Open Compute Project

DCEngine
Rack, Compute and Storage System Specification for CG-OpenRack-19
Version 1.01

Author:

Matt St Peter, Radisys Corporation

1 Revision History

Date	Name	Description
4/11/2017	Radisys Corp	Version 1.0
8/2/2017	Radisys Corp	Version 1.01 – updates per request of OCP IC. More details added on each element, improved diagrams.

2 Contents

3 License

4 Scope

DCEngine is a scalable carrier-grade rack level system designed to the OCP Approved CG-OpenRack-19 specification. It integrates high performance compute, storage and networking in several different rack sizes and configurations. All configurations are based on a set of defined modules that fit into a rack providing power and interconnects, according to the CG-OpenRack-10 Sled and Rack Specification Version 1.0.

This document provides a detailed description of the DCEngine rack and sled (compute and storage) contents as well as the associated architecture and internal components. It is a system level specification intended to provide a medium level of detail; separate specifications provide full design detail for each individual component.

Note that sleds may be customized to fit specific site environmental requirements. Therefore, parameters such as processor speed, memory size and speed, storage speed and capacity, and rack power delivery capacity are variable. The diverse range of system deployment scenarios will ultimately dictate a variety of solutions.

Additionally, overall rack configurations and the position of sleds and other major components within the rack are variable due to the flexible nature of DCEngine and CG-OpenRack-19 solutions.

5 Overview

In general, and as described in the CG-OpenRack-19 specification, a DCEngine system is made up of compute and storage sleds, Ethernet switches, a power source, and supporting rack infrastructure as shown in Figure 1. This specification addresses certain instantiations of these elements as manifested in the DCEngine system.

Rack and Power

- This system specification provides a high-level view of the rack configuration, including power supply/distribution.

Full and Half-Width Sleds

- This system specification provides details for a full width 2U storage sled, and a half-width 2U compute sled.

Top-of-Rack (ToR) Switches

- This specification does not provide specific switch details. DCEngine can support a variety of standard 19” rack mounted networking switches, typically oriented with the I/O ports facing rear (and hence back-to-front airflow). They connect to the sleds via the blind-mate interconnect system: optical fiber bundles connecting from each sled location to one or more individual switch ports.

Figure 1. Rack with Half-Width and Full-Width Sleds

6 Rack Specifications

Each 2U shelf within the rack is configured as full-width or half-width. Half-width shelves have a center partition and two position for optical blind mate connection, while the full-width shelf has no partition and has a single blind-mate optical connection site. All shelf mating surfaces, alignment and latch features, optical connections, and apertures conform to the CG-OpenRack-19 specification.

6.1 Rack and Cabinet Dimensions

There are several different size options for the overall cabinet that encases the standard 19-inch rack core within a complete rack system. For the full-height 42U rack, there are 800mm and 600mm wide cabinets, and for the shorter 16U rack, there is a 600mm wide cabinet.

The overall height of the cabinet including the rack core is from the bottom of the wheels to the top of the Rack.

Typical 42U High, 600mm Wide Rack System Dimensions:

Cabinet Height:	2114mm
Cabinet Width:	598mm
Cabinet Depth:	1196mm
Rach-core Height	1992mm
Rack-core Width:	19-inch standard
Rack-core Depth:	992mm

Typical 42U High, 800mm Wide Rack System Dimensions:

Cabinet Height:	2088mm
Cabinet Width:	798mm
Cabinet Depth:	998mm

Rach-core Height	1992mm
Rack-core Width:	19-inch standard
Rack-core Depth:	992mm

Typical 16U High, 600mm Wide Rack System Dimensions:

Cabinet Height:	935mm
Cabinet Width:	596mm
Cabinet Depth:	998mm
Rach-core Height	1992mm
Rack-core Width:	19-inch standard
Rack-core Depth:	992mm

