt, Mechanical Engineer, Meta
Cheng Chen, Thermal Engineer, Meta
Wenying Zhang, Thermal Engineer, Meta
Tanmoy Roy, Signal Integrity Engineer, Meta
Jason Liang, Hardware System Engineer, Meta

Sathish Sekar, Hardware System Engineer, Meta
Tim Chu, Hardware System Engineer, Meta
Henry Chen, Hardware System Engineer, Meta
Bob Wang, Hardware System Engineer, Meta
Maythe Baldwin, System Integration Engineer, Meta
Freddy Mercado, Compliance Engineer, Meta
Jaswanth Vutukury, Compliance Engineer, Meta
Siu Lam, Reliability Engineer, Meta
Patrick Williams, Software Engineer, Meta
Kui Ren, Manufacturing Quality Engineer, Meta
Shailja Raval, Component Quality Engineer, Meta
Amber Lu, Technical Program Manager, Meta
William Arnold, Technical Program Manager, Meta
Jenny Yu, Technical Program Manager, Meta
Kevin Lee, Technical Program Manager, Meta

Table of Contents

1. License 5
2. Compliance with OCP Tenets 6
    2.1 Openness 6
    2.2 Efficiency 6
    2.3 Impact 6
    2.4 Scale 6
    2.5 Sustainability 6
3. Change Log 6
4. Scope 7
5. Overview 8
    5.1 Description 9
6. Environmental and Regulatory Compliance 9
    6.1 Environmental 9
    6.2 Regulatory Compliance 9
7. Physical Specifications 10
    7.1 Platform Block Diagram 10
    7.2 Power Distribution Board 11
    7.5 DC-SCM 2.0 14
    7.6 E1.S Backplane 15
    7.7 OSFP Carrier Board 17
    7.7 Front IO Board 19
    7.8 Fans and Cold Plate Loop 19
    7.12 Mechanical 20
    7.13 Rack Compatibility 21
8. Electrical Requirements 21
    8.1 Input Voltage 22
    8.2 Power 22
9. Thermal Design Requirements 22
    9.1 Thermal Architecture 22
    9.2 Air Cooling Specification 22
    9.3 Fan Zone 23
    9.4 Fan Control Target 23
    9.5 Thermal Sensors 23
    9.6 Fan Control Algorithm 23

9.7 Special Fan Conditions 24
9.9 Liquid Boundary Conditions 24
9.10 Connection Interface 24
9.11 Liquid Flow Distribution 24
9.12 Leakage Handling 25
10. Interfaces 25
10.1 Signal List 25
10.2 Cables 25
11. System Firmware 26
12. Hardware Management 26
12.1 Out-Of-Band Management 26
12.2 BMC 26
12.3 Remote Upgradeability 26
12.4 Sensors 27
13. Security 32
Appendix A - Checklist for IC approval of this Specification (to be completed by contributor(s) of this Spec) 32

1. License

Contributions to this Specification are made under the terms and conditions set forth in Modified OWF-CLA-1.0.2 (As of June 1, 2023) ("Contribution License") by:

Meta Platforms, Inc.

Usage of this Specification is governed by the terms and conditions set forth in **Modified OWFa1.0.2 Final Specification Agreement (FSA) (As of June 1, 2023) ("Specification License").**

You can review the applicable Specification License(s) referenced above by the contributors to this Specification on the OCP website at http://www.opencompute.org/participate/legal-documents/. For actual executed copies of either agreement, please contact OCP directly.

Notes:

1) The above license does not apply to the Appendix or Appendices. The information in the Appendix or Appendices is for reference only and non-normative in nature.

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. Compliance with OCP Tenets

2.1 Openness

The Clemente Compute tray demonstrates openness with its design collaboration with NVIDIA's MGX. It also incorporates many other OCP contributions including the DC-SCM, OCP NIC, OCP Datacenter NVMe™ SSD Specification, and connectivity to the Rack Management Controller.

2.2 Efficiency

The heatsink for the PDB has been redesigned to greatly reduce the required fan speed. This has a large impact on the overall power draw of the compute tray.

2.3 Impact

Clemente provides a platform that enables quick time to market while reducing risk by leveraging a wide array of validation data.

2.4 Scale

The Clemente Compute tray includes high reliability features around the operation of the cooling, leak sensing, and out of band management that make it possible to scale out many trays to form large GPU clusters.

2.5 Sustainability

Clemente includes numerous commodities and contributions within OCP enabling circulatory and reuse among platforms.

3. Change Log

Date	Version #	Author	Description

13Aug2025	0.5	Matt Bowman Et al.	Initial import into OCP template from internal specification
1Sep2025	1.0	Matt Bowman Et al.	Numerous grammatical and content fixes

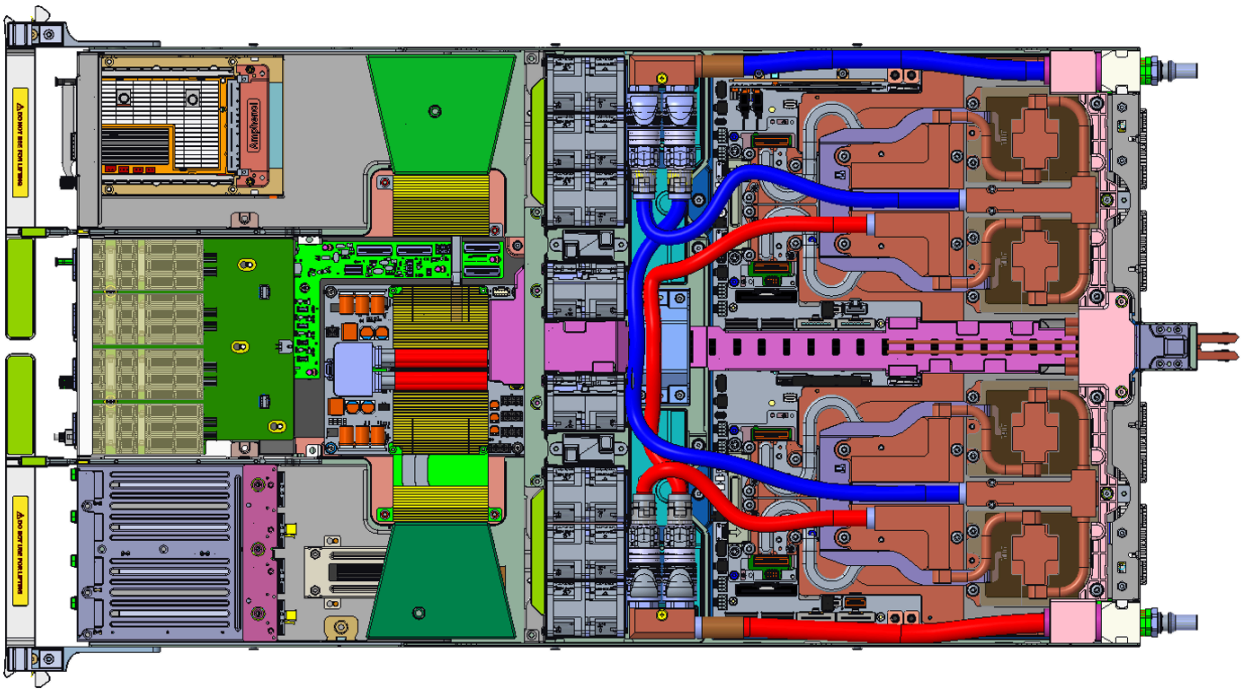
4. Scope

This specification describes the design of the Clemente compute tray which integrates NVIDIA's® GB300 Host Processor Module (HPM) into a form factor with peripherals that supports Meta's AI/ML training use cases.

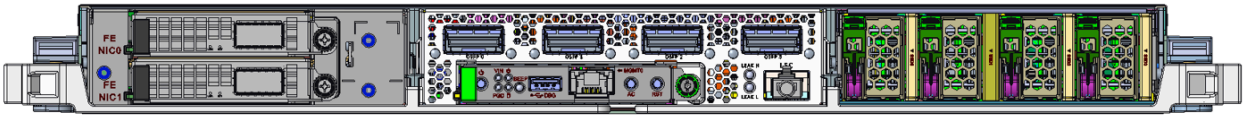
5. Overview

The Clemente Compute Tray is 1OU tall and is made up of two NVIDIA® GB300 HPM modules that are cabled together to operate with a unified memory space. Within the tray is also the power conversion, fans, network, and storage devices to enable the node to have all the necessary resources for performant and reliable operation.

This specification covers the details of the Clemente compute tray, but does not go into the design and specifics of the NVIDIA® GB300 module.



