Technical white paper

Architecture and Technologies in the HP BladeSystem c3000 Enclosure

Table of contents

Introduction

The c3000 enclosure is a wire-once, cloud-ready system architecture designed both to deploy into traditional enterprise solutions, and build on top of existing infrastructure with applications and virtualization. And when you're ready, you can move into cloud-based deployments with HP CloudSystem. HP CloudSystem is at the core of our cloud-ready Converged Infrastructure which is part of the HP Converged Cloud architecture. HP CloudSystem offers a complete system to build and manage services across public, private, and hybrid clouds. You can do all this with the HP BladeSystem infrastructure as the foundation with upgrade along the way. We realize that budgets are constrained and the underlying technologies engineered in HP BladeSystem help maximize every hour, watt, and dollar.

As a part of HP BladeSystem, the c7000 enclosures use an Intelligent Infrastructure managed by Onboard Administrator (OA) and iLO Management Engine to monitor health and dynamically allocate power and cooling resources for maximum efficiency. HP Thermal Logic technologies such as Active Cool fans, Insight Control Suite power management software, Enclosure Dynamic Power Capping, and Dynamic Power Capping can reduce energy consumption, reclaim capacity, and extend the life of your data center. With a passive high capacity NonStop signal midplane, you can meet the performance needs of your most demanding applications and their availability requirements. With the Insight Display or Insight Control manager, administrators can perform many functions locally or with remote console stations to examine a message or change a basic configuration.

The c3000 enclosure is designed for remote sites, small and medium-sized businesses, and data centers with special power and cooling constraints. It is compatible with HP ProLiant and Integrity server blades, networking interconnect modules, and HP Insight Control management software. It represents an evolution of the entire rack-mounted infrastructure consolidating and repackaging all the featured infrastructure elements—computing, storage, network, and power—into a single infrastructure-in-a-box that accelerates the integration and optimization of the data center.

This technical whitepaper gives you an overview of the HP BladeSystem c3000 enclosure This includes new features in the c3000 Platinum enclosure like the 80 Plus certified 1200W Platinum Hot Plug power supply. You will also read about new and updated ProLiant Gen8 compliant technologies for the c3000 Platinum enclosure such as Power Discovery and Thermal Logic power/cooling technologies,

For information about enclosure-level and system-level management, read the "Management architecture of HP BladeSystem c-Class Systems" technical whitepaper at: hp.com/bc/docs/support/SupportManual/c00814176/c00814176.pdf. This BladeSystem management paper includes information about BladeSystem enclosure compatibility with HP Insight Control 7.2, and RDP Migration to the new Insight Control server provisioning. You can find out more about HP Insight Control server provisioning at: hp.com/us/en/enterprise/servers/software/insight-control/deployment.aspx

HP BladeSystem design

HP BladeSystem design and engineering gives you industry-leading performance, power efficiency, and thermal control. Embedded management enhancements include Integrated Lifecycle Automation capabilities enabled by innovations such as Intelligent Provisioning for easy system set-up, Active Health for agentless hardware monitoring and alerting, and Smart Update for automated firmware and system software maintenance. The BladeSystem c-Class enclosure and ProLiant Gen8 server blades utilize HP Insight Management software to automate key management processes, including a system's physical deployment, configuration, and problem management. ProLiant Gen8 server blades enable your organization to consolidate physical servers and components while still maintaining the same workload capacity and performance.

The BladeSystem c-Class enclosure includes many blades and interconnect device options:

- ProLiant server blades using AMD Opteron™ or Intel® Xeon® x86 processors
- ProLiant workstation blades
- Integrity server blades
- HP Storage blades
- HP Tape blades/solutions
- PCI-X or PCI Express (PCIe) expansion blades

BladeSystem interconnect modules feature a variety of networking standards:

- Ethernet
- Fibre Channel
- Fibre Channel over Ethernet (FCoE)

- InfiniBand
- iSCSI
- Serial Attached SCSI (SAS)

For a full listing of compatible interconnects, see the HP BladeSystem c3000 Platinum Enclosure QuickSpecs documentation at: hp.com/products/quickspecs/12790_na/12790_na.html

The architecture of the c-Class enclosure provides a basis for broad solutions, including the HP CloudSystem (hp.com/go/cloudsystem). HP CloudSystem includes shared IT infrastructure services by integrating pools of computing, storage, and networking capabilities with management tools.

The c-Class enclosure interoperates and connects with other HP infrastructure pieces, including external storage components such as DAS (direct-attached storage), NAS (network attached storage), and SAN (storage area network) solutions (hp.com/go/blades/storage).

The design features flexibility:

- Blade form factors that can scale vertically or horizontally —half-height or full-height blades and single- or, double- wide blades (depending on whether you use the c3000 or c7000 enclosure)
- Interconnect module form factors that can scale—single-wide or double-wide modules
- Uplinks to connect up to seven enclosures
- Signal midplane that allows flexible use of I/O signals, including multiple fabrics using the same traces

c3000 enclosure physical infrastructure

While the BladeSystem c7000 enclosure[1] is optimized for enterprise data center applications, the c3000 enclosure is intended for other computing environments such as remote sites or small businesses. The enclosure design accommodates sites that may not have any special cooling capability; and can exist in environments of up to 35 degrees centigrade. The c3000 enclosure includes use of management devices such as local KVM switches for local administration.

The HP BladeSystem c3000 enclosure infrastructure has redundant signal paths between servers and interconnect modules. The c3000 Platinum enclosure includes a shared, 10 terabit per second high-speed NonStop midplane for wire-once connectivity of server blades to network and shared storage. Power is delivered through a pooled-power backplane that ensures the full capacity of the power supplies is available to all server blades for maximum flexibility and redundancy. The c3000 incorporates Intelligent Infrastructure, managing power with the 1200W Platinum Hot Plug Power Supply (80 Plus certified), Single Phase Intelligent Power Module, and Intelligent Power Distribution Units (iPDUs).

The HP BladeSystem c3000 enclosure provides the following capabilities:

- It fits into standard-size HP and third-party racks.
- It accommodates BladeSystem c-Class server blades, storage blades, and interconnect modules.
- It supplies all the power, cooling, and I/O infrastructure for the c-Class components.

You can populate the c3000 enclosure with the following components:

- Up to four full-height (FH) or eight half-height (HH) server and/or storage blades per enclosure
- Up to four interconnect modules simultaneously featuring a variety of network interconnect fabrics such as Ethernet, Fibre Channel (FC), InfiniBand (IB), Internet Small Computer System Interface (iSCSI), or Serial-attached SCSI (SAS)
- Active Cool fan kits for a maximum of six fans
- Up to six power supplies with either low-line or high-line power input[2]
- Single or dual BladeSystem Onboard Administrator (OA) management modules
- DVD drive
- Optional KVM enclosure module for connecting the c3000 to an in-rack KVM switch or HP TFT 7600 Rack Mount Keyboard/Monitor

Figures 1 and 2 show the front and rear views of the HP BladeSystem c3000 Enclosure. The 6U high enclosure accommodates servers, interconnects, mezzanine cards, storage blades, power supplies, and fans.

[1] More information on ProLiant c-Class architecture and the c7000 enclosure is available at hp.com/servers/technology
[2] Typically, only four power supplies are required, especially in sites where non redundant AC input is acceptable.

Figure 1.

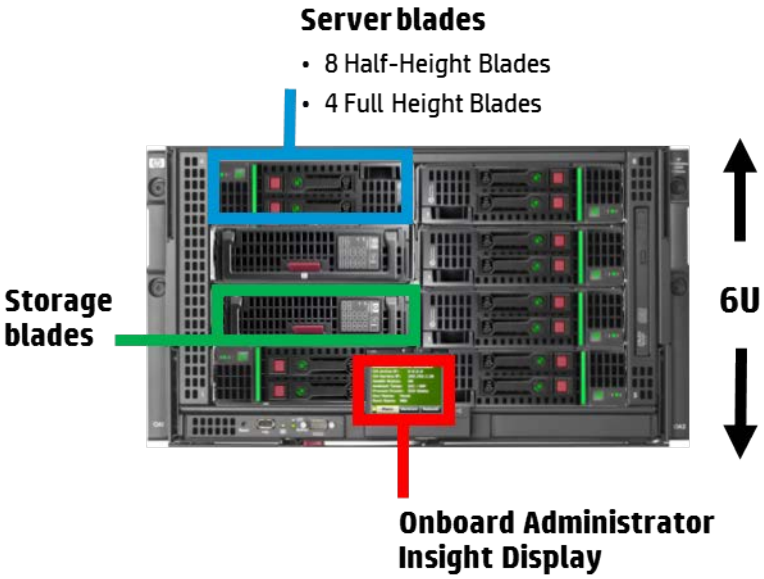

Figure 2.

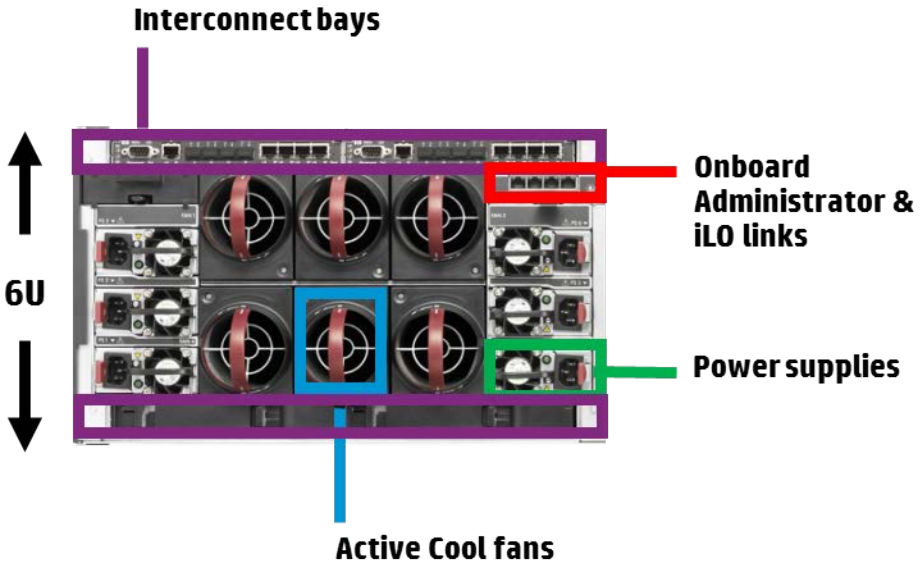

The passive NonStop signal midplane and separate power backplane in the enclosure have no active components. Separating the power delivery in the backplane from the high-speed interconnect signals in the midplane results in minimal thermal stress to the signal midplane. A pooled-power backplane delivers power and ensures that the full capacity of the power supplies remains available to all server blades and interconnects.

HP BladeSystem c3000 Platinum enclosure is available worldwide as a single-phase enclosure for flexibility in connecting to datacenter power or alternatively, in countries with low-line (100VAC to 110VAC) power outlets, the enclosure's power supplies can be connected directly to low-line wall outlets. You also have a choice of AC or DC power supplies.

BladeSystem Onboard Administrator cabling

The standard BladeSystem Onboard Administrator module is pre-installed in a front-loading tray (Figure 3) that also houses the Insight Display. The module contains a serial connector to connect to a PC with a null-modem RS232 serial cable. A USB connector is also available for future USB connectivity.

Figure 3.

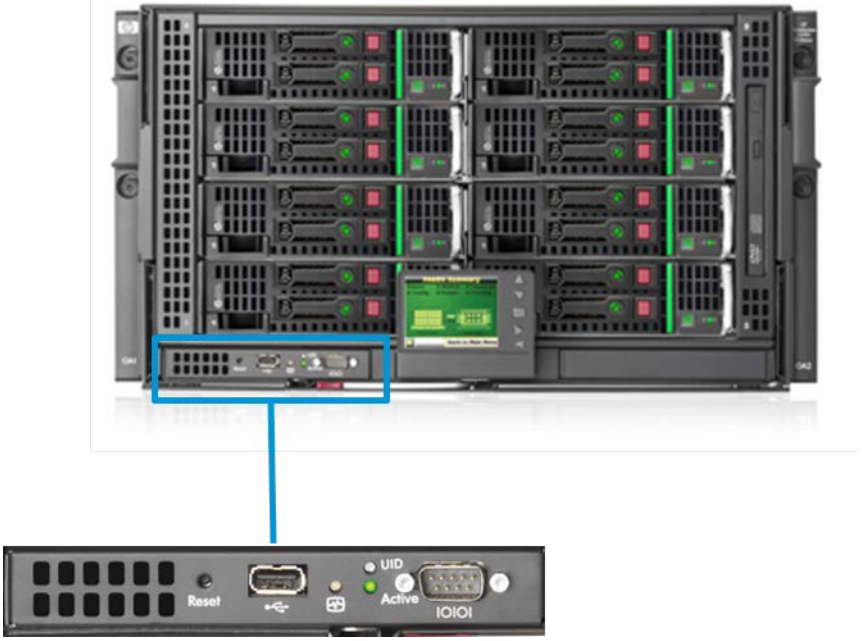

A separate rear-loading link module contains RJ-45 ports for enclosure link-up/link-down connectivity and BladeSystem Onboard Administrator network access (Figure 4).

Figure 4.