Figure 2. 42U Rack-Core, Front View

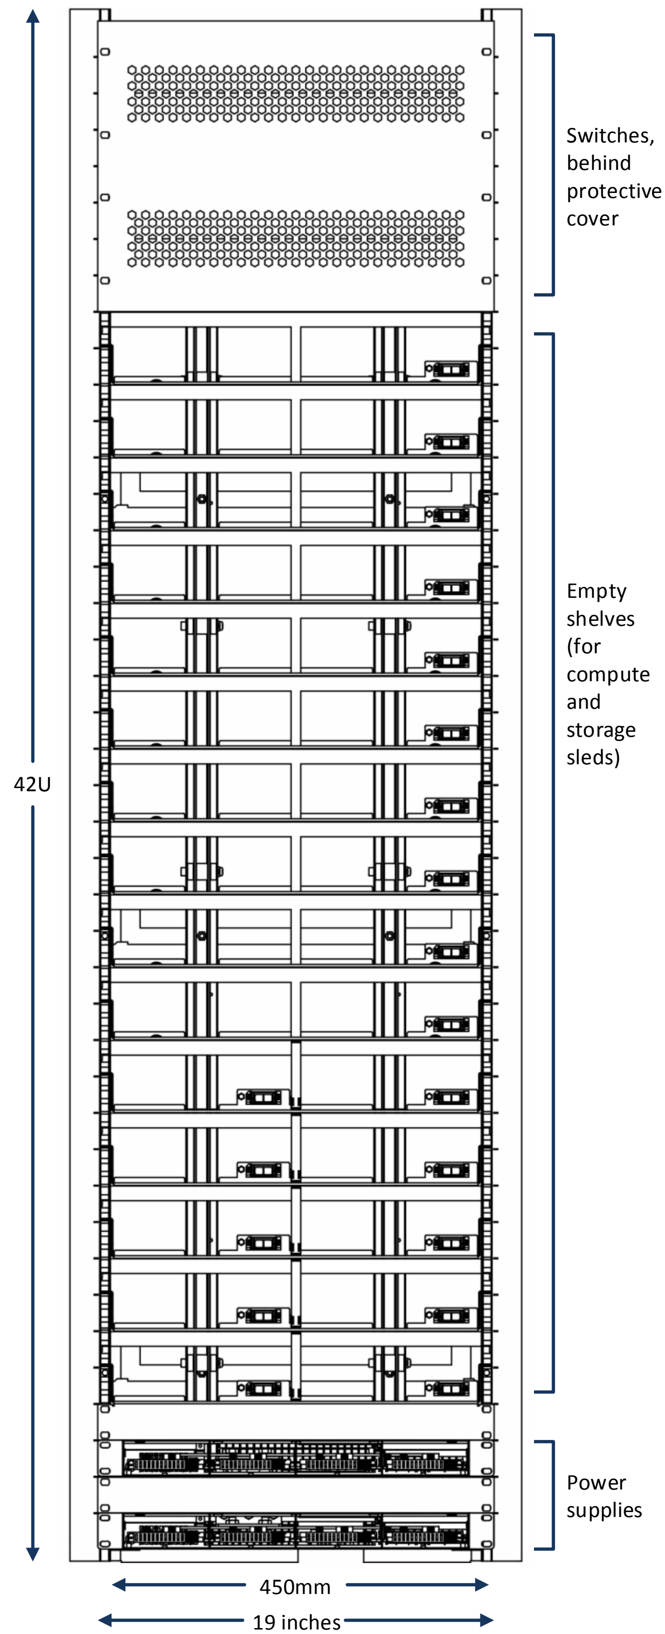

Figure 3. 42U Rack-Core, Rear View

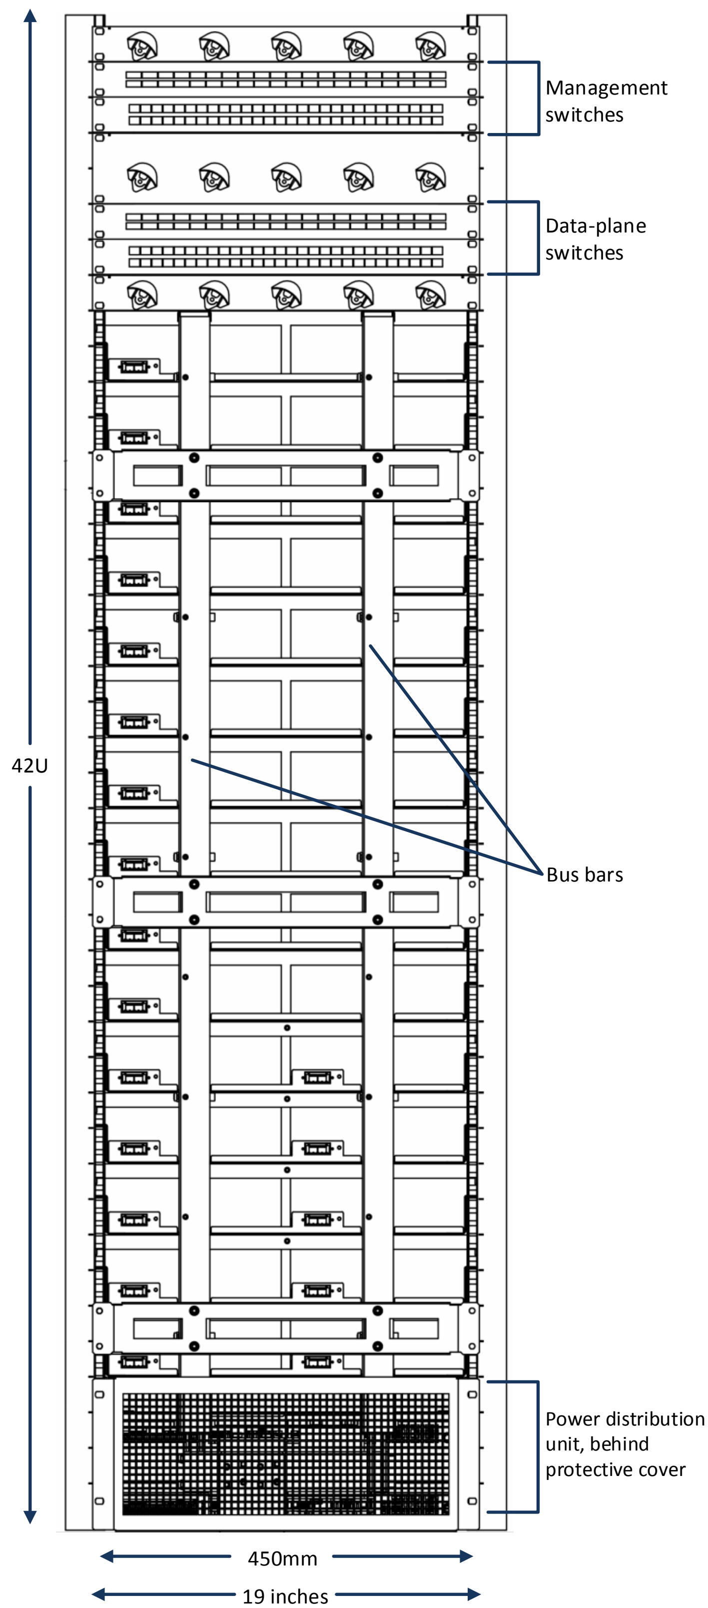

Figure 4. 42U Rack-Core, Side View

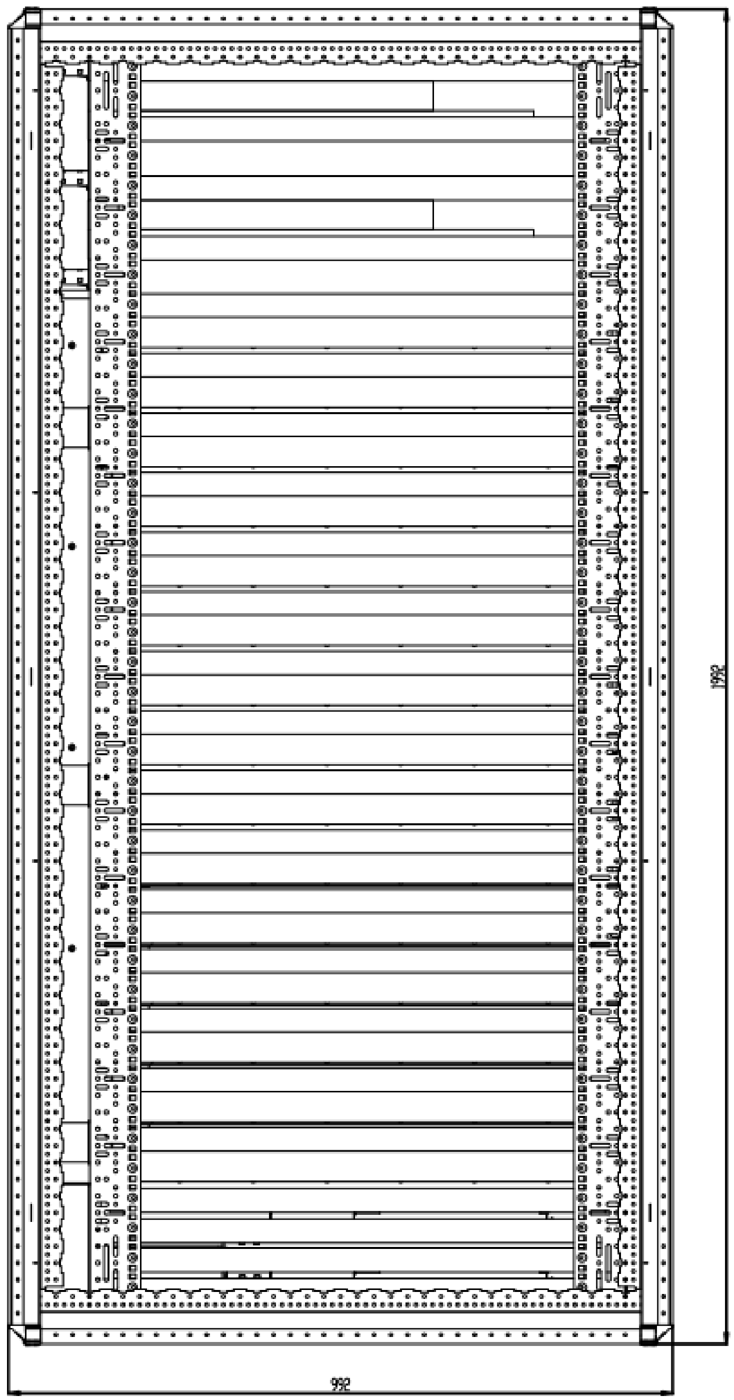

Figure 5. 42U 800mm Rack-System Front View

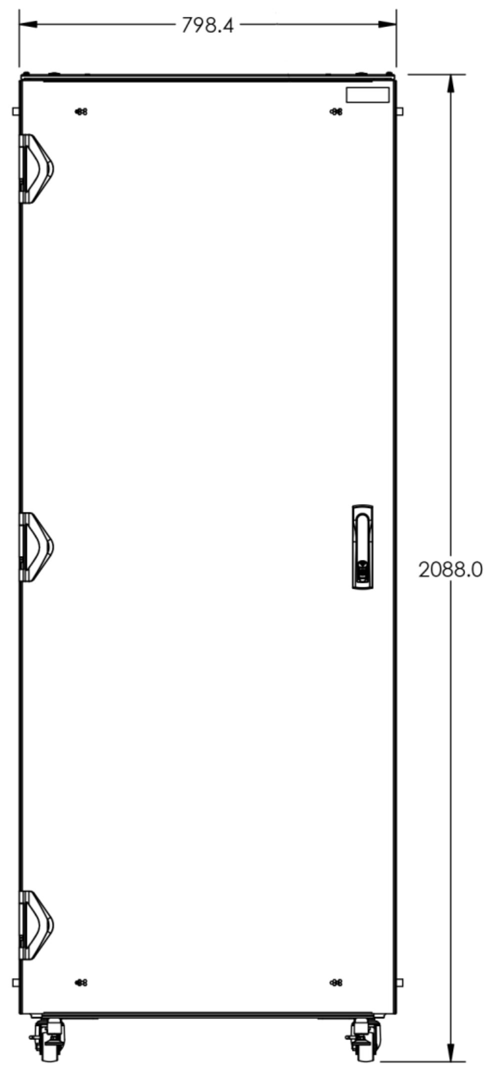

Figure 6. 42U 800mm Rack-System Top View

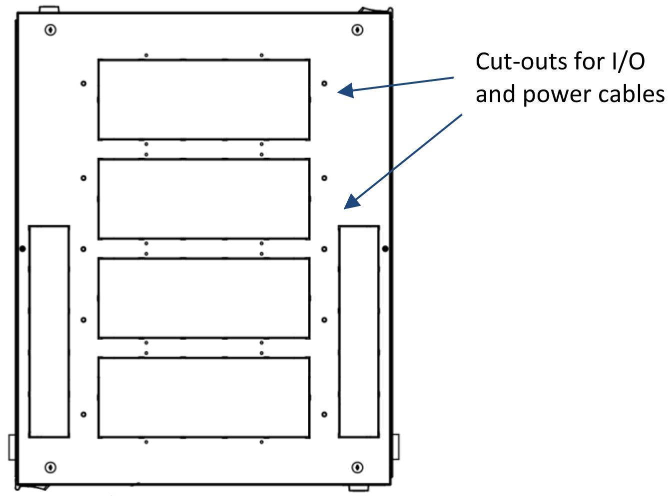

Figure 7. 42U 600mm Rack-System Front View

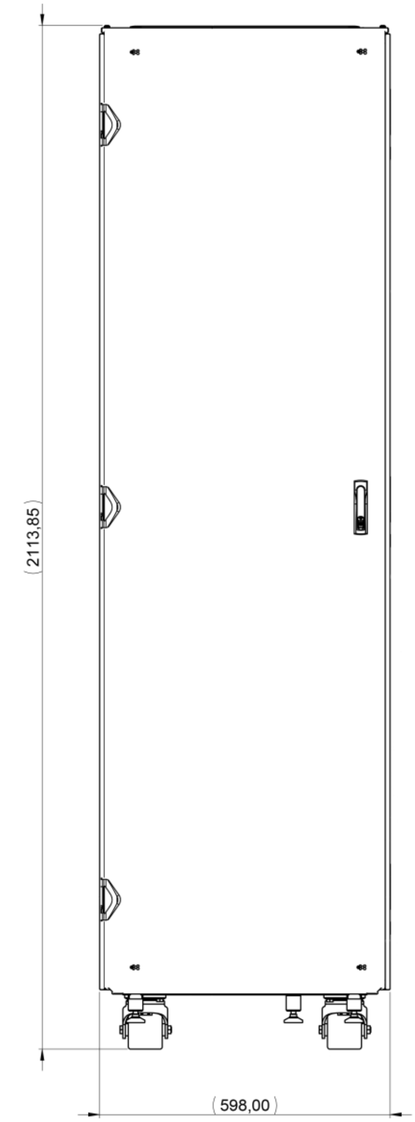

Figure 8. 16U 600mm Rack Front View

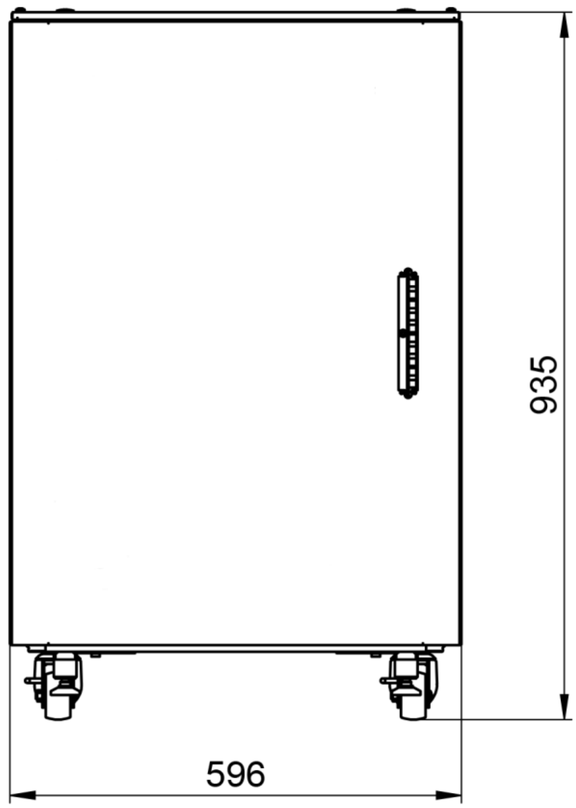

6.2 Rack Power Input Specifications

DCEngine supports several types of rack-level power, including AC (single-phase or three-phase Wye and Delta configurations) and DC (-48V and 400V).

AC 3-Phase Delta power feed requirements:

Voltage: Nominal 240VAC, 192 to 380VAC phase to phase

Current: Maximum 17.5 Amps, RMS per PSU

Plug type: 3-pole, 4-wire (incl. ground), Hubbell CS8365C, 50A

Feed qty: 6 feeds, including redundant source

Figure 9. 120/208VAC Delta Power Plug

CS6365C

AC 3-Phase Wye power feed requirements:

Voltage: Nominal 120/208VAC, 192 to 380VAC phase-to-neutral

Current: Maximum 17.5 Amps, RMS per PSU

Plug type: 4-pole, 5-wire (incl. ground), IEC 60309 4P+E, 32A

Feed qty: 6 feeds, including redundant source

Figure 10. 240VAC Wye Power Plug

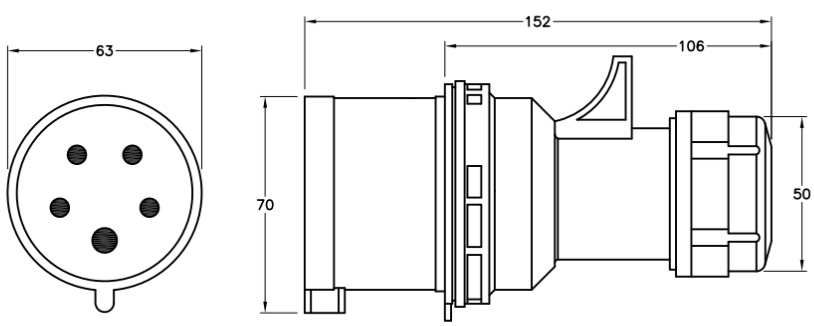

AC 3-Phase Wye power feed requirements:

Voltage:	Nominal 120/208VAC, 192 to 380VAC phase-to-neutral
Current:	Maximum 17.5 Amps, RMS per PSU
Plug type:	4-pole, 5-wire (incl. ground), IEC 60309 4P+E, 32A
Feed qty:	6 feeds, including redundant source

Figure 11. 240VAC Wye Power Plug

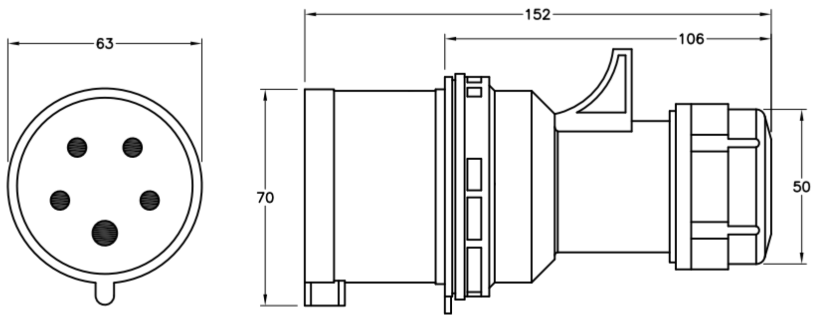

DC power feed requirements:

Voltage:	Nominal -48VDC, -40 to -72VDC
Current:	Maximum 85 Amps, DC per PSU
Plug type:	2-pole, 2-wire, Anderson SBX 75X, 110A
Feed qty:	Up to 20 feeds, depending on PSU count and redundancy

Figure 12. -48VDC Power Plug

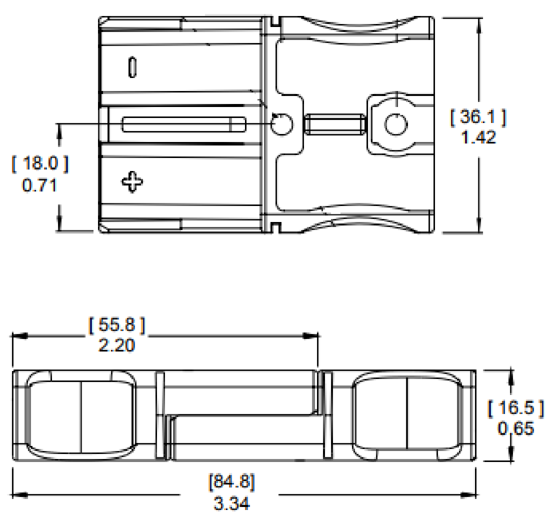

6.3 Power Distribution

All external rack power for the sleds is converted to 12VDC at the base of the rack, and distributed to the sleds through two vertical bus bars. The rack Power Distribution Unit (PDU) provides overcurrent protection and distribution of external power to the modular 12VDC power supplies at the base of the rack. The PDU also provides power to rack devices such as top-of-rack switches.

Power supply bays can be populated as needed, depending on the load in the rack. The table below provides an overview of the COTS TDK-Lambda HFE2500 series power supplies. For detailed information about the HFE2500 series, see the TDK-Lamba data sheet: http://www.us.tdk-lambda.com/ftp/specs/hfe2500.pdf

Table 1. Power supply specifications

Name	Definition
Input Voltage	85 – 265VAC
Input Frequency	47 - 63Hz
Input Current	12A Max.
Inrush Current	> 50A
Operating Temperature	-10 to 70 C
Safety Agency Certification	UL60950-1, CE Mark

A logical diagram showing the power distribution is shown below.

Figure 13. Logical Diagram of Rack Power Distribution

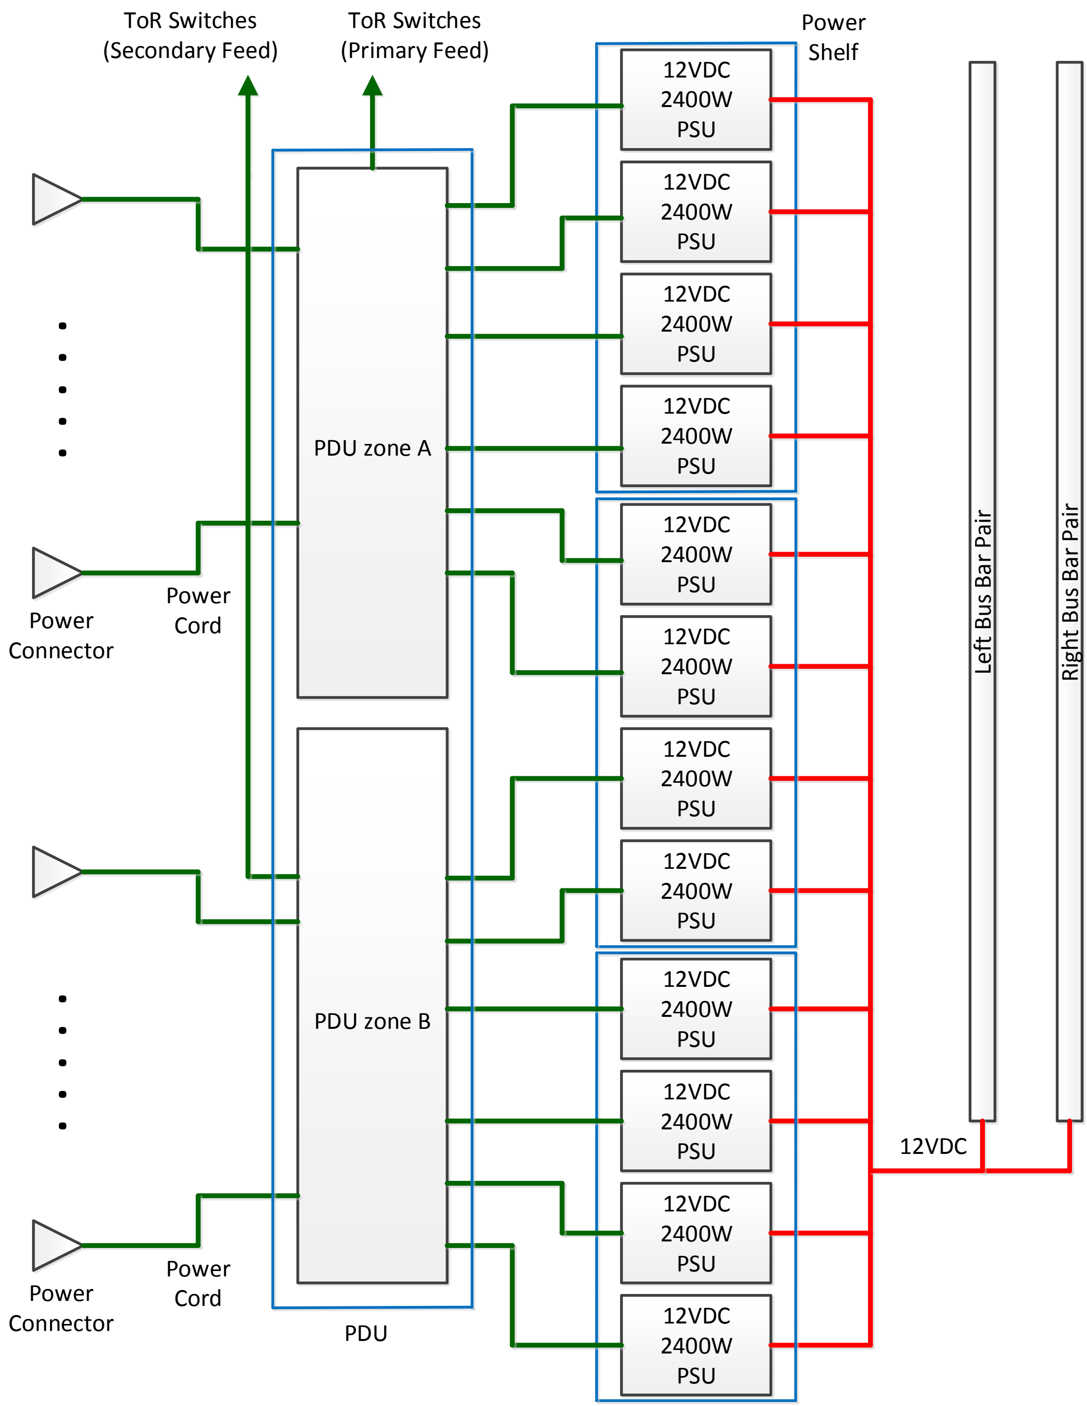

The diagram below shows a basic full size (42U) rack configuration with both full-width and half-width sleds.

Figure 14. 42U Rack Power Distribution Block Diagram

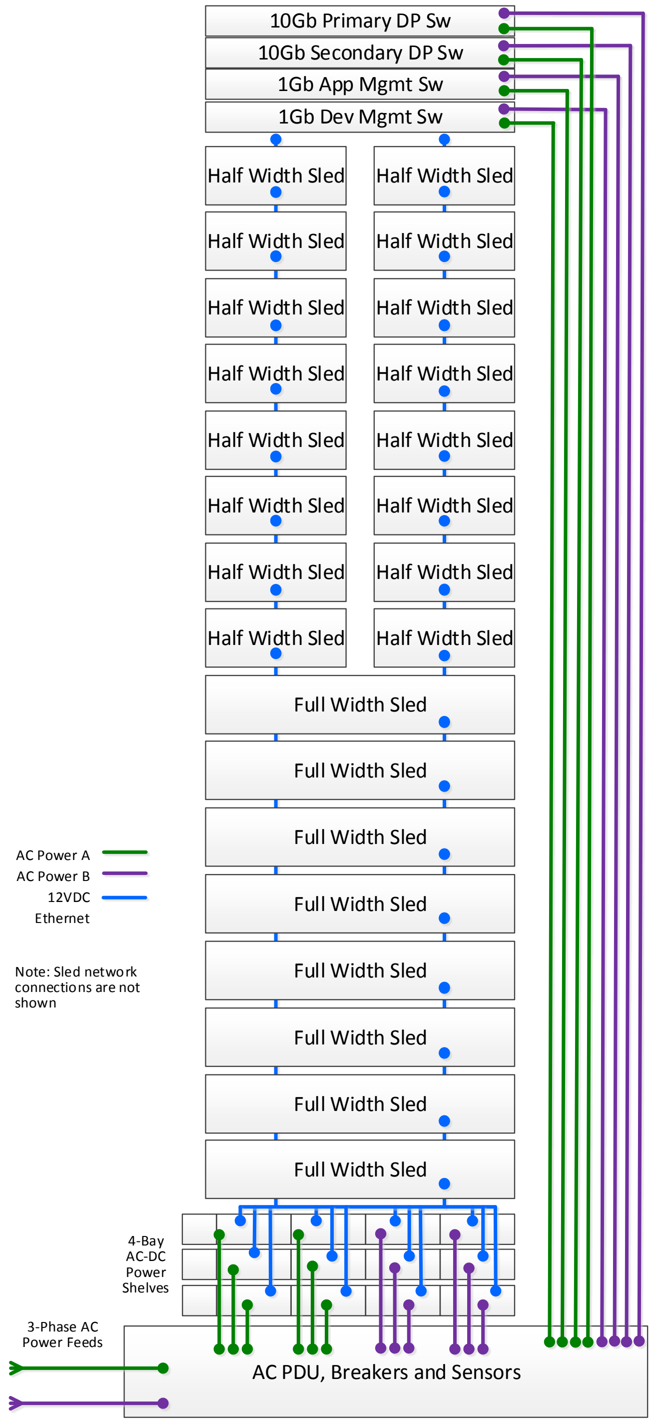

The diagram below shows a 16U rack configuration with both full-width and half-width sleds.