Overall view of the Clemente Compute Tray



Front view of the Clemente Compute Tray

5.1 Description

Clemente is Meta's next generation AI/ML rack that supports large cluster training and inference use cases. Multiple compute trays can be combined in a rack to create a large domain of GPUs to support our ever increasing workloads.

This design focuses on achieving a fast time to market, alignment with industry references, providing cutting edge performance, and seamlessly integrating into Meta's data center environments.

It also represents a milestone of our first deployment of a design that uses ORv3 HPRv3 side power racks.

6. Environmental and Regulatory Compliance

6.1 Environmental

- Gaseous Contamination: Severity Level G1 per ANSI/ISA 71.04-1985
- Ambient operating temperature range: +5°C to +35°C
- Operating and Storage relative humidity: 10% to 90% (non-condensing)
- Storage temperature range: -40°C to +70°C
- Transportation temperature range: -55°C to +85°C (short-term storage)
- Operating altitude with no de-ratings: 6000 feet

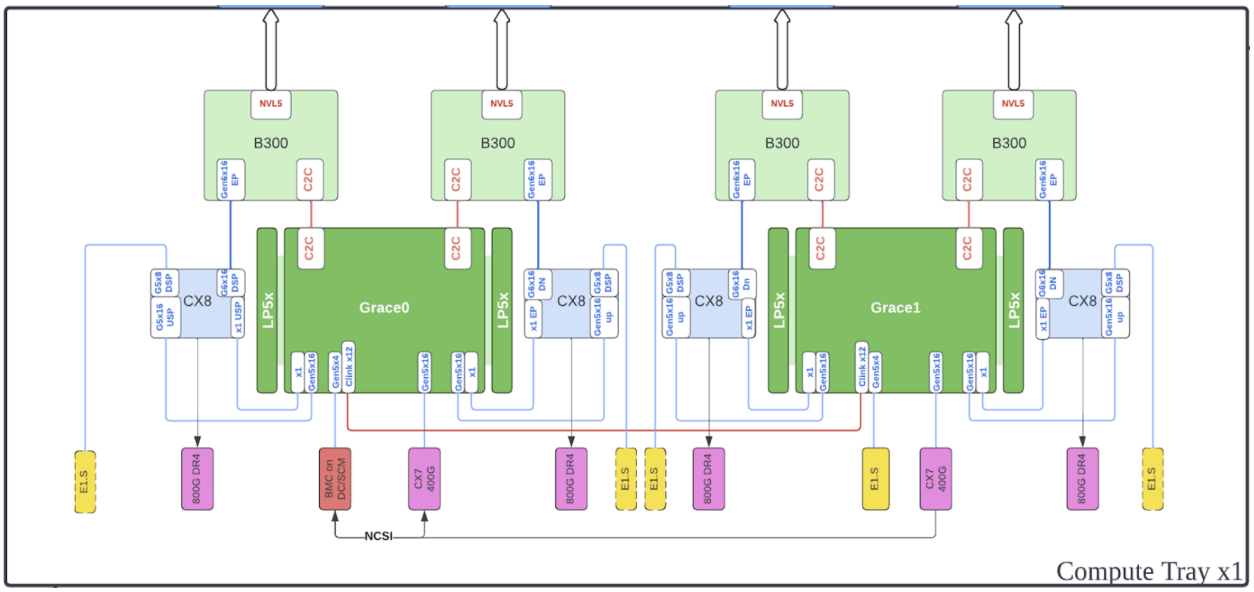
6.2 Regulatory Compliance

- CE and NRTL (UL listed) marks in accordance with the listing and surveillance requirements of the NRTL
- EU Low Voltage Directive (2014/35/EC)
- EU EMC Directive (2014/30/EC)
- EN55032/CISPR32, Class A criteria
- EN55035/CISPR35
- CISPR16
- ANSI C63.4
- EN61000-4-2
- EN61000-4-3
- EN61000-4-4
- EN61000-4-5
- EN61000-4-6
- EN61000-4-11

- FCC CFR47 Part 15, Subpart B, Class A criteria
- RoHS Directive (2015/863/EU)
- WEEE Directive (2012/19/EU)
- REACH regulation (EC 1097/2006)
- US EPA 40 CFR751
- IEC/EN 62368-1, hazard-based performance standard for audio video, and ITE

7. Physical Specifications

7.1 Platform Block Diagram



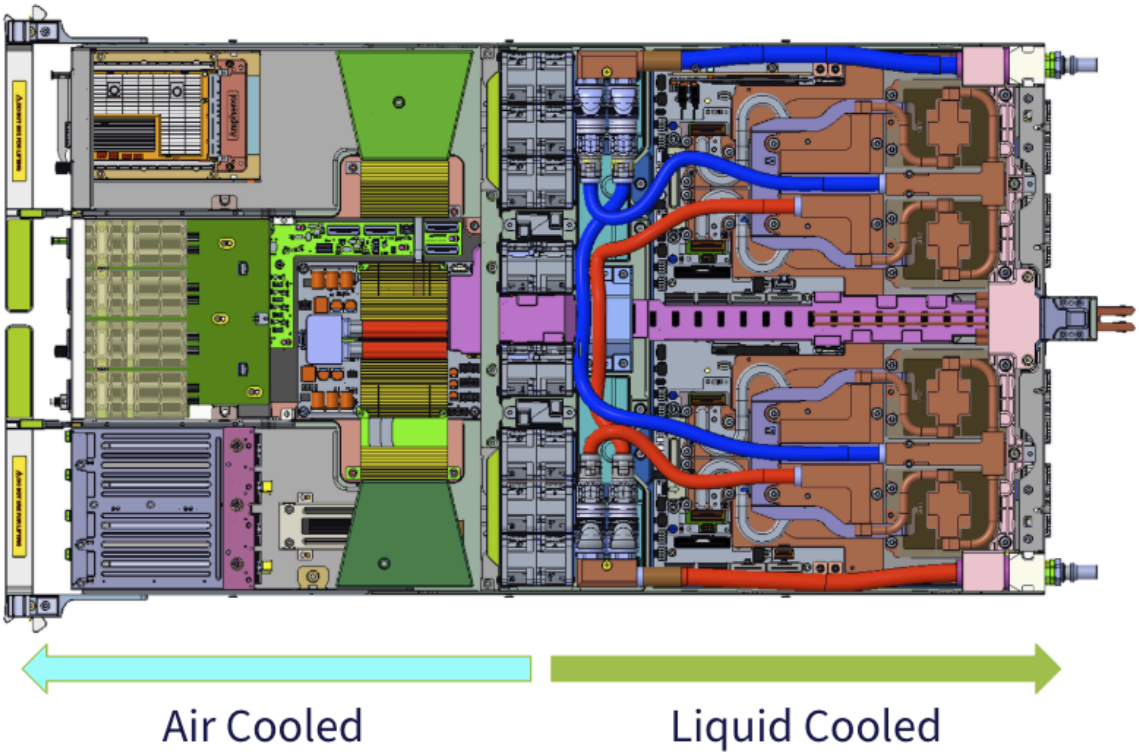
Block Diagram of the Clemente Compute Tray

Multiple components are integrated within the compute tray to provide all of the functionality. These include:

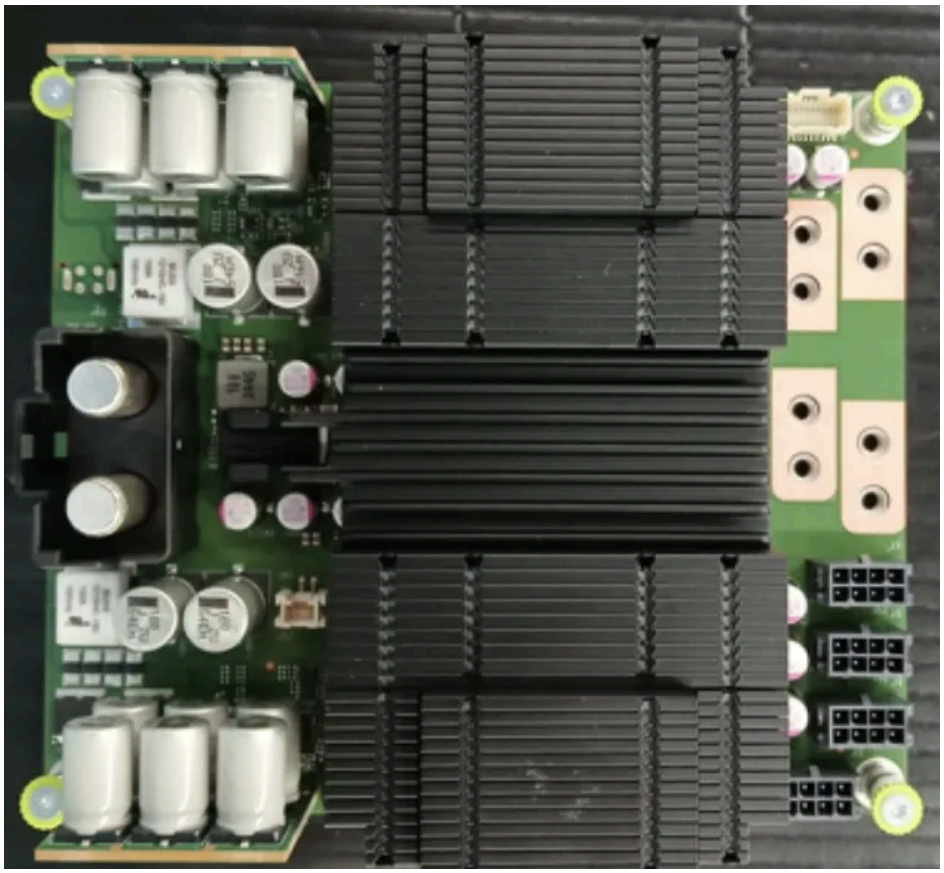
- GB300 Host Processor Module (NVIDIA® designed)
- Host Management Controller or HMC (NVIDIA® designed)
- CX8 Network Module for the backend network (NVIDIA® designed)
- Power Distribution Board or PDB (NVIDIA® designed)
- Interposer (Meta designed)
- DC-SCM 2.0 (Meta designed)
- E1.S NVMe™ Backplane (NVIDIA® designed)
- OSFP Carrier Board (NVIDIA® designed)
- Front IO Board (Meta designed)