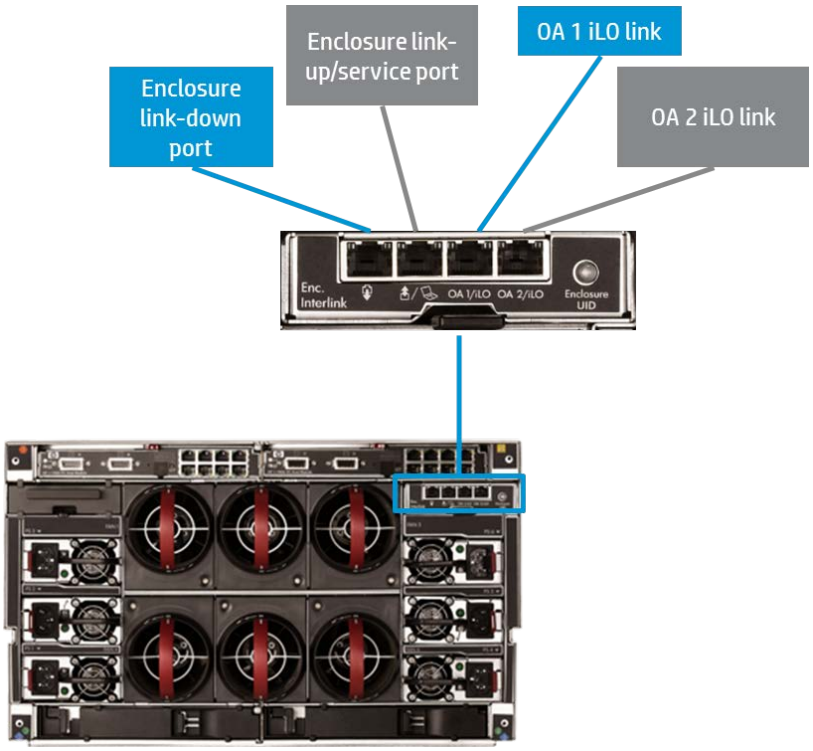

Enclosure-based DVD ROM

The HP BladeSystem c3000 Enclosure has an optional CD/DVD ROM drive that installs in the front of the enclosure. The Insight Display and BladeSystem Onboard Administrator allow system administrators to connect (or disconnect) the media device to (from) one or multiple servers at a time. In addition, a browser-based console is available through the iLO functionality of each server blade. The console gives you additional options:

- Perform local OS deployment, updates, and patches
- Update server platform ROMs

The enclosure-based CD/DVD offers local drive access to server blades by using the BladeSystem Onboard Administrator or Insight Display. When media is loaded in the enclosure-based DVD ROM, local administrators can use the Insight Display to attach the media device to a single or to multiple server blades simultaneously.

KVM Module

The optional KVM Module plugs into the rear bay adjacent to interconnect module 1 and provides a VGA connector and two additional USB connectors for the c3000 enclosure. An external VGA monitor and external USB keyboard/mouse can provide access not only to all server video consoles, but also to the OA command line interface (CLI), and/or to the Insight Display. Using "PrintScrn" as a hot key to switch consoles, the user can select a particular server console, control the server power, or connect to the enclosure DVD from the KVM menu screen (Figure 5). Besides allowing the user to select a server video console, the menu provides current server health status, power status, and DVD connect status. The OA automatically provides a server name, based on server information. From a server video console session, the user presses PrintScrn to return to the KVM menu. The CLI console provides a text screen to log in and run command-line commands to the OA. The Insight Display provides all the enclosure's Insight Display screens on the KVM monitor and uses the KVM keyboard to navigate those screens using the KVM station.

Figure 5.

Scalable blade form factors and device bays

The half-height and full-height blade form factors, which scale blades horizontally in the c3000 enclosure and vertically in the c7000 enclosure, provide several benefits. They include reduced cost, increased reliability, and improved ease-of-use. Placing full-height form factors or half-height form factors (or both in the c7000 enclosure) in the same enclosure lets you exploit its flexibility.

The thickness of a device bay provides several advantages over narrower or wider alternatives:

- It holds industry-standard components
- It holds enough blades to amortize the cost of the shared enclosure infrastructure (power supplies and fans)
- It uses cost-effective, standard-height DIMMs in the server blades

- It uses vertical, rather than angled, DIMM connectors to give better signal integrity and more room for heat sinks. The vertical DIMM connectors also allow more DIMM slots per processor and provide better airflow across the DIMMs.

Figure 6 shows BladeSystem c-Class form factors, scaling vertically with half-height and full-height blades.

Figure 6.

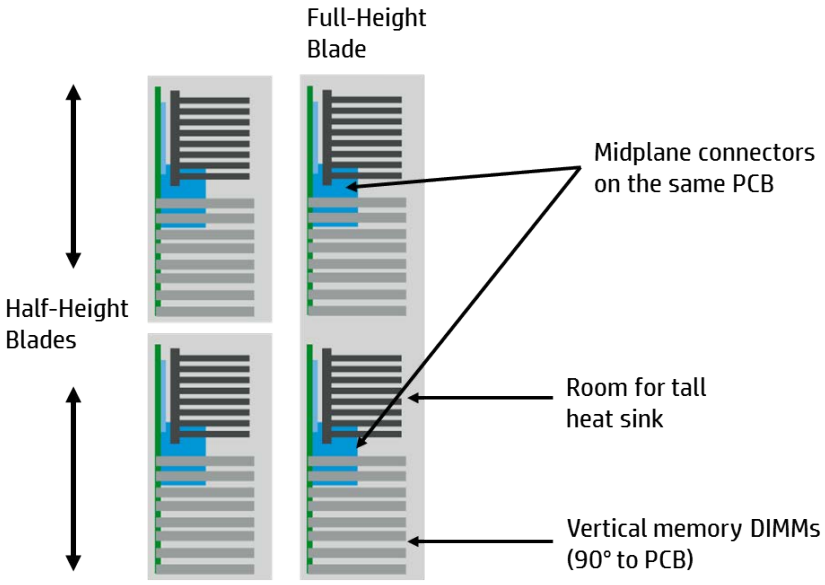

You can use two full-height bays in combination for a full-height double-wide blade form factor. For example, HP Integrity server blades can combine multiple blades to create 2-, 4-, or 8-socket systems, as shown in Figure 4. Each base server blade has a Blade Link connector. The Blade Link connector joins selected QPI (Quick Path Interconnect) ports among processors, the required clock signals, and side band signals for the system to operate as a scale-up multiprocessor system. Blade Link is available only with Integrity blades.

Figure 7 shows horizontal scaling with Integrity blades and Blade Link connections.

Figure 7.

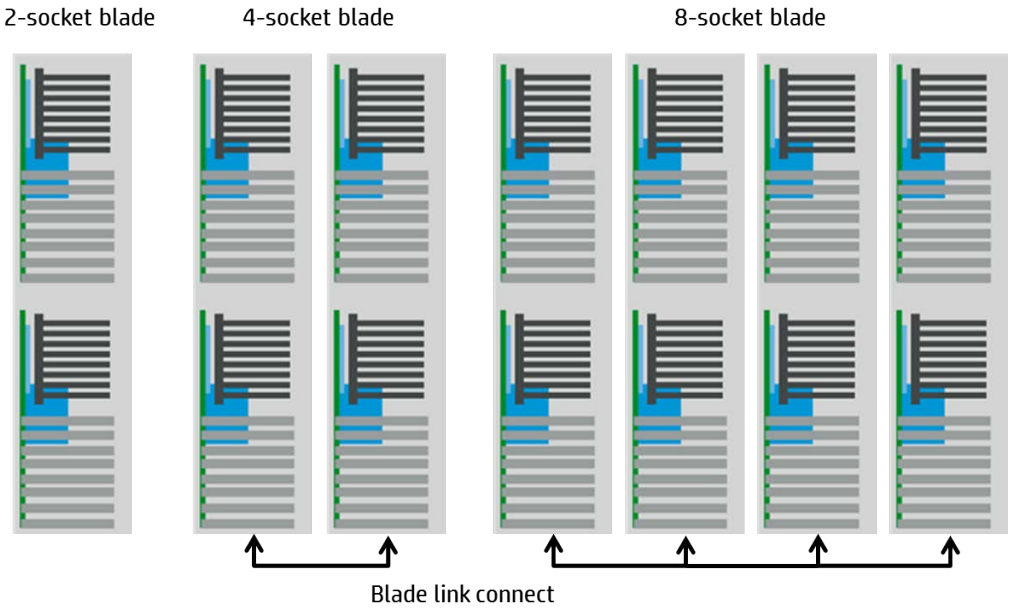

Scalable interconnect form factors

The interconnect bays offer network redundancy in the form of two single-wide or double-wide form factor slots for efficient use of space and improved performance. The eight interconnect bays are designed to accommodate two single-wide redundant interconnect modules in a scale-out configuration or a larger, higher-bandwidth double-wide interconnect module for higher performance interconnect modules (Figure 8).

Figure 8.

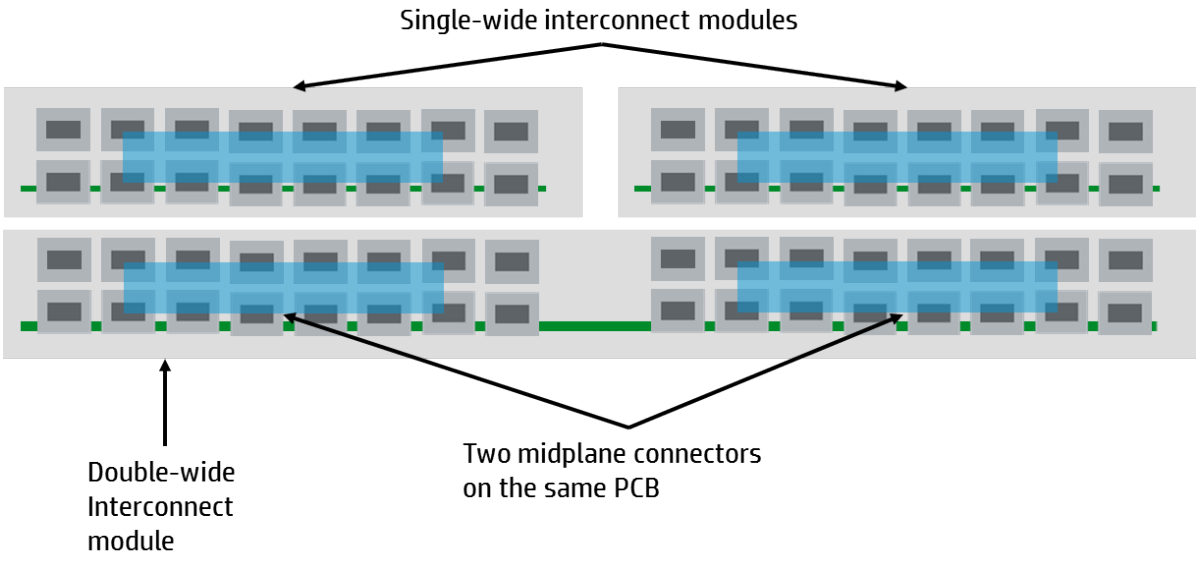

The scalable interconnect form factor provides similar advantages as the scalable device bays:

- It accommodates the maximum number of interconnect modules.
- It allows enough space in a double-wide module to include two signal connectors on the same PCB plane. This affords reliable and simple connectivity to the NonStop signal midplane.

Star topology

The device bays and interconnect bays connect in a fan-out, or star topology centered around the interconnect modules. The exact topology depends on your configuration and the enclosure. For example, if you place two single-wide interconnect modules side-by-side (Figure 9a), the architecture is a dual-star topology. Each blade has redundant connections to the two interconnect modules. If you use a double-wide interconnect module, it is a single star topology, providing more bandwidth to each server blade. Figure 9b shows the redundant configuration using double-wide interconnect modules.

Figure 9.

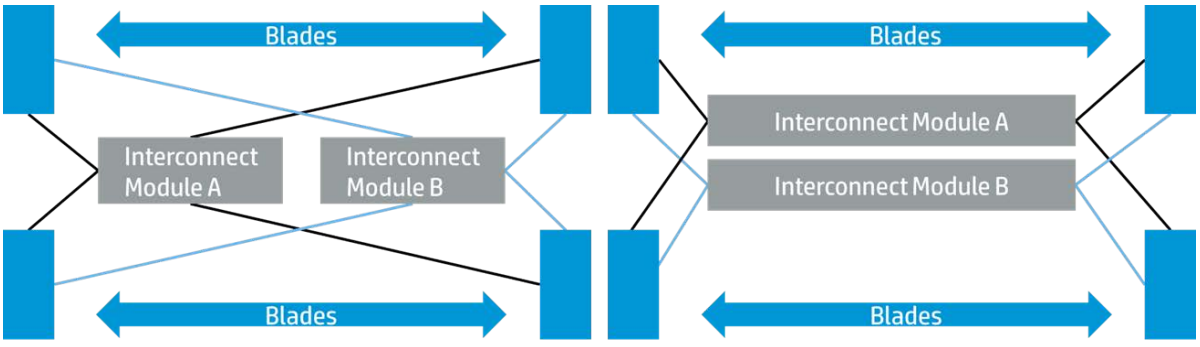

9a 9b

Signal midplane design and function

In addition to the scalable server blade and device bay form factor, the c7000 enclosure includes a high-band width NonStop signal midplane and a separate power backplane.

