OSS High Density Compute Accelerator Implementation and Usage Guide

ONE STOP SYSTEMS

www.onestopsystems.com

Table of Contents

Revisions:

1st Publication: April 11, 2014

1. Overview

About the HDCA

The **High Performance Compute Accelerator (HDCA)** is a 3U rackmount appliance that attaches to one to four servers, adding thousands of cores of computational performance to a variety of HPC applications. The appliance contains up to 16 high-performance GPUs or coprocessors in four removable canisters. The system supports up to four PCIe cable connections operating at 128Gb/s bandwidth each. A single server can operate all 16 boards or the canisters can be divided between multiple servers.

The **HDCA** is a modular system that is easy to install with three basic parts: rackmountable chassis, canisters with GPUs, and power supplies. The all-steel construction 3U rackmount chassis houses three 3,000-watt redundant power supplies, four canisters with up to four boards each, and an IPMI-based system management monitor. The IPMI system monitor tracks chassis resources including an add-in card SM bus and provides monitoring, control and alarming features. It supports command line or web interface (GUI) and SNMP v2c and v3 protocols as well as RMCP+ protocol.

GPUs and Coprocessors

Graphics Processing Units (GPUs) are used in numerous defense, intelligence, financial, medical and seismic applications today and the number is rapidly growing. The need to get the huge amounts of data transcribed and made useful through data and image processing is becoming overwhelming. The more GPUs available, the quicker the data can be used. GPU appliances supporting multiple NVIDIA® GPUs and Intel® Phi™ coprocessors, like the HDCA are quickly becoming the best and most economical way of accomplishing this tremendous feat.

NVIDIA GPUs

NVIDIA Tesla® K-Series GPU Accelerators are based on the NVIDIA Kepler™ computer architecture and powered by CUDA®, the world's most pervasive parallel computing model. They include innovative technologies like Dynamic Parallelism and Hyper-Q to boost performance as well as power efficiency and deliver record application speeds for seismic processing, biochemistry simulations, weather and climate modeling, image, video and signal processing, computational finance, computational physics, CAE, CFD and data analytics.

NVIDIA's Kepler-based GRID™ boards are specifically designed to enable rich graphics in virtualized environments. GRID boards have an optimized multi-GPU design that allows hardware virtualization of the GPU. This means multiple users can share a single GPU, improving user density while providing true PC performance and compatibility.

Intel Phi Coprocessors

Intel Xeon Phi coprocessors provide up to 61 cores, 244 threads, and 1.2 teraflops of performance, and they come in a variety of configurations to address diverse hardware, software, workload, and performance and efficiency requirements.

PCIe Connectivity to Servers

The HDCA has four PCIe x16 Gen3 connections available at the rear of the chassis to support up to four host servers. As the system powers on it automatically selects the number of servers attached and maps the GPUs to the appropriate server connections.

See Appendix 1 and the One Stop Systems' website for the latest list of tested and supported servers. http://www.onestopsystems.com/documents/1u-2u-3u-nvidia-ca-tested-server-matrix.pdf

1.1 Safety Instructions

Please read this section carefully and observe the instructions for your own safety and correct use of the device.

⚠ Observe the warnings and instructions on the device and in the manual.

One Stop Systems products leave the factory in a perfectly safe condition. In order to maintain this condition and ensure safe operation, the user must observe the instructions and warnings contained in this manual.

- The device must be used in accordance with the instructions for use.
- The equipment must be installed in accordance with the National Standards and/or Electrical Codes of the country in question.
- The electrical installations in the room must correspond to the requirements of the respective regulations.
- Take care that there are no cables, particularly mains cables, in areas where persons can trip over them.
- Use only the cable supplied by One Stop Systems.
- Do not set up the device where it is exposed to direct sunlight, in the proximity of heat sources or in a damp location. Make sure the device has adequate ventilation.
- Only devices and components may be connected to the interfaces of the system, which fulfill the requirements of a SELV circuit (Safety extra Low Voltage) in accordance with EN60950.
- It must be assumed that safe operation is no longer possible:
    - If the device has visible damage.
    - If the device no longer functions.
- In these cases the device must be shut down and secured against unintentional operation.
- If extensions are made to the device the legal stipulations and the device specifications must be observed.
- Assembly, extensions, new settings, alterations or repairs may be carried out only as authorized by One Stop Systems.
- Adjustment, maintenance or repair on the open device may be carried out only as authorized by One Stop Systems.
- Only original accessories approved by One Stop Systems may be used.

- Top cover must always be properly installed before power is applied.
- The inside of the enclosure is considered a “Service Access Area.”

 Warning! Hazardous voltages exist.

1.2. Unpacking Instructions

1. If the carton is damaged when you receive it, request that the carrier's agent be present when you unpack and inspect the equipment.
2. After unpacking, verify that all items listed in the packing list are present.
3. Observe all proper ESD safety procedures.
4. Inspect the equipment for shipping damage.
5. Save all packing material for storage or return shipment of the equipment.
6. For repairs or replacement of equipment damaged during shipment, contact One Stop Systems, Inc. to obtain a Return Materials Authorization (RMA) number and further shipping instructions.

2. System Components

2.1. Enclosure

The HDCA enclosure is a 3U rack mountable chassis that is all steel construction. The chassis is 38” deep. The enclosure contains four accelerator canisters and three 3,000-watt power supplies. The HDCA supports up to sixteen double-wide PCIe cards. It also contains an IPMI 2.0 system monitor and alarming module.

Figure 2.1 Chassis (with power supplies)

The top cover of the chassis is attached with standard screws. The chassis has a snap-on front bezel with a replaceable air filter and status LED’s on the chassis which provides alarm data and chassis identification. LEDs are bi-color: amber and green. There are four power switches, one on each canister and four rear-mounted fans attached with standard screws.

The HDCA can be installed into a rack forward mounted, with the bezel (front) facing forward, or reversed, with the PCIe connections (rear) facing forward. The two configurations require different internal parts which will be installed prior to shipment depending on each customer’s configuration requirements. For the purposes of this manual, when the front is referenced, this will always mean the part of the chassis with the bezel. When the rear is referenced, this will always mean the part of the chassis with the PCIe cable connectors.

The enclosure base also houses the main system backplane. This backplane is a fixed part of the enclosure and not a user serviceable item. It contains the four PCIe x16 Cable connectors to the host system(s), three PCIe 3.0 switches that automatically configure the 16 accelerator slots depending on what cables are plugged into the system and four PCIe x16 connections to the four accelerator canisters. The main backplane also serves as the power distribution system to the canisters, mechanical alignment for the canisters and contains the system monitor modules. See the Appendix for more information and diagrams of the main backplane.

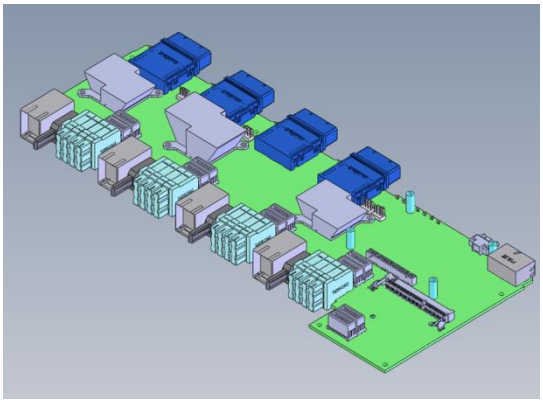

Figure 2.2 Main Backplane

Table 2.1 Specifications

Enclosure	• Dimensions: 17”W x 5.25” H x 38”D • Supports 16 full-length, full-height, 2-slot PCIe x16 GPU/PHIs • All 16 boards face the rear of the chassis (no IO bracket access) • Removable front bezel with air filter • Front panel LEDs • Four rear panel PCIe x16 Gen3 cable interfaces • 4 individually-removable rear fans and 4 canister mounted fans • Weight: 92lbs when fully loaded with 16 accelerators
Main backplane	• Four PCIe x16 cable inputs to rear of enclosure • Four PCIe x16 high-density connectors to each canister • 1x PLX PEX 8796 and 2 x PLX PEX 8749 PCIe 3.0 switches manage PCIe cross connects from cables to canisters • 2x RJ45 connectors for IPMI v2.0 System Monitor • 1x HD DB-9 serial port for IPMI network configuration • Optional RJ45 for basic SYSMON2 chassis monitor (not required when using IPMI System Monitor) • Supports bus-bar power distribution to the canisters through 8 high-power bladed connectors (2 per canister) • On board IPMI System Monitor & SYSMON2 connectors
Canister backplane	• 4x PCIe 3.0 x16 double-spaced slots in 2 ranks of 2 GPUs each • 4x 8-Pin 12V power connectors for GPU/PHI AUX power cables • PLX PEX 8796 PCIe 3.0 switch
Power	• 6000W redundant power subsystem • Three 3U 3,000-watt front removable, hot-swap supplies • Each supply measures 1U (1.65”) x 2.7” x 28.5” • 2+1 redundant with full current sharing operation • 3,000W each at 208-277VAC, 15A max input • 1,500W each at 90-124VAC, 15A max input • 20A breaker and IEC C14 power input at rear for each supply • +12V and +12V standby voltage outputs • All +12V power rails shared on copper bus bar delivery system
Power Cords	• 110V power cord for PDUs and Wall receptacles o OSS Part number: OSS-CBL-PWR-5-15-C13-15A-6 o NEMA 5-15 to IEC C13, Straight, 14AWG, 15A, 6' • 240V power cord for PDUs o OSS Part number: OSS-CBL-PWR-C14-C13-15A-6 o IEC C14 to IEC C13, Straight, 14AWG, 15A, 6' • 240V power cord for US Wall receptacles o OSS Part number: OSS-CBL-PWR-6-15-C13-15A-6 o NEMA 6-15 to IEC C13, Straight, 14AWG, 15A, 6'
System monitoring/alarming	Fully IPMI v2.0 compliant monitoring, control & alarming system *Temperature* • Monitors inlet & exhaust temps • Fan speed auto adjusts by temp • Alarm set-points for over temp *Fans*

	• Monitors all system fan tachs • PWM fan speed control • Alarms for slow or failed fans *Power* • Monitors supply telemetry • Monitors output voltage rails • Alarms for voltages out of range • Alarms for supply failure *Add-in Cards* • Monitors add-in card I2C SM bus • Alarms for abnormal card telemetry *Interface* • CLI or web GUI • Supports SNMP and RCMP • Remote chassis and canister LED tagging
Air filter	• 30 ppi open cell polyfoam • Die-cut, removable and replaceable
Cooling	• Four 80 x 80 x 38mm fans on the rear of the enclosure • One 80 x 80 x 38mm fan on the front of each canister • All fans are 143CFM each in push-pull configuration • All fans PWM monitored and speed controlled by the IPMI system monitor • Rear fans hot-swap from rear of the chassis • Power supplies separately cooled from internal 25mm fans

Operating environment	• Temperature range:
	  • Operating: 10°—35°C
	  • Storage: -40°—85°C
	• Humidity range:
	  • Operating: 20% to 80% relative (non-condensing)
	  • Non-operating: 5% to 95% relative (non-condensing)
	• Altitude range:
	  • Operating: 0 to 10,000 ft.
	  • Storage: 0 to 50,000 ft.

Agency compliance	Designed to meet the following agency certifications with testing currently pending: • FCC - Part 15 of the FCC Rules, Class A, 47CFR • Canada ICES-003, issue 4, Class A • UL/IEC 60950-1 • Canada: CSA C22.2 No. 60950-1 • Argentina: IEC60950-1 • Japan: VCCI, Class A • Australia/New Zealand AS/NZS CISPR 22, Class A • IEC 60950-1 (CB Certificate and CB Test Report) • CE Mark (EN55022 Class A, EN60950-1, EN55024, EN61000-3-2, EN61000-3-3) • CISPR 22, CISPR 24, Class A • CE Emissions 2004-108EC • RoHS compliance (Directive 2002/95/EC) • CCN NWGQ, NWGQ7

2.2 Power Bus Bar

The solid copper alloy bus bars provide a direct link from the power system to the main system backplane. From there, the 12V DC power output and ground is distributed directly to where the power is needed using minimal wiring. The reduction in wiring allows fewer blockages and the most efficient cooling of high-powered accelerators.

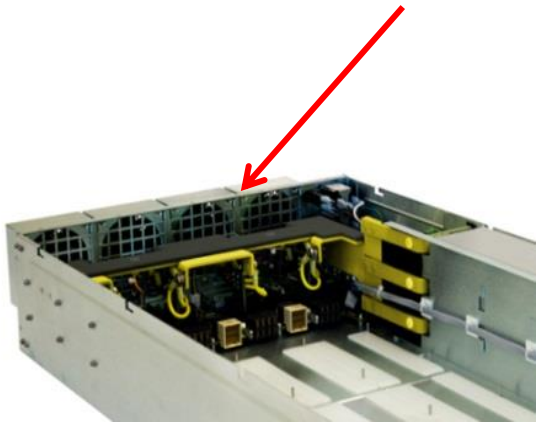

Figure 2.3 Power Bus Bar

There is a power drop from the bus bars to the main backplane directly behind each canister power connector thereby minimizing any voltage drop to each canister. The bus bars are covered with an UL approved coating, separated by a non-conductive gap material, secured with nylon washers, inserts and fasteners and painted the common yellow and black, representing 12V and ground.

2.3 Power Supplies

Figure 2.4 Power Supply

Depending on the configuration of HDCA ordered, the HDCA contains either two or three 1U x 2.7" x 28.5" 3,000-watt power supplies. The power supplies are fully load sharing and hot-swappable when removed one at a time. Systems that ship with 8 or less accelerators installed can operate with two power supplies in 1+1 redundant mode. Three redundant power supplies provide ample power for nine to sixteen GPUs or coprocessors. Due to continuous load sharing of the supplies under normal 3-supply operations, even under the 6000W maximum power capacity of the HDCA, each supply is never under more than 66% load thereby prolonging the life of each supply.

Since the HDCA can be mounted in forward or reverse airflow directions, there are two possible SKUs for additional or replacement supplies that have the proper airflow. When the bezel side is facing the "cold aisle" air intake, the standard airflow power supply is required. If the PCIe Cable connectors are facing the "cold aisle" air intake, the reverse airflow power supply is required.

HDCA Power Supply Options:
OSS-PS-3000W-AC – Standard air flow
OSS-PS-3000W-AC-R – Reverse air flow

Each power supply has a separate IEC 60320 C14 filtered power input connector on the rear of the system. Each power input has a separate circuit breaker for system protection. The appropriate 15A 14AWG minimum power cord should be used depending on the system installation requirements of the HDCA including the voltage supplied to the system, wall power outlet, power strip or rack Power Distribution Unit (PDU). If the installation requires 110V nominal input power, each power supply will be able to deliver 1500W maximum to the HDCA. This configuration can supply up to 3000W total in 2+1 redundant mode or supply up to 4500W total maximum power with no supply redundancy. In 208-

277V operation, the max power output is limited to 6000W so a three supply HDCA system will always operate in 2+1 redundant mode. Due to the myriad of power cable needs in various installations around the world, customers can use locally sourced power cables or order an appropriate power cable from OSS using the table of common power cables below.