- CX7 OCP NIC 3.0 for the front end NICs (purchased commodity)

Fans within the tray are used to cool components at the front of the tray while the high powered components utilize liquid cooling at the rear.



7.2 Power Distribution Board

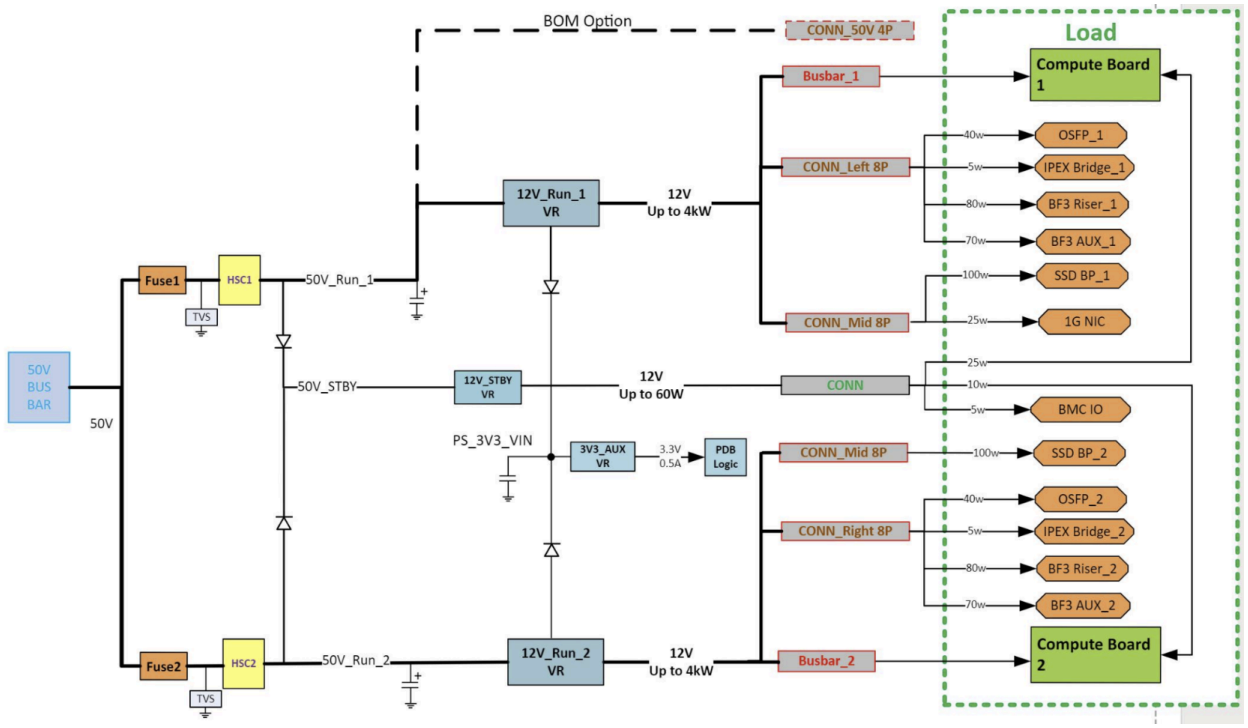


Top view of the PDB

The Power Distribution Board (PDB) has a few important functions in the Clemente Compute Tray:

- Convert the 48V ORv3 bus bar power to 12V, and distribute it to the other components in the system
- Provide hot swap functionality for the compute tray to be hot installed/removed from the rack.
- The compute tray fans get their power and control signals from the PDB
- The management and control CPLD for the compute tray resides on the PDB to utilize its broad connectivity to the other components.

The overall power delivery design of the compute tray is shown below:



Clemente Compute Tray Power Delivery Block Diagram

Key settings for the hot swap controllers:

- Over Voltage Protection (OVP) set to 60.4V-rising/59.3V-falling
- Under Voltage Protection (UVP) set to 46.1V-rising/39.6V-falling
- Over Current Protection (OCP) set to 120A

The firmware implementation provides two different types of power cycling options for the compute tray.

- **“DC Cycle”** instructs the primary 48V -> 12V converters to toggle. This cycles power to the GB300 modules, E1.S drives, and back end NICs. However, the management 12V “P12V_STBY” remains on so the DC-SCM and front end NICs remain powered and accessible.

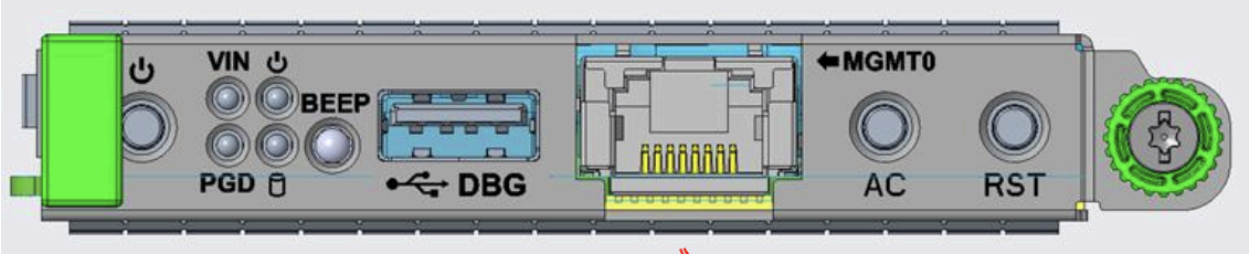
- “**AC Cycle”** instructs all of the 48V -> 12V to toggle. This results in everything including the management devices (e.g. BMC, OOB path, etc) to power cycle.

7.5 DC-SCM 2.0

The DC-SCM follows the DC-SCM v2.0 specification, with pinout exceptions permitted to include features not covered in the base specification.