NonStop signal midplane scalability

A key component of the BladeSystem c3000 and c7000 enclosures is the I/O infrastructure, a NonStop signal midplane that provides the internal wiring between the server or storage blades and the interconnect modules. The NonStop signal midplane is an entirely passive board that uses serializer/deserializer (SerDes) technology to incorporate multiple protocols and provide point-to-point connectivity between device bays and interconnect bays. The term passive means there are no active electrical components on the board. On one side of the board are the connectors for the server or storage blades. Internal traces link them to the eight connectors on the other side of the board for the interconnect modules (Figure 10).

Figure 10.

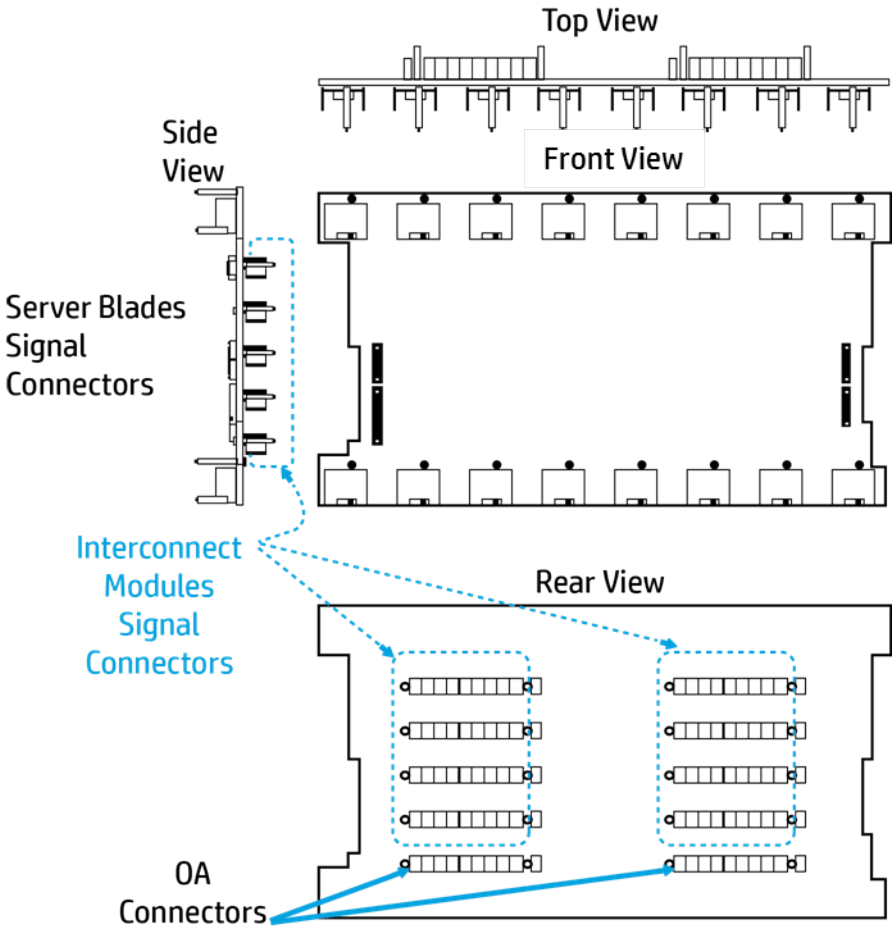

The NonStop signal midplane features signal rates of up to 10 gigabits per second (Gb/s) per lane (each lane consists of four SerDes transmit/receive traces). Each half-height server bay has 16 lanes delivering the cross-sectional bandwidth to conduct up to 160 Gb/s per direction.

In the c3000 enclosure, the aggregate bandwidth between device bays and interconnect bays is 2.5 Terabits per second (Tb/s) across the NonStop signal midplane. The aggregate bandwidth does not include additional traffic capacity between interconnect modules or cross-bay connections. Current half-height blades use 12 of the 16 lanes to the interconnect bays.

One of the areas of improved design focus was high-speed signal integrity. Getting this level of bandwidth between bays required special attention to high-speed signal integrity:

- Using general best practices for signal integrity to minimize end-to-end signal losses across the signal midplane
- Moving the power into an entirely separate backplane to independently optimize the NonStop signal midplane
- Providing a method to set optimal signal waveform shapes in the transmitters, depending on the topology of the end-to-end signal channel

The signal midplane also includes the management signals from each bay to the OA modules. These management signals are isolated from the high-speed server-to-interconnect signals. The OA is the terminating point for all interconnect bays. An interconnect module cannot use the connection to the OA to communicate with another interconnect module. Interconnect module pairs (side-by-side) have cross-connect capabilities.

By taking advantage of the similar four-trace, differential SerDes transmit and receive signals, the NonStop signal midplane is compatible with either network semantic protocols (such as Ethernet, Fibre Channel, and InfiniBand) or memory semantic protocols (PCI Express), using the same signal traces.

Figure 11 illustrates how you can logically overlay the physical lanes onto sets of four traces. Interfaces such as Gigabit Ethernet (1000 base-KX) or Fibre Channel need a 1x lane, or a single set of four traces. Higher bandwidth interfaces, such as InfiniBand DDR, use up to four lanes (4x). Each half-height bay connects with 16 lanes to the midplane to provide each bay a design limit of 160 Gbps in each direction to the interconnects. Current half-height server blade offerings use up to 120 Gbps with LOMs and Mezzanine cards.

Figure 11.

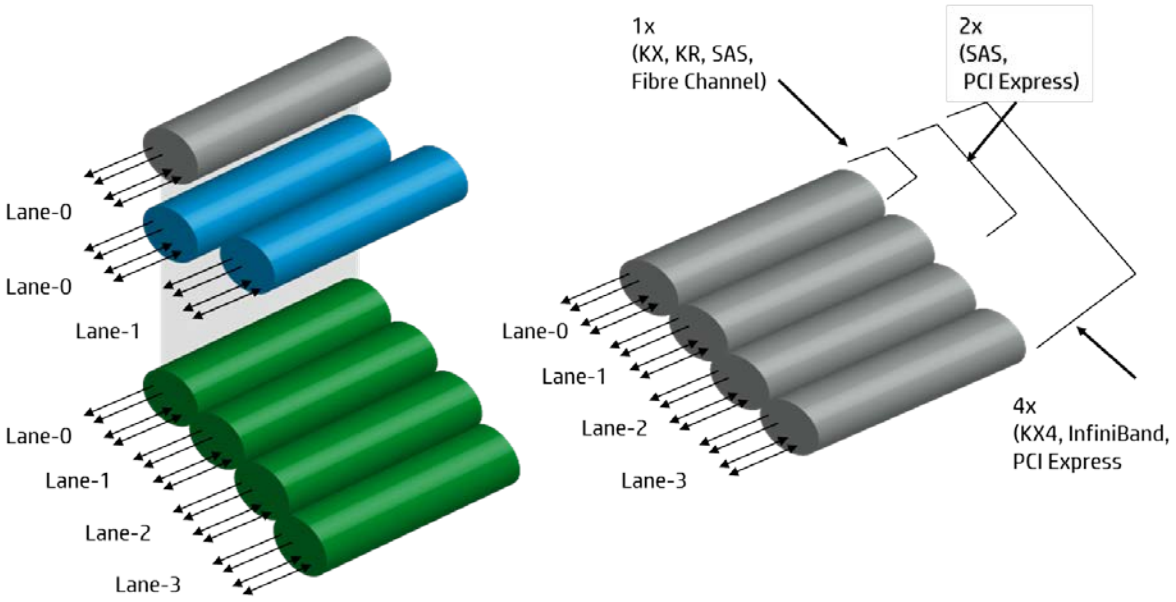

Note

Network-semantic interconnect protocols use network addresses in the packet headers to exchange data between two nodes such as, MAC addresses and IP addresses for Ethernet, world-wide port name for FC, or GUID for InfiniBand. Memory-semantic interconnect protocols use memory addresses in the packet headers to deposit or retrieve data where these addresses can be memory-mapped registers of a chip or system memory location.

The NonStop signal midplane has eight 200 pin connectors to include eight individual switches, four double bay switches, or a combination of the two. It provides the flexibility of 1x, 2x, or 4x connections from the server blade mezzanine cards, which connect the interconnect bays. The rear of the enclosure includes eight interconnect bays that can accommodate eight single or four redundant interconnect modules. All interconnect modules plug directly into these interconnect bays. Each c-Class enclosure requires two interconnect switches or two pass-thru modules, side-by-side, for a fully redundant configuration.

Best practices

To ensure high-speed connectivity among all blades and interconnect modules, we leveraged our many years of experience in designing HP Superdome computers. Specifically, our engineers paid special attention to

- Controlling the differential signal impedance along each end-to-end signal trace across the PCBs and through the connector stages
- Using a ground plane to isolate, receive, and transmit signal pins (see Figure 12)
- Keeping signal traces short to minimize losses
- Routing signals in groups to minimize signal skew

Figure 12.

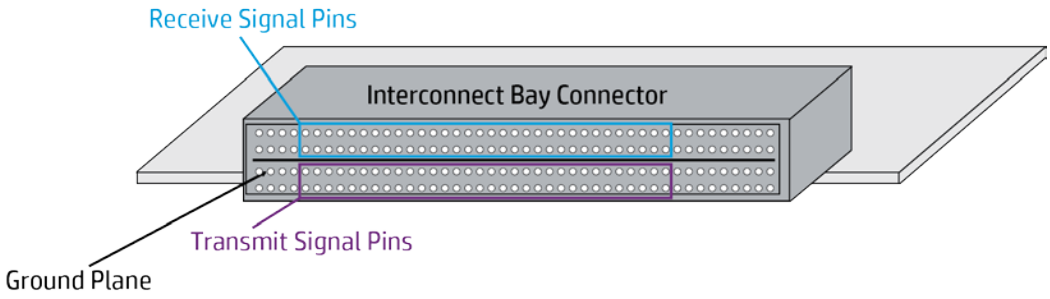

Channel topology and equalization settings

Even when using best practices, insertion and reflection losses can degrade high-speed signals transmitted across multiple connectors and long PCB traces. Insertion losses, such as conductor and dielectric material losses, increase at higher frequencies. Impedance discontinuities, primarily at connectors cause reflection losses. To compensate for these losses, shape the transmitter's signal waveform by selecting signal equalization settings. But a transmitter's equalization settings depend on the end-to-end channel topology and the type of component sending the signal. Both topology and the transmitting component can vary in the BladeSystem c-Class because of the flexible architecture and the use of mezzanine cards and NICs or other embedded I/O devices. As shown in Figure 13, the topology for device 1 on server blade 1 (a b c) is different from the topology for device 1 on server blade 4 (a-d-e). So, a link configuration mechanism in the Onboard Administrator (assisted by iLO on each server blade) identifies the channel topology for each device and configures the proper equalization settings for that device.

Figure 13.

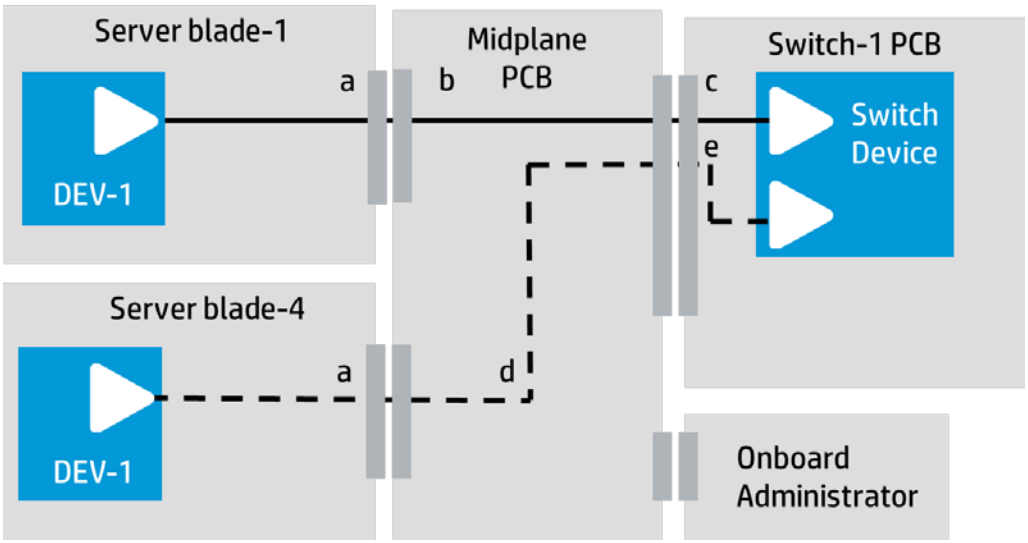

Signal midplane provides reliability

To provide high reliability, we designed the NonStop signal midplane as a completely passive board. It has no active components along the high-speed signal paths. The PCB consists primarily of traces and connectors. While there are a few components on the PCB, they are limited to passive devices that are unlikely to fail. The only active device is an EEPROM (Electrically Erasable Programmable Read-Only Memory) utility, which the Onboard Administrator uses to get information such as the midplane serial number. If this device fails, it does not affect the NonStop signal midplane.

Power backplane scalability and reliability

To minimize losses while distributing power, the power backplane has solid copper plates and integrated power delivery pins (Figure 14). Solid copper plates reduce voltage drops, produce high current density, and provide high reliability.

Figure 14.