Figure 15. 16U Rack Power Distribution Block Diagram

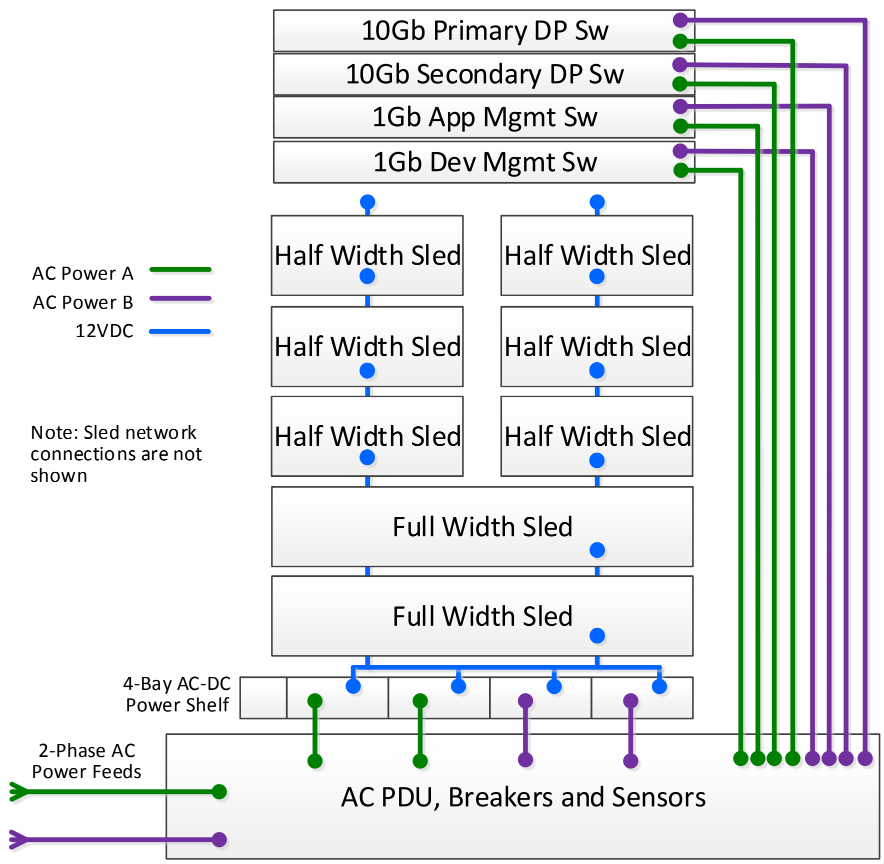

6.4 Bus Bar Locations

The rack has two vertical 12V bus bar pairs at the rear of all shelves.

Each right-hand (facing front) vertical bus bar pair is positioned 108mm from the right side of the sled bay wall. Each left-hand vertical bus bar is centered 229mm from the right-hand pair.

6.5 Rack Management

Optional basic rack level management through a single IP address may be achieved by adding a management agent on the device management switch. This agent provides the following functions:

- Discovery and inventory of the servers and switches present in the rack
- Health status monitoring and alerting of each server
- Power and boot control of each server.

This functionality is optional in that it does not need to be present for the rack elements to function, but it adds a convenient level of abstraction and aggregation for network level system management of the constituent elements. It is not required by the current CG-OpenRack-19 specification, but it is an expected area of evolution in parallel to direct node management.

DCEngine uses the Intel® Rack Scale Design (RSD) reference implementation (available at https://github.com/01org/intelRSD) as a basis for adding basic rack level management capabilities to the rack. The Pooled System Management Engine (PSME) of the reference implementation has been ported to the Linux environment running on the device management switch as shown in Figure 16.

Figure 16. Rack Management using Intel PSME Agent

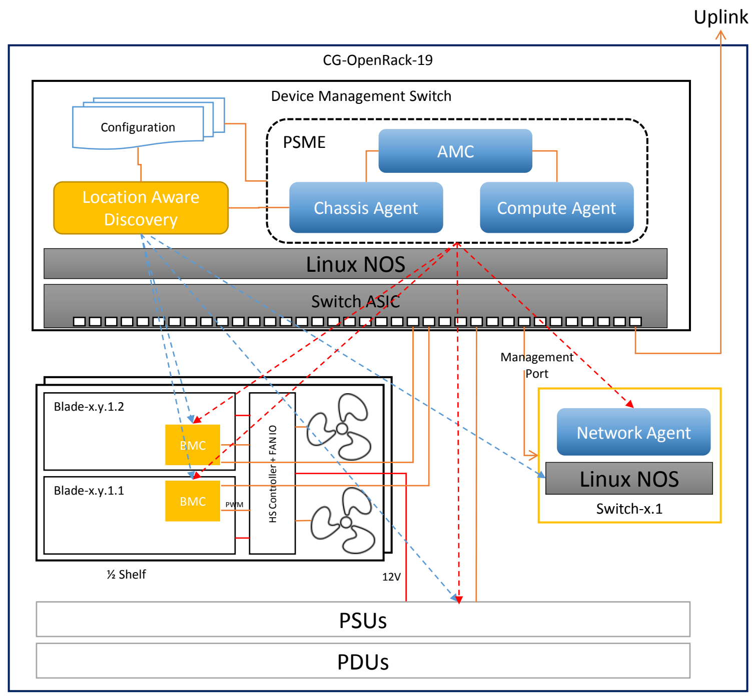

The RSD PSME reference code contains multiple agents for managing each of the asset types in the rack, i.e. compute, chassis, network and storage. The agents register with a common Asset Management Controller (AMC) that provides the northbound REST API to external management applications. The chassis and compute agents have been updated to support the management of the sleds and servers in the rack.

A location aware discovery agent was also added to the device management switch to discover and determine the locations of each of the sleds and servers in the rack. This agent monitors the switch ports to determine the presence/absence of the devices in the rack. It uses a Port-to-Device Location mapping configuration file to map learned MAC addresses to Blade & Switch locations (i.e. MAC → Port → Location). The agent provides a set of APIs that allow the chassis agent to learn the location, management IP address, and presence of devices in the rack.

Once the devices are discovered through the discovery agent, the chassis and compute agents are able to access inventory information and health status from the server BMCs using IPMI messaging. This information is then made available through the Redfish REST APIs supported by the AMC, essentially acting as an aggregation + proxy function for the nodes. In the future native Redfish nodes could similarly be aggregated at the rack level and/or be independently managed.

7 Sled Specifications

The sleds are based on a uniform form factor that allows interoperability with the rack, as specified by CG-OpenRack-19. This includes the external dimensions, external connections, and features designed to interoperate with the rack.

Sleds are positioned at a 2RU vertical pitch. The number of shelves within a rack is flexible, and can be defined by the required density or application. The vertical bus bars are designed to cover the sled area.

All dimensional tolerances are per DIN ISO 2768-mK, except sled height, width, and depth, which is +/-0.20mm.

Each sled has locally controlled cooling that allow the in-sled airflow to be set according to its own needs within the rack system. This is done with two fans in the half-width sleds and four fans in the full-width sleds. Sled airflow is front-to-back for all sleds.

The two sled types described in this document are based on the Intel S2600TP Server Board, which in turn is based on the Xeon processor family. Of course, many other possible sled variations and SKUs may also be compatible with CG-Openrack-19 and the DCEngine rack system.

7.1 Compute Sled

7.1.1 Detailed Description

The "Compute Sled" is 2U high and occupies half the width of a rack shelf. Each sled has two Intel S2600TP Server Boards.

Compute sled features

- Designed to CGOpenRack-19 specifications
- 2U half-width sled for DCEngine hyperscale rack systems
- 1 or 2 Intel S2600TP dual-socket Xeon server boards
- Supports Intel E5-26xx class processors, including Broadwell (v4) up to 20 cores, with long lifecycle and low power options
- 16x DIMM sockets, DDR4 288-pin PC4-2133, 1.2V with ECC (up to 512GB)
- Integrated SSD boot/OS/App drive
- Up to 3x 2.5" SATA HDDs or SSDs per server
- 2x 10G dataplane ports per server to rear optical interconnect
- 2x 1Gb IPMI/management plane ports per server to rear optical interconnect
- Integrated 80mm fans reduce noise levels compared to 1U servers
- Blind mate rear IO simplifies sled replacement

All of the sleds are powered through the vertical 12VDC bus bars and rack level power supplies.

The compute sled is shown in Figure 17. The internal block diagram showing power feed management, storage connections, and external network connections to the backplane connector is shown in Figure 18.

Figure 17. Compute Sled View witout Top Cover

Figure 18. Compute Sled Internal Connections

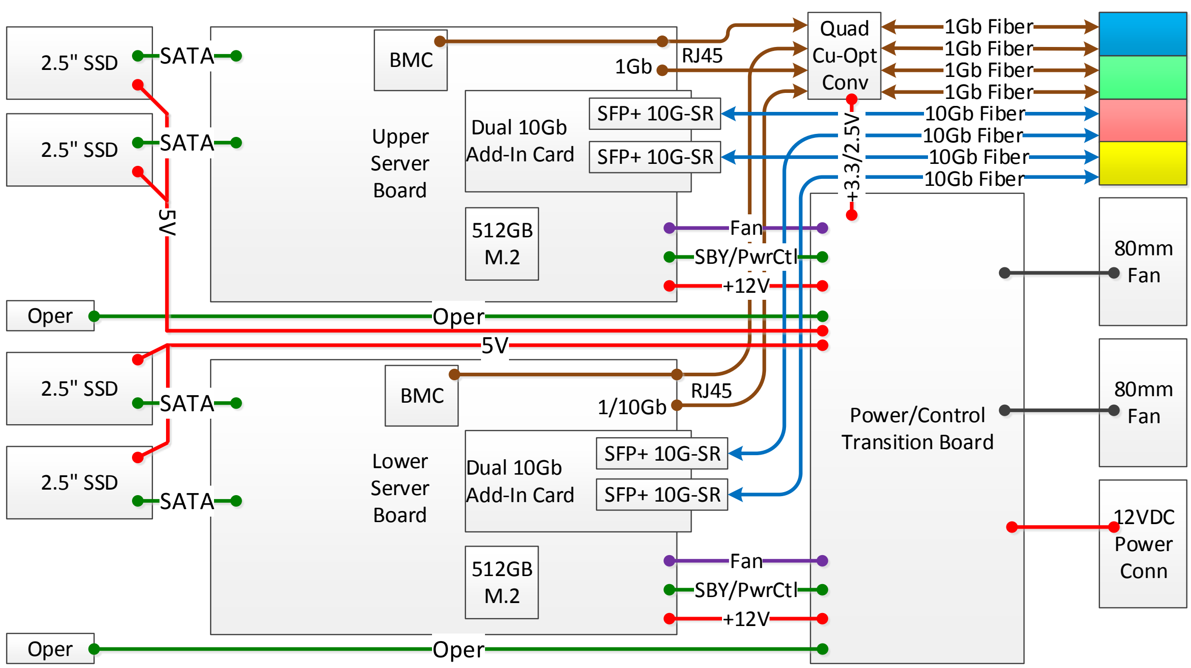

7.1.2 Serviceability

The compute sled can be removed and replaced quickly, without the need to power down the entire rack system.

Although it is recommended that internal compute sled components are replaced by a certified technician, the following components can be replaced:

- Server board
- Power board
- Media converter board
- Dual in-line memory module (DIMM)
- Solid state drive (SSD)
- Peripheral component interconnect express (PCIe) plug-in card
- Riser card
- M.2 flash boot drive
- Optical transceivers
- External connectors

7.1.2.1 Compute Sled Replacement Procedure

This section provides brief user-centric instructions on how to replace a compute sled, for illustrative purposes.

<u>Remove sled:</u>

Note: Sleds can be hot swapped; they do not need to be powered down for electrical or safety reasons before removal. However, sleds may be powered down by pushing the power button on the right side of the sled. The OK LED next to the power button indicates when the power is off. Compute sleds containing dual server nodes have two power buttons.

1. Push the sled release latch toward the middle of the sled and pull the sled from the rack using the front handle.

2. Avoid hitting the connectors at the rear of the sled by ensuring the sled is supported when removing it from the shelf. Keep the sled level; do not allow the sled to drop or rotate in the rack.

Install replacement sled:

1. Remove the new sled from its packaging.
2. Remove the yellow protective cover.
3. Insert the replacement sled into the rack, using mechanical assistance from a lift if necessary. Make sure the sled is level and not twisted in the rack. The sled should slide easily, latching with an audible click.

7.2 Storage Sled

7.2.1 Detailed Description

The "Storage Sled" is also 2U high and uses the full width of a rack shelf. Each sled has a single Intel S2600TP Server Board along with sixteen 3.5" hard drives connected to a disk controller through an expansion card.