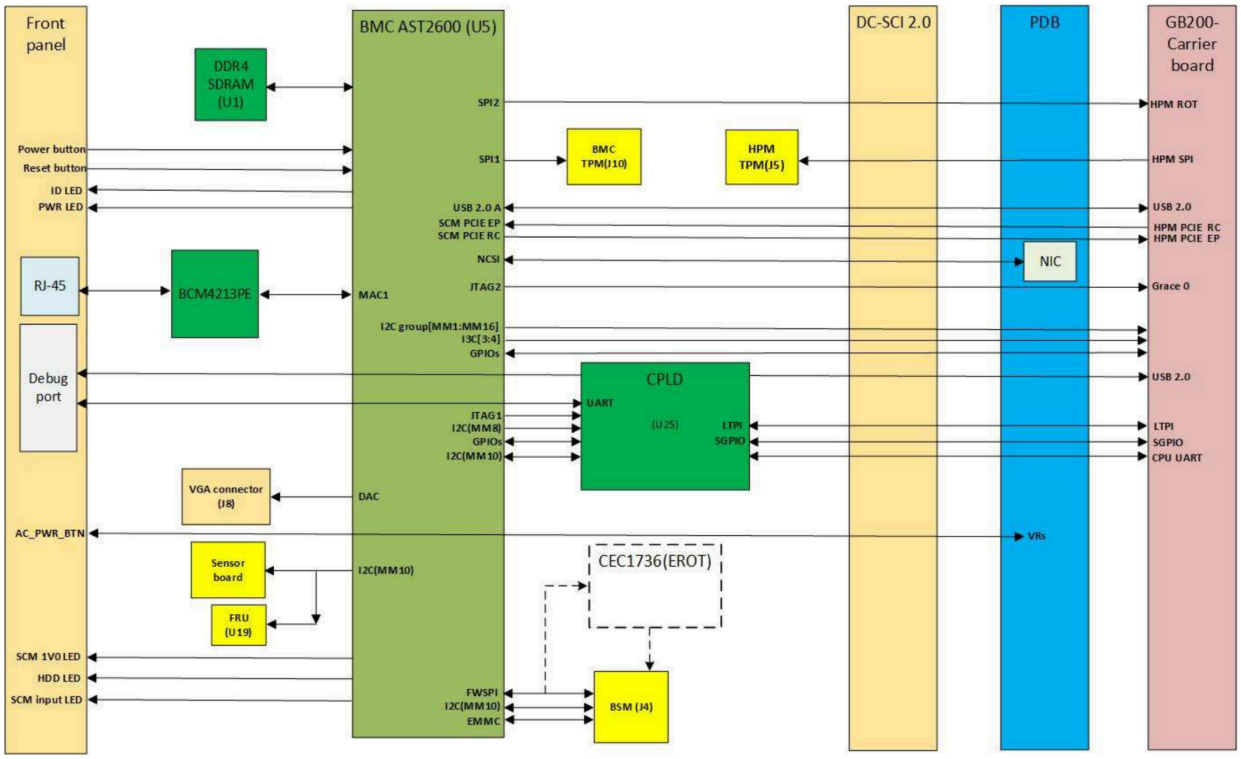
The system uses an ASPEED AST2600 BMC for baseboard management with the following connectivity:

- Single x16 8Gb DDR4 DRAM memory interface
- NCSI/RMII network interface on NIC0 to allow out-of-band management
- I2C management of all system sensors for power, thermal and environment
- I2C and USB connectivity to the HMC
- JTAG to CPU subsystem (with physical access to system)
- RJ45 for management network access



Front View of the DC-SCM

The BMC operates independently of the host system and any other programmable devices within the system with the exception of NIC0 required for network access. The BMC is capable of recovering any firmware in the system regardless of the state of the host system.
The SCM contains removable TPM2.0 for both the HPM and the BMC through SPI interface, a SPI flash for BIOS, and a removable board storage module (BSM) containing all BMC non-volatile storage.



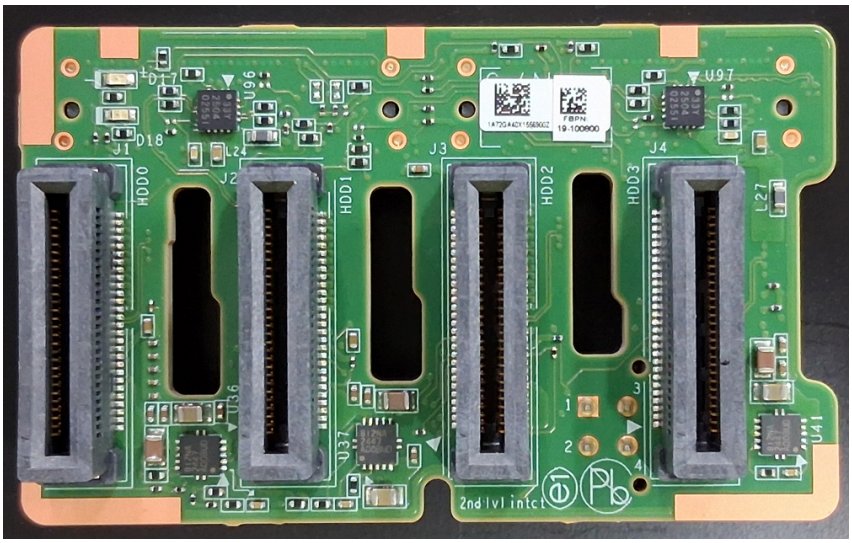
Block Diagram of the DC-SCM

The DC-SCM follows the DC-SCM v2.0 specification. The front panel of it has the following:

- Power-on pushbutton
- Power LED
- HDD LED
- Fault/ID LED
- USB 3.0 OCP debug port
- SCM voltage status LED
- SCM power good status LED
- AC power hard-cycle pushbutton
- Reset pushbutton
- RJ45 Management Port

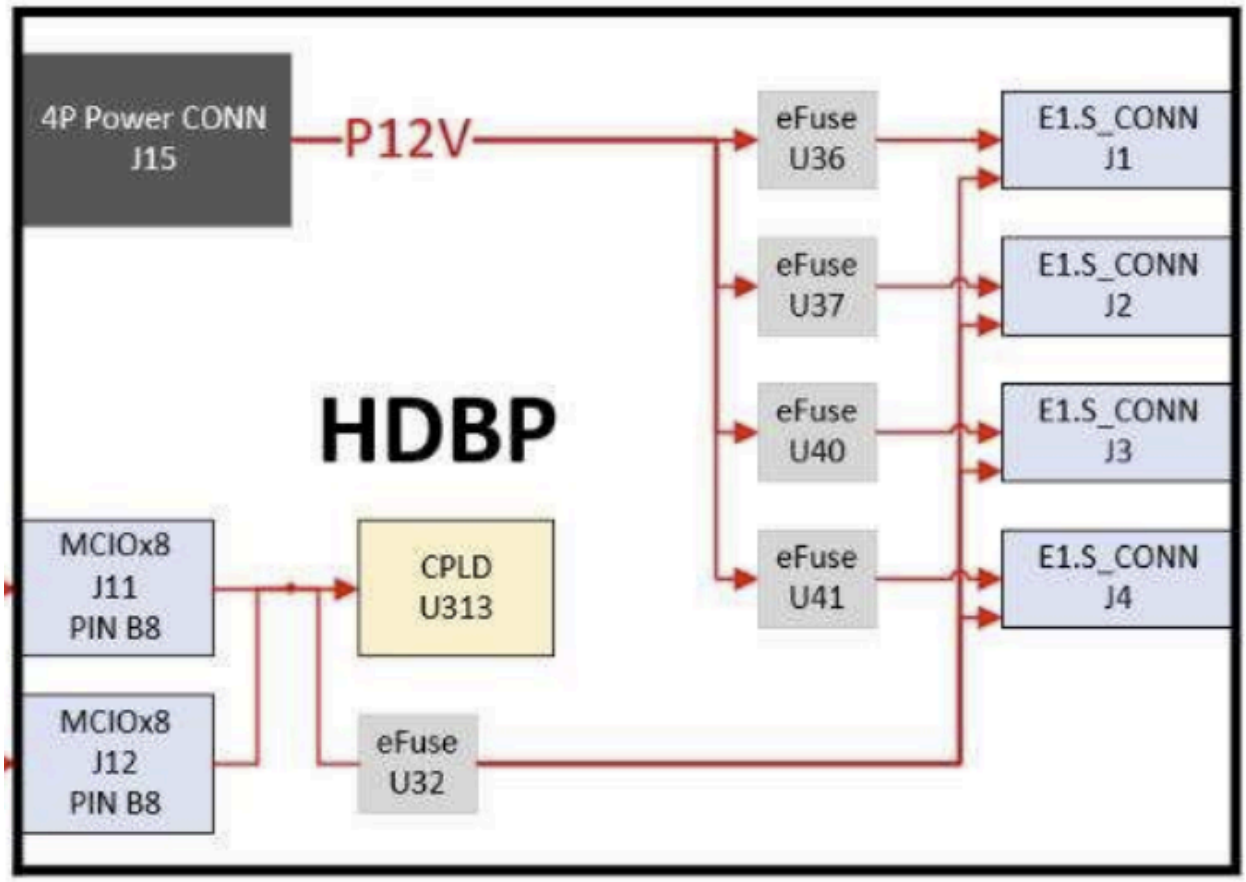
7.6 E1.S Backplane

The E1.S backplane is used to provide direct attached storage for the compute tray's boot and data drives. All power, management, and high speed routing for four 25mm drives are implemented on the E1.S backplane.