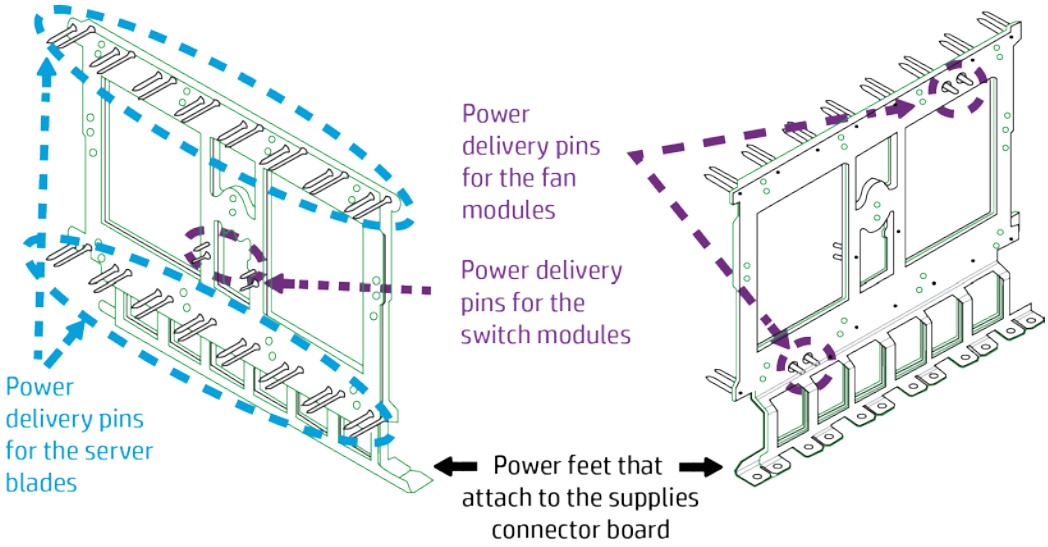

Enclosure connectivity and interconnect fabrics

HP Virtual Connect lets you virtualize your network, your storage and your power infrastructure so that resources can be dynamically allocated based on server workloads. The c3000 enclosure includes up to four different interconnect fabrics (such as Ethernet, Fibre Channel, InfiniBand, iSCSI, SAS) simultaneously within the enclosure.

HP Virtual Connect

HP Virtual Connect technology removes traditional infrastructure constraints by abstracting servers from their uplinks, simplifying management between networks and servers. HP Virtual Connect allows the creation of pools of network, storage, and compute resources. With wire-once technology you can add, move, or change servers in minutes. Virtual Connect includes the following components:

- HP Virtual Connect Flex-10 10Gb Ethernet Module
- HP Virtual Connect Flex-10D 10Gb Ethernet Module
- HP Virtual Connect 8Gb 20-Port Fibre Channel Module
- HP Virtual Connect 8Gb 24-Port Fibre Channel Module
- HP Virtual Connect Enterprise Manager

Virtual Connect implements server-edge virtualization so that you can add, replace, or move server blades without having to make changes to LAN and SAN connections. We recommend using Virtual Connect to reduce cabling and management overhead and improve flexibility of connection management. Virtual Connect modules slide into the interconnect bays of c-Class enclosures. When using Virtual Connect Fibre Channel modules, the enclosure must have at least one Virtual Connect Ethernet or FlexFabric module. The Virtual Connect Manager software runs on those modules. You can also use HP Virtual Connect Enterprise Manager software for centralized management of Virtual Connect domains with up to 16,000 server blades. After you have established the LAN and SAN connections to the pool of servers, you can use Virtual Connect Manager to define a server connection profile for each server. Instead of using the default media access control (MAC) addresses for all NICs and default World Wide Names (WWNs) for all host bus adapters (HBAs), the Virtual Connect Manager can assign unique MAC addresses and WWNs to these profiles from a pre-allocated pool of addresses.

You can establish all LAN and SAN connections once during initial deployment. If you later add, deploy, or change a server, you do not need the change a LAN and SAN connections because Virtual Connect keeps the profile for that server bay constant.

HP Virtual Connect includes Flex-10 technology, which allows you to fully utilize 10 GbE connection bandwidth. Using Flex-10, you can allocate the bandwidth of a single 10 Gb pipeline into multiple network connections called FlexNICs. With Virtual Connect FlexFabric modules and to FlexFabric Adapters, each 10 Gb port can be allocated across four network connections,

or FlexNICs, or one of those 4 connections can be a FlexHBA connection to use Fibre Channel over Ethernet or iSCSI traffic. In addition, you can dynamically allocate the bandwidth for each FlexNIC and FlexHBA connection by setting it to a user-defined portion of the total 10 Gb connection. You can set speed from 10 Mb per second to 10 Gb per second in 100 Mb increments.

You'll find advantages to partitioning a 10 GbE pipeline:

- More network connections per server, which is especially important in a virtual machine environment
- Ability to match bandwidths to the network function, such as management console or production data

You can use the HP Virtual Connect Flex-10 10Gb Ethernet Module or HP Virtual Connect FlexFabric 10Gb/24-port Module to manage the 10GbE (Flex-10) server connections to the data center network. The 10Gb Flex-10 and FlexFabric adapter LOMs and mezzanine cards are CNAs (Converged Network Adapters), each with two 10 Gb ports. Each 10 Gb port can be configured from one to a maximum of four individual FlexNICs. The server ROM and the operating system or hypervisor recognize each FlexNIC and FlexHBA as an individual NIC or HBA, respectively.

You can find the latest Virtual Connect Flex-10 and FlexFabric hardware components here: hp.com/go/vc

Interconnect modules

The BladeSystem c3000 enclosure offers a variety of interconnect options, including:

- Pass-thru modules
- Ethernet and Fibre Channel switches
- High-bandwidth fabrics such as InfiniBand
- SAS BL switches

Switches offer a traditional approach to administering a network. The primary value in switches is cable consolidation through high-speed uplinks and the shared power and cooling infrastructure.

Ethernet and Fibre Channel pass-thru modules are available when you require direct one-to-one connections between servers and either a LAN or a SAN. HP Ethernet and Fibre Channel Pass-Thru Modules provide 16-port, transparent, 1-to-1 port connectivity between the server and an external switch.

Interconnect modules in the BladeSystem c3000 enclosure are available in two widths:

- Single-wide modules provide sixteen internal ports, each connects to a separate device bay in the front of the enclosure.
- Double-wide modules provide sixteen double-wide internal ports providing connectivity to DDR InfiniBand and other 4-lane, high speed interconnects.

Each interconnect module also provides external connectors that vary based on its particular design.

Note
The c-Class Ethernet Pass-Thru Module is only compatible with fixed-speed gigabit Ethernet. Because the server, storage, or other optional blades are connected through SerDes to the interconnect bays, and SerDes Ethernet does not have an auto-negotiation protocol, you must use a switch to connect to 10/100 networks outside of the enclosure. The 10 Gb Ethernet Pass-Thru Module is compatible with 1 Gb or 10 Gb connections to the server. Note that this is a limitation of the 1 GB Ethernet Pass-Thru Module only. The Fibre Channel Pass-Thru Module ports have the auto-negotiate protocol.

Server blades

We build server blades for the BladeSystem c3000 enclosure according to c-Class standard form-factors referred to as half-height and full-height. The enclosure holds either full-height or half-height server blades or a combination of the two. We preconfigure the enclosure with device bay dividers to house half-height server or storage blades. To accommodate full-height server blades, you must remove dividers.

For connectivity, every server blade ships with at least two built-in Ethernet connections. You can install optional mezzanine cards for additional interconnect fabric connections such as 10 Gb Ethernet, InfiniBand, and Fibre Channel.

You can configure up to four different interconnect fabrics without sacrificing redundancy or performance. Here are some options:

- Half-height server blades typically have two embedded Gigabit NICs and two c-Class PCIe mezzanine option connectors. A half-height server configured with one dual-port Gigabit NIC mezzanine card and one quad-port NIC mezzanine card provides eight independent NICs.
- Full-height server blades typically have four embedded Gigabit NICs and three c-Class PCIe mezzanine option connectors. A full-height server configured with one dual-port and two quad-port Gigabit NIC mezzanine cards provides 14 independent Gigabit NICs.

Storage options inside the BladeSystem enclosure

Storage options inside the enclosure provide an alternative to local disk drives or SAN connectivity. The BladeSystem c3000 enclosure is compatible with several types of storage solutions. Tape and PCI option blades are also available for c-Class. Each of these blades increases configuration flexibility by adding options that do not fit inside the server blade.

Direct attach storage blades

We deliver direct attach storage for c-Class server blades with the following storage blades:

- The D2220sb Storage Blade: It features up to 12 hot plug small form factor (SFF) SAS or SATA Midline hard disk drives or SAS/SATA SSDs. The D2220sb Storage Blade features an onboard Smart Array P420i controller with 2GB flash-backed write cache, for increased performance and data protection.
- HP D2200sb Storage Blade: It includes up to 12 hot plug small form factor (SFF) SAS or SATA hard disk drives and solid state drives. For increased performance and data protection, the D2200sb features 1 GB flash-backed write cache and an embedded Smart Array P410i controller.

Note

You need the mezzanine Pass-Thru Option Kit to support the HP D2200sb PCIe Storage Blade with HP c-Class BladeSystem full-height server blades. Up to two storage blades are supported per server mix and match.

For mechanical compatibility, both direct attach storage blades use the same half-height form factor as server blades. The enclosure backplane provides a PCIe connection from the storage blade to the adjacent c-Class server blade. The design enables high performance storage access without any additional cables.

You must pair the direct attach storage blade with an adjacent server blade in the same cooling zone. That's because the physical connection between the direct attach storage blade and its adjacent server blade is a dedicated x4 PCIe connection across the NonStop midplane that connects the adjacent bays. You must pair the half-height server blade with a direct attach storage blade in specific bays. Refer to the appendix for bay number details. You do not need a mezzanine card to connect a half-height server blade to an adjacent direct attach storage blade.

Some full-height server blades include up to two direct attach storage blades. You might need a mezzanine card to connect specific full-height server blades to one or two direct attach storage blades. See the documentation that ships with the mezzanine card and the server blade for installation requirements.

Before installing a direct attach storage blade with a full-height server blade in a BladeSystem c3000 enclosure, install the half-height blade shelf on the direct attach storage blade first. Refer to the storage-blade installation guide for instructions. If you are installing two direct attach storage blades with one partner full-height server blade in an HP BladeSystem c3000 enclosure, use the mini divider instead of the half-height blade shelf. See the HP BladeSystem c3000 Enclosure Quick Setup Instructions for more information.

Note

When you use direct attach storage blades with a full-height server blade, the first direct attach storage blade must be in the bottom bay. To maintain essential airflow, you must install a blank above the storage blade to block the empty upper bay. Or, you must insert a half-height server blade or second direct attach storage blade into the upper bay. For this configuration, you should install the first storage blade before installing the half-height server blade or second storage blade. And you should remove the half-height server blade or second storage blade before removing the storage blade.

Shared storage

With the HP S P4000 Virtual SAN Appliance (VSA) Software, you can use the D2220sb or D2200sb as an iSCSI SAN for use by all servers in the enclosure and any server on the network. You must install LeftHand VSA software in a virtual machine on a VMware ESX host server adjacent to the storage blades. LeftHand P4000 VSA software features storage clustering for scalability, network RAID for storage failover, thin provisioning, snapshots, remote replication, and cloning. You can expand

capacity within the same enclosure or to other BladeSystem enclosures by adding additional D2220sb or D2200sbstorage blades and LeftHand VSA software licenses.

External SAS connectivity with direct connect SAS storage for HP BladeSystem
With direct connect SAS storage for HP BladeSystem, you can build local server storage with zoned storage. Alternatively, you can enable low-cost shared storage within the rack with high performance 3 Gb/s or 6 Gb/s SAS architecture. Keep in mind the following considerations when using external SAS storage:

- Each HP server blade requires an HP Smart Array P700m, P711m, or P712m controller installed in a mezzanine slot for access to external storage.
- You must install single or redundant HP StorageWorks 3 Gb or 6 Gb SAS BL switches in the interconnect bays of the enclosure.
- Depending on application requirements, you must connect the switches through an SAS cable to external storage.

NAS/SAN Gateway
The X3800sb Network Storage Gateway Blade is a ready-to-deploy SAN gateway solution that has Microsoft® Windows® Storage Server 2008 R2 Enterprise x64 Edition pre-installed. You can use the X3800sb to access Fibre Channel, SAS, or iSCSI SAN storage. You can also use it to translate file data from the server into blocks for storage to provide consolidated file, print, and management hosting services clustered together.

Storage server
The X1800sb Network Storage Blade comes with Microsoft Windows Storage Server 2008 R2 Standard x64 Edition pre-installed with Microsoft iSCSI Software Target and HP Automated Storage Manager Software included. You can pair the X1800sb with the D2220sb or D2200sb storage blade to create shared storage and file serving inside the BladeSystem enclosure. You can also use the X1800sb as a low cost gateway to external Fibre Channel, SAS, or iSCSI storage.

Mezzanine cards

We offer a variety of mezzanine card options to provide connectivity to networks and storage. HP ProLiant c-Class server blades use two types of mezzanine cards to connect to the various interconnect fabrics such as Fibre Channel, Ethernet, serial-attached SCSI, or InfiniBand.