Storage sled Features

- Designed to CG-OpenRack-19 specifications
- 2U full-width sled for DCEngine hyperscale rack systems
- 1 Intel S2600TP dual-socket Xeon server board
- Supports Intel E5-26xx class processors, including Broadwell (v4) up to 20 cores, with long lifecycle and low power options
- 16x DIMM sockets, DDR4 288-pin PC4-2133, 1.2V with ECC (up to 512GB)
- Integrated SSD boot/OS/App drive
- Up to 4x 2.5" SATA HDDs or SSDs
- Up to 16x 3.5" or 2.5" SAS/SATA HDDs/SSDs, configurable as one or more RAID groups
- Supports HW RAID (on-controller MegaRAID) or SW RAID
- 2x 10G dataplane ports to rear optical interconnect
- 2x 1Gb IPMI/management plane ports to rear optical interconnect

- Integrated 80mm fans reduce noise levels compared to 1U servers
- Blind mate rear IO simplifies sled replacement

The storage sled is shown in Figure 19. The internal block diagram showing power feed management, storage connections, and external network connections to the backplane connector is shown in Figure 20.

Figure 19. Storage Sled View witout Top Cover

Figure 20. Storage Sled Internal Connections

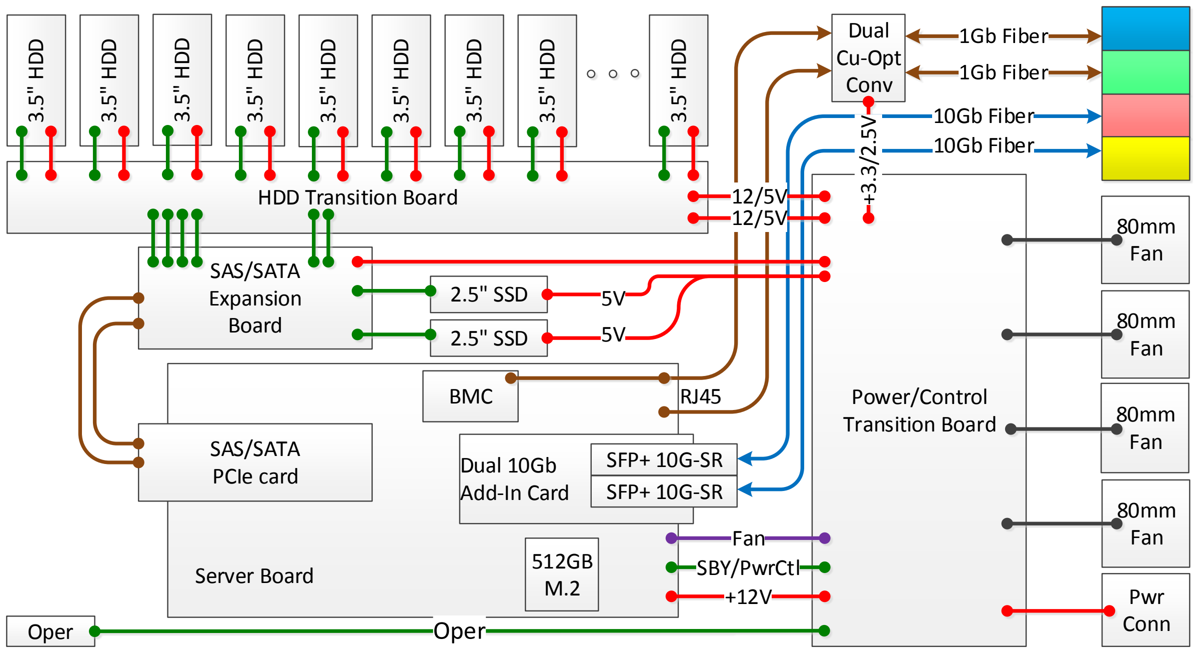

7.2.2 Serviceability

The storage sled can be removed and replaced quickly, without the need to power down the entire rack system.

Although it is recommended that internal storage sled components are replaced by a certified technician, the following components can be replaced:

- Server board
- Power board
- Media converter board
- Dual in-line memory module (DIMM)
- Solid state drive (SSD)
- Hard disk drive (HDD)
- Peripheral component interconnect express (PCIe) plug-in card
- Riser card
- M.2 flash boot drive
- Serial attached SCSCI (SAS) expander card

- Optical transceivers
- External connectors

7.2.2.1 Storage Sled Replacement Procedure

This section provides brief user-centric instructions on how to replace a storage sled, for illustrative purposes.

Remove sled:

Note: Sleds can be hot swapped; they do not need to be powered down for electrical or safety reasons before removal. However, sleds may be powered down by pushing the power button on the right side of the sled. The OK LED next to the power button indicates when the power is off.

1. Push the sled release latch toward the middle of the sled and pull the sled from the rack using the front handle.
2. Avoid hitting the connectors at the rear of the sled by ensuring the sled is supported when removing it from the shelf. Keep the sled level; do not allow the sled to drop or rotate in the rack.

Install replacement sled:

1. Remove the new sled from its packaging.
2. Remove the yellow protective cover.
3. Insert the replacement sled into the rack, using mechanical assistance from a lift if necessary. Make sure the sled is level and not twisted in the rack. The sled should slide easily, latching with an audible click.

7.3 Sled Mechanical Dimensions

7.3.1 Common Sled Dimensions

All DCEngine sleds and shelves comply with the CG-OpenRack-19 physical specifications. The significant shelf and sled dimensions are described here, but refer to the CG-OpenRack-19 document for all physical requirements and further details.

All sled shelves within the rack are set up in the rack at a vertical spacing of two Rack Units (RU), equivalent to 3.50 inches (88.9mm).

All sleds have a vertical height of 84.0mm ±0.20mm as measured from the bottom face to the top face and a horizontal depth of 780.0mm ±0.20mm as measured from the front face to the rear face.

Figure 21. Sled Vertical Height (front view)

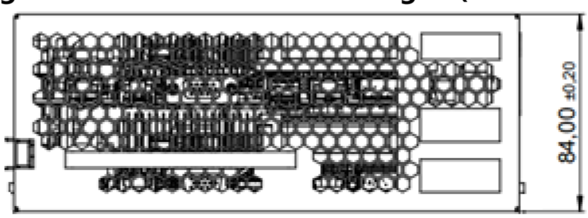

Figure 22. Sled Horizontal Depth (side view)

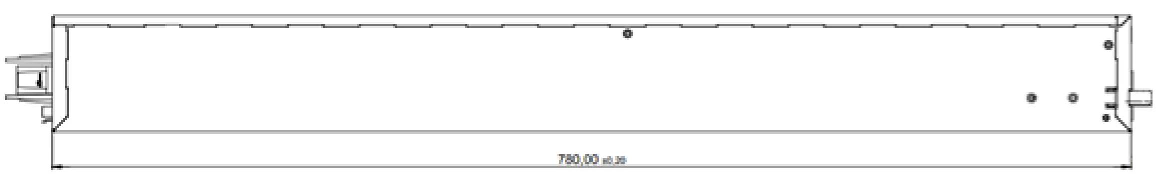

All sled front handles are centered 22.00mm from the bottom of the sled.

Figure 23. Sled handle vertical position (side view)

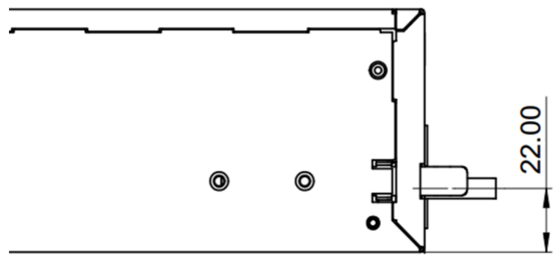

7.3.2 Half-Width Sled Dimensions

The half-width sleds share a common height and depth with full-width sleds. They occupy half of a shelf (left or right side), and fit two-across. They have a horizontal width of 215.0mm.

Figure 24. Half-width sled horizontal width (front view)

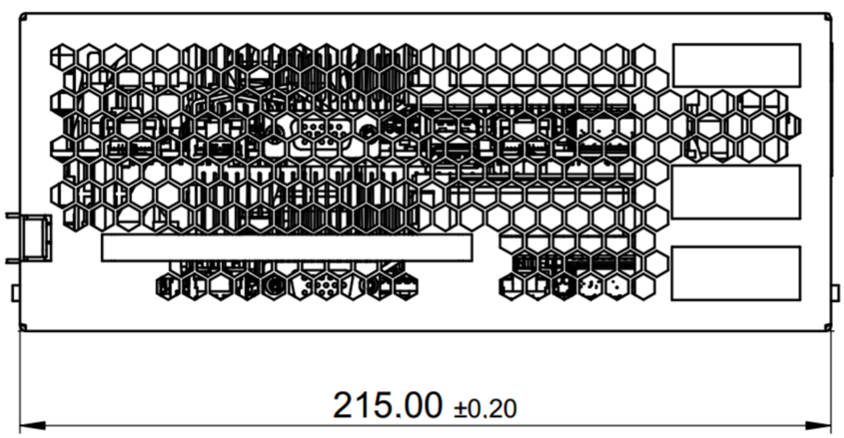

Half-width sled front handles are centered 144.25mm from the right side of the sled.

Figure 25. Half-width sled handle horizontal position (top-down view)

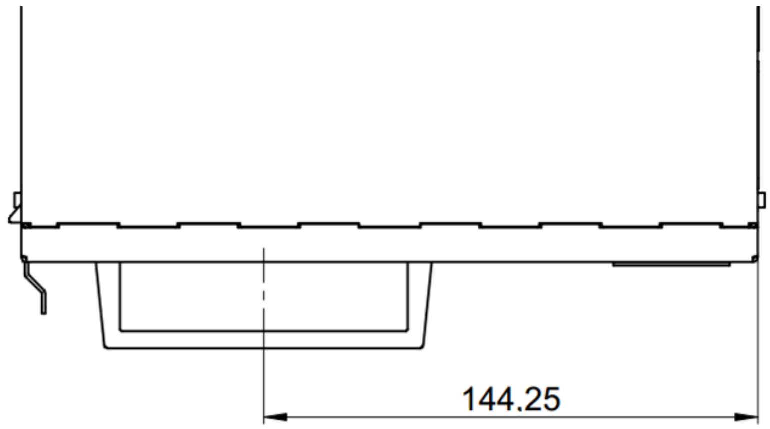

7.3.3 Full-Width Sled Dimensions

The full-width sleds share a common height and depth with half-width sleds. They occupy the full width of a shelf. They have a width of 444.00mm. The sled front handles are centered 143.25mm from the right side of the sled.

Figure 26. Full-width sled horizontal width (front view)

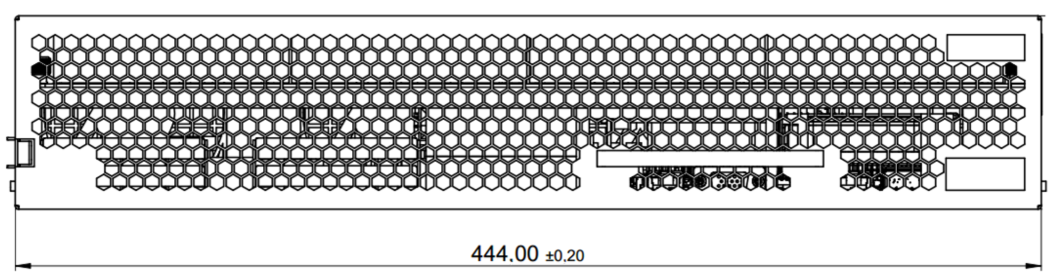

Figure 27. Full-width sled handle horizontal position (top-down view)

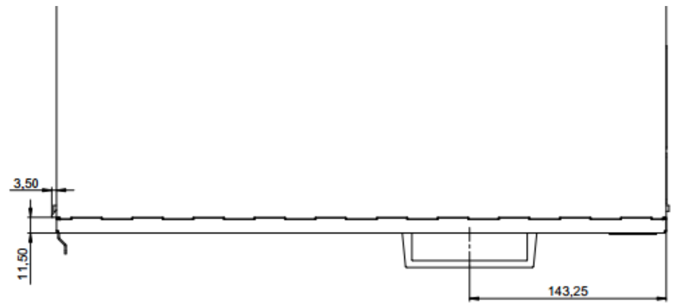

7.4 Sled Optical and Power Connections

The DCEngine rack system follows the CG-Openrack-19 specification in supporting an optical interconnect network (similar to a "backplane" function) connecting server and storage sleds to top of rack switches. This network consists of individual fiber optic interfaces multiplexed into multi-channel fiber cable connecting at the rear of the rack. The interface to each sled is a multi-channel blind-mate optical connector.