Overall view of the E1.S Backplane

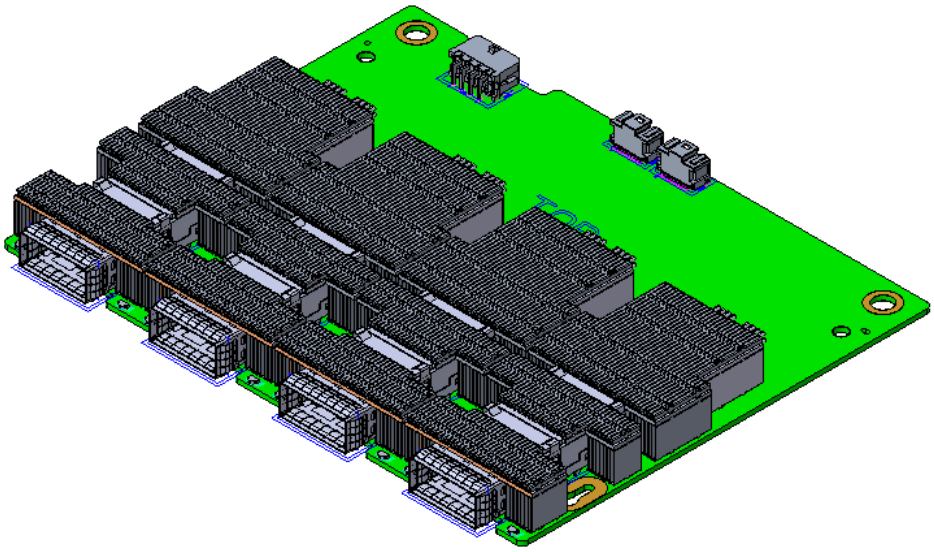
The E1.S drives are aircooled using the tray fans so the backplane includes PCB cutouts to facilitate airflow. Electrical hot add/remove is implemented enabling drives to be serviced on a live system. A CPLD is included on the backplane to monitor and act upon the drive management signals. Each drive can operate up to a 25W TDP.



High-level block diagram of the E1.S Backplane

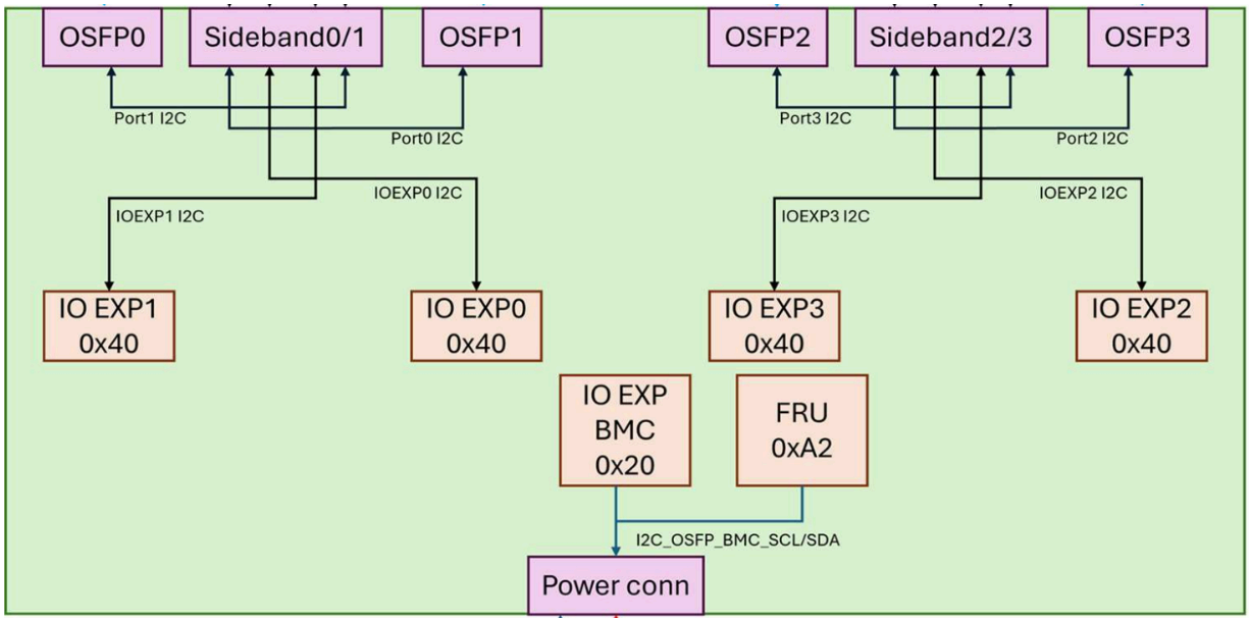
7.7 OSFP Carrier Board

The OSFP board is used to provide the backend network connections to the CX8 NIC Modules connected to the GB300 HPMs.



OSFP Carrier

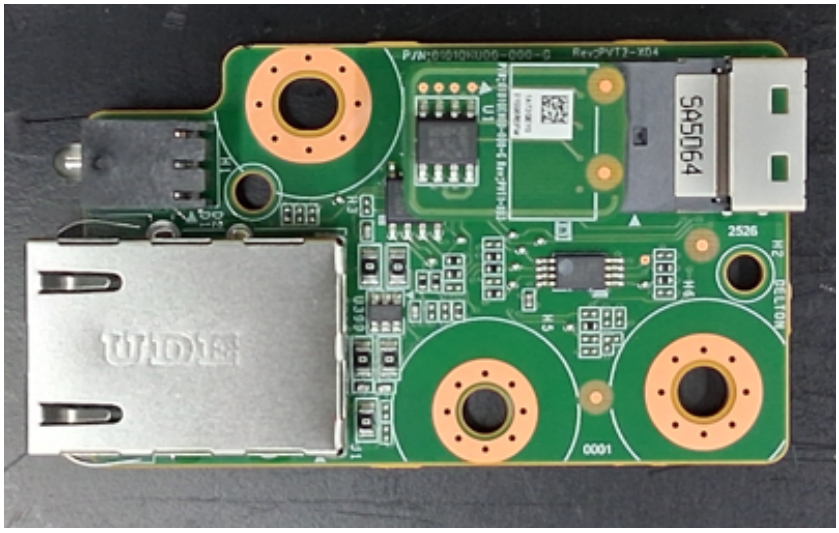
Fly over cables are used to connect the high speed routing directly from the CX8 NIC modules to the OSFP ports. No high speed routing is directly on the PCB which simplifies the design. IO expanders are controlled by the management circuitry to discover and configure the board.



OSFP Carrier Block Diagram

7.7 Front IO Board

The leak sensing and control (LSC) connection from the compute tray to the rack management controller (RMC) is done with an RJ45 cable connected to the front IO (FIO) board. The pinout is custom, and includes power management and I2C communication. For details of the LSC and RMC please refer to their OCP specification “Rack Management Controller v1 Specification”.



The pinout of the RJ45 was done to ensure connecting a standard ethernet cable would not result in any damage.

7.8 Fans and Cold Plate Loop

Within the compute tray there are eight fans. These fans are:

- 46-52V operating voltage
- Dual rotor design
- Size 4056
- Have a common PWM signal
- Independent TACH for monitoring each fan rotor

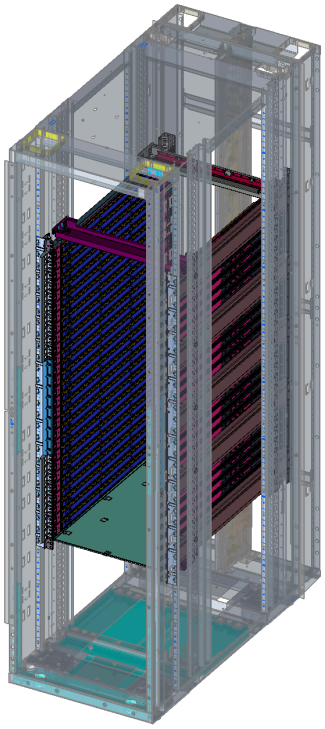
These fans provide air cooling for the E1.S drives, front end CX7 OCP NICs, the DC-SCM, and the power conversion circuitry on the PDB. The thermal design is resilient, and can continue to operate with a single fan rotor failure.

The cold plate loop is used for liquid cooling of the GB300 and CX7 backend NIC modules. The compute trays are validated to work with air assisted liquid cooling with PG25, and facility water.