Table 2.2 OSS Power Cables

OSS Part Number	Description
OSS-CBL-PWR-5-15-C13-15A-6	110V US Power wall cable, NEMA 5-15 to IEC C13, Straight, 14 AWG, 15A, 6'
OSS-CBL-PWR-C14-C13-15A-6	240V Power PDU cable, IEC C4 to IEC C13, Straight, 14 AWG, 15A, 6'
OSS-CBL-PWR-6-15-C13-15A-6	240V US Power wall cable, NEMA 6-15 to IEC C13, Straight, 14 AWG, 15A, 6'
Other configurations or lengths	Call OSS

Table 2.3 HDCA Power Requirements

Configuration	HDCA Power Requirement Estimate (Max @240V Input)
Base configuration (no GPUs)	740W
HDCA with 16 K20 GPUs (225W max)	4600W
HDCA with 16 K20X/K40 GPUs (235W max)	4775W
HDCA with 16 Intel Phi 7120P (300W max)	5880W

The power system of the HDCA is intended to always be 'on' or in 'standby' mode waiting for the host system to power on and initialize the accelerators in the HDCA. Once on and connected to one or more operating host systems, turning off power to the HDCA or a canister in operation is likely to cause a host system operating system fault as accelerators in the HDCA operate the same as if they were plugged into a PCIe slot in a host system. Therefore, there is no main power switch for the HDCA. All 3 breakers can be switched off or all three power cables can be removed from the HDCA to cut all power to the expansion system. Alternatively, all 4 power switches on the front of the canisters can be turned to the off position which will inhibit DC power to the canisters with standby power remaining available to the system. See section 3 for information on setting up the HDCA for first use or after a maintenance cycle.

2.4 Canisters

The four canisters in the HDCA contain a rear-mounted PCIe signal connector, canister alignment receptacle and an 8 bladed power connector. The canister backplane consists of four GPU power auxiliary connectors and four x16 PCIe slots (up to 350W each) spaced for double-slot accelerator add-in boards. Each canister has a ring handle ejector, DC inhibit power switch and one dedicated front mounted 80mm x 38mm high output fan. Each canister has a six LED display visible from with the front bezel off or via light pipes on the front filtered bezel when it is installed. See section 4.1 for more information about canister LEDs. The canisters are designed for passive accelerators available from OSS with special brackets and power harnesses. Diagrams and connector

Figure 2.5 Canister

information on the canister can be found in the Appendix.

2.5 IPMI System Monitor

The Intelligent Platform Management Interface (IPMI) is a standardized computer system interface used by system administrators for out-of-band management of computer systems and monitoring of their operation. It is a way to manage a computer that may be powered off or otherwise unresponsive by using a network connection to the hardware rather than to an operating system or login shell. The HDCA has a fully IPMI version 2.0 compliant system monitor installed that helps the user monitor and control system resources such as fans as well as perform alarming functions for system failures. The IPMI System Monitor also provides advanced datacenter features such as remote "tagging" of an HDCA making finding the chassis as easy as searching for the lit blue LED on the front or rear of the chassis marked for maintenance. Other datacenter features include external interface to higher level protocols such as SNMP, three user settable alarm levels, external alarm notifications via email and SMS and tools to detect out of the norm fan and power supply operation to fix potential issues before they become an critical.

Because IPMI is an open standard and there are hundreds of commands available to the user, OSS has pre-set the most common operational parameters as defaults and provides a separate IPMI System Monitor manual with all of the available monitoring, alarming and control commands for the user to customize. These commands are available using a simple command line interface (CLI) or a web browser interface included with the HDCA system. Once a configuration is determined by the user to be optimal, OSS can provide a custom HDCA part number for the user with the user settings pre-loaded into the IPMI System Monitor. The base IPMI manual can be found on our web site at: http://onestopsystems.com/documents/ShelfManagerUG_3_2_0_20130731.pdf.

Alternatively, the user can interface to the IPMI System Monitor using a higher level protocol such as SNMP or RCMP used to monitor and control a large number of systems in an entire datacenter such as servers and multiple HDCAs. A separate external interface specification manual is available for the IPMI System Monitor at our web site at:
http://onestopsystems.com/documents/ShelfManagerEIR_3_2_0_20130531.pdf

2.6 NVIDIA Tesla and Grid K-Series GPU's

The HDCA contains up to 16 K-Series GPU Accelerators supplied with the system, depending on the SKU ordered. Available GPU Options are Tesla K10, K20, K20X or K40 and Grid K1 or K2. Since the HDCA provides expansion to standard PCIe 3.0 x16 slots, the number and type of Tesla or Grid GPUs can be mixed and matched in the system depending on the application requirements. For special configurations mixing various GPUs, please contact OSS sales for a configured part number. All GPUs installed in the system are fully tested, supported by OSS and nVidia and have the appropriate brackets, firmware, wiring harnesses and VBIOS to support this high density system.

Figure 2.6 Four NVIDIA GPUs installed in a canister

Tesla GPUs have two DMA engines for bi-directional PCIe communication. They also have large on-board memory. For more information on the Tesla GPU product line see the Tesla table below or you can access the datasheet at: http://www.nvidia.com/content/tesla/pdf/NVIDIA-Tesla-Kepler-Family-Datasheet.pdf.

GRID GPUs are specifically designed to enable rich graphics in virtualized environments. GRID boards feature the NVIDIA Kepler architecture that allows hardware virtualization of the GPU. This means multiple users can share a single GPU, improving user density while providing true PC performance and compatibility. For more information on the GRID GPU product line see the GRID table below or you can access the datasheet at: http://www.nvidia.com/content/grid/resources/8556_NVIDIA_GRID_DS_Feb14_US_HR.pdf.

Table 2.4 NVIDIA Tesla K-Series GPU Specifications

	Peak Double Precision Performance	Peak Single Precision Performance	# of GPUs	# of CUDA Cores	Memory Capacity	Memory Bandwidth	GPU Computing Applications	Architecture Features
Tesla K10	0.19 teraflops	4.58 teraflops	2 x GK104s	2 x 1536	8GB	320GB/s	Seismic, image, signal	SMX
Tesla K20	1.17 teraflops	3.52 teraflops	1 x GK110	2496	5GB	208GB/s	CFD, CAE, Financial computing, Computational Processing, video analytics, chemistry and physics, data analytics, satellite imaging, weather modeling	SMX, Dynamic Parallelism, Hyper-Q
Tesla K20X	1.31 teraflops	3.95 teraflops	1 x GK110	2688	6GB	250GB/s		
Tesla K40	1.4 teraflops	4.0 teraflops	1 x GK110B	2880	12GB	288GB/s		

Table 2.5 NVIDIA GRID GPU Specifications

	# of GPUs	# of CUDA Cores	Memory Capacity	Memory Clock	GPU Virtualization Applications	Architecture Features
GRID K1	4 x GK107s 850MHz	768	16GB	891 MHz	Maximum number of concurrent users	H.264 Encoding, SMX
GRID K2	2 x GK104s	3072	8GB	2.5GHz	Maximum user density for graphic-intensive applications	H.264 Encoding, SMX

2.7 Intel Xeon Phi Coprocessors

The HDCA also supports up to sixteen Intel Xeon Phi Coprocessors including the 3100 family, 5100 family and 7100 family. OSS supplies the passive versions of the required coprocessor depending on the system SKU ordered. The number and type of Phi coprocessors can be mixed and matched in the system depending on the application requirements. Phis and GPUs can also be mixed in the HDCA. Contact OSS sales for special configured product part numbers to meet your acceleration needs.

Figure 2.7 Intel Xeon Phi Coprocessors

Intel Xeon Phi coprocessors provide up to 61 cores, 244 threads, and 1.2 teraflops of performance, and they come in a variety of configurations to address diverse hardware, software, workload, performance, and efficiency requirements. The Phi 3100 family provides outstanding parallel performance. It is an excellent choice for compute-bound workloads, such as MonteCarlo, Black-Scholes, HPL, LifeSc, and many others. The Phi 5100 family is optimized for high-density computing and is well-suited for workloads that are memory-bandwidth bound, such as STREAM, memory-capacity bound, such as ray-tracing, or both, such as reverse time migration (RTM). These coprocessors have the lowest thermal design power (TDP) of the Intel Xeon Phi product family. The Phi 7100 family provides the most features and the highest performance and memory capacity of the Intel Xeon Phi product family. This family supports Intel® Turbo boost Technology 1.0, which increases core frequencies during peak workloads.

Table 2.6 Intel Xeon Phi Coprocessor Specifications

Model	Peak Double Precision Performance	Board TDP	# of Cores	Core Freq.	Memory Capacity	Memory Bandwidth	Computing Applications	Architecture Features
3120P	1.003 teraflops	300W	57	1.1 GHz	6GB	240GB/s	Monte Carlo, Black-Scholes, HPL, LifeSC	Maximum Value
5110P	1.011 teraflops	225W	60	1.05 GHz	8GB	320GB/s	STREAM, ray-tracing, RTM	Power Efficient
7120P	1.208 teraflops	300W	61	1.24 GHz	16GB	352GB/s	Seismic Imaging Processing, Molecular Dynamics, WRF	Intel Turbo Boost 1.0 1.33GHz

2.8 Cooling

The HDCA has eight 80 x 36mm, 143CFM, high-output fans, one on the front of each canister and four on the rear of the base enclosure. The fans act in concert to both push air into the HDCA and pull air through a densely populated system thereby allowing up to sixteen 350W accelerators to be adequately cooled across the system operating temperature range. This push-pull airflow design provides redundancy by allowing for continuous system operation even with one fan failure. The HDCA has no HDDs, processors, and other peripherals blocking fresh airflow directly to the GPUs and Coprocessors allowing an excellent environment to produce maximum computing results without sacrificing performance due to server slot power restrictions.

The fans on the inside of each canister are removed with the canister and can be serviced by a service technician as detailed in Section 5. Each canister fan gets power from the canister backplane, reports fan speeds and can be controlled by the IPMI system monitor using pulse width modulation (PWM) signals through the backplanes. This minimizes wiring in the system and provides remote monitoring, speed control and alarming capability via the IPMI System Monitor.

The fans on the rear of the enclosure are hot-swap modules providing the same level of monitoring, control and alarming capability as the front fans. These fan modules are easily removed from the rear of the enclosure with a standard screwdriver and hot-swapped while the system and all of the accelerators are under operation. The procedure for fan replacement can be found in Section 5. Upon any fan failure, canister removal, critical temperature alarm or add-in card alarm from a supported accelerator, the fans will accelerate to 100% fan speed to protect the high-powered GPUs and Coprocessors.

2.9 Agency Certifications

Designed to meet the following agency certifications with testing currently pending:

- FCC - Part 15 of the FCC Rules, Class A, 47CFR
- Canada ICES-003, issue 4, Class A
- UL/IEC 60950-1
- Canada: CSA C22.2 No. 60950-1
- Argentina: IEC60950-1
- Japan: VCCI, Class A
- Australia/New Zealand AS/NZS CISPR 22, Class A
- IEC 60950-1 (CB Certificate and CB Test Report)
- CE Mark (EN55022 Class A, EN60950-1, EN55024, EN61000-3-2, EN61000-3-3)
- CISPR 22, CISPR 24, Class A
- CE Emissions 2004-108EC
- RoHS compliance (Directive 2002/95/EC)
- CCN NWGQ, NWGQ7

3

Setting up the HDCA

3.1 System Set-Up

The HDCA is shipped pre-configured, tested and guaranteed to operate with the server supplied with the system or with a server from the published compatibility list available at www.onestopsystems.com. The HDCA comes preloaded with NVIDIA Tesla GPUs, Grid GPUs or Intel Xeon Phi Coprocessors which are fully supported by OSS. OSS customer service should be consulted prior to initial operation of any customer supplied boards added to the system or any workstations or servers used with the system not fond on the published compatibility list. Even though the HDCA is designed to be plug-and-play, unfortunately, not all host system BIOS are designed to handle the vast amount of accelerator resources the HDCA provides to the system designer.

How to read the One Stop Systems SKU for a homogenous accelerator system:

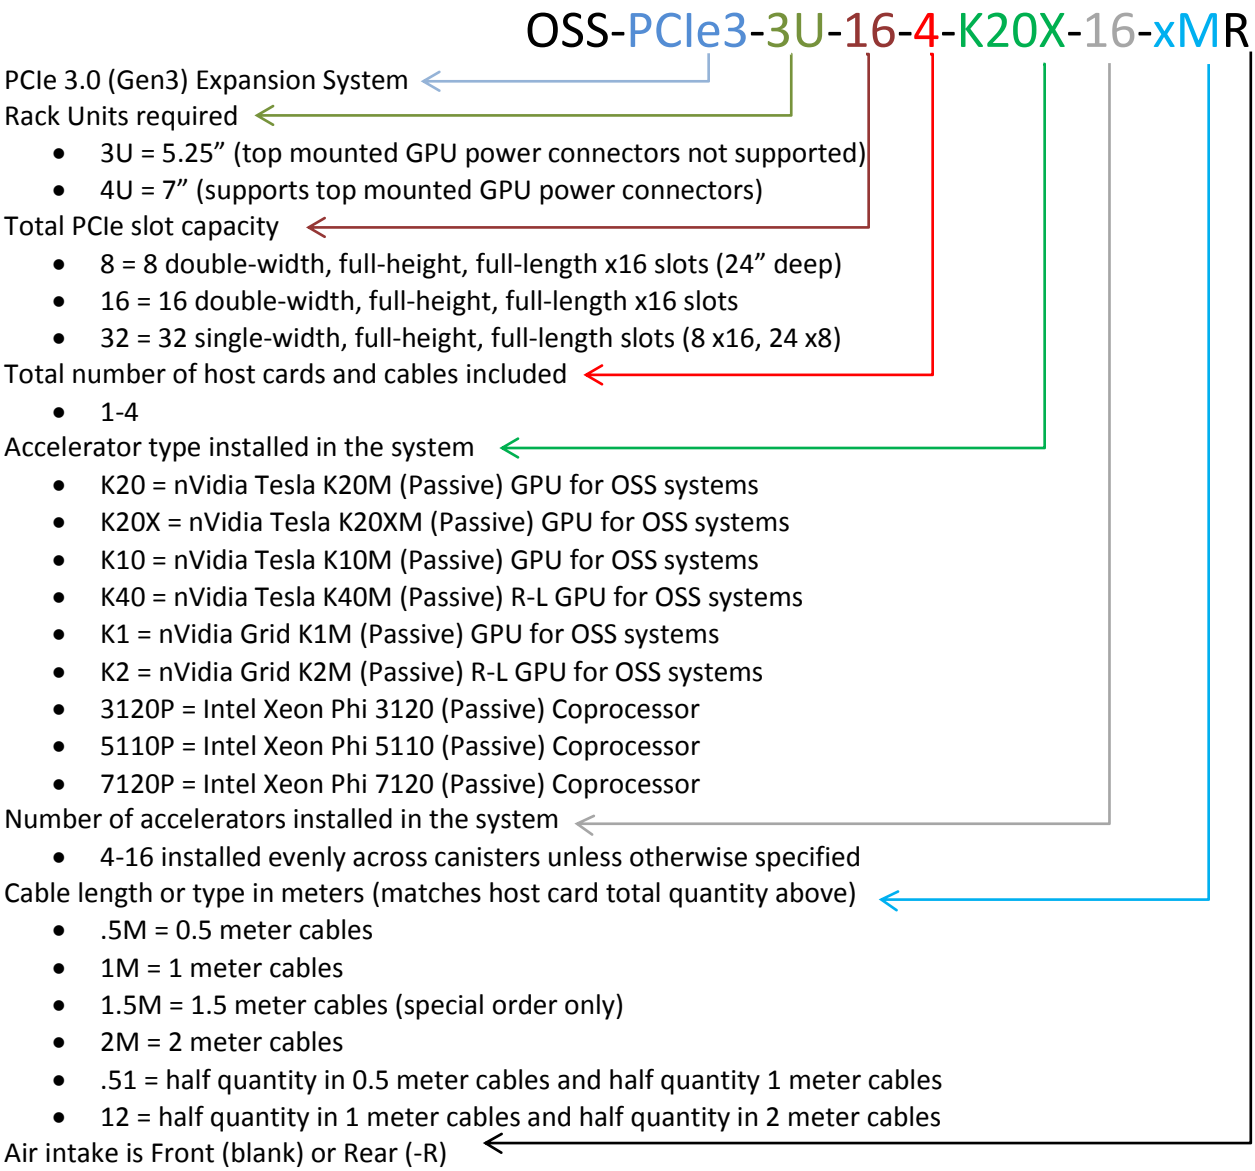

PCIe 3.0 (Gen3) Expansion System

Rack Units required

- 3U = 5.25” (top mounted GPU power connectors not supported)
- 4U = 7” (supports top mounted GPU power connectors)