Type I and Type II mezzanine cards differ in the power allocated to them by the server and in the physical space they occupy on the server blade. Type I mezzanine cards have slightly less power available and are slightly smaller. You can use Type I mezzanine cards with all ProLiant c-Class server blades in all mezzanine connectors (Table 1). You can use Type II mezzanine cards with Mezzanine 2 or 3 connectors in full-height c-Class server blades. You can also use Type II mezzanine cards with Mezzanine 2 connectors in half-height c-Class server blades.

Table 1.

	c-Class server blades	Type I mezz card	Type II mezz card
Mezz connector 1	Full height	yes	yes
	Half height	yes	no
Mezz connector 2	Full height	yes	yes
	Half height	yes	yes
Mezz connector 3	Full height	yes	yes

For the most up-to-date information about the c Class mezzanine card options, go to the HP website: hp.com/products/blades/components/c-class-interconnects.html.

Fabric connectivity and port mapping

Each enclosure requires interconnects to provide network access for data transfer. The interconnects reside in interconnect bays located on the rear of the enclosure (Figure 15). The server blades and enclosure include up to three independent interconnect fabrics, such as Ethernet, Fibre Channel, InfiniBand, and Virtual Connect modules.

Figure 15.

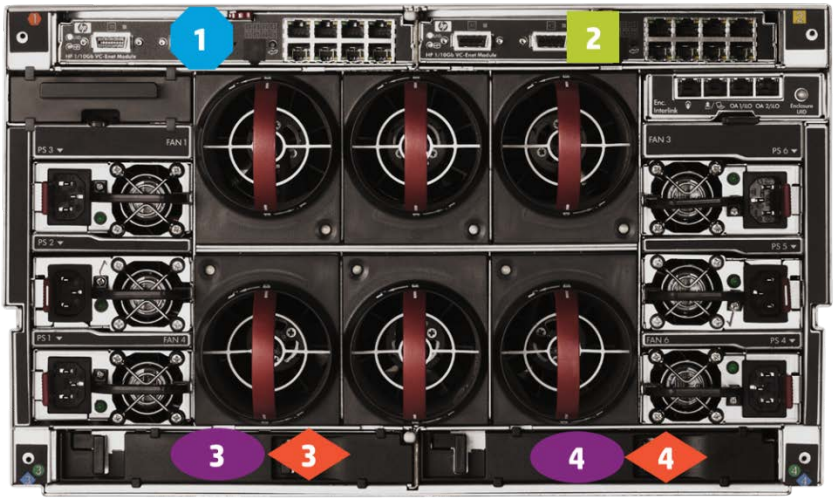

For interconnect bay mapping purposes, it does not matter in which device bay a server blade is installed. The mezzanine connectors always connect to the same interconnect bays. Because the connections between the device bays and the interconnect bays are hard-wired through the NonStop signal midplane, the server mezzanine cards must be matched to the appropriate type of interconnect module. For example, a Fibre Channel mezzanine card must be placed in the mezzanine connector that connects to an interconnect bay holding a Fibre Channel switch.

Embedded NICs and adapters are attached by single-wide interconnects in interconnect bay 1. Mezzanine 1 devices are attached by single-wide interconnects in internconnect bay 2 (Table 2). Mezzanine 2 and 3 can be attached by single-wide or double-wide interconnects such as InfiniBand or 10 Gb Ethernet devices in interconnect bays 3 and 4.

Table 2.

Server blade signal	**Interconnect bay**	**Interconnect label**
NIC 1, 2, 3, 4 (embedded)	1	Blue octagon
Mezzanine 1	2	Green square
Mezzanine 2	3 and 4	Purple oval
Mezzanine 3	3 and 4	Orange diamond

An internal connection on the midplane between interconnect bays 1 and 2 and an additional connection between interconnect bays 3 and 4 provide an internal link for use as a crosslink port between interconnect bays 1 and 2 or interconnect bays 3 and 4. NIC teaming can be configured between embedded NICs and Mezzanine 1 NICs using the internal crosslinks between the switches through this internal connection.

Several port types are referenced in Figures 16 and 17:

- Examples of 1x ports are 1-Gb Ethernet (1-GbE) pass-thru modules and Fibre Channel interconnect modules.
- An example of a 2x port is a SAS interconnect module.
- Examples of 4x ports are 10-GbE pass-thru modules and InfiniBand interconnect modules.

A full-height server blade plugs into two device bay connectors and has 32 lanes available to the 4 interconnect bays. Interconnect bay 1 is reserved for Ethernet interconnects. It connects embedded Ethernet NICs to the internal facing ports on the Ethernet interconnect. Depending on the configuration requirements, additional mezzanine cards and interconnects can be employed:

- Mezzanine 1 and Interconnect Bay 2
- Mezzanine 2 and Interconnect Bays 3 and 4
- Mezzanine 3 and Interconnect Bays 3 and 4

The full-height server blade has four embedded NICs and can accept up to three mezzanine cards. Each embedded NIC and optional mezzanine port is mapped through the signal midplane to specific ports on interconnect bays. A full-height server blade installed in device bay 1 would have NICs mapped in the following manner:

- NIC 1 (PXE default) — Interconnect bay 1 port 5
- NIC 2 — Interconnect bay 1 port 13
- NIC 3 — Interconnect bay 1 port 1
- NIC 4 — Interconnect bay 1 port 9

Figure 16. Port mapping for HP BladeSystem c3000 full-height server blades to interconnect bays

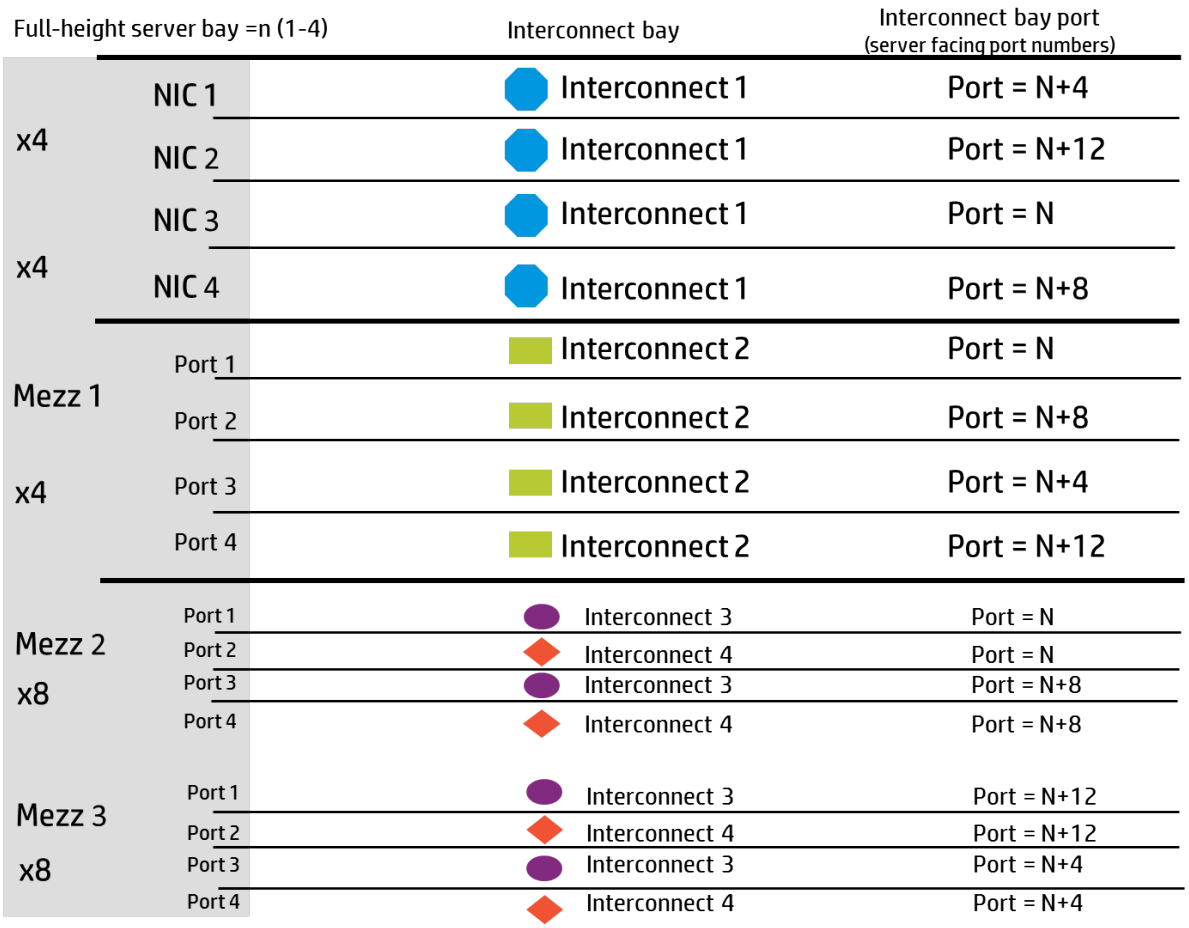

Half-height server blades connect to a single power and signal connector on the NonStop signal midplane. The remaining signal connector is allocated to the adjacent device bay (that is, device bays 1 and 5). As a result, half-height server blades do not use four-port mezzanine cards on connector 1, and they do not contain a Mezzanine 3 connector. The extra lanes on the NonStop signal midplane are allocated to the adjacent device bay. A four-port PCIe x8 mezzanine card installed in connector 2 PCIe x8 can send x2 signals to interconnect bays 3 and 4. Figure 17 lists the available configurations for half-height devices installed in device bay N (1–8).

Figure 17. Port mapping for HP BladeSystem c3000 half-height server blades to interconnect bays

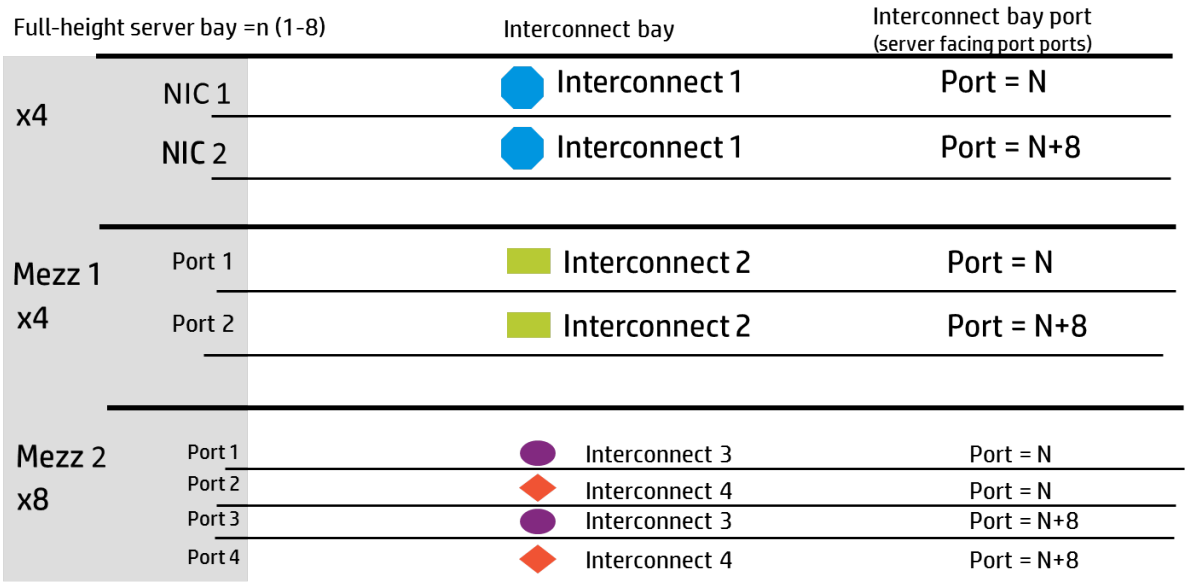

Port mapping differs slightly between full-height and half-height server blades due to the additional mezzanine cards on the full-height version. The OA and HP Systems Insight Manager (HP SIM) software include tools that simplify the process of mapping mezzanine ports to switch ports. For specific port-mapping details, see the HP BladeSystem Onboard Administrator User Guide at: hp.com/bc/docs/support/SupportManual/c00705292/c00705292.pdf.

c3000 enclosure bay-to-bay crosslinks

For bay-to-bay communication, the c3000 midplane provides four-trace SerDes signals between adjacent bays..

Device bay crosslinks

Device bay crosslinks are wired between adjacent horizontal device bay pairs as indicated by the arrows in the c3000 enclosure front view (Figure 18). For half-height server blades, these signals are used for four-lane PCIe connection to a partner blade such as a tape blade or PCI expansion blade. For full-height server blades, these signals are used for PCIe connection to a partner blade in the lower adjacent bay and require a PCIe pass-thru mezzanine card installed in mezzanine connector 3. The BladeSystem Onboard Administrator disables the device bay crosslinks in instances where they cannot be used, such as when two server blades reside in adjacent device bays.

Figure 18.