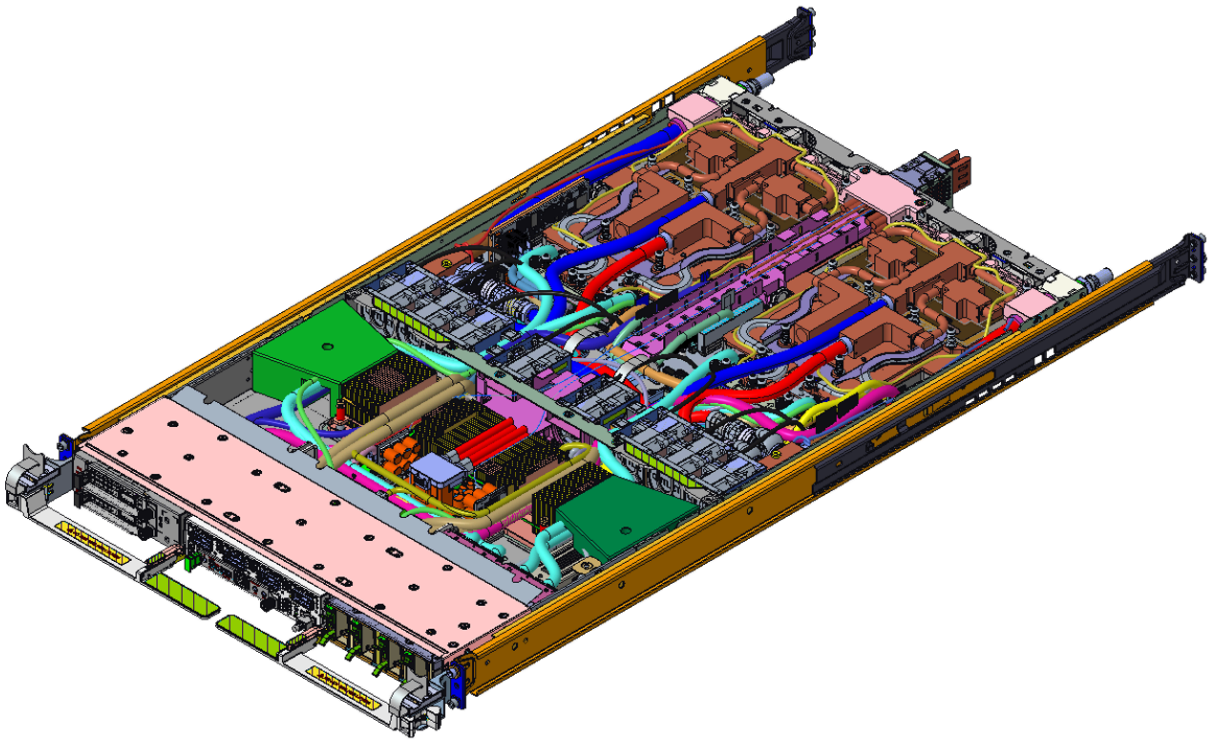
Overall view of the Compute Tray fans and cold plate loop

7.12 Mechanical

The Clemente compute tray is 1RU tall, and fits into a standard 19” rack width. A canister is used to adapt the RU height and 19” width dimensions to be compatible with a standard Open Rack V3 HPR frame.



Canister Assembly to adapt 19” RU into 21” OU



- Dimensions: 43.6 mm (H) x 498 mm (W) x 766 mm (D)
- Weight: Approximately 29.6 kg / 65.3 lb

All server and rack components are easily accessible from the front side of the rack (cold aisle), with the exception of low-failure rate redundant components such as fans, which require the tray to be removed from the rack.

Trays and other system components are field-replaceable and tool-less with green touch points and free of sharp edges.

7.13 Rack Compatibility

The Clemente compute tray is designed for use in the NVIDIA® canister installed in an Open Rack V3 (ORv3) HPR rack.

8. Electrical Requirements

The Clemente Compute Tray is compliant with ORv3 HPR specifications.

8.1 Input Voltage

The Clemente Compute Tray is compliant with the ORv3 HPR specification for 51VDC input voltage.

The system will operate normally with an input voltage tolerance range between 46V and 52V.

The absolute maximum voltage range is 40V to 59.5V.

8.2 Power

The Clemente Compute Tray power consumption depends on the implementation and workload. The power distribution for the compute tray is capable of delivering a maximum TDP power of ~4200W for all components, and up to a peak power consumption of ~7740W for short duration.

9. Thermal Design Requirements

9.1 Thermal Architecture

The platform includes both air cooled and liquid cooled components. CPU, GPU and Switch packages will be liquid cooled, with remaining air cooled. This section covers requirements for both air cooling and liquid cooling.

For details not stated in this document, please treat the NVIDIA® reference design as a starting point.

9.2 Air Cooling Specification

	Normal Operation	**Corner Case**	**Remain Power on**
Inlet Temperature	Up to 30°C	Up to 35°C	Up to 45°C
Altitude (no inlet temp deration)	Up to 6000 ft	Up to 6000 ft	Up to 6000 ft
Cold/Hot Aisle Δp	Within +/- 0.005 InH2O	Within +/- 0.2 InH2O	Within +/- 0.2 InH2O
Relative Humidity	10% ~ 90%	10% ~ 90%	10% ~ 90%
Fan Status	All Healthy	Two fan rotor failure (N+2)	Two fan rotor failure (N+2)
Air side ΔT*	More than 12.5°C (22°F)	NA	NA

(CFM/W)	Less than 0.145 CFM/W at sea level [Max loading condition only]		
Thermal Throttling	Not allowed	Not allowed	Allowed
Minimum Thermal Margin	7% of component temperature spec as a baseline.	NA	NA
Fan Power Ratio*	Below 5% of air side IT load [Max loading condition only]	NA	NA

*: may be waived if the corresponding fan duty becomes too low that impairs fan speed stability or system robustness upon boundary condition oscillation.

9.3 Fan Zone

Channel Island adopts individual fan arrays and control algorithms for each tray. 40mm counter rotating fans are expected for 1OU trays.

9.4 Fan Control Target

Under normal operation conditions, minimum margin (critical component) shall have at least 7% of component temperature spec, while fan power consumption shall be less than 5% of system total power (account for max loading power only). Exceptions may apply as stated above.
For components with large sensor discrepancies, or not included in PID control, the target thermal margin will be discussed separately.

9.5 Thermal Sensors

The BMC monitors all applicable temperature, power, and fan speed sensors, and adopts selective temperature readings for fan speed control algorithms when applicable. The maximum allowable tolerance of the inlet temperature sensor is ±2°C. The maximum allowable tolerance of other thermal sensors is ±5°C.

9.6 Fan Control Algorithm

Look up tables are applied to the inlet temperature sensor, and temperature sensors of all components with thermal risk under all possible operation conditions, if the fans operate at the minimum duty specified (20%). PID control tables should be applied to temperature sensors of all components that may be thermally critical under all possible operation conditions.

9.7 Special Fan Conditions

For devices/components/parts that do not have built-in temperature sensors or cannot be included in fan speed control, their thermal solution design should enable the inlet-temperature based FSC to ensure their cooling at idling/low loading conditions. The PID control FSCs should ensure their cooling at heavy load conditions.

Special fan modes shall be designed for the following scenarios:

- AC Cycle
- DC Power Off/Cycle
- Single rotor Failure
- Fan Board Failure
- Sensor failure or exceed LCR/UCR limits

9.9 Liquid Boundary Conditions

At rack level, the Channel Island liquid cooling solution design should be able to enable full TDP performance robustly, with:

- Coolant Type: PG25 based fluid, for instance Dow Frost LC-25
- Coolant Supply Temp: 40 degC nominal, 42 degC excursion
- Coolant Flow Rate and **Δ**p: up to 140 LPM and 15 psi

In nominal operation with max loading, the liquid side temperature difference should stay between 10~12 degC, i.e. 1.25 ~ 1.5 LPM/kW.

9.10 Connection Interface

Tray level QDs: Blind Mate Quick Connects/Disconnects (QD/QC) will be adopted for Compute Tray and Switch Tray. The pressure drop of each pairing is desired to be around 1psi or below, at flow rate of 4 LPM. UQD04, UQDB04, and ORV3 BMQC options are known to be capable of meeting such requirements.

9.11 Liquid Flow Distribution

In operation, the Channel Island platform is expected to operate at fixed liquid flow rate and pressure drop. The liquid loop will remain stable unless otherwise specified.

The liquid flow rate allocation shall be designed aiming at equivalent liquid side temperature difference across all trays. Flow rate variation across the same type of trays, i.e. comparing across all Compute Trays, or across all Switch Trays, shall be kept within +/- 10%.

9.12 Leakage Handling

The Channel Island platform is expected to have following functionalities:

- Leakage Detection: leak detection sensors (wire or spot type) that can capture leakage events, and communicate with system BMC/CPLD to raise the alarm to RMC. Desired solutions should have leakage detection at component, tray and rack level, allowing adjustment of sensitivity and granularity.
- Leakage Containment: mechanical design features that can contain the liquid within localized spaces in the tray, or drip pan.
- Leakage Response: being able to communicate with the necessary components/devices to shut down power delivery and turn off coolant supply, upon leakage detection. The detailed timing, sequence, blast radius and recovery mechanism will be optimized based on NPI validation and field operation outcomes, aiming at maximizing hardware availability and healthiness.