Total PCIe slot capacity

- 8 = 8 double-width, full-height, full-length x16 slots (24” deep)
- 16 = 16 double-width, full-height, full-length x16 slots
- 32 = 32 single-width, full-height, full-length slots (8 x16, 24 x8)

Total number of host cards and cables included

- 1-4

Accelerator type installed in the system

- K20 = nVidia Tesla K20M (Passive) GPU for OSS systems
- K20X = nVidia Tesla K20XM (Passive) GPU for OSS systems
- K10 = nVidia Tesla K10M (Passive) GPU for OSS systems
- K40 = nVidia Tesla K40M (Passive) R-L GPU for OSS systems
- K1 = nVidia Grid K1M (Passive) GPU for OSS systems
- K2 = nVidia Grid K2M (Passive) R-L GPU for OSS systems
- 3120P = Intel Xeon Phi 3120 (Passive) Coprocessor
- 5110P = Intel Xeon Phi 5110 (Passive) Coprocessor
- 7120P = Intel Xeon Phi 7120 (Passive) Coprocessor

Number of accelerators installed in the system

- 4-16 installed evenly across canisters unless otherwise specified

Cable length or type in meters (matches host card total quantity above)

- .5M = 0.5 meter cables
- 1M = 1 meter cables
- 1.5M = 1.5 meter cables (special order only)
- 2M = 2 meter cables
- .51 = half quantity in 0.5 meter cables and half quantity 1 meter cables
- 12 = half quantity in 1 meter cables and half quantity in 2 meter cables

Air intake is Front (blank) or Rear (-R)

3.2 PCIe Cable Configuration

Once the system is powered on with the desired cable configuration, that configuration will persist until the entire HDCA is powered down. This is true as long as at least one connected host remains active from the original power-on sequence. This allows servicing hosts or canisters while the others remain operational. The HDCA is capable of supporting hot swap of canisters but, unfortunately, most PCIe add-in card vendors do not have drivers to support hot swap of their boards. To reset the cable configuration of the HDCA, all hosts should be powered down, the cables reconfigured and the system restarted with the procedure in section 4.

Figure 3.1 Rear View of HDCA

Table 3.1 Cable Configurations: The table below shows the configuration selection with ■ representing a cable present and □ representing a cable not present on the rear of the enclosure.

□□□□ Each box represents one canister. For reference, they will be referred to as canister 1234 from left to right.

Configuration	Cables Installed	Description
0	□□□□	Reserved
1	□□□■	One cable controls canister 4 and the canisters 1-3 are unused
2	□□■□	One cable controls all canisters 1-4
3	□□■■	One cable (left) controls canisters 1-3, one cable (right) controls canister 4
4	□■□□	One cable controls all canisters 1-4
5	□■□■	One cable (left) controls canisters 1-3, one cable (right) controls canister 4
6	□■■□	Two cables control 2 canisters each. Left cable controls canisters 1-2. Right cable controls canisters 3-4.
7	□■■■	One cable (left) controls canisters 1-2, two cables (middle and right) control 1 canister each. Middle cable controls canister 3 and right cable controls canister 4.
8	■□□□	One cable controls canister 1 and canisters 2-4 are unused
9	■□□■	Two cables control 1 canister each, two canisters unused. Left cable controls canister 1 and right cable controls canister 4. 2 and 3 are unused.
10	■□■□	One cable (right) controls canisters 2-4, one cable (left) controls canister 1.
11	■□■■	Two cables (left and right) control 1 canister each. Left controls canister 1, and right controls canister 4. One cable (middle) controls canisters 2 and 3.
12	■■□□	One cable (right) controls canisters 2-4, one cable (left) controls canister 1.
13	■■□■	Two cables (left and right) control 1 canister each. Left cable controls canister 1, right cable controls canister 4. One cable (middle) controls canisters 2 and 3.
14	■■■□	One cable (right) controls canisters 3 and 4, two cables (left and middle) control 1 canister each. Middle cable controls canister 2 and left cable controls canister 1.
15	■■■■	Four cables control 1 canister each. Each cable controls its corresponding canister, 1 to 1, 2 to 2, etc.

Operating

4

4.1 Powering Up the System

To set up the HDCA for first operation:

- Do not modify the shipped configuration prior to first operation. Install and test the SKU as shipped to maximize support effectiveness if required.
- Check the packaging for all ordered items including shipping damage. If there appears to be shipping damage, please contact your carrier to file a claim. Packaging should include the HDCA unit with accelerators, a separate box with host interfaces and cables.
- Remove all canisters from the chassis to inspect for damage, that the proper number of accelerators are included and to reduce the weight of the system prior to mounting. It is recommended to note the slots each canister came from in order to re-install the canisters into the same slots they are removed from, but is not required. The procedure for removing the canisters can be found in section 5.2.2 below.
- Install the HDCA in an ETSI compatible 19” rack with front and rear 19” rack rails using both the front and rear rack mounting provisions. A two person install is recommended. Mount in the rack as instructed below within the cable distance to your host system. Cables are available in 0.5, 1 and 2 meter lengths depending on the SKU ordered. Install per the airflow direction of the chassis SKU ordered and labeled on the side of the unit since there are front-to-rear and rear-to-front airflow SKUs to serve both blade and traditional servers. The main enclosure is meant to be fixed mounted to the rack rails while all of the customer serviceable items are easily removed from the front or rear of the system.
- Re-install the canisters into their respective slots. Turn the power switch on the front of each canister to the OFF (O) position.
- Install the power cables according to the rack or lab power infrastructure requirements as shown in Figure 4.1. Note there is no main power switch for the HDCA. When the power cables are plugged in and the PCIe cables are not, the system is in standby mode awaiting a power-on signal from the host system via the PCIe cables. If total power down is preferred, the power input cables must be removed from the rear of the unit. The enclosure uses standard 15 Amp IEC C14 male power connectors. Each power supply can operate up to 1500W power output each using 110V nominal voltage input or 3000W output each using 208-277V voltage input using the appropriate power cords. Some of the standard power cords available from OSS to support the HDCA include:
    - 110V power cord for PDUs and Wall receptacles
        - OSS Part number: OSS-CBL-PWR-5-15-C13-15A-6
        - NEMA 5-15 to IEC C13, Straight, 14AWG, 15A, 6'
    - 240V power cord for PDUs
        - OSS Part number: OSS-CBL-PWR-C14-C13-15A-6
        - IEC C14 to IEC C13, Straight, 14AWG, 15A, 6'
    - 240V power cord for US Wall receptacles
        - OSS Part number: OSS-CBL-PWR-6-15-C13-15A-6
        - NEMA 6-15 to IEC C13, Straight, 14AWG, 15A, 6'

Figure 4.1 Connecting Power Cords

- Install the host interfaces into an available PCIe add-in board slot in the server or workstation according to the host manufacturer instructions. If OSS supplied the server, the boards are pre-installed in the optimal slot(s). Note the OSS HIB is a PCIe 3.0 x16 board and should be installed in a server slot that can accommodate a x16 physical slot. The HIB will link train to the highest possible bandwidth available in the slot, up to x16 PCIe 3.0.
- No drivers are required for the OSS HIB since they act as a simple PCIe bus pass-through.
- Before connecting the host servers to the HDCA, it is recommended to check the BIOS version and OSS recommended BIOS settings for your host system. OSS posts the latest certified servers, BIOS revisions and recommended BIOS settings on the web site. If your host is not on the certified list make sure your system has the latest BIOS installed and the following settings are enabled in the BIOS for best results or contact OSS for testing.
    - Memory IO Map set above 4GB
    - PCIe 64bit is enabled
- With the host(s) powered down, install the desired number of PCIe cables in the desired configuration per Table 3.1. Simply connecting the cables will allocate the accelerators to the desired configuration upon power-on. Suggestions for optimal cable configurations can be found in section 3.2

Figure 4.2 Connecting PCIe

- Ensure the main chassis top is properly installed before powering on the enclosure. The system will not cool properly without the main chassis top in place.
- Turn on the power switches on the front of all canisters.

Figure 4.3 Powering up the Canisters

- The HDCA is in standby mode waiting for a host system to power on
- When any of the connected hosts are powered on, the HDCA will power on and accelerator enumeration with the hosts will commence. The PCIe cables present on the rear of the system at the time the first host is powered on will determine the configuration, not the power state of the host connected to the other end of the PCIe cable.

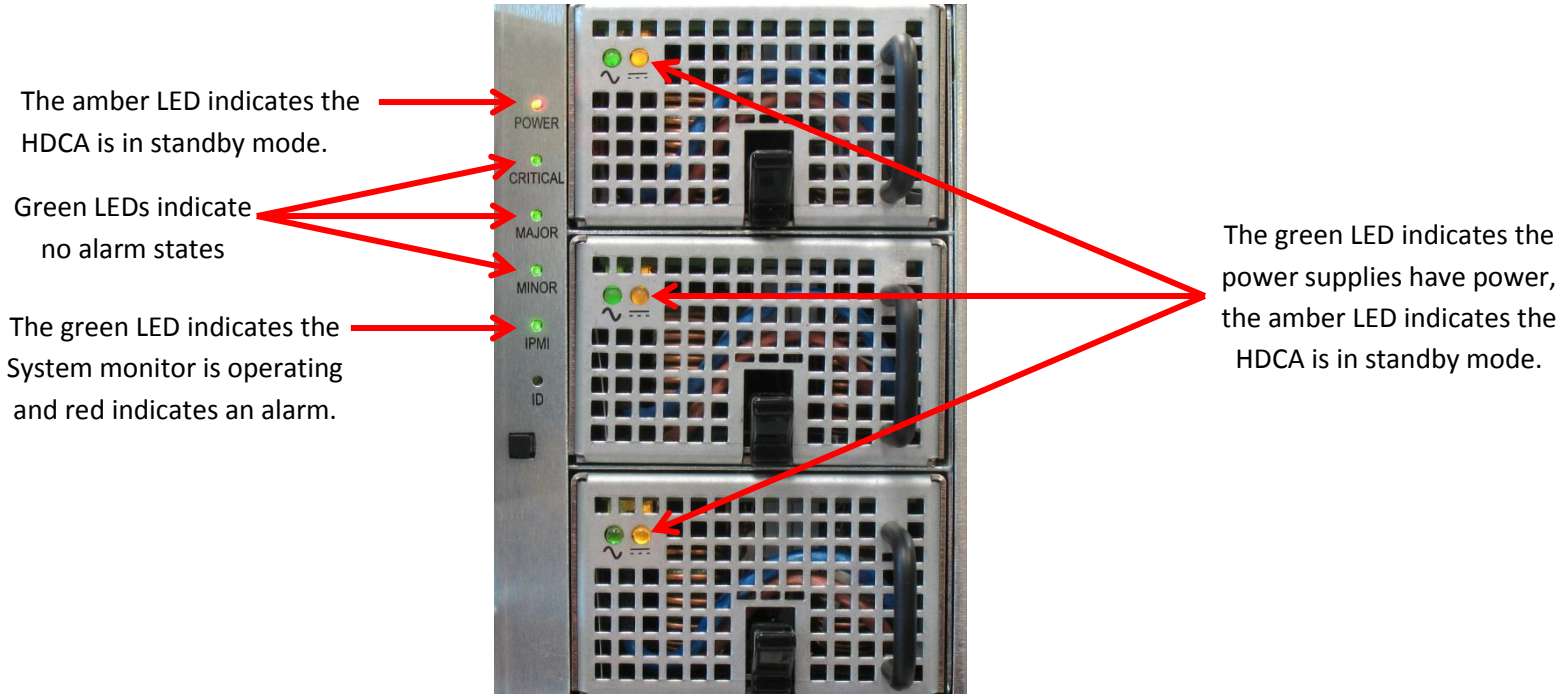

Figure 4.4 Standby Mode

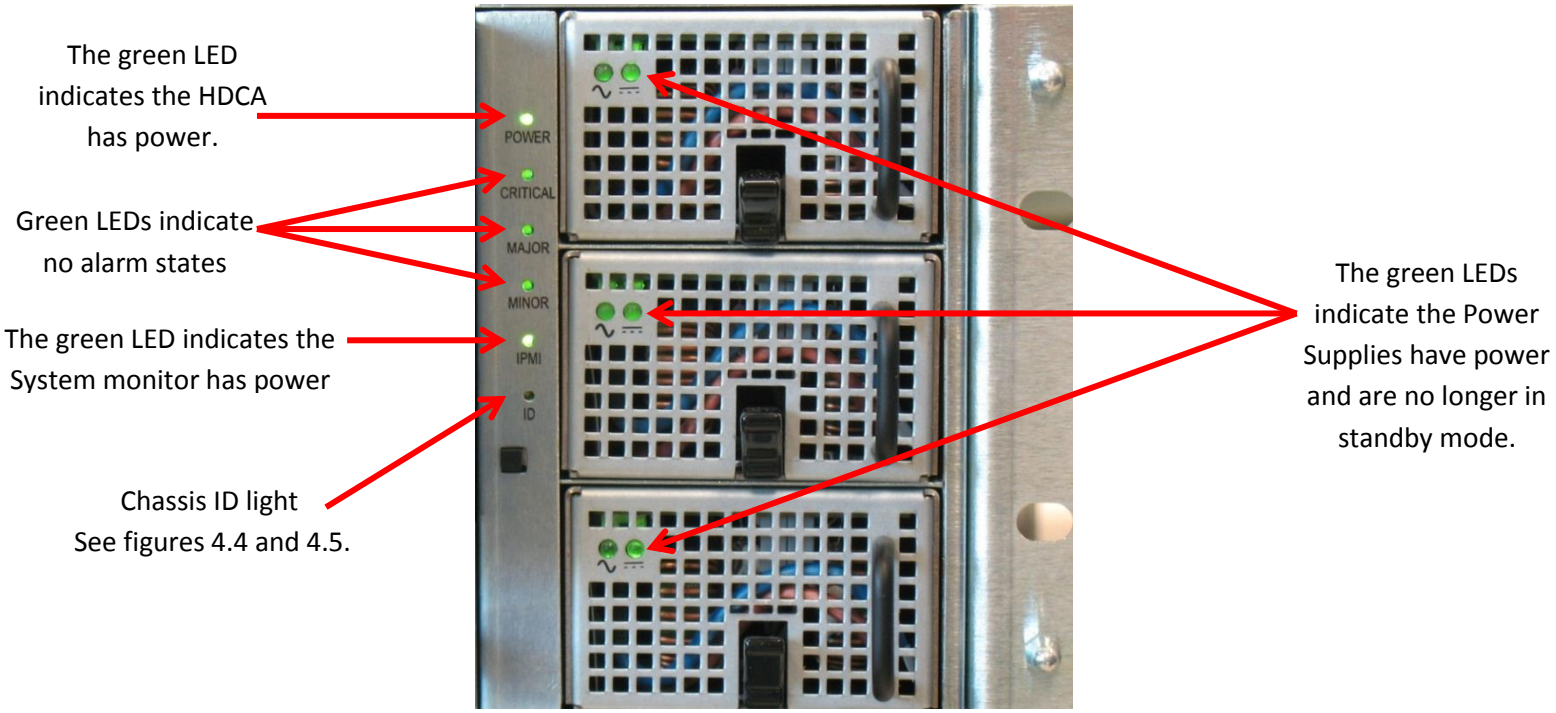

Figure 4.5 System LEDs under Normal Operation

- The following recommended power-up sequences for the most common configurations are as follows.
  - **Four host servers connected to one HDCA (See Figure 4.7)**
    - Setup: 4 hosts each with 1 HIB connected by PCIe cable to each of the 4 HDCA PCIe cable connectors. All 4 PCIe cables are connected at both ends, all 4 canister power connectors are on, all 3 circuit breakers are on. The HDCA is in standby mode.