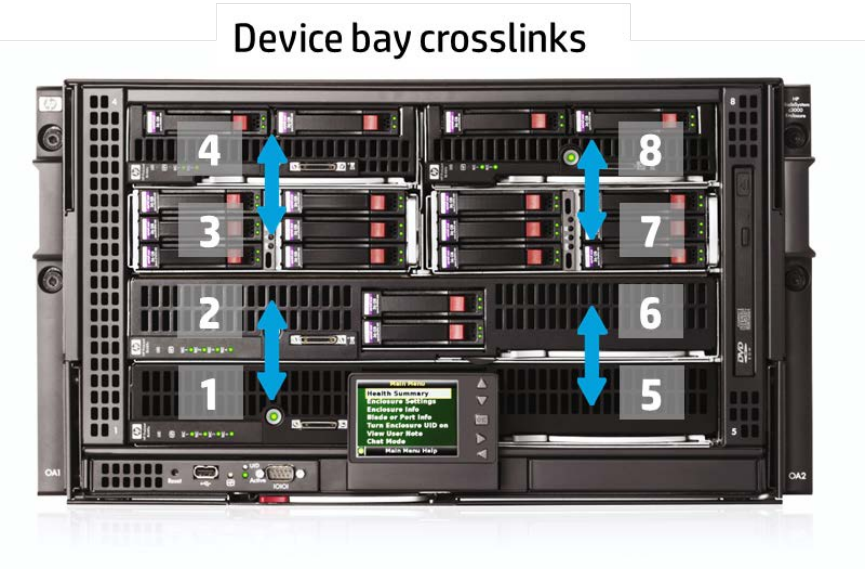

Interconnect bay crosslinks

Interconnect bay crosslinks are wired between adjacent interconnect bay pairs as indicated by the arrows in the c3000 enclosure rear view (Figure 19). These signals can provide module-to-module connections (such as Ethernet crosslink ports between matching switches), or they can be used by Virtual Connect modules as stacking links. The BladeSystem Onboard Administrator disables the interconnect bay crosslinks in instances where they cannot be used, such as two different modules residing in adjacent horizontal interconnect bays.

Figure 19.

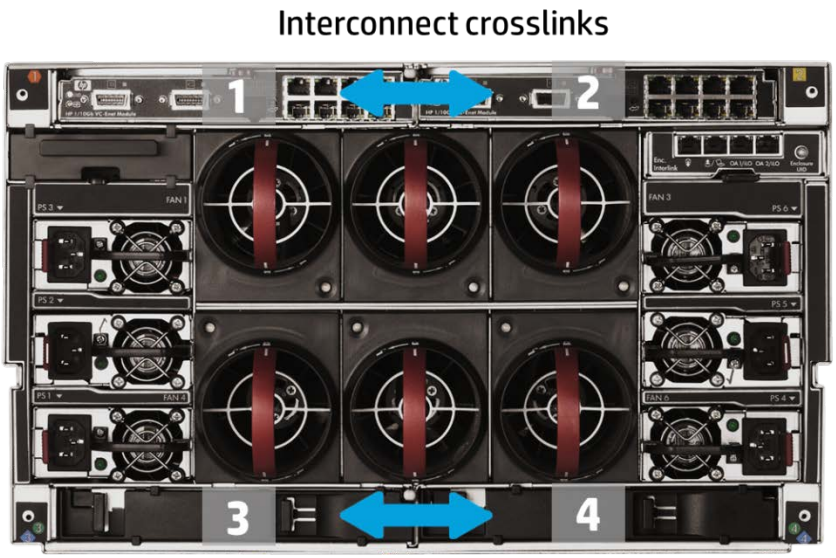

Blade-to-blade connectivity

The NonStop signal midplane allows you to use more modular components than previous generations of blade systems. We can develop components in the blade form factor and connect them across the NonStop signal midplane—front-to-back or side-to-side. The architecture includes front-to-back modularity by connecting mezzanine cards in the server blades at the front of the enclosure to the matching interconnect modules in the rear of the enclosure.

There are four dedicated lanes between each side-by-side pair (odd/even) of device bays. This lets you connect optional storage or I/O devices to a server blade through the midplane. The Onboard Administrator disables these side-to-side links when they are unusable, such as when two server blades reside in adjacent device bays.

Some examples of storage and I/O device options using the dedicated links to a server blade include:

- StorageWorks D2220sb or D2200 Storage Blade that consist of a RAID controller and additional drives
- StorageWorks Ultrium Tape Blades, or LTO-4 Ultrium tape cartridges
- BladeSystem PCI Expansion Blade that holds two off-the-shelf PCI-X or PCIe cards

HP Thermal Logic technologies

The BladeSystem c3000 enclosure uses several HP Thermal Logic technologies, including mechanical design, built-in power and thermal monitoring, and control capabilities. Thermal Logic technologies yield significant power and cooling savings compared to traditional rack-and-tower-based servers. Thermal Logic technologies also provide an instant view of power use and temperature at the server, enclosure, or rack level. They automatically adjust power and thermal controls to minimize power usage while maintaining adequate cooling for all devices and ensuring high availability.

HP Thermal Logic technologies include the following elements and capabilities:

- Active Cool fans
- Parallel Redundant Scalable Enclosure Cooling (PARSEC) design
- Platinum Efficiency Power Supplies
- Instant power and thermal monitoring
- Pooled power for a variety of power redundancy modes
- Dynamic Power Saver mode

- Power Regulator
- Dynamic Power Capping

Active Cool fans

Quite often, dense, full-featured, small form-factor servers use very small fans for localized cooling in the specific areas. Because the fans generate low airflow (in cubic feet per minute, or CFM) at medium backpressure, a single server often requires multiple fans to ensure adequate cooling. If each server blade contains several fans, installing many server blades together in an enclosure can result in a significant cost and space overhead.

A second solution for cooling is to use larger, blower-style fans for an entire enclosure. The fans generate high airflow, but they typically require higher power input and more space. They are loud and designed for the maximum load in an enclosure. As a result, designers may have to sacrifice server features to allow large, high-power fans to fit in the enclosure. Even then, ensuring adequate airflow to all the servers without leakage, over provisioning, or bypass is a challenge.

To overcome these issues in the BladeSystem c3000 enclosure, our engineers designed a new type of fan that delivers high airflow and high pressure in a small form factor that can scale to meet future cooling needs. We have 20 patents pending for the Active Cool fan technology and implementation.

HP Active Cool fans can cool eight server blades using as little as 100 watts of power. Active Cool fans use ducted fan technology with a high-performance motor and impeller to deliver high CFM at high pressure (Figure 20). The fan includes a bell mouth inlet with a specially designed impeller and a stator section that also provides cooling fins for the motor and acoustic treatments at the rear of the fan. The fan's unique shape generates high-volume, high-pressure airflow at even the slowest fan speeds, with low noise levels and minimal power consumption.

Figure 20.

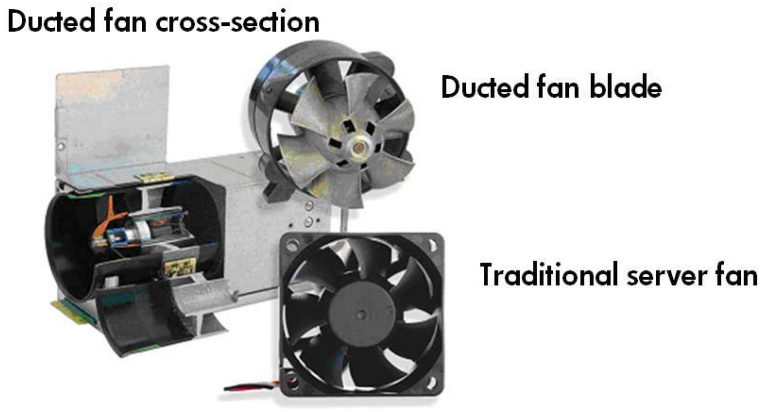

The OA controls Active Cool fans, ramping cooling capacity either up or down based on system needs. Along with optimizing the airflow, the control algorithm optimizes the acoustic levels and power consumption. As a result, the BladeSystem c3000 enclosure requires less airflow than traditional rack-mount servers to properly cool the server blades within the enclosure.. As we introduce new Active Cool fans, they may not interoperate with earlier Active Cool fans. If the OA detects mixed fans that are not interoperable in an enclosure, it will indicate a fan compatibility problem.

HP PARSEC architecture

The BladeSystem c3000 enclosure uses PARSEC architecture—parallel, redundant, scalable, enclosure-based cooling.

Parallel

In this context, "parallel" means that fresh, cool air flows over all the blades (in the front of the enclosure) and all the interconnect modules (in the back of the enclosure). Our designers divided the enclosure into four cooling zones with fans in each. The Active Cool fans supply cooling for their own zone and redundant cooling for the rest of the enclosure. To ensure scalability, we designed both the fans and the power supplies with enough capacity to meet the needs of compute, storage, and I/O components well into the future.

To optimize thermal design, we developed a relatively airtight center air plenum, or air chamber. In the BladeSystem c3000 enclosure, all device bays include a shutoff door, normally closed, to prevent air leakage. When a server blade is inserted, it seals into the center air plenum docking collar, and the server shutoff door opens to allow airflow across the server blade.

Similarly, the fan seals into the center air plenum docking collar. Each fan bay includes louvers that open automatically when a fan begins operating. If a fan is not functional, the pressure distribution around the fan changes. This pressure change causes the louvers to close, ensuring that cooling air does not flow through the inoperative fan (Figure 21).

Figure 21.

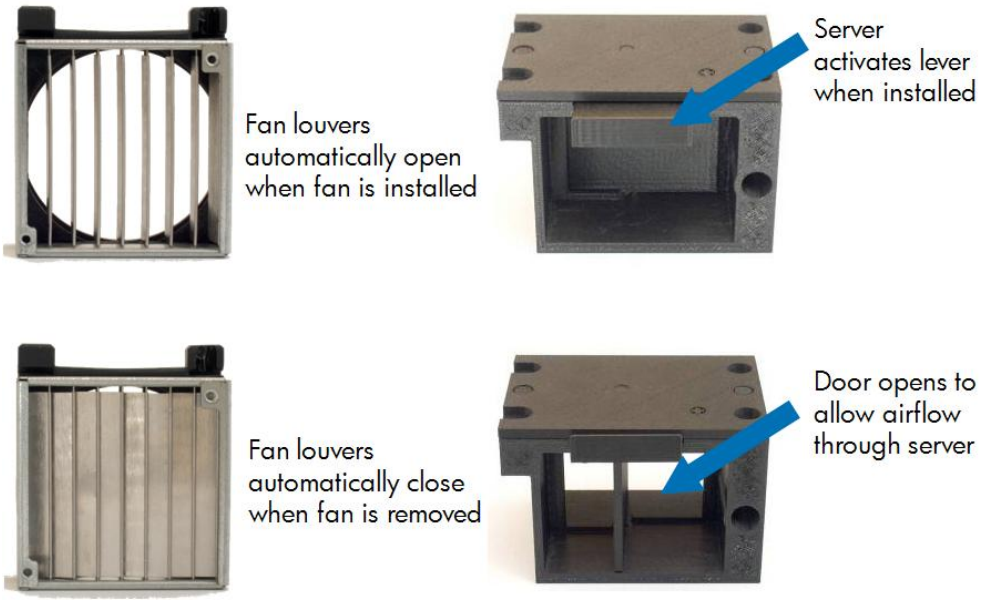

The enclosure and the components within it optimize the cooling capacity through unique mechanical designs. Managed airflow through the enclosure ensures the following:

- Every device gets cool air
- No device stands in the path of hot exhaust air from another device
- Air goes only where necessary for cooling

Redundant and scalable

BladeSystem c3000 enclosures ship with four installed fans that provide redundancy and support up to four half-height devices in device bays 1, 2, 5, and 6, or two full-height server blades in device bays 1 and 2.

In a four-fan configuration, the BladeSystem Onboard Administrator prevents server and storage blades installed in device bays 3, 4, 7, and 8 from powering on until two additional fans are added into fan bays 1 and 3. To install blade devices in all eight device bays, it is necessary to install all six Active Cool fans.

Thermal Logic for the server blade

Precise ducting on ProLiant server blades manages airflow and temperature based on the unique thermal requirements of all the critical components. The airflow is tightly ducted to ensure that no air bypasses the server blade and to obtain the most work from the least amount of air.

This concept allows more flexibility in heat sink design. The heat sink design closely matches the server blade and processor architecture requirements. For example, in the HP BladeSystem BL460c server blade using Intel Xeon processors, we were able to use a smaller, high-power processor heat sink than in rack-mount servers. These heat sinks have vapor chamber bases, thinner fins, and tighter fin pitch than previous designs. This creates the largest possible heat transfer surface in the smallest possible package (Figure 22). The smaller heat sink allows more space on the server blades for DIMM slots and hot-plug hard drives.

Figure 22.