10. Interfaces

10.1 Signal List

The system includes the following general signals:

- High-speed signals
    - PCIe® 5.0 to the front end NICs and E1.S drives
    - PCIe® 6.0 between the CX8 mezzanine and the GPUs
    - Clink as the interconnect between the CPUs
    - C2C interface between the CPU and GPUs
    - Clocks
- Low-speed signals
    - I2C/I3C®/SPI for management, sensors, I/O expanders, and other components
    - USB 3.0 for debug connections
    - UART
    - GPIO
    - JTAG

10.2 Cables

Various cables (high speed, low speed, and power) are used to connect up the components within the compute tray. This is done to improve the SI loss budget on the high speed signals,

provide placement flexibility for the components within the tray, and enable more granular serviceability options during repairs.

11. System Firmware

Clemente uses the following open source firmware:

- Linux Foundation BMC, an open-source firmware stack for the BMC on the SCM:

12. Hardware Management

12.1 Out-Of-Band Management

The Clemente Compute Tray has a dedicated 200G NIC for out-of-band (OOB) management (NIC0). OOB access is always online as long as standby power is provided. The NIC attaches to the CPU Tray and is OCP NIC 3.0 compliant.

The BMC can update NIC firmware through OOB to ensure NIC update is available in case the OS is not accessible and in-band update fails.

12.2 BMC

This system uses an ASPEED AST2600 BMC with one 8Gb x16 DDR4 SDRAM and 65MB flash for various platform management services. The BMC interfaces with hardware, BIOS, and host firmware. It is part of the Linux Foundation OpenBMC:

https://github.com/openbmc/openbmc/tree/master/meta-facebook/meta-catalina/meta-clemente

The BMC resides on the SCM board, which follows the DC-SCM v2.0 specification (with board pinout exceptions).

The BMC is a standalone system in parallel to the host. The health status of the host system does not affect the normal operation and network connectivity of BMC.

The BMC is remotely upgradable in an out-of-band through the front end NIC.

12.3 Remote Upgradeability

The system is capable of remotely upgrading the VRs, CPLDs, BIC, switches, and FRU EEPROMs. The BMC can access each device, update code from BMC to the device, and verify the code.

The BMC is capable of recovering all CPLD and BIC in the system from image corruption or an accidental erase. The BMC does not rely on proper operation of the BIC or CPLD for booting or to be accessible remotely.

12.4 Sensors

BMC has access to all analog sensors directly or through host management. Below is a table of all sensors:

Sensor Name	Sensor Name
FAN0_PWM	HGX_CPU1_CORE36_UTIL_PCT
FAN0_TACH_IL_SPEED_RPM	HGX_CPU1_CORE37_UTIL_PCT
FAN0_TACH_OL_SPEED_RPM	HGX_CPU1_CORE38_UTIL_PCT
FAN1_PWM	HGX_CPU1_CORE39_UTIL_PCT
FAN1_TACH_IL_SPEED_RPM	HGX_CPU1_CORE3_UTIL_PCT
FAN1_TACH_OL_SPEED_RPM	HGX_CPU1_CORE40_UTIL_PCT
FAN2_PWM	HGX_CPU1_CORE41_UTIL_PCT
FAN2_TACH_IL_SPEED_RPM	HGX_CPU1_CORE42_UTIL_PCT
FAN2_TACH_OL_SPEED_RPM	HGX_CPU1_CORE43_UTIL_PCT
FAN3_PWM	HGX_CPU1_CORE44_UTIL_PCT
FAN3_TACH_IL_SPEED_RPM	HGX_CPU1_CORE45_UTIL_PCT
FAN3_TACH_OL_SPEED_RPM	HGX_CPU1_CORE46_UTIL_PCT
FAN4_PWM	HGX_CPU1_CORE47_UTIL_PCT
FAN4_TACH_IL_SPEED_RPM	HGX_CPU1_CORE48_UTIL_PCT
FAN4_TACH_OL_SPEED_RPM	HGX_CPU1_CORE49_UTIL_PCT
FAN5_PWM	HGX_CPU1_CORE4_UTIL_PCT
FAN5_TACH_IL_SPEED_RPM	HGX_CPU1_CORE50_UTIL_PCT
FAN5_TACH_OL_SPEED_RPM	HGX_CPU1_CORE51_UTIL_PCT
FAN6_PWM	HGX_CPU1_CORE52_UTIL_PCT
FAN6_TACH_IL_SPEED_RPM	HGX_CPU1_CORE53_UTIL_PCT
FAN6_TACH_OL_SPEED_RPM	HGX_CPU1_CORE54_UTIL_PCT
FAN7_PWM	HGX_CPU1_CORE55_UTIL_PCT
FAN7_TACH_IL_SPEED_RPM	HGX_CPU1_CORE56_UTIL_PCT

FAN7_TACH_OL_SPEED_RPM	HGX_CPU1_CORE57_UTIL_PCT
FIOBOARD_TEMP_C	HGX_CPU1_CORE58_UTIL_PCT
HDDBOARD1_SSD0_TEMP_C	HGX_CPU1_CORE59_UTIL_PCT
HDDBOARD1_SSD1_TEMP_C	HGX_CPU1_CORE5_UTIL_PCT
HGX_CPU0_CORE0_UTIL_PCT	HGX_CPU1_CORE60_UTIL_PCT
HGX_CPU0_CORE10_UTIL_PCT	HGX_CPU1_CORE61_UTIL_PCT
HGX_CPU0_CORE11_UTIL_PCT	HGX_CPU1_CORE62_UTIL_PCT
HGX_CPU0_CORE12_UTIL_PCT	HGX_CPU1_CORE63_UTIL_PCT
HGX_CPU0_CORE13_UTIL_PCT	HGX_CPU1_CORE64_UTIL_PCT
HGX_CPU0_CORE14_UTIL_PCT	HGX_CPU1_CORE65_UTIL_PCT
HGX_CPU0_CORE15_UTIL_PCT	HGX_CPU1_CORE66_UTIL_PCT
HGX_CPU0_CORE16_UTIL_PCT	HGX_CPU1_CORE67_UTIL_PCT
HGX_CPU0_CORE17_UTIL_PCT	HGX_CPU1_CORE68_UTIL_PCT
HGX_CPU0_CORE18_UTIL_PCT	HGX_CPU1_CORE69_UTIL_PCT
HGX_CPU0_CORE19_UTIL_PCT	HGX_CPU1_CORE6_UTIL_PCT
HGX_CPU0_CORE1_UTIL_PCT	HGX_CPU1_CORE70_UTIL_PCT
HGX_CPU0_CORE20_UTIL_PCT	HGX_CPU1_CORE71_UTIL_PCT
HGX_CPU0_CORE21_UTIL_PCT	HGX_CPU1_CORE7_UTIL_PCT
HGX_CPU0_CORE22_UTIL_PCT	HGX_CPU1_CORE8_UTIL_PCT
HGX_CPU0_CORE23_UTIL_PCT	HGX_CPU1_CORE9_UTIL_PCT
HGX_CPU0_CORE24_UTIL_PCT	HGX_CPU1_ENERGY_J
HGX_CPU0_CORE25_UTIL_PCT	HGX_CPU1_EXHAUST_TEMP_C
HGX_CPU0_CORE26_UTIL_PCT	HGX_CPU1_INLET0_TEMP_C
HGX_CPU0_CORE27_UTIL_PCT	HGX_CPU1_INLET1_TEMP_C
HGX_CPU0_CORE28_UTIL_PCT	HGX_CPU1_PWR_W
HGX_CPU0_CORE29_UTIL_PCT	HGX_CPU1_TEMP_C
HGX_CPU0_CORE2_UTIL_PCT	HGX_CPU1_THERMAL_MARGIN_TEMP_C
HGX_CPU0_CORE30_UTIL_PCT	HGX_CPU1_VR_CPU_PWR_W
HGX_CPU0_CORE31_UTIL_PCT	HGX_CPU1_VR_CPU_VOLT_V
HGX_CPU0_CORE32_UTIL_PCT	HGX_CPU1_VR_SOC_PWR_W