Figure 4.6 Rear View of HDCA

    - Host 1 connected to cable Slot 1, host 2 connected to cable Slot 2 and so on.
    - Each host will control 1 canister of 4 slots each per the block diagram 4.7 where each color path represents each host connection to the appropriate canister
    - For best results it is recommended to power on host 2 and host 3 at the same time.
    - The HDCA will power on at that time and configure for 4-connection operation.
    - Power on host 1 and host 4.

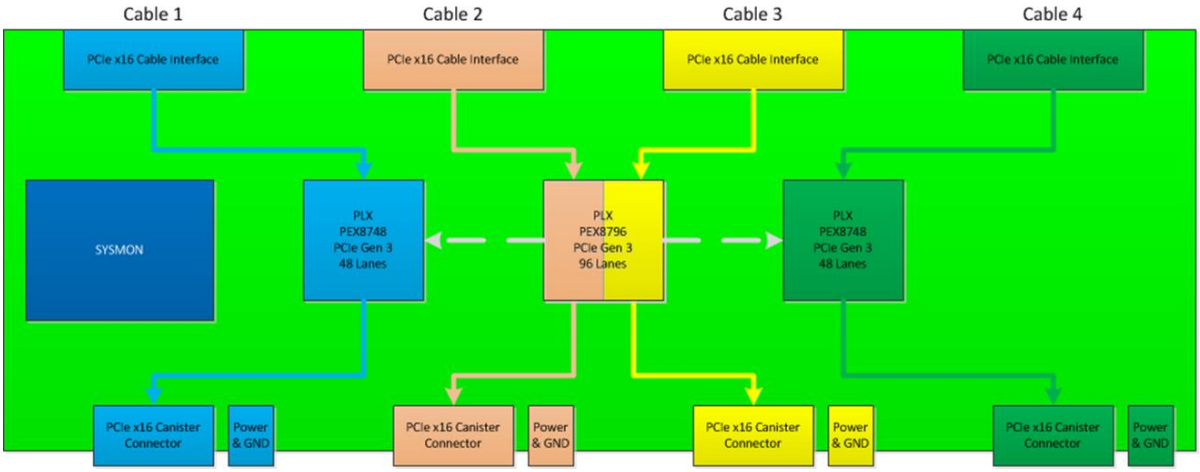

Figure 4.7 Main Backplane Connection Diagram for 4 Host Connections

- **Two host servers connected to one HDCA using two HIBs each (Also Figure 4.7)**
  - Setup: 2 hosts each with 2 HIBs connected by PCIe cable to each of the 4 HDCA PCIe cable connectors. All 4 PCIe cables are connected at both ends, all 4 canister power connectors are on, all 3 circuit breakers are on. The HDCA is in standby mode.
  - To minimize latency and switch hops through the main backplane, it is recommended to connect host 1 to cable Slots 1 & 4 (salmon & yellow paths in figure 4.7) and connect host 2 to cable Slots 2 & 3 (blue and green paths).
  - For best results it is recommended to power on host 2 first.
  - The HDCA will power on at that time and configure for 4-connection operation.
  - Power on host 1.
- **Two host servers connected to one HDCA using one HIB each (See Figure 4.8)**
  - Setup: 2 hosts each with 1 HIB connected by PCIe cable to two HDCA PCIe cable connectors. Both PCIe cables are connected at both ends, all 4 canister power connectors are on, all 3 circuit breakers are on. The HDCA is in standby mode.
  - Connect host 1 to cable Slot 2 (salmon path in figure 4.8) and connect host 2 to cable Slot 3 (yellow path).
  - For best results it is recommended to power on both hosts at the same time.
  - The HDCA will power on at that time and configure for 2-connection operation.

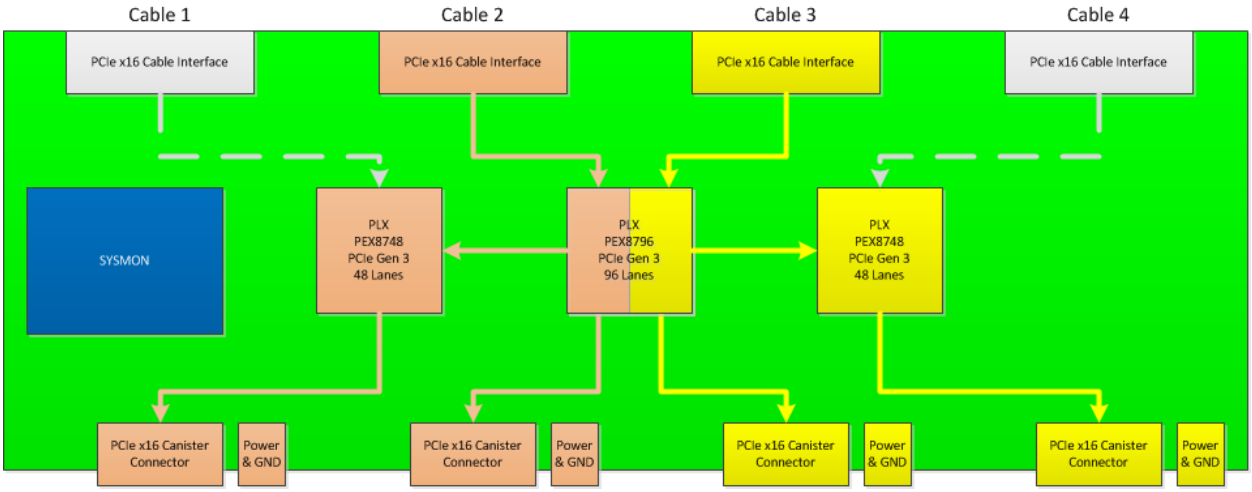

Figure 4.8 Main Backplane Connection Diagram for 2 Host Connections

- **One host server connected to one HDCA using two HIBs (Also Figure 4.8)**
  - Setup: 1 host with 2 HIBs connected by PCIe cable to two HDCA PCIe cable connectors. Both PCIe cables are connected at both ends, all 4 canister power connectors are on, all 3 circuit breakers are on. The HDCA is in standby mode.
  - Connect the HIBs to cable Slots 2 & 3.
  - Power on the host server.
  - The HDCA will power on at that time and configure for 2-connection operation.
- **One host server connected to one HDCA using one HIB (See Figure 4.9)**
  - Setup: 1 host with 1 HIB connected by PCIe cable to one HDCA PCIe cable connector. The PCIe cable is connected at both ends, all 4 canister power connectors are on, all 3 circuit breakers are on. The HDCA is in standby mode.
  - Connect the HIBs to cable Slot 2.
  - Power on the host server.
  - The HDCA will power on at that time and configure for 1-connection operation.

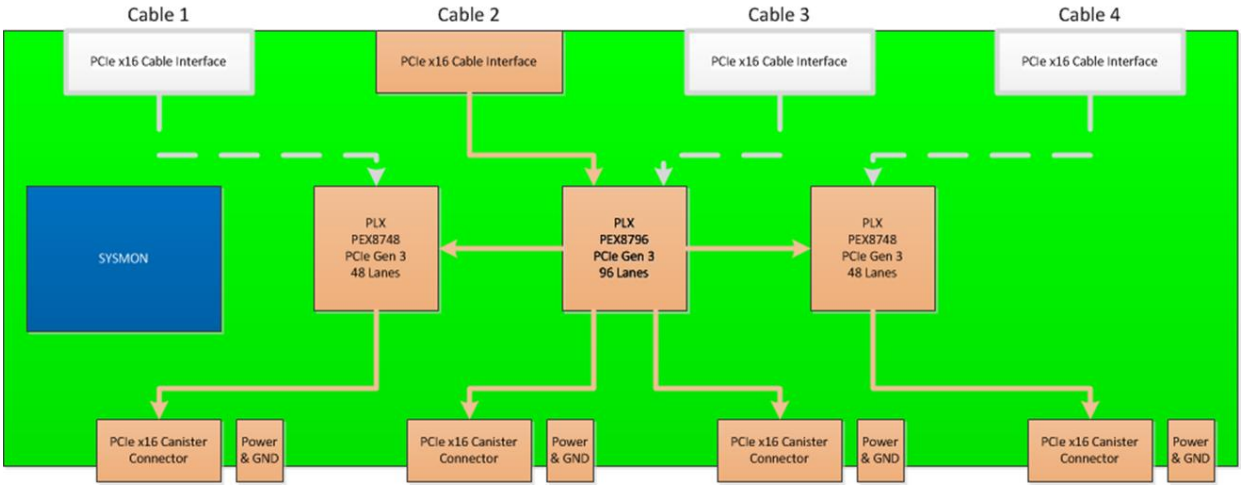

Figure 4.9 Main Backplane Connection Diagram for 1 Host Connection

For any questions on this procedure contact the OSS customer support team.

4.2 LED Features

To easily identify the chassis from the back in a datacenter rack environment, use the ID button or IPMI system monitor to turn on the Chassis ID light. The ID LED is also controlled by the IPMI System monitor. The ID LED will light blue. A matching blue LED will light on the back for easy identification.

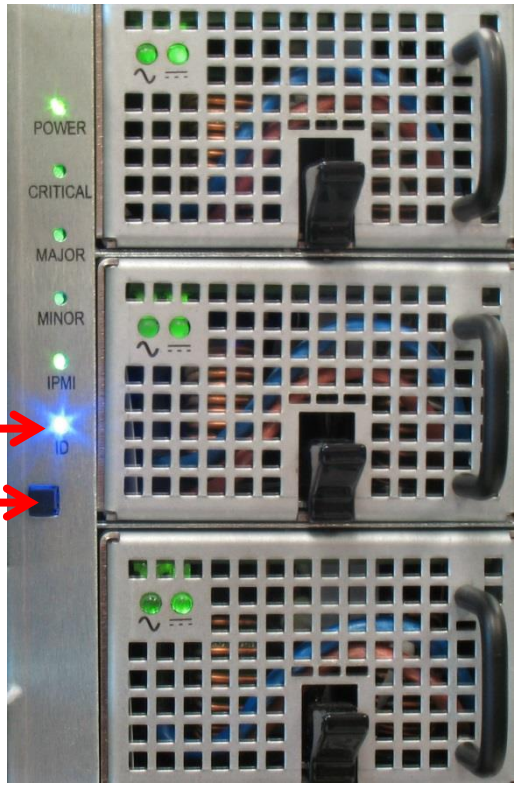

Figure 4.10 Chassis ID LED Location on the front

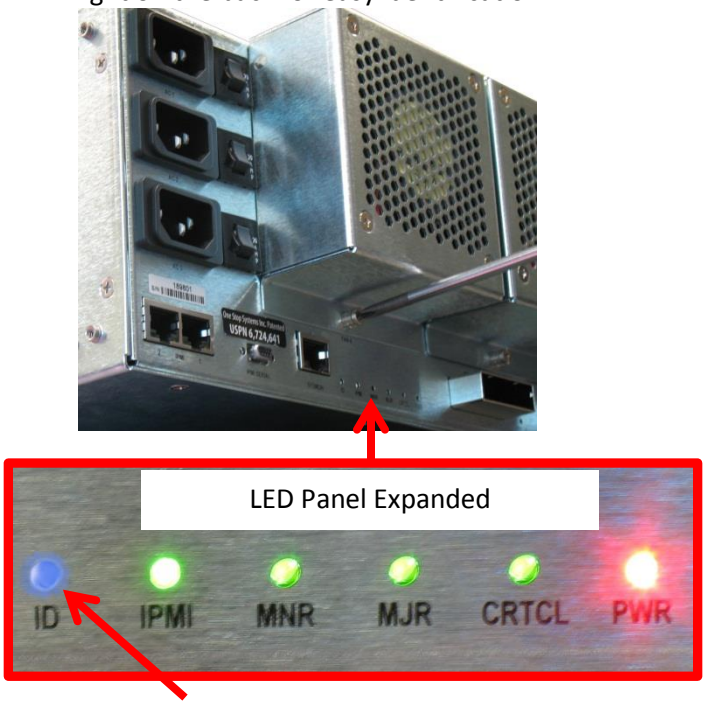

Figure 4.11 Chassis ID LED Location on the Rear: Arrow points to the chassis ID LED on the rear of the Enclosure. This feature can be used while Chassis is in standby mode (as pictured) or when fully powered up.

1. The PWR LEDs for all four canisters are green.
2. The C/ID LED is lit for the canister(s) that are connected to the host system.