The following diagrams show how individual Compute and Storage sleds, which both have the same blind-mate optical connector interface, connect to rack switches (management and data plane) via the fiber optic blind-mate interconnect system. Note that other variations with different speeds and levels of multiplexing are possible. See also section 8 for a full system connectivity view.

Figure 28. End-to-End Management Plane Optical Connections

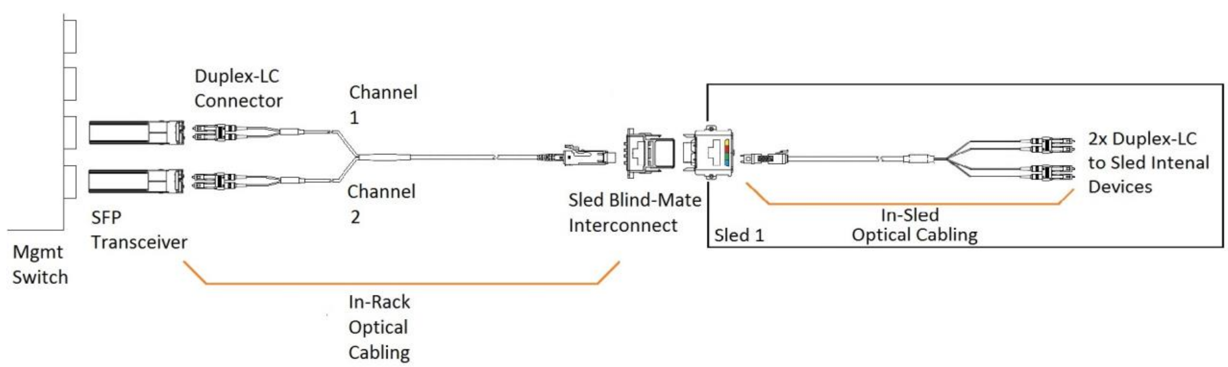

Figure 29. End-to-End Data Plane Optical Connections

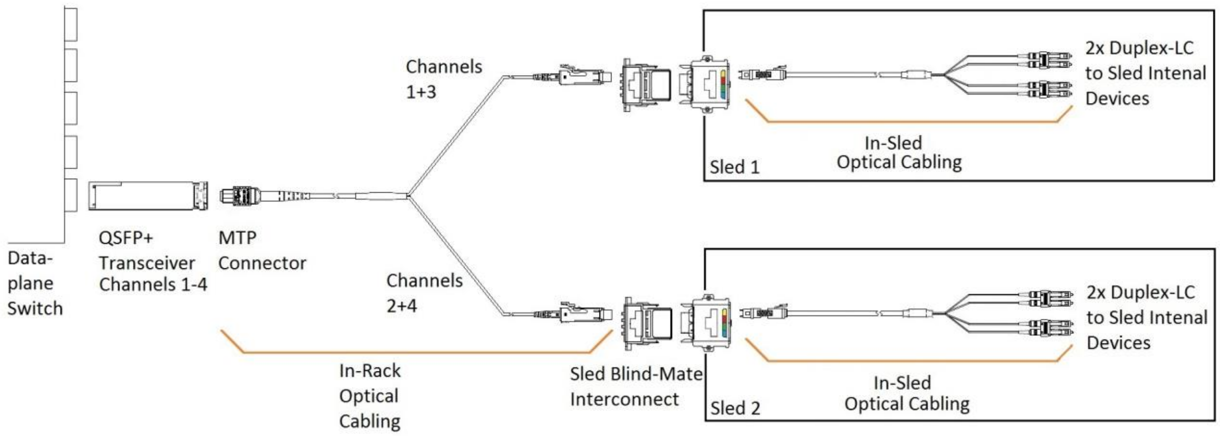

The physical placement of the power connector and blind-mate optical connector on the back of the sleds follows the CG-Openrack-19 specification. They are both located in the lower left side of the sled as shown in the following diagrams. The full-width Storage Sled has the same connector placement as the half-width Compute Sled except the sled itself extends to the full rack width; it does not add additional connectors.

Figure 30. Sled Power & Optical connector horizontal position (top-down view)

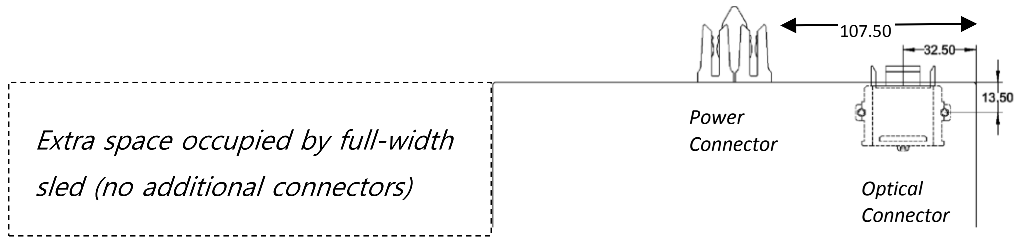

Figure 31. Sled optical connector vertical position (rear view)

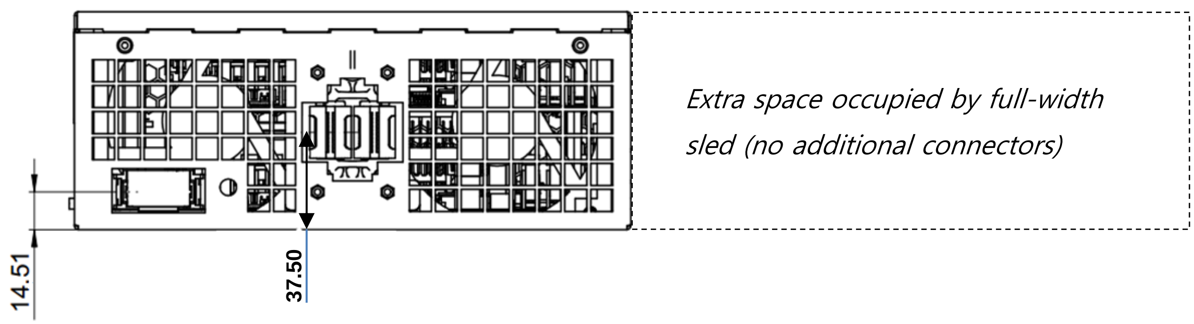

7.5 Mechanical Retention and Handle Features

All sleds include retention latches and alignment features as described in the CG-OpenRack-19 specification.

7.6 Sled Front Interface

The front panel of each sled follows the CG-OpenRack-19 recommendations. All sleds include one operator panel for each separately managed server node. The dimensions of each panel are 34mm width and 14mm height.

Figure 32. Sled Operator Panel Art and Dimensions

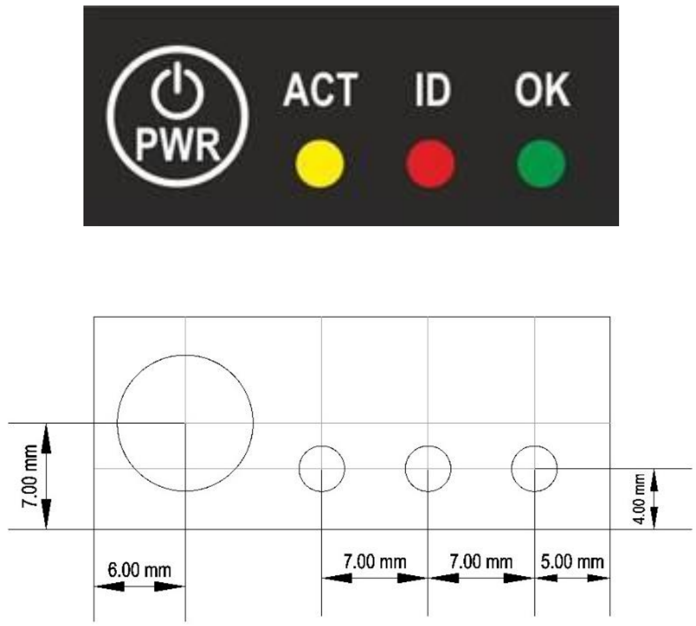

7.7 Sled Electrical Specifications

The following specifications apply to both Compute and Storage sled types.

7.7.1 Input Voltage, Power, and Current

Table 2 lists the nominal, maximum, and minimum values for the sled input voltage, in accordance with the CG-Openrack-19 specification.

The maximum and minimum voltages include the effects of connector temperature, age, noise/ripple, and dynamic loading.

Table 2. Input Voltage Range

Nominal voltage	Maximum voltage	Minimum voltage
12.25V DC	12.95V DC	11.65V DC

The maximum amount of power allowed per sled is defined during system power allocation. The number of sleds in a rack might be limited by the capacity of the AC power cord or by the cooling capacity of the deployment environment. Table 3 lists the input power allocation for a low-power sled (which may be either half- or full-width).

Table 3. Input Power Allocation and Maximum Current

Connector	Nominal voltage	System power allocation (in W)	Maximum current (in A)
Methode Mini-MCC Bus Connector	12.25VDC	1000W	105A

All sleds include over-current protection and in-rush current control from the 12V bus, with a linear ramp from 0A to the full load current at less than 20A/ms. All power and logic grounding is via 12V power return as well as through the sled enclosure rack mechanical connection as part of a chassis grounding method.

7.7.2 Electromagnetic Interference Mitigation

Sleds, as well as rack-level systems have been tested to meet FCC Class A limits as part of a fully loaded rack-level configuration, including power distribution, power conversion, switching networks, and a maximum population of sleds.

7.8 Sled Management

The “Compute Sled” has two Intel S2600TP Server Boards and the “Storage Sled” has a single Intel S2600TP board. The sled management is supported by several hardware and software components integrated on the server board and power supply board that work together to support the following.

- Control system functions – power system, ACPI, system reset control, system initialization, front panel interface, system event log

- Monitor various board and system sensors, regulate platform thermals and performance in order to maintain (when possible) server functionality in the event of component failure and/or environmentally stressed conditions
- Monitor and report system health
- Provide an interface for Server Management Software applications

Refer to Figure 18 and Figure 20 for block diagrams of each of the sled types (Compute and Storage, respectively) and also Figure 35 showing how the sled management entities are connected in the system. Also see Figure 16 in section 6.5 (Rack Management) which shows how sled level management relates to (optional) rack level management.

This section provides a high-level overview of the sled management features and functionality implemented on the server board and power supply board. The S2600TP Technical Product Specification provides a detailed description of the management features and the list of sensors supported by the server board.

Note that while the current node designs support an IPMI interface for the information described below, the CG-Openrack-19 spec and future DCEngine products are open to adoption of a native Redfish interface supporting similar functionality. This would conform to the forthcoming "OCP Hardware Management with Redfish" specification and remain compatible with the concept of an optional rack management agent (as described in section 6.5) while also normalizing and enhancing the node interface for usage by other levels of network management.