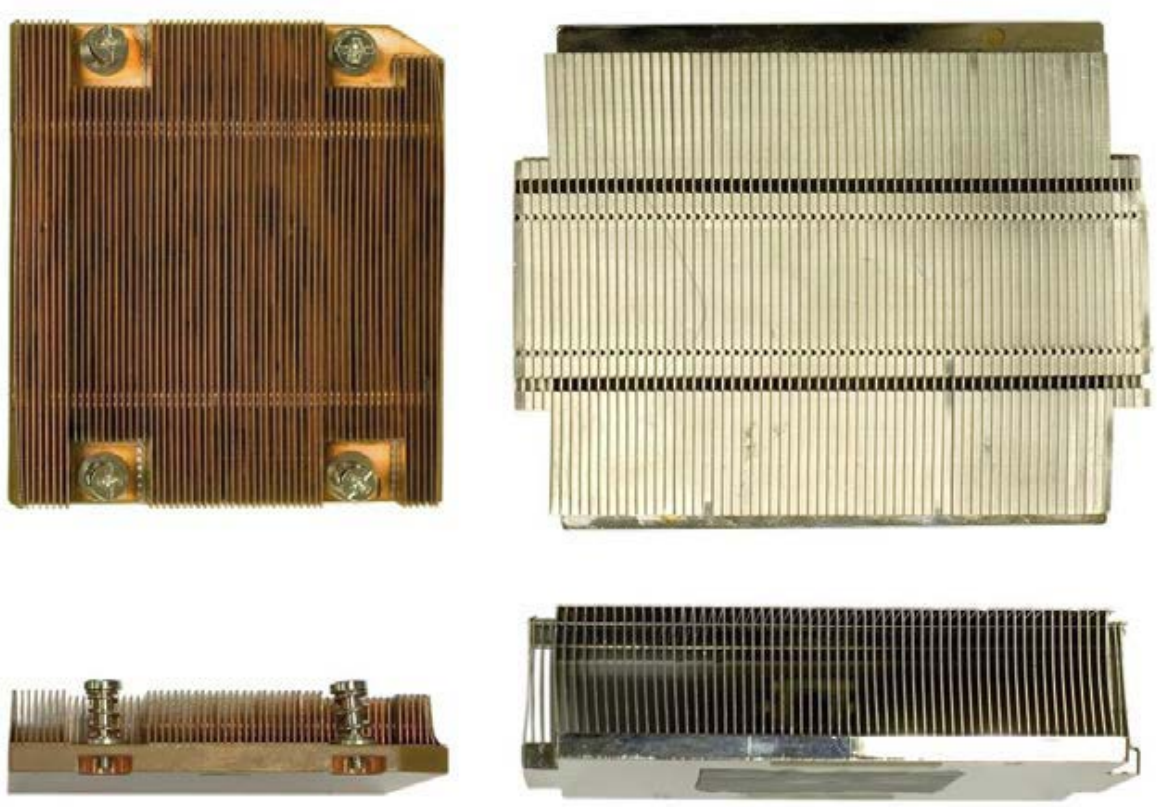

Instant thermal monitoring provides a real-time view of heat, power, and cooling data. The OA retrieves thermal information from all server blades, storage blades, and interconnect modules in the enclosure to ensure an optimal balance between cooling, acoustic levels, and power consumption.

The Thermal Logic technology in the OA keeps fan and system power at the lowest level possible. If the thermal load in the enclosure increases, the Thermal Logic feature instructs the fan controllers to increase fan speeds to accommodate the additional demand. If high temperature levels occur, the iLO processors on server blades and the OA modules provide alerts to various management tools such as HP Insight Control Environment and HP Systems Insight Manager. In addition, built-in failsafe procedures shut down devices in the enclosure if temperature levels exceed specified parameters. This prevents permanent damage to any devices within the enclosure.

HP Thermal Logic includes sophisticated algorithms in each BladeSystem ROM, ProLiant OA, iLO, and BladeSystem OA. In combination, these algorithms minimize the power and cooling that is necessary to maintain the proper BladeSystem environment.

3D Sea of Sensors

Servers use internal fans to exhaust heat into the data center to keep their components operating within a safe temperature range. Previous generations of servers used temperature sensors and a “fan curve” to set the speed of the fans to a preset value based on the measured temperature.

The 3D Sea of Sensors technology in ProLiant Gen8 servers uses an array of sensors to map the server’s temperature profile more accurately. Instead of using a simple fan curve, the iLO processor monitors multiple temperature sensors throughout the server and then employs a proportional–integral–derivative (PID) control feedback algorithm to set and control the speed of each server fan. With the ProLiant Gen8 servers, we have added sensors to most HP option cards so that stand-up PCI cards, backplanes, and mezzanine cards can be effectively monitored. Thermal Logic adjusts fan speed in the enclosure cooling zones to minimize power consumption and maximize cooling efficiency. The 3D Sea of Sensors is part of the Thermal Logic portfolio of features designed for thermal and power management. Read more about these technologies in the “Power monitoring and management” section later in this paper.

Enclosure power components and capabilities

The HP BladeSystem c3000 Enclosure ships with two power supplies. However, up to six power supplies can be installed, depending on the AC redundancy level required and the number of devices installed in the enclosure. BladeSystem c3000 single-phase power supplies automatically switch between low-line (120 VAC) and high-line (240 VAC) to address both environments. A pooled power backplane delivers power to the enclosure. This ensures that the full capacity of the power supplies is available to all server blades.

Moving the power supplies into the enclosure reduced the transmission distance for DC power distribution allowing the use of an industry-standard 12V infrastructure. Using a 12V infrastructure eliminated several power-related components and improved power efficiency on the server blades and in the infrastructure. The control circuitry was stripped and put on the management board and fans.

High efficiency HP c3000 power supplies provide greater than 90% efficiency in AC to DC conversion. These power supplies are in the ProLiant universal form factor so they can also be used in other ProLiant servers. Each power supply ships with a standard power distribution unit (PDU) power cord (C13 to C14), and each enclosure includes c13 to c20 power cords for different types of PDUs. By purchasing proper wall outlet cords, users can connect the power supplies to standard wall outlets.

Caution
Wall outlet power cords should only be used with low-line power sources. If high-line power outlets are required, safety regulations require either a PDU or a UPS between the c3000 enclosure power supplies and wall outlets

---

Power Discovery and management
The HP Common Slot (CS) Platinum Power supplies and Intelligent Power Distribution Units (iPDUs) are key components of Power Discovery Services and Intelligent Power Management. Power Discovery allows BladeSystem Enclosures to communicate information to iPDUs that automatically track enclosure power connections to the specific iPDU outlets to ensure redundancy and prevent downtime. Together they communicate through power line communication (PLC) technology embedded into HP CS Platinum and Platinum Plus power supplies. PLC allows the power supplies to share data with the intelligent power module, and with the HP Insight Control power management software. Power Discovery gives you accurate, automated, error-free power mapping, monitoring, and control. IPD services are available with ProLiant G7 and later servers. OA 3.5 is required. You can read more about Power Discovery at: hp.com/go/ipdu.

Power supplies and power configuration
The enclosure can contain up to six 1200-watt self-cooled power supplies. Most typically, the c3000 enclosure is deployed at sites that would not normally have data center AC redundancy to racks. Therefore, the c3000 has been configured so that only four power supplies are needed in a Power Supply Redundant (N+1) mode, where the enclosure would be connected to a UPS, to a single PDU, or directly into 110V wall outlets. If there is a need for dual AC power feeds and data center-like AC redundancy, six power supplies can be configured to connect to a pair of PDUs (three connected to each PDU). A variety of PDUs are available, as indicated in the c3000 QuickSpecs: http://h18004.www1.hp.com/products/quickspecs/12790_div/12790_div.html. The HP Power Advisor is a tool to assist in estimating power consumption and selecting proper components for the c3000 storage and server configuration.

HP expects that in many of the markets targeted for the c3000 enclosure (midmarket and remote sites), the c3000 will be connected to an uninterruptible power supply (UPS) for power backup instead of to a PDU. HP recommends using HP Power Advisor to determine the number of power supplies needed in the c3000 and to determine the UPS capacity requirement. You can access the HP Power Advisor at: hp.com/go/HPPowerAdvisor.

Note
The rack-mountable HP R5500 UPS (5000VA/4500W) includes four power supplies in the power supply redundant (N+1) power mode.

The c3000 Platinum Enclosure delivers a 5% improvement in system power efficiency provided by new ProLiant Gen8 94% efficient Common Slot Platinum Plus power supplies. Total power savings include both server power and the power required for server cooling. You also have the choice of low-line or high-line AC Power supplies. The c3000 Platinum Enclosure includes these power supply options:

- 1200W Platinum (94%) Hot Plug Power Supply, Improved efficiency, saving power per enclosure (80 Plus certified)
- Three-phase enclosure with a pair of international power cords with IEC 309, 5-Pin, 16A power connectors
- –48V DC Input Modules using a screw down terminal lug (45DG 4AWG 1/4 2H)

Moving the power supplies into the enclosure reduces the transmission distance for DC power distribution and employs an industry-standard 12V infrastructure for the enclosure. By using a 12V infrastructure, we eliminated several power-related components and improved power efficiency on the server blades and in the infrastructure. We put the control circuitry on the management board and fans.

Pooled power

All the power in the BladeSystem c3000 enclosure is part of a single power pool that any server blade within the enclosure can access. This provides maximum flexibility when configuring the power in the system so that you can choose the level of redundancy.

Because this power design has no zones, it facilitates both N+N and N+1 power modes, to accommodate higher power requirements in the future, if needed.

The BladeSystem c3000 enclosure has three redundancy modes:

- AC or DC redundant
- Power supply redundant
- No redundancy mode

You can use the OA or the Insight Display to select the power redundancy mode.

Since all power supplies share in delivering pooled power to the entire enclosure, the OA report shows that all power supply outputs track each other closely. This distributes the load evenly, which is particularly important when using three-phase AC power.

If you change the power mode of the enclosure to AC redundant, the OA ensures that equal numbers of power supplies are available on the A and B sides as indicated in Figure 23. This logical grouping ensures that when the enclosure is configured with the three-phase AC input module, all the power supplies on one side maintain enclosure power if AC input power is lost on the other side. If you configure the enclosure power mode to N+1 or not redundant, you can install the power supplies in any bay.

Figure 23.

Connecting to PDUs with AC redundancy to each rack

In an N+N AC redundancy configuration, the total power available equals the amount from the A or B side, whichever contains fewer power supplies. In this configuration, N power supplies are used to provide power and the same number are used to provide redundancy, where N can equal 1, 2, or 3. Any number of power supplies from 1 to N can fail without causing the enclosure to lose power. When correctly wired with redundant AC line feeds, this configuration will also ensure that a single AC line feed failure will not cause the enclosure to power off.

Typical power configuration connecting to an uninterruptible power supply (UPS)
In a configuration with N+1 power supply redundancy connecting to a UPS, the total power available equals the total power available less one power supply. In this configuration, there can be a maximum of six power supplies and one of them is always available to provide redundancy. In the event of a single power supply failure, the redundant power supply will take over the load of the failed power supply.

Connecting with no power redundancy configured
In a configuration with no power redundancy, the total power available equals the power available from all installed power supplies. Any power supply or AC line failure causes the system to power off if the remaining power supplies are unable to handle the full load.

The OA manages power allocation rules of various components and can limit overall power capacity for the enclosure.

Dynamic Power Saver mode
Most power supplies operate more efficiently when heavily loaded and less efficiently when lightly loaded. Dynamic Power Saver mode provides power-load shifting for maximum efficiency and reliability. Dynamic Power Saver technology maximizes power supply efficiency to reduce operating costs. Power supply efficiency is simply a measure of DC watts output divided by AC or DC watts input. At 50% efficiency, 2000W input would yield 1000W output. The difference is costly wasted energy that generates unnecessary heat.

Dynamic Power Saver mode is active by default because it saves power in most situations. When enabled, Dynamic Power Saver runs the required power supplies at a higher use rate and puts unneeded power supplies in standby mode. A typical power supply running at 20% load could have an efficiency rating as low as 60%. But at 50% load, the efficiency rating could be up to 94%, providing a significant savings in power consumption.

The OA module enables Dynamic Power Saver. When enabled, this feature monitors the total power consumed by the enclosure in real-time and automatically adjusts for changes in demand.

Note
In redundant environments, at least two power supplies are always active, and the maximum load reached on any power supply is 50%. When the load reaches 50%, another two power supplies activate, ensuring redundancy at all times.

Power Regulator
HP ProLiant Power Regulator provides iLO-controlled speed stepping for Intel x86 and recent AMD Opteron™ processors. The Power Regulator feature improves server energy efficiency by giving processors full power when they need it and reducing power when they do not. This power management feature allows ProLiant servers with policy-based power management to control processor power states. You can configure Power Regulator for continuous, static low power mode or for Dynamic Power Savings mode, which automatically adjusts power to match processor demand.

Basic Power Capping for each server blade
iLO firmware can limit power consumption on ProLiant server blades. You can set a limit in watts or BTUs per hour. The limit constrains the amount of power consumed, which reduces heat output into the data center. The iLO firmware monitors server power consumption and checks it against the power cap goal. If necessary, iLO adjusts server performance to maintain an average power consumption that is less than or equal to the power cap goal. This functionality is available on all ProLiant server blades using Intel or recent AMD processors.

Insight Control power management lets you set power caps on groups of included servers and statically allocates the group power cap among the servers in the group. Insight Control power management allocates the group cap equitably among all servers in the group based on a calculation using the idle and maximum measured power consumption of each server. You can also track and record over time the actual power use of groups of servers and enclosures. This provides your data center facilities managers with measured power consumption for various periods, reducing the need to install monitored PDUs to measure actual power use in data centers.

HP Dynamic Power Capping
HP launched Dynamic Power Capping in December 2008. Compatible servers contain an internal hardware circuit that monitors server power consumption on a sub-second basis. If consumption approaches the power cap set in iLO, the internal hardware circuit will limit power consumption quickly enough to protect PDU-level circuits from over-subscription and prevent power-related server outages. At the enclosure level, HP Enclosure Dynamic Power Capping lets you set an enclosure-level power cap. The OA manages the power distribution to the server blades and will change server caps over time as workloads change.