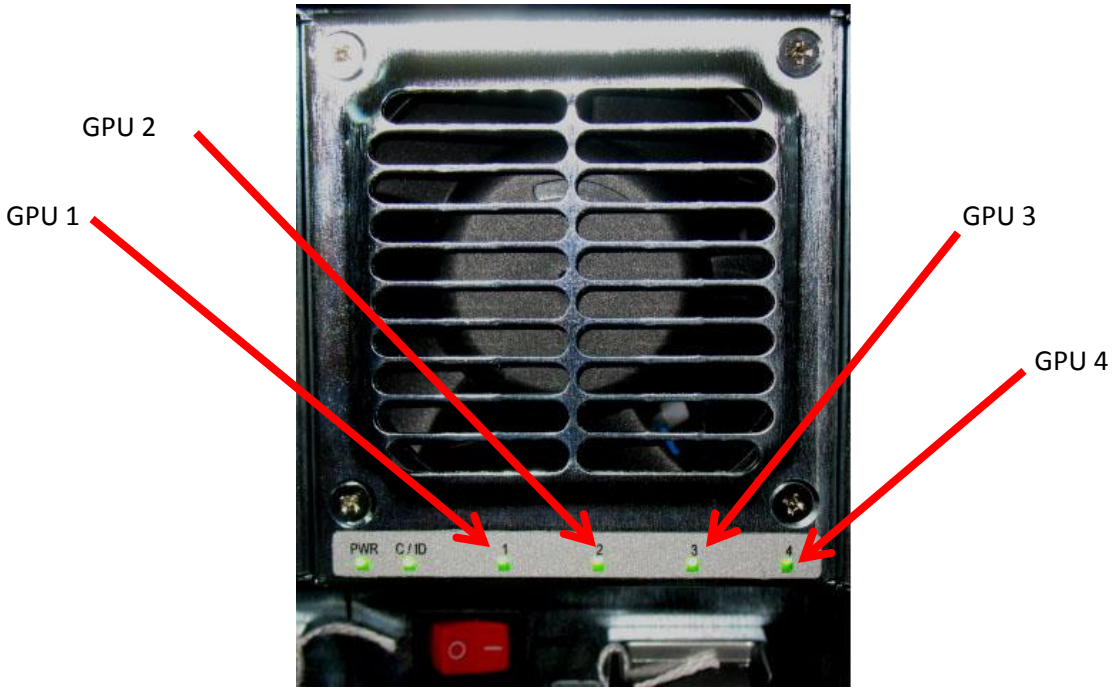

Figure 4.12 LED Panel on Canisters: Lights 1-4 are lit, indicating all slots in the canister are full.

3. LEDs 1-4 correspond to the four GPUs as shown in Figure 4.2.
4. Slot Status LED Blink Pattern
   - 4.1. Off – Link is Down
   - 4.2. Blinking, 512ms On, 512ms Off (1Hz) – Link is Up, 2.5 GT/s (PCIe Gen 1)any negotiated lane width
   - 4.3. Blinking, 256ms On, 256ms Off (2Hz) – Link is Up, 5 GT/s (PCIe Gen 2)any negotiated lane width
   - 4.4. On – Link is Up, 8.0 GT/s (PCIe Gen 3), Any negotiated lane width

NOTE: Some GPUs may idle and enter power saving mode at Gen 1 or Gen 2 even though they support higher speeds when fully active.

Figure 4.13 GPU Arrangement Pertaining to LEDs: The GPU number corresponds to the number on the LED panel. (Image shown is with NVIDIA Tesla GPUs. Arrangement is the same with Intel Xeon Phi coprocessors).

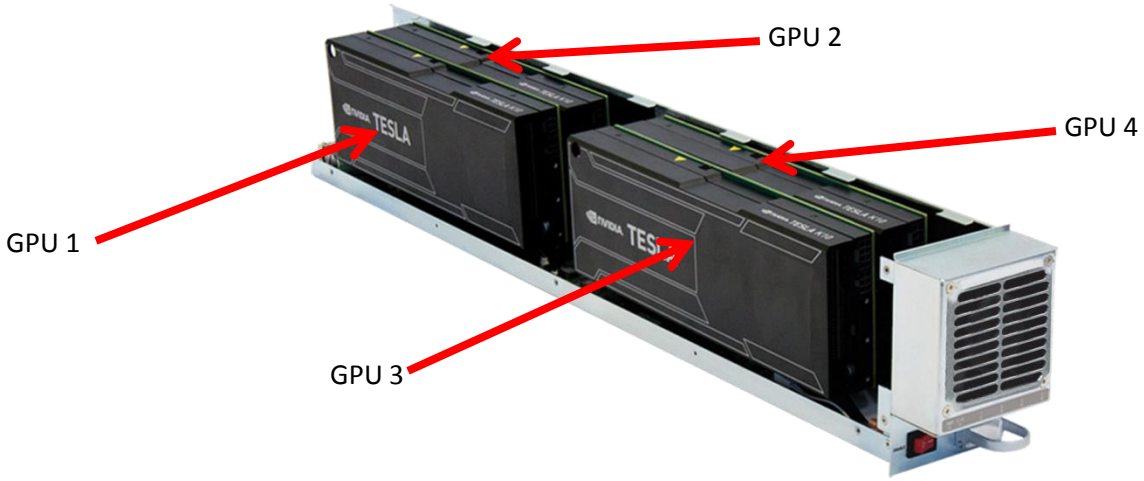

4.3 Using the Enclosure with Windows and Linux

1. The HDCA can be used with either Windows or Linux. OS compatibility is determined by the accelerators installed in the system since there are no drivers required for the HDCA.
2. For Windows users, open device manager to ensure proper connection to the system. If OSS supplied the host system and the accelerator, drivers are pre-loaded on the server. If another host is used, install the necessary accelerator drivers.

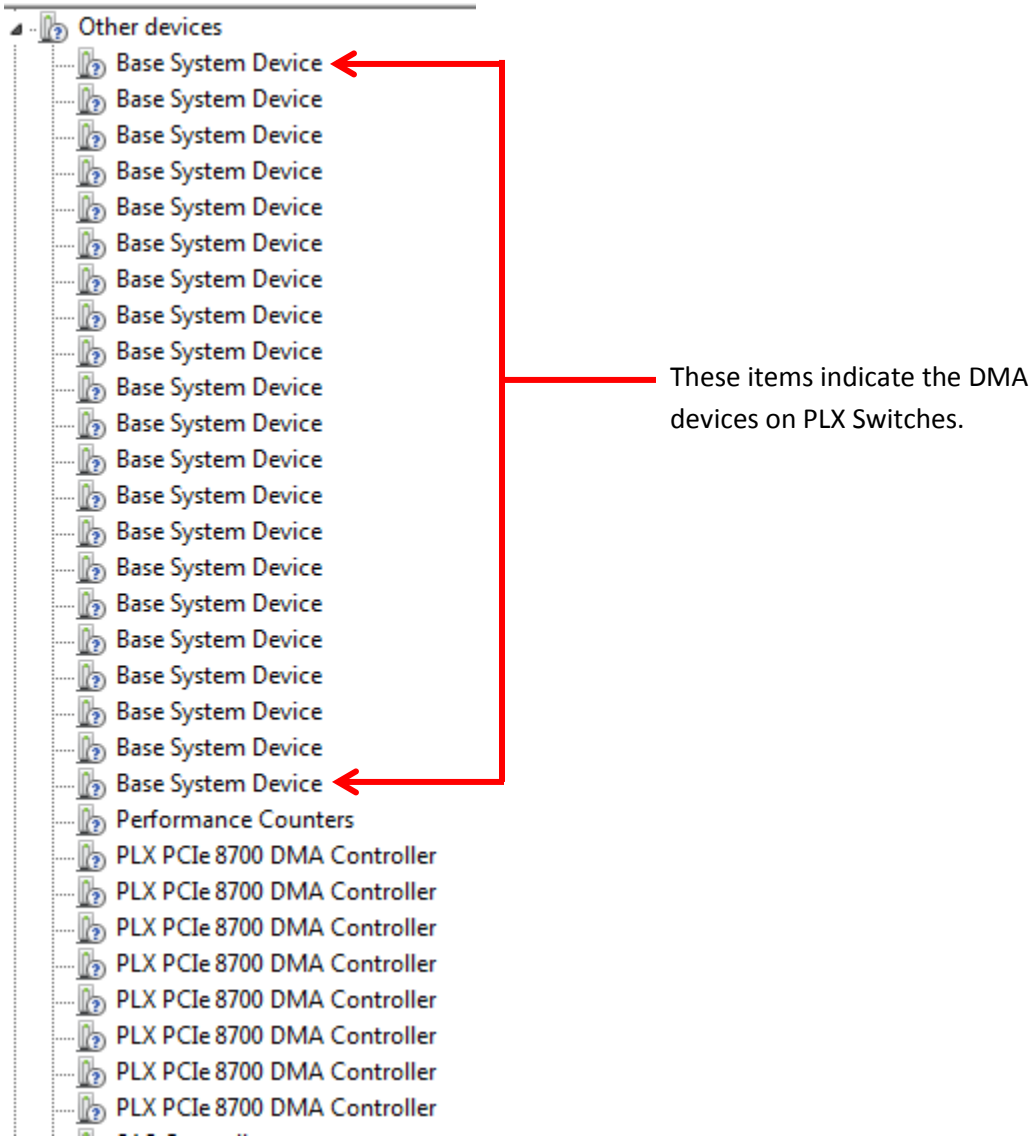

Figure 4.14 Using the Enclosure with Windows: This is an example of what will be shown in device manager (example configuration shows a system with no accelerators installed, accelerators will appear below each Base System Device).

3. For Linux users, use LSPCI to ensure a proper connection to the system. Install the drivers for the GPUs or Coprocessors.

```
04:00.0 PCI bridge: PLX Technology, Inc. Device 8733 (rev ca)
04:00.1 System peripheral: PLX Technology, Inc. Device 87d0 (rev ca)
04:00.2 System peripheral: PLX Technology, Inc. Device 87d0 (rev ca)
04:00.3 System peripheral: PLX Technology, Inc. Device 87d0 (rev ca)
04:00.4 System peripheral: PLX Technology, Inc. Device 87d0 (rev ca)
05:08.0 PCI bridge: PLX Technology, Inc. Device 8733 (rev ca)
06:00.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
07:00.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
07:08.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
07:0c.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
07:10.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
07:14.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
08:00.0 PCI bridge: PLX Technology, Inc. Device 8749 (rev ca)
08:00.1 System peripheral: PLX Technology, Inc. Device 87d0 (rev ca)
08:00.2 System peripheral: PLX Technology, Inc. Device 87d0 (rev ca)
08:00.3 System peripheral: PLX Technology, Inc. Device 87d0 (rev ca)
08:00.4 System peripheral: PLX Technology, Inc. Device 87d0 (rev ca)
09:08.0 PCI bridge: PLX Technology, Inc. Device 8749 (rev ca)
0a:00.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
0b:00.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
0b:04.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
0b:0c.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
0b:10.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
10:00.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
11:00.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
11:04.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
11:0c.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
11:10.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
17:00.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
18:00.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
18:04.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
18:0c.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
18:10.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
1d:00.0 PCI bridge: PLX Technology, Inc. Device 8749 (rev ca)
1e:00.0 PCI bridge: PLX Technology, Inc. Device 8749 (rev ca)
1e:08.0 PCI bridge: PLX Technology, Inc. Device 8749 (rev ca)
20:00.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
21:00.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
21:04.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
21:0c.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
21:10.0 PCI bridge: PLX Technology, Inc. Device 8796 (rev aa)
```

Figure 4.15 Using the Enclosure with Linux: This is an example of what will be shown after using LSPCI (example configuration is with an empty enclosure).

4. It is recommended to shut down any host in the event you need to remove or service a canister attached to that host, unless the accelerator boards support hot swap. At this time, NVIDIA Tesla, Grid and Intel Phi accelerators do not support hot swap. Canisters attached to other hosts attached to the same HDCA will continue to run while a canister is being serviced.

5

Service Support

There are only a few customer serviceable parts on the HDCA. They include the power supplies, the front and rear fans and the GPUs or coprocessors. If there are issues with any other parts of the system, One Stop Systems will replace parts as per the One Stop Systems Warranty. The power supplies and rear fans are hot swappable. The front fans and GPUs or coprocessors cannot be replaced while the canister is in operation, as it requires the removal of the canister.

5.1 Replacing the Power Supplies:

NOTE: It is recommended that you keep spare power supplies on site for maintenance operations. These are easy customer field replaceable units (FRUs).

1. Power supplies can be removed with the system powered on or off. If replacing a power supply while the system is on, only replace one at a time.
2. Lift the latch and pull firmly on the handle to remove the power supply. Removing the power supply will require approximately 35lbs of pressure. Ensure the enclosure is securely in place, then lift the latch and pull the handle.

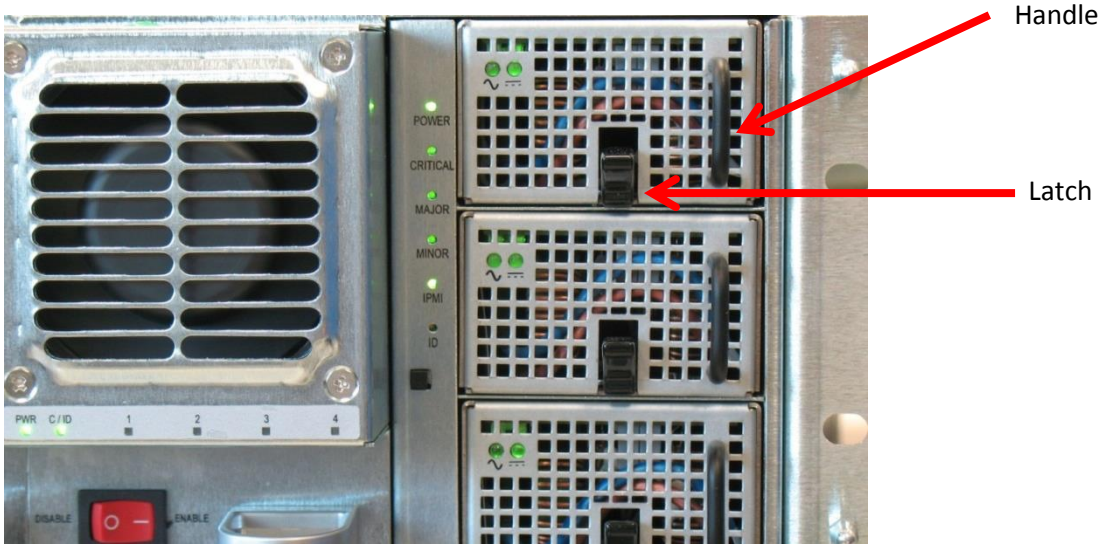

Figure 5.1: Removing the Power Supply

3. Slide the power supply out carefully.

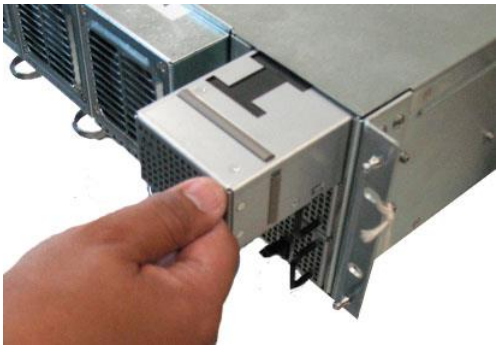

Figure 5.2: Removing the Power Supply

4. Replace the defective power supply with the new one.
5. Return any defective power supplies to OSS for replacement if under warranty.

5.2 Removing and Replacing Fans

5.2.1 Rear Fans:

NOTE: It is recommended that you keep spare rear fan modules on site for maintenance operations as damage could occur to GPUs if not properly cooled for extended periods. These are easily replaced from the rear of the enclosure while the system is operating.