The integrated Baseboard Management Controller (BMC) firmware supports the following IPMI 2.0 features:

- IPMI Watchdog timer
- Messaging support, including command bridging and user/session support
- Chassis device functionality, including power/reset control and BIOS boot flags support
- Event receiver device: The BMC receives and processes events from other platform subsystems
- Field Replaceable Unit (FRU) inventory device functionality: The BMC supports access to system FRU devices using IPMI FRU commands

- System Event Log (SEL) device functionality: The BMC supports and provides access to a SEL including SEL Severity Tracking and the Extended SEL
- Sensor Data Record (SDR) repository device functionality: The BMC supports storage and access of system SDRs
- Sensor device and sensor scanning/monitoring: The BMC provides IPMI management of sensors. It polls sensors to monitor and report system health.
- IPMI interfaces
  - Host interfaces include system management software (SMS) with receive message queue support, and server management mode (SMM)
  - IPMB interface
  - LAN interface that supports the IPMI-over-LAN protocol (RMCP, RMCP+)
- Serial-over-LAN (SOL)
- ACPI state synchronization: The BMC tracks ACPI state changes that are provided by the BIOS.
- BMC self-test: The BMC performs initialization and run-time self-tests and makes results available to external entities.

The integrated Baseboard Management Controller (BMC) firmware also supports the following non-IPMI features:

- In-circuit BMC firmware update
- Fault resilient booting (FRB): FRB2 is supported by the watchdog timer functionality.
- Fan speed control with SDR
- Fan redundancy monitoring and support
- Enhancements to fan speed control.
- Power supply redundancy monitoring and support
- Acoustic management: Support for multiple fan profiles
- Signal testing support: The BMC provides test commands for setting and getting platform signal states.
- The BMC generates diagnostic beep codes for fault conditions.
- System GUID storage and retrieval
- Front panel management: The BMC controls the system status LED and chassis ID LED. It supports secure lockout of certain front panel functionality

and monitors button presses. The chassis ID LED is turned on using a front panel button or a command.

- Power state retention
- Power fault analysis
- Intel® Light-Guided Diagnostics
- Power unit management: Support for power unit sensor. The BMC handles power good dropout conditions.
- DIMM temperature monitoring: New sensors and improved acoustic management using closed-loop fan control algorithm taking into account DIMM temperature readings.
- Address Resolution Protocol (ARP): The BMC sends and responds to ARPs (supported on embedded NICs).
- Dynamic Host Configuration Protocol (DHCP): The BMC performs DHCP (supported on embedded NICs).
- Platform environment control interface (PECI) thermal management support
- E-mail alerting
- Support for embedded web server UI in Basic Manageability feature set.
- Enhancements to embedded web server
    - Human-readable SEL
    - Additional system configurability
    - Additional system monitoring capability
    - Enhanced on-line help
- Integrated KVM
- Enhancements to KVM redirection
    - Support for higher resolution
- Integrated Remote Media Redirection
- Lightweight Directory Access Protocol (LDAP) support
- Intel® Intelligent Power Node Manager support
- Embedded platform debug feature which allows capture of detailed data for later analysis.
- Provisioning and inventory enhancements:
    - Inventory data/system information export (partial SMBIOS table)
- DCMI 1.5 compliance.

- Management support for PMBus* rev1.2 compliant power supplies
- BMC Data Repository (Managed Data Region Feature)
- System Airflow Monitoring
- Exit Air Temperature Monitoring
- Ethernet Controller Thermal Monitoring
- Global Aggregate Temperature Margin Sensor
- Memory Thermal Management
- Smart Ride Through (SmaRT)/ Closed Loop System Throttling (CLST)
- BMC FW reliability enhancements:
  - Redundant BMC boot blocks to avoid possibility of a corrupted boot block resulting in a scenario that prevents a user from updating the BMC.
  - BMC System Management Health Monitoring

8 System Interconnect

Both compute and storage sleds have a rear-mounted blind-mate optical housing that supports data-plane and management plane connections to the top of rack switches, according to the CG-OpenRack-19 Rack and Sled specification.

8.1 System Interconnect Topology

DCEngine interconnection between sleds and switches follows the general concept and examples shown in the CG-OpenRack specification.

- Device Management Switch: 1GbE downlink management traffic interfaces (e.g., to BMCs) and 10GbE uplink interfaces.
- Application Management Switch: 1GbE downlink management traffic interfaces (e.g., to payload processors) and 10GbE uplink interfaces.
- Primary Dataplane Switch: 10bE downlink data plane traffic interfaces (e.g., to payload processors) and 40bGE uplink interfaces.

- Secondary Dataplane Switch: 10GbE downlink data plane traffic interfaces (e.g., to payload processors) and 40GbE uplink interfaces.

Note that the CG-Openrack-19 spec does not specify connection speeds, and higher speeds and different configurations are possible.

As shown in Figure 18, for each Compute sled there is:

- 4x10Gb data plane connections, 2 (primary + secondary) to each server
- 4x1Gb management connections, 2 (application + device) to each server

As shown in Figure 20, for each Storage sled there is:

- 2x10Gb data plane connections, primary + secondary, to the single server
- 2x1Gb management connections, application + device, to the single server

The first diagram below shows the basic connectivity concept and the next two show how the compute and storage sleds connect to the switches in one particular example system configuration with 8 storage sleds, 16 compute sleds, and 4 top of rack switches, using 10Gb connections for the data planes and 1Gb connections for the management planes.

Figure 33. Optical Interconnects between Sleds and Switches

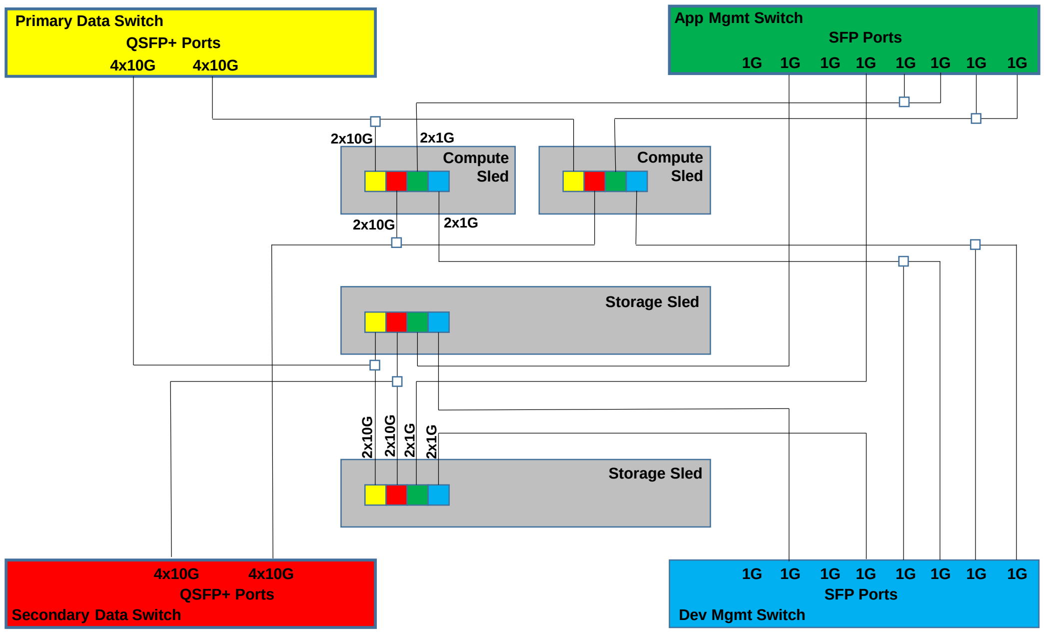

In the following two diagrams, the switches are shown on top, and the sleds below: 2 half-width compute sleds per shelf and 1 full-width storage sled per shelf.