For more information on HP Insight Control power management features, as well as the full range of HP BladeSystem Management architecture, please refer to the "Management architecture of HP BladeSystem c-Class Systems" technical whitepaper at: hp.com/bc/docs/support/SupportManual/c00814176/c00814176.pdf

For a listing of servers that support Insight Control power management please see the Insight Management Support Matrix at: http://www.hp.com/go/insightsoftware/docs. You can read more about Insight Control power capping software at: hp.com/go/powercapping.

HP Power Advisor

In the latest generation of HP BladeSystem servers, power is both critical and flexible. The HP Power Advisor is a tool to assist in estimating power consumption and selecting proper components—including power supplies—at a system, rack, and multi-rack level. Administrators can measure power usage when they configure servers and racks with this downloadable tool. It produces a condensed bill of materials, a cost of ownership calculator, and a power report. See the Resources section at the end of this document for a link to the HP Power Advisor.

Conclusion

The BladeSystem c3000 enclosure serves as the foundation of a modular computing architecture that consolidates and simplifies infrastructure, reduces operational cost, and delivers IT services more effectively. Thermal Logic technologies supply the mechanical design features, built-in monitoring, and control capabilities that enable your IT administrators to optimize their power and thermal environments.

The c3000 Platinum enclosure includes Intelligent Power Management with Power Discovery for automated, error-free power mapping, The optional CS 1200W Platinum Hot Plug power supply is 80 Plus certified and part of the Power Discovery technology. The shared, high-speed, NonStop midplane and pooled-power backplane in the enclosure accommodate higher bandwidths and new technologies. The signal midplane increases its aggregate bandwidth to 10Tb/s and includes FDR InfiniBand capability for demanding workloads.

The c3000 enclosure is designed for remote sites, small and medium-sized businesses, and data centers with special power and cooling constraints. The OA supplies the infrastructure to provide essential power and cooling information and help to automate infrastructure management. The HP BladeSystem c3000 enclosure provides all the power, cooling, and I/O infrastructure required by c-Class modular servers, interconnects, and storage components.

Appendix: Fan, power supply, and device bay population guidelines

For correct operation, you must install fans and server blades in the correct bays of the HP BladeSystem c3000 enclosure. The OA ensures that you have correctly placed the fans and blades before it allows systems to power on.

Fan bay numbering

. HP BladeSystem c3000 Enclosure – Fan population guidelines. For correct operation, fans and server blades must be installed in the correct fan bays. The BladeSystem Onboard Administrator will ensure that fans and server/storage blades are correctly placed before allowing systems to power on..

Figure A-1.

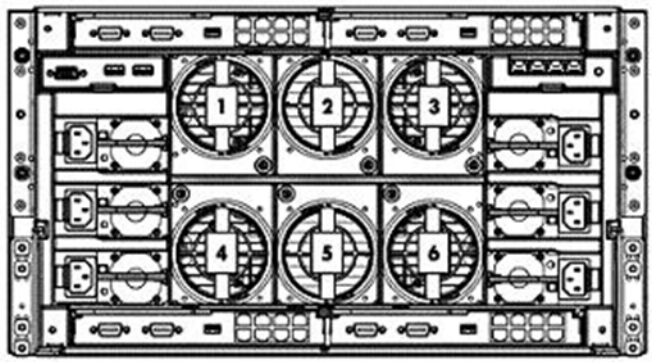

BladeSystem c3000 enclosures ship with four fans installed, serving up to four half-height devices or two full-height server blades. Adding two additional fans to the enclosure allows population with eight half-height or four full-height devices:

- Four-fan configuration requires population of fan bays 2, 4, 5, and 6.
- Six-fan configuration enables population of all fan bays.

In a four-fan configuration, the BladeSystem Onboard Administrator prevents blade devices in device bays 3, 4, 7, and 8 from powering on and identifies the fan subsystem as degraded. To incorporate blade devices in these device bays, install six Active Cool fans.

Figure A-2.

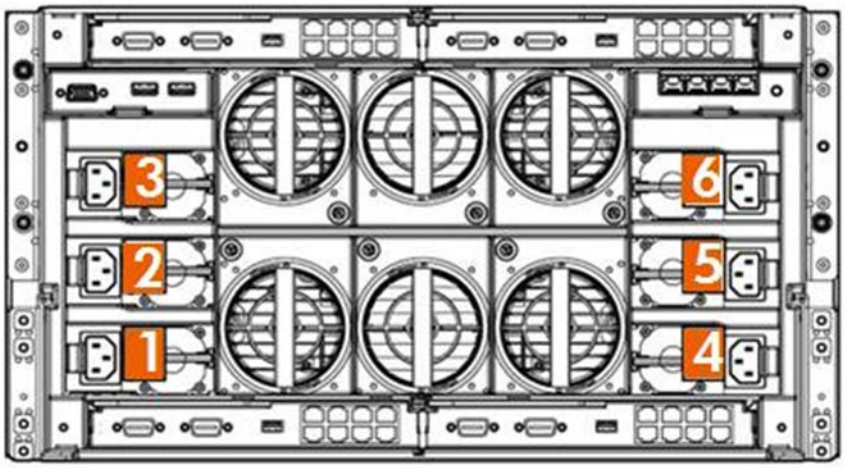

Table A-1. Power supply placement

Number of power supplies	Power supply bays used
2	1 and 4
4	1, 2, 4, and 5
6	All power supply bays filled

Table A-2. **Power supply redundancy options**

Number of power supplies	Power supply bays used
1+1	1 and 4
2+1	1, 4, and 2
3+1	1, 4, 2, and 5
4+1	1, 4, 2, 5, and 3
5+1	Populate all power supply bays

Table A-3. **AC redundancy options**

Number of power supplies	Power supply bays used
1+1	1 and 4
2+2	1, 2, 4, and 5
3+3	Populate all power supply bays

Device bay numbering

HP BladeSystem c3000 Enclosure – Full-height server blade device bay numbering. Full-height servers should be populated from bottom to top when viewing from the front of the enclosure. With four fans, only the bottom or left two device bays can be used; with six fans, all device bays can be used.

Figure A-3.

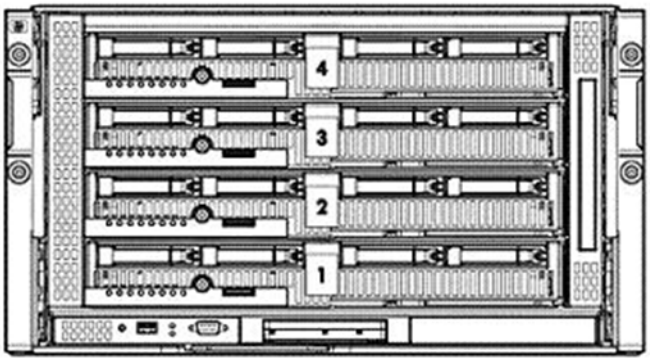

HP BladeSystem c3000 Enclosure – Half-height server blade device bay numbering. Half--height servers should be populated in the following order: Device bays 1, 5, 2, 6, 3, 7, 4, 8.

. **Figure A-4.**

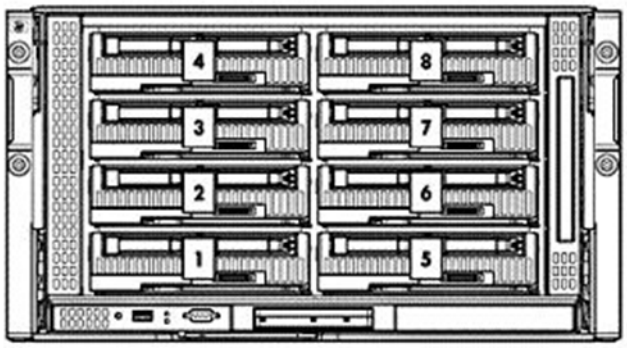

Important
When looking at the rear of the enclosure, device bay numbering is reversed.

Caution
To prevent improper cooling or thermal damage, do not operate the server blade or the enclosure unless all device bays are populated with either a component or a blank.

Enclosure blade zones

The c3000 enclosure is divided by sheet metal panels into two full-height zones. Zone 1 and Zone 2 are divided in half by a removable shelf to accommodate a maximum of 8 half-height device bays per enclosure. These zones reflect the PCIe bus mapping in the signal midplane and limit placement of the server blade/storage blade combination. The signal midplane has a direct PCIe link connecting adjacent paired device bays. If installing two direct-attach storage blades with one partner full-height server blade in an HP BladeSystem c3000 Enclosure; use the mini divider instead of the half-height blade shelf. See HP BladeSystem c3000 Enclosure Quick Setup Instructions for more information.

. **Figure A-5.**

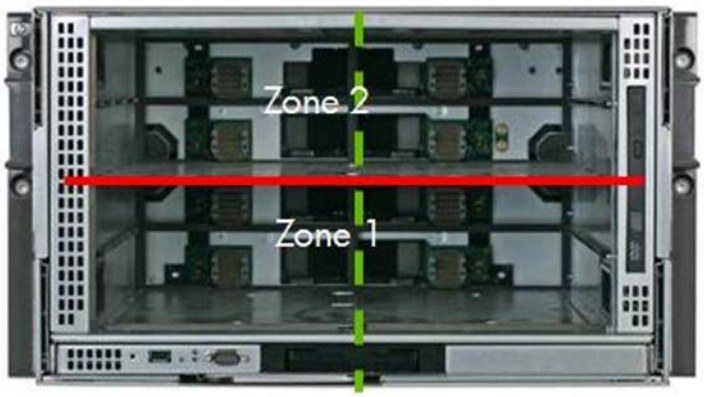

Important
The server blade/storage blade relationship cannot extend beyond the removable dividers between full height device bays, nor can it span the removable bay shelf dividing the zone into half-height device bays.

The enclosure comes preinstalled with removable full-height dividers between the four device bays in Zone 2 and the four device bays in Zone 1. In addition, a half-height divider is available for use between device bays 4 and 8 if the full-height divider is removed. Using these combinations of dividers, the following combinations of server blades can be installed:

- 8 half-height server blades with both full-height dividers installed
- 4 full-height server blades with both full-height dividers removed
- 4 half-height server blades in Zone 1 with 1 full-height divider installed
- 2 full-height server blades in Zone 1 with 1 full-height divider removed
- 4 half-height server blades in Zone 2 with 1 full-height divider installed
- 2 full-height server blades in Zone 2 with 1 full-height divider removed
- 1 full-height server blade and 2 half-height server blades in Zone 2 with 1 full-height divider removed and the half-height divider installed

Important
The server blade/storage blade relationship cannot extend beyond the removable dividers between full height device bays, nor can it span the removable bay shelf dividing the zone into half-height device bays.

Removing the full-height divider in Zone 1 allows only full-height server blades to be installed in Zone 1.

Removing the full-height divider in the Zone 2 requires either installing only full-height server blades in Zone 2 or installing the half-height divider between device bays 4 and 8. With the half-height divider installed, two half-height devices (two server blades, one companion blade and one server blade, or one blade blank and one companion blade or server blade) can be installed in device bays 4 and 8 and one full-height server blade in device bay 3/7.

A companion blade (HP StorageWorks D2220sb or D2200sb Storage Blade, HP PCI Expansion Blade, or HP StorageWorks Ultrium 448c Tape Blade) can be installed in either of the paired device bays (1/2, 3/4, 5/6, or 7/8) with a half-height server blade installed in the other paired device bay.

To install a companion blade with a full-height server blade, the companion blade must be installed in device bay 8 with the full-height server blade installed in device bay 3/7. The half-height divider must be installed between device bays 4 and 8, and either a blade blank or a half-height server blade can be installed in device bay 4.

Resources, contacts, or additional links

HP BladeSystem c3000 Enclosure QuickSpecs
hp.com/products/quickspecs/12790_div/12790_div.HTML

HP BladeSystem c-Class documentation
hp.com/enterprise/cache/316735-0-0-0-121.html

HP BladeSystem c3000 Enclosure Setup and Installation Guide
hp.com/bc/docs/support/SupportManual/c01167165/c01167165.pdf

HP BladeSystem Onboard Administrator User Guide
hp.com/bc/docs/support/SupportManual/c00705292/c00705292.pdf

HP BladeSystem c-Class interconnects
hp.com/go/bladesystem/interconnects

General HP BladeSystem information
hp.com/go/bladesystem/

Technical white papers and podcasts about HP BladeSystem
hp.com/products/servers/technology/whitepapers/proliant-servers.html

HP Power Advisor
hp.com/go/HPPowerAdvisor

"Management architecture of HP BladeSystem c-Class Systems" technology brief
hp.com/bc/docs/support/SupportManual/c00814176/c00814176.pdf

HP c-Class mezzanine card options
hp.com/products/blades/components/c-class-interconnects.html

"Technologies in HP ProLiant Gen8 c-Class server blades" technology brief.
hp.com/bc/docs/support/SupportManual/c01136096/c01136096.pdf

iLO 2 firmware updates
hp.com/go/ilo

"HP Power Capping and Dynamic Power Capping for ProLiant servers" technology brief
hp.com/bc/docs/support/SupportManual/c01549455/c01549455.pdf?jumpid=reg_R1002_USEN

Power supply efficiency
80plus.org/manu/psu/psu_detail.aspx?id=41&type=1

Sign up for updates
hp.com/go/getupdated

Share with colleagues

Rate this document

4AA4-8129ENW, August 2013