1. Remove the screw at the bottom of the defective fan to separate it from the enclosure.

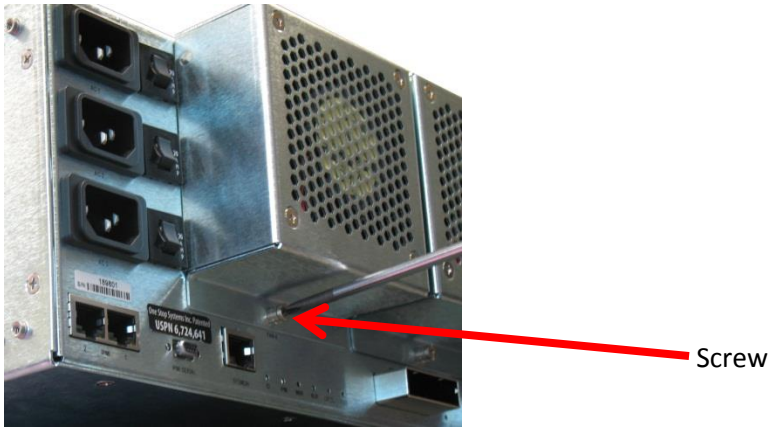

Figure 5.3: Removing the Rear Fan Enclosure

2. Carefully lift up on the fan and pull out.

Figure 5.4: Removing the Rear Fan Enclosure

3. The fan can be removed within the 3U space.
4. Replace the defective fan enclosure with the new one.
5. Replace the screw.
6. Return any defective fan modules to OSS for replacement if under warranty.

5.2.2 Front Fans:

NOTE: It is recommended that you keep a spare canister or spare front fans on site for maintenance operations as damage could occur to GPUs if not properly cooled for extended periods. Replacing the canister fan is an advanced customer or service person FRU.

1. Either leave the unit securely in the rack or place the HDCA on a steady surface with the bezel facing forward.
2. Remove the bezel by pulling toward you and slightly up, down or to one side, as the installation allows, to disengage the Tinnerman clips from the ball studs.
3. To remove the canister, push down on the latch slightly. The latch will depress ¼ of an inch, which will release the latch. Do not use excessive force to push down on the latch, or it will break.

Figure 5.5 Removing the Canister

4. Using the handle, pull out the canister.

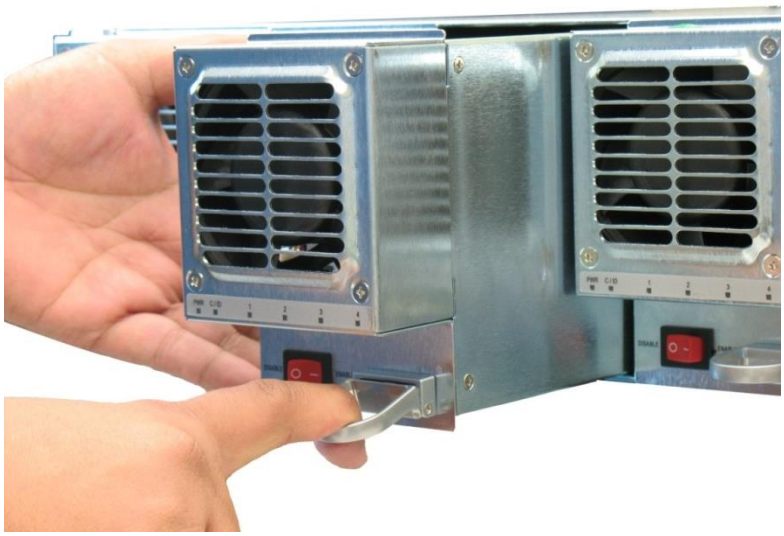

Figure 5.6 Removing the Canister

This should only be done in an ESD safe area using standard ANSI/ESD 20.20 practices. GPU and or system damage could occur if rules are not properly followed.

5. Remove 8 screws from the canister and remove the top.

Figure 5.7 Removing the Top of the Canister: The only screws that need to be removed to access the canister are marked by arrows etched into the metal.

Figure 5.7.1 Removing the Top of the Canister: Unscrew 8 screws.

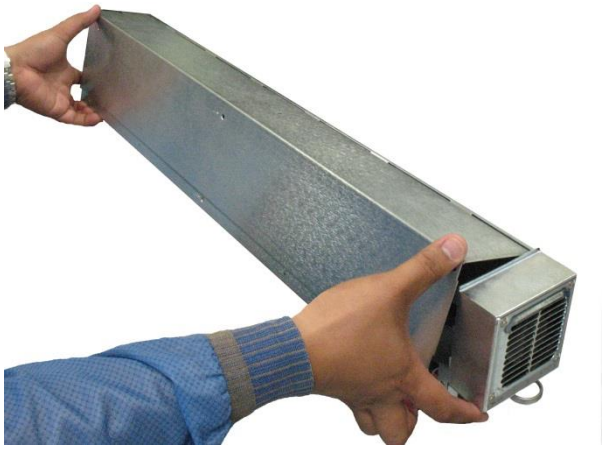

Figure 5.7.2 Removing the Top of the Canister: Carefully lift the top from the canister. The edge of the top is slightly sharp. Take care when removing the top.

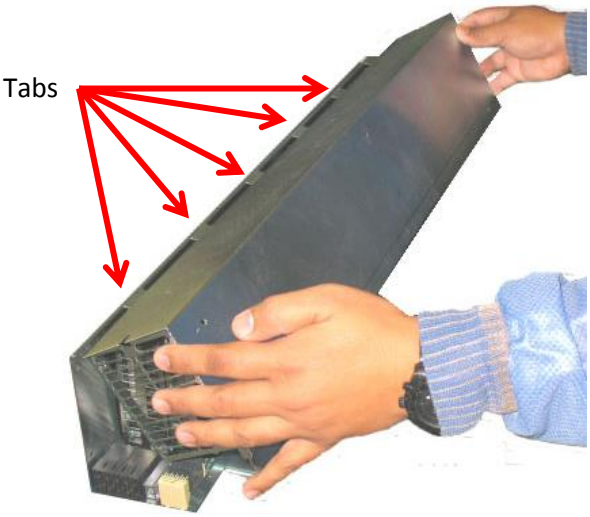

Figure 5.7.3 Removing the Top of the Canister: There are tabs holding the top of the canister in place. Lift and tilt the top to free it from the tabs.

6. Unplug the fan power connector

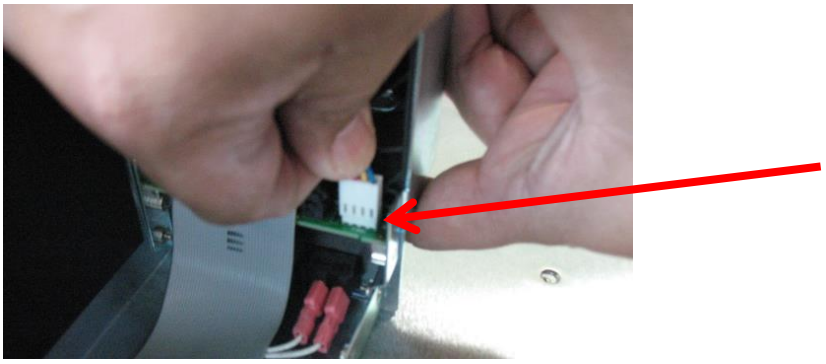

Figure 5.8: Removing the Front Fan Enclosure

7. Unlatch the ribbon cable by pressing down on the ejection tabs on both sides. Remove the ribbon cable.

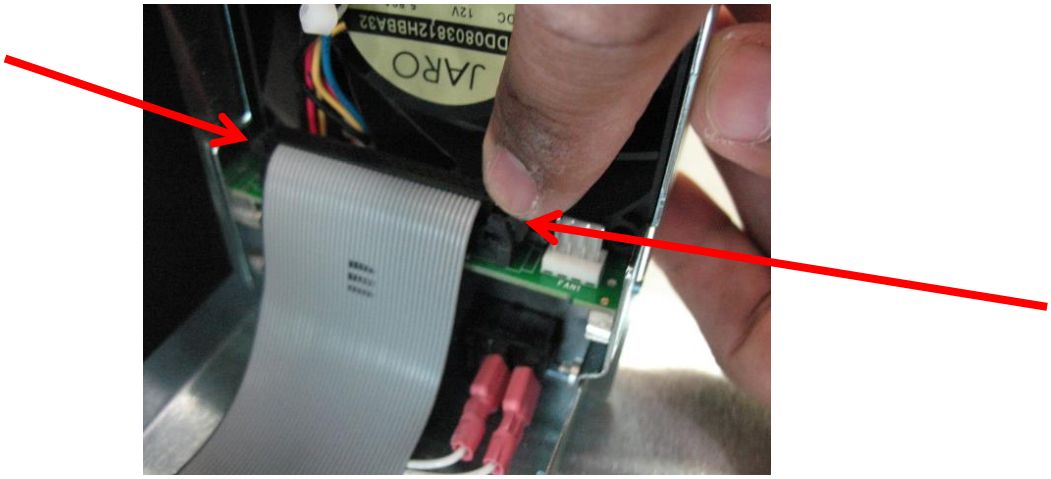

Figure 5.9: Removing the Front Fan Enclosure

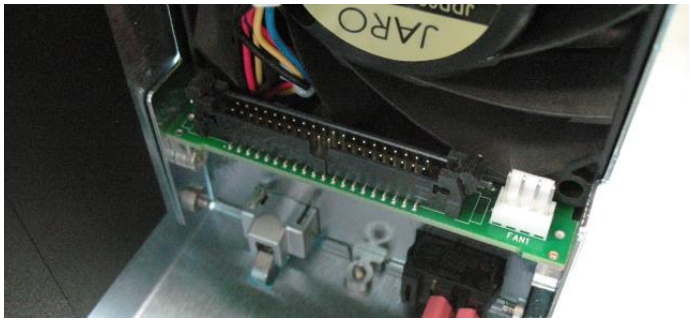

Figure 5.10: Removing the Front Fan Enclosure: The picture shows the power connector and ribbon cable removed.

8. Remove 10 screws total. After the first 5 screws (A), the fan enclosure will separate from the canister.
9. Remove the final 5 screws (B).

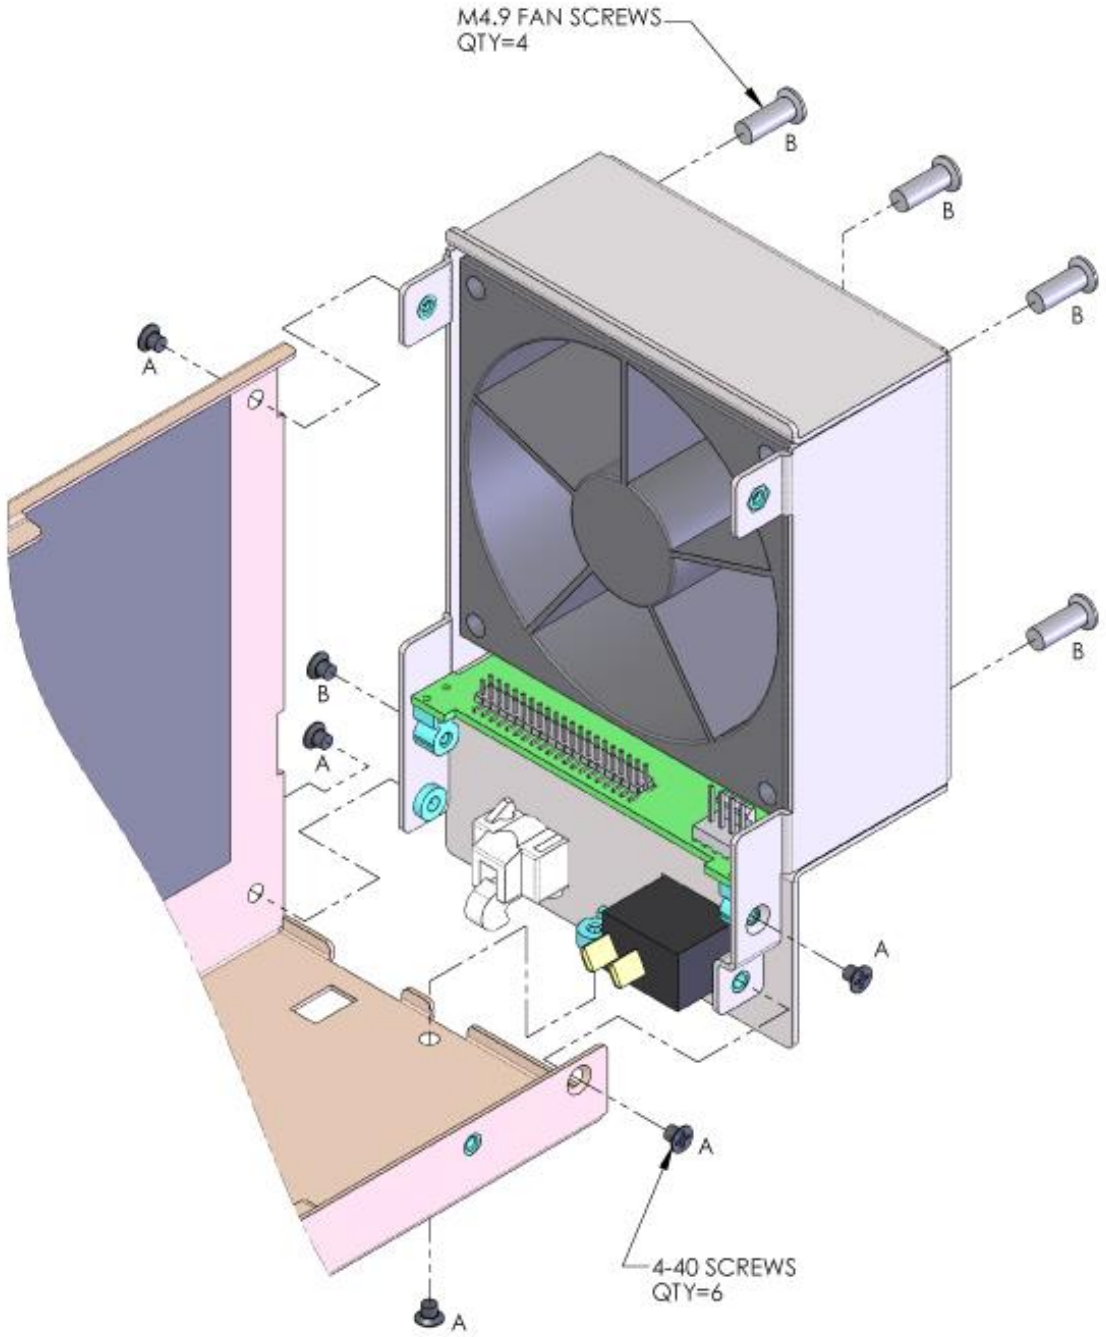

Figure 5.11: Removing the Front Fan Enclosure: The diagram shows the location of the 10 screws that need to be removed in order to replace the front fan. Label A indicates screws that need to be removed to detach the fan enclosure from the canister. Label B indicates screws that need to be removed to remove the fan from the enclosure.