Figure 34. Example Rack System Config: Data Plane Connections

Fiber Optic Cables between ToR switches and Sleds Rear View in Rack

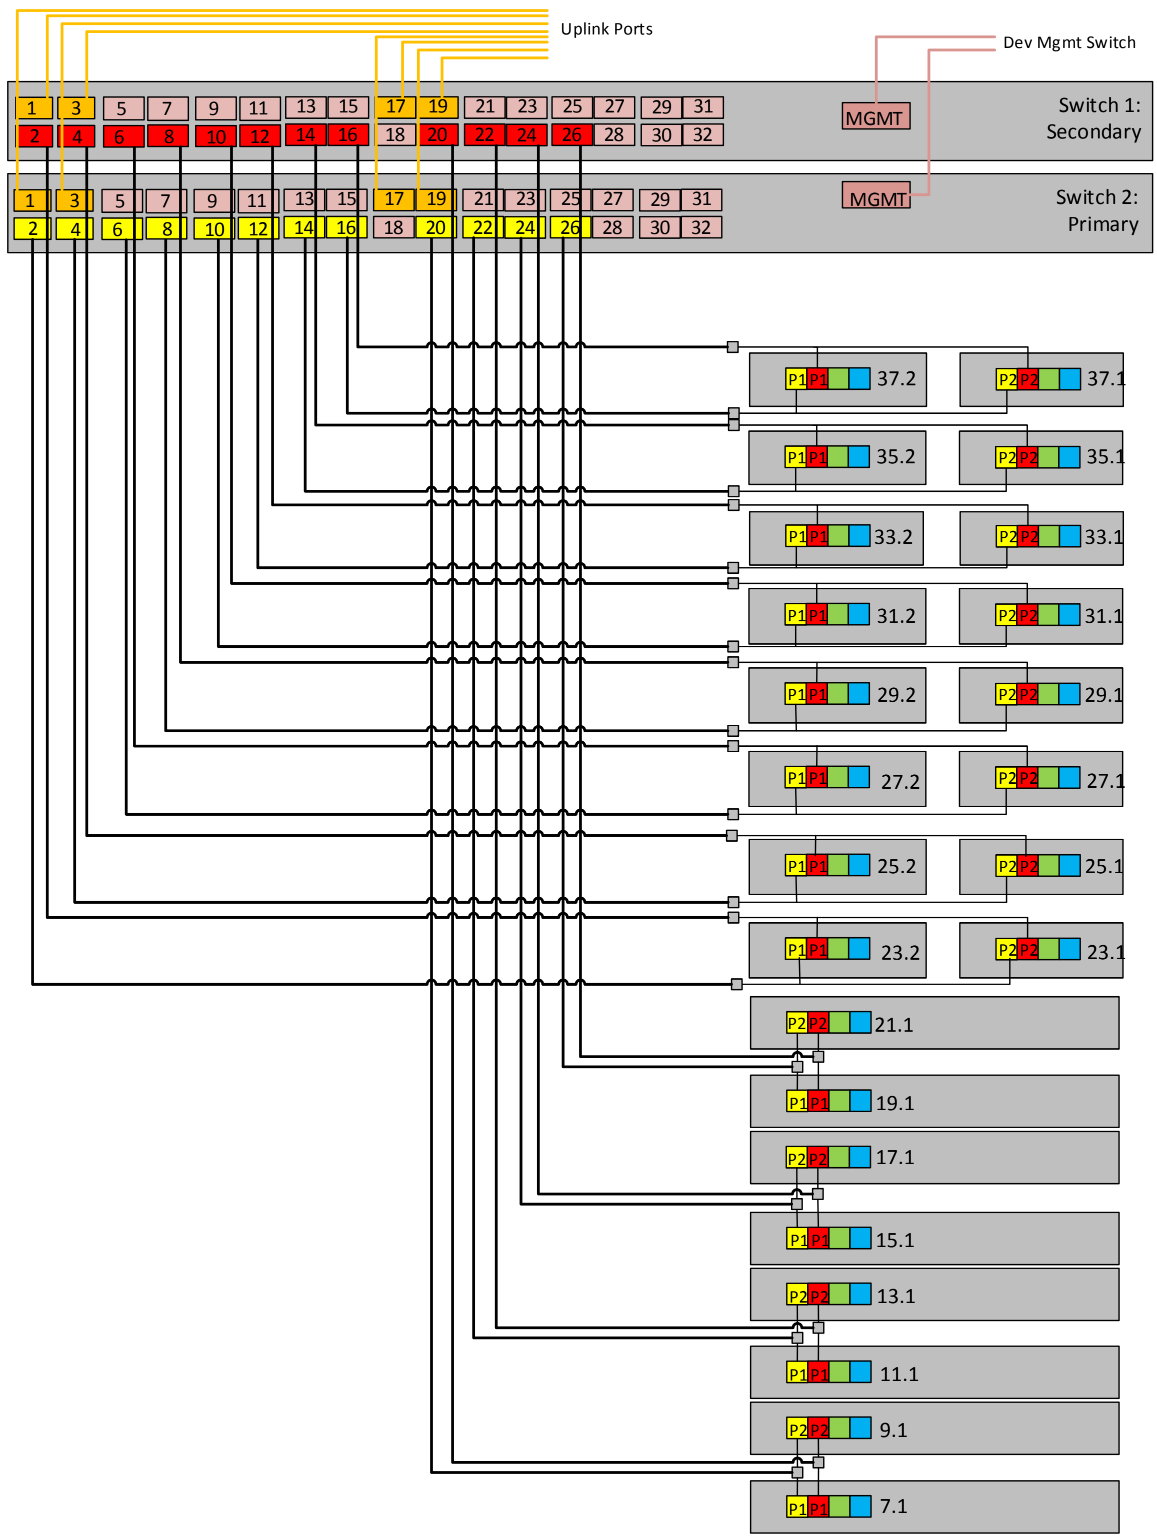

Figure 35. Example Rack System Config: Management Plane Connections

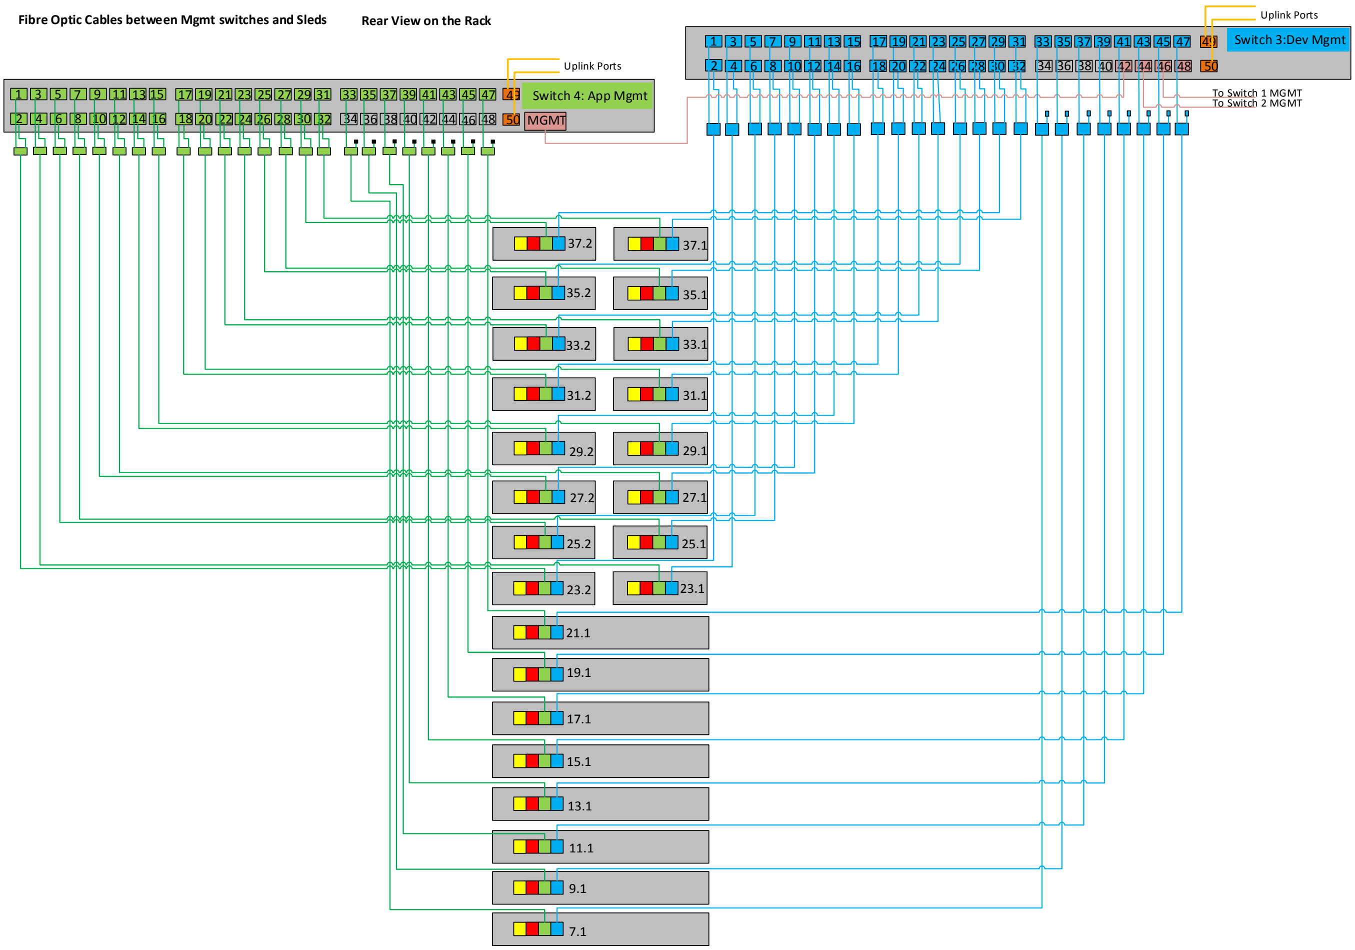

8.2 Sled Common Data Port Interfaces

Each DCEngine compute or storage shelf has two or one "backplane" housings (respectively) that will blind-mate with the "daughtercard" version of the housing that is part of the installed sled.

All interfaces between server nodes and the rack infrastructure pass through the blind-mate optical connector at the rear of the compute and storage sleds.

The housing has four sections, each of which is capable of holding a single fiber optic ferrule. The sections are color-coded.

Figure 36. Sled blind-mate section assignments (rear view)

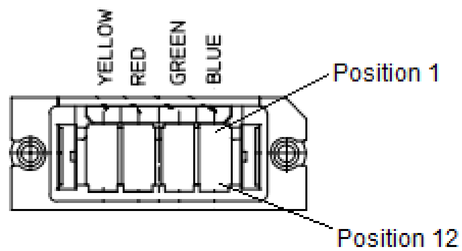

DCEngine compute and storage sleds follow the CG-OpenRack-19 specification suggestions:

Green → Application Management

Blue → Device Management

Yellow → Primary Data-plane

Red → Secondary Data-plane

8.3 Sled to Shelf Interconnects

The fiber positions in the rear connector match the defined positions of the shelf connector. Sleds have a rear-mounted blind-mate "daughtercard" optical housing that mate with a compatible "backplane" housing on the rear of the DCEngine shelves.

The optical blind-mate connector housing supports four ferrules. The typical number of fiber positions per ferrule in DCEngine is 12, with 8 fibers are loaded on each ferrule. The ferrule positions are color-coded, with each color corresponding to one of four network types.

The sled to shelf interconnects follow the CG-OpenRack-19 specification.

Table 4. Sled Interconnect Components

Assembly mounting	Connector description	Manufacturing part number
Sled	HBMT daughtercard housing	(Molex) 106105-2100
Sled	Power rail connector	(Methode) 5313-07458-00107

Table 5. Shelf Interconnect Components

Assembly mounting	Connector description	Manufacturing part number
Shelf	HBMT backplane housing	(Molex) 106105-2000

8.4 Sled Data Connections

Each sled connects to the rack-level networks via a set of up to four optical ferules in a blind-mate housing system. This connector system is positioned at the rear face of each sled and at the back of each shelf. When the sled is installed on the shelf, all of the ferules automatically align and connect, which ultimately connects the on-sled network interfaces to the rack-level network switches.

The optical blind-mate system and the in-rack optical cables provide the end-to-end network connections to all installed sleds.

Figure 37. Sled Blind-Mate Optical Connections, Front View

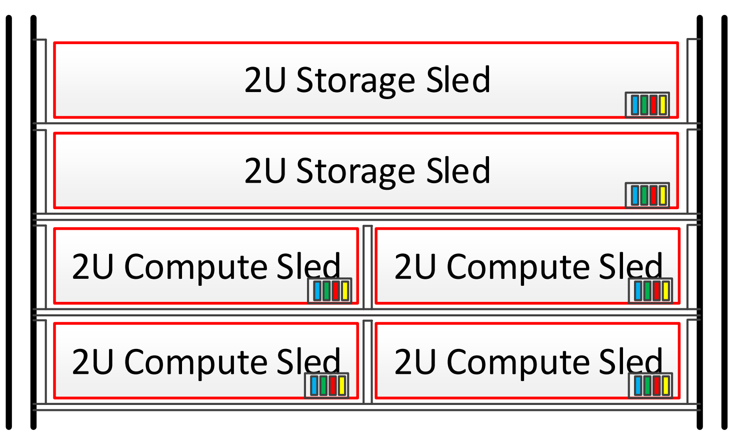

8.4.1 Signal Definitions

Table 6 below shows the fiber naming and assignments for the blind-mate optical interconnect.

Table 6. Optical Connector Asignments

Signal	Housing Section	Fiber Location	Description	Example Use (Switch Connections)
ENET_SW1.1_TX	Blue	1	1GbE (1000BASE-SX) lane	1Gb Device Management Plane connection (BMC IPMI link)
ENET_SW1.1_RX	Blue	12		
ENET_SW1.2_TX	Blue	2	1GbE (1000BASE-SX) lane	1Gb Device Management Plane connection (BMC IPMI link)
ENET_SW1.2_RX	Blue	11		
ENET_SW2.1_TX	Green	1	1GbE (1000BASE-SX) lane	1Gb Application Management Plane connection
ENET_SW2.1_RX	Green	12		
ENET_SW2.2_TX	Green	2	1GbE (1000BASE-SX) lane	1Gb Application Management Plane connection
ENET_SW2.2_RX	Green	11		
ENET_SW3.1_TX	Yellow	1	10GbE (10GBASE-SR) lane	10Gb Primary Dataplane connection
ENET_SW3.1_RX	Yellow	12		
ENET_SW3.2_TX	Yellow	2	10GbE (10GBASE-SR) lane	10Gb Primary Dataplane connection
ENET_SW3.2_RX	Yellow	11		
ENET_SW4.1_TX	Red	1	10GbE (10GBASE-SR) lane	10Gb Secondary Dataplane connection
ENET_SW4.1_RX	Red	12		
ENET_SW4.2_TX	Red	2	10GbE (10GBASE-SR) lane	10Gb Secondary Dataplane connection
ENET_SW4.2_RX	Red	11		

Note: “Tx” and “Rx” are as referenced from the rack-level switch ports

Note: The fiber signals are presently defined at 1Gb/s and 10Gb/s speeds; future speed increases are possible.

Radisys Corporation

www.radisys.com

5435 NE Dawson Creek Drive

Hillsboro, OR 97124