HGX_CPU0_CORE33_UTIL_PCT	HGX_CPU1_VR_SOC_VOLT_V
HGX_CPU0_CORE34_UTIL_PCT	HGX_GPU0_DRAM_PWR_W
HGX_CPU0_CORE35_UTIL_PCT	HGX_GPU0_DRAM_TEMP_C
HGX_CPU0_CORE36_UTIL_PCT	HGX_GPU0_ENERGY_J
HGX_CPU0_CORE37_UTIL_PCT	HGX_GPU0_PWR_W
HGX_CPU0_CORE38_UTIL_PCT	HGX_GPU0_TEMP_C
HGX_CPU0_CORE39_UTIL_PCT	HGX_GPU0_THERMAL_MARGIN_TEMP_C
HGX_CPU0_CORE3_UTIL_PCT	HGX_GPU0_VOLT_V
HGX_CPU0_CORE40_UTIL_PCT	HGX_GPU1_DRAM_PWR_W
HGX_CPU0_CORE41_UTIL_PCT	HGX_GPU1_DRAM_TEMP_C
HGX_CPU0_CORE42_UTIL_PCT	HGX_GPU1_ENERGY_J
HGX_CPU0_CORE43_UTIL_PCT	HGX_GPU1_PWR_W
HGX_CPU0_CORE44_UTIL_PCT	HGX_GPU1_TEMP_C
HGX_CPU0_CORE45_UTIL_PCT	HGX_GPU1_THERMAL_MARGIN_TEMP_C
HGX_CPU0_CORE46_UTIL_PCT	HGX_GPU1_VOLT_V
HGX_CPU0_CORE47_UTIL_PCT	HGX_GPU2_DRAM_PWR_W
HGX_CPU0_CORE48_UTIL_PCT	HGX_GPU2_DRAM_TEMP_C
HGX_CPU0_CORE49_UTIL_PCT	HGX_GPU2_ENERGY_J
HGX_CPU0_CORE4_UTIL_PCT	HGX_GPU2_PWR_W
HGX_CPU0_CORE50_UTIL_PCT	HGX_GPU2_TEMP_C
HGX_CPU0_CORE51_UTIL_PCT	HGX_GPU2_THERMAL_MARGIN_TEMP_C
HGX_CPU0_CORE52_UTIL_PCT	HGX_GPU3_DRAM_PWR_W
HGX_CPU0_CORE53_UTIL_PCT	HGX_GPU3_DRAM_TEMP_C
HGX_CPU0_CORE54_UTIL_PCT	HGX_GPU3_ENERGY_J
HGX_CPU0_CORE55_UTIL_PCT	HGX_GPU3_PWR_W
HGX_CPU0_CORE56_UTIL_PCT	HGX_GPU3_TEMP_C
HGX_CPU0_CORE57_UTIL_PCT	HGX_GPU3_THERMAL_MARGIN_TEMP_C
HGX_CPU0_CORE58_UTIL_PCT	HGX_HMC_TEMP_C
HGX_CPU0_CORE59_UTIL_PCT	HGX_HPM0_ALTITUDE_PRESSURE_PA
HGX_CPU0_CORE5_UTIL_PCT	HGX_HPM1_ALTITUDE_PRESSURE_PA

HGX_CPU0_CORE60_UTIL_PCT	HGX_TOTALGPU_PWR_W
HGX_CPU0_CORE61_UTIL_PCT	Interposer_CABLE_TSENSE_NEGATIVE_CURR_A
HGX_CPU0_CORE62_UTIL_PCT	Interposer_CABLE_TSENSE_NEGATIVE_TEMP_C
HGX_CPU0_CORE63_UTIL_PCT	Interposer_CABLE_TSENSE_POSITIVE_CURR_A
HGX_CPU0_CORE64_UTIL_PCT	Interposer_CABLE_TSENSE_POSITIVE_TEMP_C
HGX_CPU0_CORE65_UTIL_PCT	Interposer_P12V_NIC0_CURR_A
HGX_CPU0_CORE66_UTIL_PCT	Interposer_P12V_NIC0_PWR_W
HGX_CPU0_CORE67_UTIL_PCT	Interposer_P12V_NIC0_TEMP_C
HGX_CPU0_CORE68_UTIL_PCT	Interposer_P12V_NIC0_VOLT_V
HGX_CPU0_CORE69_UTIL_PCT	Interposer_P12V_NIC1_CURR_A
HGX_CPU0_CORE6_UTIL_PCT	Interposer_P12V_NIC1_PWR_W
HGX_CPU0_CORE70_UTIL_PCT	Interposer_P12V_NIC1_TEMP_C
HGX_CPU0_CORE71_UTIL_PCT	Interposer_P12V_NIC1_VOLT_V
HGX_CPU0_CORE7_UTIL_PCT	Interposer_P12V_SCM_CURR_A
HGX_CPU0_CORE8_UTIL_PCT	Interposer_P12V_SCM_PWR_W
HGX_CPU0_CORE9_UTIL_PCT	Interposer_P12V_SCM_TEMP_C
HGX_CPU0_ENERGY_J	Interposer_P12V_SCM_VOLT_V
HGX_CPU0_EXHAUST_TEMP_C	NIC0_TEMP_C
HGX_CPU0_INLET0_TEMP_C	NIC1_TEMP_C
HGX_CPU0_INLET1_TEMP_C	PDB_HSC1_CURR_A
HGX_CPU0_PWR_W	PDB_HSC1_PWR_W
HGX_CPU0_TEMP_C	PDB_HSC1_TEMP_C
HGX_CPU0_THERMAL_MARGIN_TEMP_C	PDB_HSC1_VIN_VOLT_V
HGX_CPU0_VR_CPU_PWR_W	PDB_HSC1_VOUT_VOLT_V
HGX_CPU0_VR_CPU_VOLT_V	PDB_HSC2_CURR_A
HGX_CPU0_VR_SOC_PWR_W	PDB_HSC2_PWR_W
HGX_CPU0_VR_SOC_VOLT_V	PDB_HSC2_TEMP_C

HGX_CPU1_CORE0_UTIL_PCT	PDB_HSC2_VIN_VOLT_V
HGX_CPU1_CORE10_UTIL_PCT	PDB_HSC2_VOUT_VOLT_V
HGX_CPU1_CORE11_UTIL_PCT	PDB_INLET_TEMP_C
HGX_CPU1_CORE12_UTIL_PCT	PDB_MAX_VR_TEMP_C
HGX_CPU1_CORE13_UTIL_PCT	PDB_VR_P12V_N1_CURR_A
HGX_CPU1_CORE14_UTIL_PCT	PDB_VR_P12V_N1_PWR_W
HGX_CPU1_CORE15_UTIL_PCT	PDB_VR_P12V_N1_TEMP_C
HGX_CPU1_CORE16_UTIL_PCT	PDB_VR_P12V_N1_VOLT_V
HGX_CPU1_CORE17_UTIL_PCT	PDB_VR_P12V_N2_CURR_A
HGX_CPU1_CORE18_UTIL_PCT	PDB_VR_P12V_N2_PWR_W
HGX_CPU1_CORE19_UTIL_PCT	PDB_VR_P12V_N2_TEMP_C
HGX_CPU1_CORE1_UTIL_PCT	PDB_VR_P12V_N2_VOLT_V
HGX_CPU1_CORE20_UTIL_PCT	SCM_P12V_VOLT_V
HGX_CPU1_CORE21_UTIL_PCT	SCM_P1V0_VOLT_V
HGX_CPU1_CORE22_UTIL_PCT	SCM_P1V2_VOLT_V
HGX_CPU1_CORE23_UTIL_PCT	SCM_P1V8_VOLT_V
HGX_CPU1_CORE24_UTIL_PCT	SCM_P2V5_VOLT_V
HGX_CPU1_CORE25_UTIL_PCT	SCM_P3V3_VOLT_V
HGX_CPU1_CORE26_UTIL_PCT	SCM_P3V_BAT_VOLT_V
HGX_CPU1_CORE27_UTIL_PCT	SCM_P5V_VOLT_V
HGX_CPU1_CORE28_UTIL_PCT	SCM_PGPPA_VOLT_V
HGX_CPU1_CORE29_UTIL_PCT	SCM_TEMP_C
HGX_CPU1_CORE2_UTIL_PCT	SYSTEM_AIRFLOW_CFM
HGX_CPU1_CORE30_UTIL_PCT	VIRTUAL_FIOBOARD_TEMP_C
HGX_CPU1_CORE31_UTIL_PCT	
HGX_CPU1_CORE32_UTIL_PCT	
HGX_CPU1_CORE33_UTIL_PCT	
HGX_CPU1_CORE34_UTIL_PCT	
HGX_CPU1_CORE35_UTIL_PCT	

13. Security

The Clemente platform includes the following security features:

- SCM with Linux Foundation BMC
    - Verified boot with signed firmware
- NVIDIA® based HPM
    - External root of trust

Appendix A - Checklist for IC approval of this Specification (to be completed by contributor(s) of this Spec)

Complete all the checklist items in the table with links to the section where it is described in this spec or an external document .

Item	Status	Link to detailed explanation
Has this contribution been presented to an OCP Project group during a project call or engineering workshop?	**Yes**	If “No”, please state the reason.
Approval by Project Leads		If “No”, please state the reason.
Is this contribution entered into the OCP Contribution Portal?		If “No”, please state the reason.
Was it approved in the OCP Contribution Portal?		If “No”, please state the reason.