10. Lift the fan out of the enclosure.

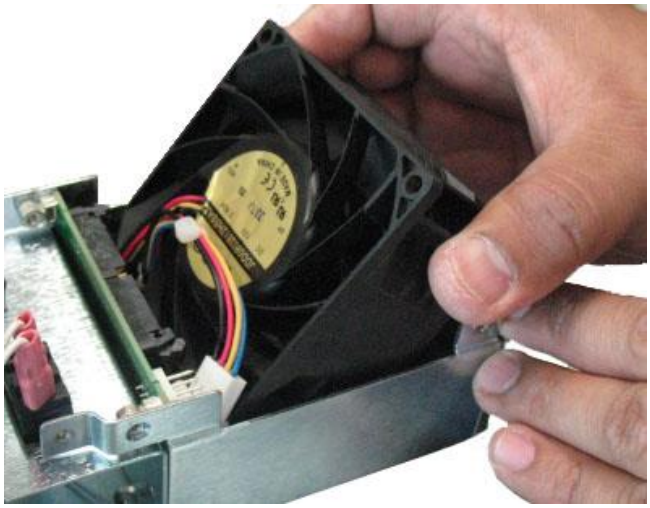

Figure 5.12: Removing the Front Fan Enclosure.

11. Install the new fan.
12. Replace 5 screws (labeled B in Figure 5.11).
13. Reattach the fan enclosure to the canister.
14. Replace the remaining 5 screws (labeled A in Figure 5.11).
15. Plug in the ribbon cable.
16. Plug in the fan power connector.
17. Replace the top cover (including 8 screws marked by arrows etched into the metal)

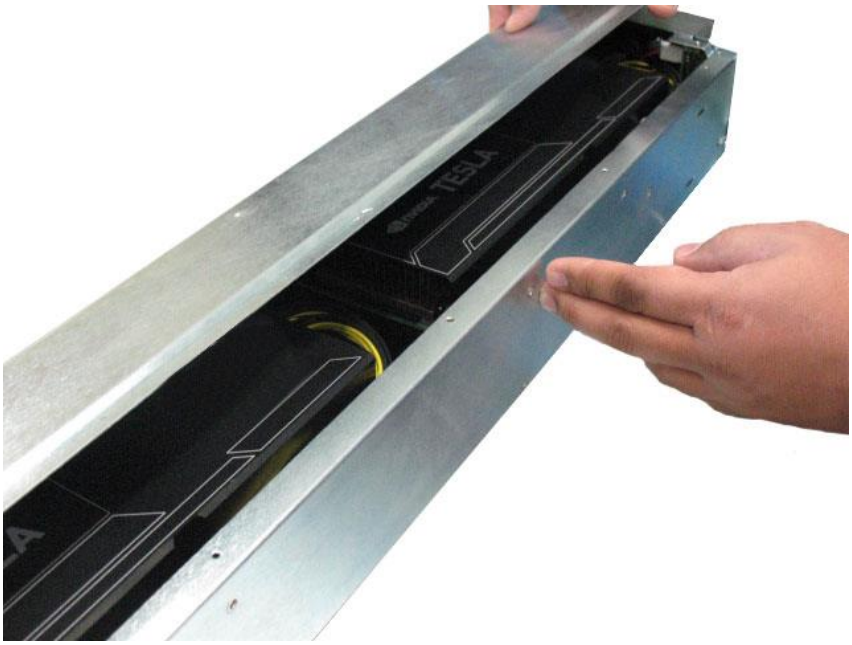

Figure 5.13 Replacing the Top of the Canister: Ensure the tabs on the canister line up with the slots of the top. Tilt and close the canister.

18. Replace 8 screws.

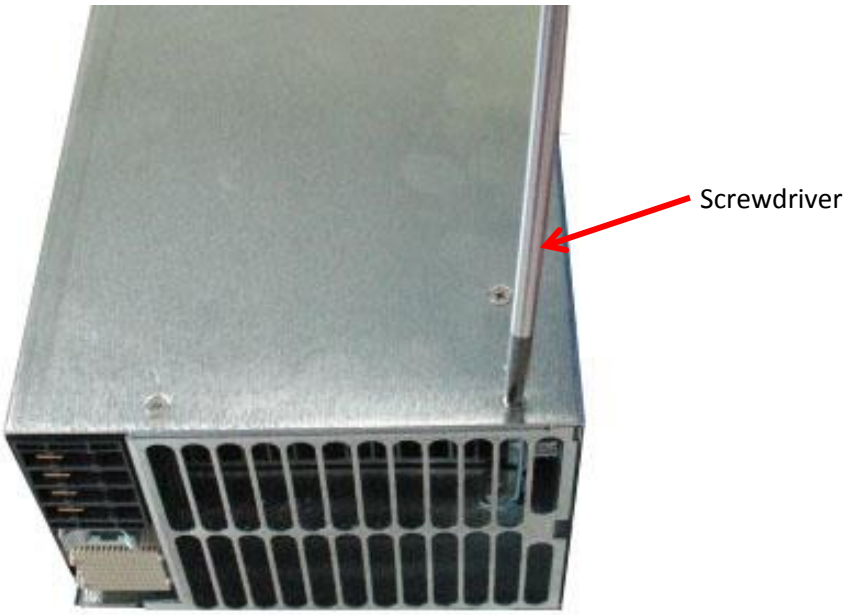

Figure 5.14 Replacing the Top of the Canister: Replace 8 screws, marked by arrows on the canister.

19. Reinsert the canister into the enclosure.

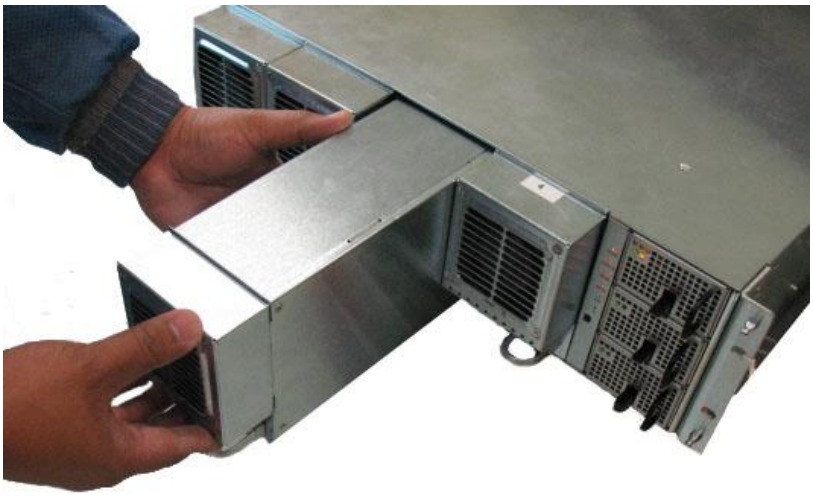

Figure 5.15 Reinserting the Canister

20. Turn on the canister.
21. Then turn on the host system attached to that canister.
22. Return any defective fans to OS for replacement if under warranty.

5.3 Removing and Replacing GPUs or Coprocessors

NOTE: Images in Section 5.3 show process with GPUs as the example. All GPUs should be purchased from One Stop Systems due to custom mounting brackets, wire harnesses, firmware and cooling considerations. Using non-OSS GPUs could result in incompatibility, mounting difficulties, power, warranty and support issues.

This should only be done in an ESD safe area using standard ANSI/ESD 20.20 practices. GPU and or system damage could occur if rules are not properly followed.

1. To replace the GPUs, remove the top of the canister. See Figures 5.7 through 5.7.3 for more information.
2. Unscrew the two screws from the top bracket. Set top bracket aside.

Figure 5.16: Removing the Top Bracket.

3. Slowly remove the GPU from the slot. Press the latch on the cable connector and remove the auxiliary power cables from the GPU.
   Note: If the GPU being removed is next to the metal wall of the chassis, the GPU away from the wall needs to be removed first.

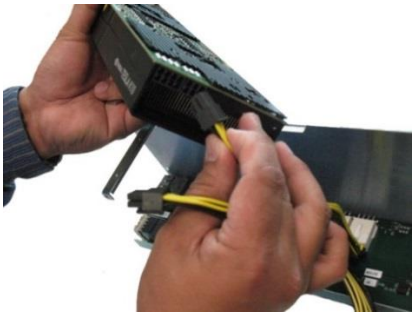

Figure 5.17: Removing the Auxiliary Power Cables from the GPU

4. Insert the auxillary power cables into the replacement GPU.

Figure 5.18: Inserting the Auxiliary Power Cables into the GPU

5. Insert the replacement GPU by carefully aligning the connectors and pushing the board firmly into the slot.

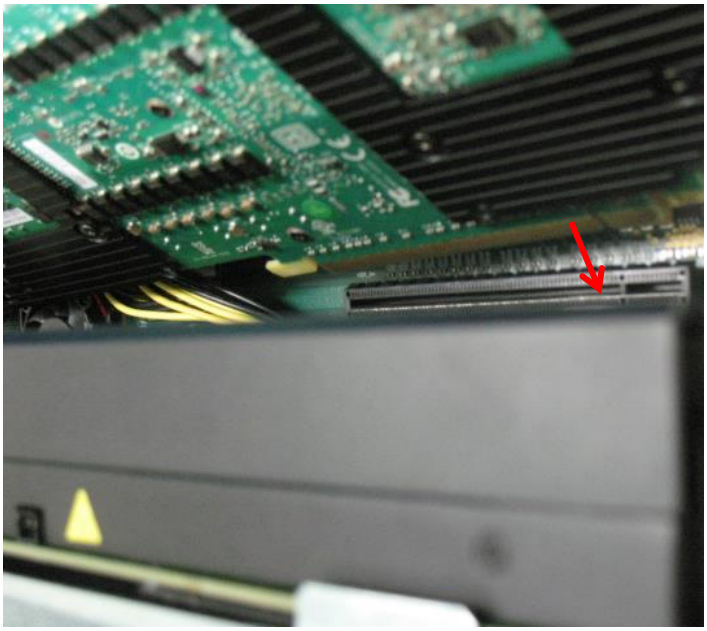

Figure 5.19: Installing the Replacement GPU

6. Ensure the power cables are routed around not over the fan on the backplane switch chip after insertion.

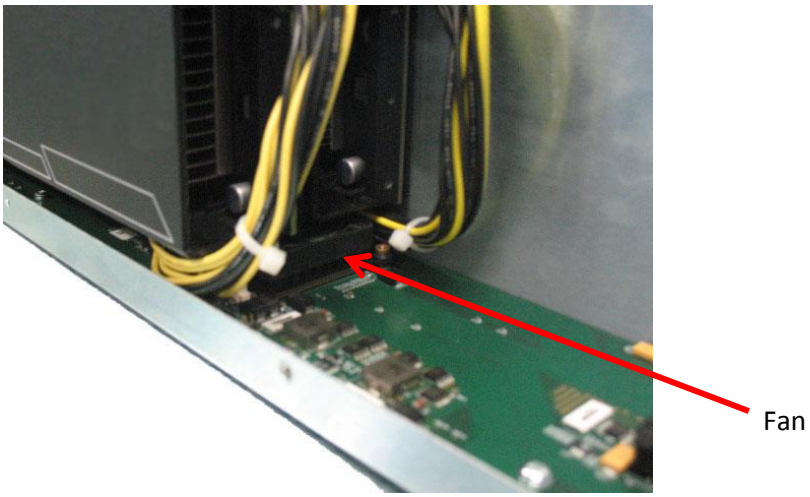

Figure 5.20: Ensuring Auxiliary Power Cables are Properly Installed: Picture shows two GPUs with auxiliary power cables routed around the fan, not over it.

7. Replace the top bracket and screws.
8. Replace the top of the canister and reinsert it into the enclosure. For more information see Figures 5.13-5.15.

For issues not mentioned in this User Manual, email support@onestopsystems.com or call 877-438-2724.

Appendix 1

A1.1 BIOS Settings

Table 6.1: BIOS Settings for Various Servers and Server Boards

DELL PowerEdge C6220	
Bios	1.1.19 (Current) **(1.1.9 previous)**
Build Date:	2/25/2013
CPU Cfg	Virtualization Technology = Disable
Memory Cfg	Memory Mapped I/O 64 Bit
DELL PowerEdge R720	
Bios	1.6.0
System CPLD Version	103
CPU Cfg	Virtualization Technology = Disable
System BIOS Settings - SATA Settings	Embedded SATA = AHCI Mode
Memory Cfg	Memory Mapped I/O 64 Bit
	Memory Operating Mode = Optimizer Mode
	Node Interleaving = Enable
Integrated Devices	I/OAT DMA Engine = Enable
	Memory Mapped I/O above 4GB = enable
Systems Profile Settings	System Profile = Custom
	Turbo Boost = Enable
	C1E = Enable
	C States = Enable
	Memory Patrol Scrub = Disable
Super Micro (All SuperMicro)	
Bios	1.0b
Build Date:	11/19/1012
Memory Cfg	Memory Mapped I/O 64 Bit = Enable
SATA Configuration	Windows OS = IDE
	Linux OS = AHCI
IBM 3650	
Bios	
CPU Cfg	Virtualization Technology should be disable
Memory Cfg	Memory Mapped I/O set to 3G (Max)
	Memory Mapped I/O 64 Bit - Enable
Intel S2600CO Server Board	
Bios Version	SE5C600.86B.01.08.0003
Build Date	2/26/2013

CPU Cfg	Intel ® Virtualization Technology = Disable
	Intel ® VT for Directed I/O = Disable
	Intel ® Turbo Boost Technology = Enable
Power & Performance Policy	CPU Power & Performance Policy = Performance
Memory Configuration	Patrol Scrub = Disable
	Demand Scrub = enable
Memory RAS and Performance Configuration	NUMA Optimized = enable
	Force 2x Refresh = Disable
Mass Storage Controller Configuration	AHCI Capable SATA Controller = Enhanced
	SAS/SATA Capable Controller = Disable
PCI Configuration	Maximize Memory below 4GB = Disable
	Memory Mapped I/O above 4GB = Enable
	Memory Mapped I/O Size = 1024G
Server Management Menu	Assert NMI on SERR = Disable
	Assert NMI on PERR = Disable
	PCIe AER Support = Disable

Appendix 2

A2.1 Canister Connector Pin-outs

1.1 Power-On (25-072-xxx-RC)

Pin	Signal
1	PS_ON#
2	GND

1.2 PCIe Auxiliary Power Connector RA (25-072-317-RC)

Pin	Signal	Pin	Signal
1	GND	5	+12V
2	GND	6	+12V
3	GND	7	+12V
4	GND	8	+12V

1.3 PCIe Connector

Pin #	Name	Pin #	Name	
B1	+12V	A1	PRSNT1#	
B2	+12V	A2	+12V	
B3	+12V	A3	+12V	
B4	GND	A4	GND	
B5	SMCLK	A5	TCK	
B6	SMDAT	A6	TDI	
B7	GND	A7	TDO	
B8	+3.3V	A8	TMS	
B9	TRST#	A9	+3.3V	
B10	3.3Vaux	A10	+3.3V	Mechanical Key
B11	WAKE#	A11	PERST#	
B12	RSVD	A12	GND	
B13	GND	A13	REFCLK+	
B14	PETp0	A14	REFCLK-	
B15	PETn0	A15	GND	
B16	GND	A16	PERp0	End of the x1 Connector
B17	PRSNT2#	A17	PERn0	
B18	GND	A18	GND	
B19	PETp1	A19	RSVD	
B20	PETn1	A20	GND	
B21	GND	A21	PERp1	
B22	GND	A22	PERn1	
B23	PETp2	A23	GND	End of the x4 Connector
B24	PETn2	A24	GND	
B25	GND	A25	PERp2	

B26	GND	A26	PERn2	
B27	PETp3	A27	GND	
B28	PETn3	A28	GND	
B29	GND	A29	PERp3	
B30	RSVD	A30	PERn3	
B31	PRSNT2#	A31	GND	
B32	GND	A32	RSVD	
B33	PETp4	A33	RSVD	
B34	PETn4	A34	GND	
B35	GND	A35	PERp4	
B36	GND	A36	PERn4	
B37	PETp5	A37	GND	
B38	PETn5	A38	GND	
B39	GND	A39	PERp5	
B40	GND	A40	PERn5	
B41	PETp6	A41	GND	
B42	PETn6	A42	GND	
B43	GND	A43	PERp6	
B44	GND	A44	PERn6	
B45	PETp7	A45	GND	
B46	PETn7	A46	GND	
B47	GND	A47	PERp7	End of the x8 Connector
B48	PRSNT2#	A48	PERn7	
B49	GND	A49	GND	
B50	PETp8	A50	RSVD	
B51	PETn8	A51	GND	
B52	GND	A52	PERp8	
B53	GND	A53	PERn8	
B54	PETp9	A54	GND	
B55	PETn9	A55	GND	
B56	GND	A56	PERp9	
B57	GND	A57	PERn9	
B58	PETp10	A58	GND	
B59	PETn10	A59	GND	
B60	GND	A60	PERp10	
B61	GND	A61	PERn10	
B62	PETp11	A62	GND	
B63	PETn11	A63	GND	
B64	GND	A64	PERp11	
B65	GND	A65	PERn11	
B66	PETp12	A66	GND	
B67	PETn12	A67	GND	
B68	GND	A68	PERp12	
B69	GND	A69	PERn12	
B70	PETp13	A70	GND	
B71	PETn13	A71	GND	
B72	GND	A72	PERp13	End of the x16 Connector
B73	GND	A73	PERn13	
B74	PETp14	A74	GND	

B75	PETn14	A75	GND	
B76	GND	A76	PERp14	
B77	GND	A77	PERn14	
B78	PETp15	A78	GND	
B79	PETn15	A79	GND	
B80	GND	A80	PERp15	
B81	PRSNT2#	A81	PERn15	
B82	RSVD	A82	GND	

A2.2 Backplane Connector Pin-outs

PCIe x16 Cable Connector

Pin #	Name	Pin #	Name	Pin #	Name	Pin #	Name
A1	GND	B1	GND	C1	GND	D1	GND
A2	PERp1	B2	PERp0	C2	PETp1	D2	PETp0
A3	PEPn1	B3	PERn0	C3	PETn1	D3	PETn0
A4	GND	B4	GND	C4	GND	D4	GND
A5	PERp3	B5	PERp2	C5	PETp3	D5	PETp2
A6	PERn3	B6	PERn2	C6	PETn3	D6	PETn2
A7	GND	B7	GND	C7	GND	D7	GND
A8	PERp5	B8	PERp4	C8	PETp5	D8	PETp4
A9	PERn5	B9	PERn4	C9	PETn5	D9	PETn4
A10	GND	B10	GND	C10	GND	D10	GND
A11	PERp7	B11	PERp6	C11	PETp7	D11	PETp6
A12	PERn7	B12	PERn6	C12	PETn7	D12	PETn6
A13	GND	B13	GND	C13	GND	D13	GND
A14	PWR	B14	PWR_RTN	C14	PWR	D14	PWR_RTN
A15	PWR	B15	PWR_RTN	C15	PWR	D15	PWR_RTN
A16	PWR	B16	PWR_RTN	C16	PWR	D16	PWR_RTN
A17	uSB_RTN	B17	uPWRON	C17	dSB_RTN	D17	dPWRON
A18	GND	B18	uWAKE#	C18	GND	D18	dWAKE#
A19	uREFCLKp	B19	uPRSNT#	C19	dREFCLKp	D19	dPRSNT#
A20	uREFCLKn	B20	uPERST#	C20	dREFCLKn	D20	dPERST#
A21	GND	B21	GND	C21	GND	D21	GND
A22	PERp9	B22	PERp8	C22	PETp9	D22	PETp8
A23	PERn9	B23	PERn8	C23	PETn9	D23	PETn8
A24	GND	B24	GND	C24	GND	D24	GND
A25	PERp11	B25	PERp10	C25	PETp11	D25	PETp10
A26	PERn11	B26	PERn10	C26	PETn11	D26	PETn10
A27	GND	B27	GND	C27	GND	D27	GND
A28	PERp13	B28	PERp12	C28	PETp13	D28	PETp12
A29	PERn13	B29	PERn12	C29	PETn13	D29	PETn12
A30	GND	B30	GND	C30	GND	D30	GND
A31	PERp15	B31	PERp14	C31	PETp15	D31	PETp14
A32	PERn15	B32	PERn14	C32	PETn15	D32	PETn14
A33	GND	B33	GND	C33	GND	D33	GND
A34	RSVD	B34	RSVD	C34	RSVD	D34	RSVD

1.1 Power Entry Module

Pin	Signal
1	+12V
2	GND

1.2 Power Supply Signal Connector (25-072-898-RC)

Pin	Signal	Pin	Signal
1	PSON#	7	+12V Sense_N
2	PSU_OK#	8	+12V Sense_P
3	LED_AUX	9	+Vaux
4	GND	10	+Vaux
5	GND	11	N/C
6	GND	12	N/C

*+Vaux is +12Vaux in the 3U system and +5Vaux in the 4U system

1.3 Fan Header (25-072-049)

Pin	Signal	Wire Color
1	GND	Black
2	+12V	Red
3	TACH	Yellow
4	PWM	Blue

1.4 System LED Header (25-072-894-RC)

Pin	Signal	Pin	Signal
1	+3.3Vaux	2	N/C
3	MINOR_ALARM#	4	MINOR_ALARM
5	MAJOR_ALARM#	6	MAJOR_ALARM
7	CRITICAL_ALARM#	8	CRITICAL_ALARM
9	POWER_GOOD#	10	POWER_GOOD
11	BLUE_LED_SW#	12	BLUE_LED#
13	IPMI_RED#	14	IPMI_GREEN#
15	GND	16	GND

SYSMON II Pinout (25-072-410-RC)

Pin #	SYSMON II	Connection
1	TEMP1	Temp Sensor 1
2	TEMP2	Temp Sensor 2
3	TEMP3	Temp Sensor 3
4	TEMP4	Temp Sensor 4
5	TEMP5	Temp Sensor 5
6	TEMP6	Temp Sensor 6

7	TEMP7	Temp Sensor 7
8	TEMP8	Temp Sensor 8
9	VCC12_IN	+12V
10	VCC3_IN	+3.3V
11	VCC5_IN	+5Vaux
12	VCC-12_IN	N/C
13	GP_IN2	C1_PGOOD (Canster 1)
14	GP_IN1	C0_PGOOD (Canster 0)
15	GP_IN4	C3_PGOOD (Canster 3)
16	GP_IN3	C2_PGOOD (Canster 2)
17	GP_IN6	IPMI Input A
18	GP_IN5	IPMI Input B
19	GP_IN8	N/C
20	GP_IN7	PWR_GOOD
21	GP_OUT2	Major Alarm
22	GP_OUT1	Minor Alarm
23	GP_OUT4	N/C
24	GP_OUT3	Critical Alarm
25	GP_OUT6	N/C
26	GP_OUT5	N/C
27	GP_OUT8	N/C
28	GP_OUT7	N/C
29	FANTACH2	Fan Tachometer 2
30	FANTACH1	Fan Tachometer 1
31	FANTACH4	Fan Tachometer 4
32	FANTACH3	Fan Tachometer 3
33	FANTACH6	CAN1_Fan Tachometer 6
34	FANTACH5	CAN0_Fan Tachometer 5
35	FANTACH8	CAN3_Fan Tachometer 8
36	FANTACH7	CAN2_Fan Tachometer 7
37	VCC5	+5Vaux
38	VCC5	+5Vaux
39	GND	Ground
40	GND	Ground

1.5 Power Supply I2C Connector (25-072-896-RC)

Pin	Signal	Pin	Signal
1	SDA	3	SCL
2	GND	4	INT#

1.6 IPMI MICRO-SERIAL PORT (25-072-879-RC)

Pin	Signal	Pin	Signal
1	Data Carrier Detect	6	Data Set Ready
2	Received Data	7	Request to Send
3	Transmitted Data	8	Clear to Send

4	Data Terminal Ready	9	Ring Indicator
5	Signal Ground		

1.7 SYSMON Configuration (Preliminary)

Sensor	Gauge Enabled	Alarm Enabled	Alarm Type	Input Polarity	Fan Tach Divisor	Low Limit	High Limit	Notes
Fan 1	Enabled	Enabled	Critical	High	2	600	N/A	
Fan 2	Enabled	Enabled	Critical	High	2	600	N/A	
Fan 3	Enabled	Enabled	Critical	High	2	600	N/A	
Fan 4	Enabled	Enabled	Critical	High	2	600	N/A	
Fan 5	Enabled	Enabled	Critical	High	2	600	N/A	
Fan 6	Enabled	Enabled	Critical	High	2	600	N/A	
Fan 7	Enabled	Enabled	Critical	High	2	600	N/A	
Fan 8	Enabled	Enabled	Critical	High	2	600	N/A	
Temp 1	Enabled	Enabled	Critical	Low	2	0	+50	
Temp 2	Enabled	Enabled	Critical	Low	2	0	+60	
Temp 3	Enabled	Enabled	Critical	Low	2	0	+50	
Temp 4	Enabled	Enabled	Critical	Low	2	0	+60	
Temp 5	Enabled	Enabled	Critical	Low	2	0	+50	
Temp 6	Enabled	Enabled	Critical	Low	2	0	+60	
Temp 7	Disabled	Disabled	Critical	Low	2	0	+50	
Temp 8	Disabled	Disabled	Critical	Low	2	0	+60	
Input 1	Enabled	Enabled	Minor	Low	1	N/A	N/A	
Input 2	Enabled	Enabled	Minor	Low	1	N/A	N/A	
Input 3	Enabled	Enabled	Minor	Low	1	N/A	N/A	
Input 4	Enabled	Enabled	Minor	Low	1	N/A	N/A	
Input 5	Enabled	Disabled	Minor	Low	1	N/A	N/A	
Input 6	Enabled	Disabled	Minor	Low	1	N/A	N/A	
Input 7	Enabled	Disabled	Minor	Low	1	N/A	N/A	
Input 8	Disabled	Disabled	Minor	Low	1	N/A	N/A	
+12V	Enabled	Enabled	Minor	N/A	N/A	11.4	12.59	
+5V	Disabled	Disabled	Minor	N/A	N/A	4.751	5.25	
+5Vaux	Enabled	Enabled	Minor	N/A	N/A	4.751	5.25	
+3.3V	Enabled	Enabled	Minor	N/A	N/A	3.135	3.466	
-12V	Disabled	Disabled	Disabled	N/A	N/A	-11.4	-12.6	

Appendix 3:

A3.1 Disclaimer

The information contained in this manual may be subject to technical alteration as a result of constant upgrading by One Stop Systems products. The attached documentation does not entail any guarantee on the part of One Stop Systems with respect to technical processes described in the manual or any product characteristics set out in the manual.

One Stop Systems does not accept any liability for any printing errors or other inaccuracies in the manual unless it can be proven that One Stop Systems is aware of such errors or inaccuracies or that One Stop Systems is unaware of these as a result of gross negligence and One Stop Systems has failed to eliminate these errors or inaccuracies for this reason. One Stop Systems expressly informs the user that this manual only contains a general description of technical processes and instructions, which may not be applicable in every individual case. In cases of doubt, please contact One Stop Systems.

Note on the Guarantee: Due to their limited service life, parts which by their nature are subject to a particularly high degree of wear (wearing parts) are excluded from the guarantee beyond that provided by law. This applies, for example, to batteries.

Exemption from the Guarantee Obligation: One Stop Systems shall be exempted from its guarantee obligations if the user fails to observe the safety instructions. One Stop Systems can only guarantee the safety, reliability and performance of the device if all of the safety instructions are observed.

Exclusion of Accident Liability Obligation: One Stop Systems shall be exempted from the statutory accident liability obligation if the user fails to observe the safety instructions.

Liability Limitation: In the event of damage to the device caused by failure to observe the safety regulations, One Stop Systems shall not be required to honor the guarantee even during the guarantee period and shall be exempted from the statutory accident liability obligation.

Appendix 4

A4.1 Limited Warranty

One Stop Systems warrants this product to be free of defects in material and workmanship p for an initial period of two years from date of delivery to the original purchaser. During this period, One Stop System s will, at its option, repair or replace this product at no additional charge to the purchaser, except as set forth in this warranty agreement.

One Stop Systems will, at its option, repair or replace this product at no additional charge to the purchaser, if the defect is related to One Stop Systems' manufactured product, such as a power supply, backplane, other chassis components or CPUs. One Stop Systems is not liable for any defects in material or workmanship of any peripherals, products or parts, which One Stop Systems does not design or manufacture. However, One Stop Systems will honor the original manufacturer's warranty on these products.

One Stop Systems will analyze the defective component and the customer will be charged in the following instances:
• No problem found: $125 (U.S. dollars).
• Damage: Parts and labor at $125 per hour with a $100 minimum charge (U.S. dollars). Receipt of damaged goods voids the One Stop Systems warranty.

Repair parts and replacement products will be furnished on an exchange basis and will be either new or reconditioned. All replacement parts and products shall become the property of One Stop Systems, if such parts or products are provided under this warranty agreement. In the event a defect is not related to the One Stop Systems manufactured product, One Stop Systems shall repair or replace the defective parts at the purchaser's cost and deliver the defective parts to the purchaser.

This limited warranty shall not apply if the product has been misused, carelessly handled, defaced, modified or altered, or if unauthorized repairs have been attempted by others. The above warranty is the only warranty authorized by One Stop Systems and is in lieu of any implied warranties, including implied warranty of merchantability and fitness for a particular purpose.

In no event will One Stop Systems be liable for any such damage as lost business, lost profits, lost savings, downtime or delay, labor, repair or material cost, injury to person or property or any similar or dissimilar consequential loss or damage incurred by the purchaser, even if One Stop Systems has been advised of the possibility of such losses or damages.

In order to obtain warranty service, the product must be delivered to One Stop Systems' facility, or to an authorized One Stop Systems service representative, with all included parts and accessories as originally shipped, along with proof of purchase and a Returned Merchandise Authorization number.

The RMA number is obtained, in advance, from One Stop Systems' Customer Service Department and is valid for 30 days. The RMA number must be clearly marked on the exterior of the original shipping container. Purchaser will be responsible and liable for any missing or damaged parts. Purchaser agrees

to pay for shipping charges one way, and to either insure the product or assume the liability for loss or damage during transit.

Ship to:
One Stop Systems
ATTENTION: RMA REPAIR DEPARTMENT
RMA ####
2235 Enterprise Street, Suite 110
Escondido, CA 92029