



Technologies in the HP BladeSystem c7000 Enclosure

technology brief, 4th edition

Abstract.....	3
Overview of HP BladeSystem c7000 Enclosure	3
Enclosure management.....	5
BladeSystem Onboard Administrator.....	5
Detecting component insertion and removal.....	6
Identifying components	6
Managing power and cooling	6
Controlling components.....	7
Redundant enclosure management	8
User interfaces for the BladeSystem Onboard Administrator	8
Security.....	8
Role-based user accounts.....	9
Enclosure linking	9
ProLiant Onboard Administrator (Integrated Lights-Out 2) for ProLiant server blades	10
Insight Display	10
HP Insight Control suite.....	11
Interconnect options and infrastructure.....	12
Interconnect modules	13
Server blades	14
Storage options inside the BladeSystem enclosure.....	14
Direct attach storage blades	14
Shared storage	15
External SAS connectivity with direct connect SAS storage for HP BladeSystem.....	15
NAS/SAN Gateway	15
Storage server.....	15
Mezzanine cards	16
Virtual Connect.....	16
Fabric connectivity and port mapping.....	17
c7000 bay-to-bay crosslinks	19
Device bay crosslinks.....	19
Interconnect bay crosslinks	19
HP Thermal Logic technologies	20
Active Cool 200 fans	20
HP PARSEC architecture.....	21
Parallel.....	21
Redundant.....	23
Scalable.....	24
Thermal Logic for the server blade.....	24
Power supplies and enclosure power subsystem.....	25
Pooled power	27

Dynamic Power Saver mode	29
Power Regulator	30
Basic Power Capping for each server blade	30
HP Dynamic Power Capping	30
HP BladeSystem Power Sizer	31
Conclusion.....	31
Appendix: Fan and server population guidelines	32
Fan bay numbering	32
Server blade bay numbering	32
Enclosure blade zones	33
For more information.....	35
Call to action	35

Abstract

The HP BladeSystem c7000 Enclosure is an evolution of the entire rack-mounted infrastructure. It consolidates and repackages all the supporting infrastructure elements—compute, storage, network, and power—into a single “infrastructure-in-a-box” that accelerates the integration and optimization of the data center. This technology brief provides an overview of the HP BladeSystem c7000 Enclosure, including Thermal Logic power and cooling technologies and interconnect options.

It is assumed that the reader is familiar with HP ProLiant server technology and has some general knowledge of BladeSystem architecture. More information about the infrastructure components and BladeSystem technologies is available on the HP website: www.hp.com/go/bladeSystem.

Overview of HP BladeSystem c7000 Enclosure

The HP BladeSystem c7000 Enclosure was the first enclosure implemented using the BladeSystem c-Class architecture.¹ It is optimized for enterprise data center applications. It fits into standard size HP and third-party racks; accommodates BladeSystem c-Class server blades, storage blades, and interconnect modules; and provides all the power, cooling, and I/O infrastructure needed to support them throughout the next several years.

The c7000 enclosure can be populated with the following components:

- Up to 8 full-height (FH) or 16 half-height (HH) server, storage, or other option blades per enclosure
- Up to eight interconnect modules simultaneously supporting 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)
- Up to 10 Active Cool 200 fan kits
- Up to six power supplies
- Redundant BladeSystem Onboard Administrator (OA) management modules (optional active-standby design)

Figures 1 and 2 show front and rear views of the c7000 enclosure. The c7000 enclosure and the c3000 enclosure support many of the same critical components such as servers, interconnects, mezzanine cards, storage blades, power supplies, and fans.

¹ HP also offers the HP BladeSystem c3000 Enclosure, which is optimized for remote sites or small businesses. More information about the c3000 enclosure is in the technology brief titled “HP BladeSystem c3000 Enclosure technologies”:

<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01508406/c01508406.pdf>.

Figure 1. HP BladeSystem c7000 Enclosure – front view

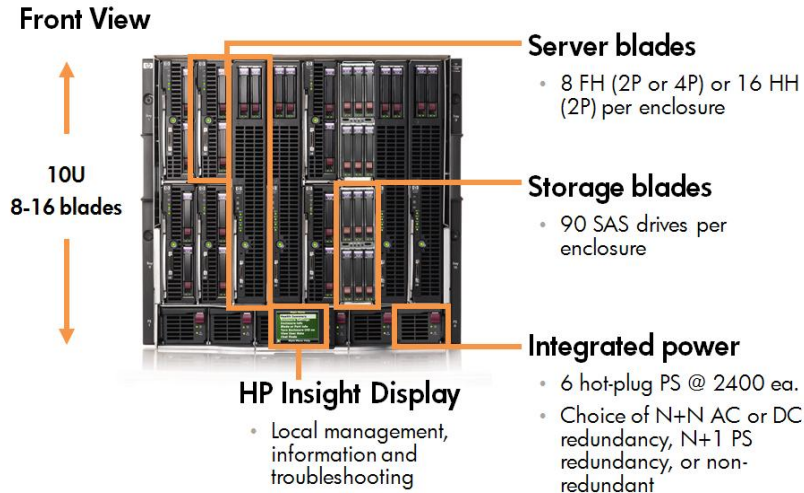
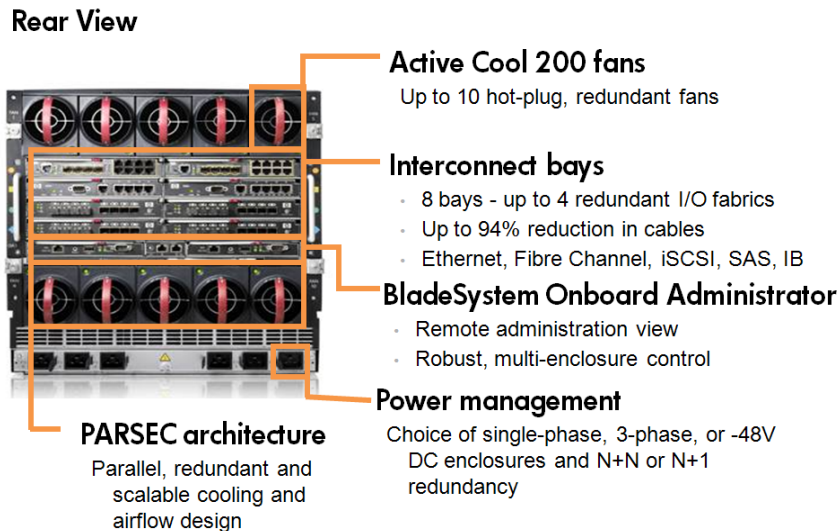


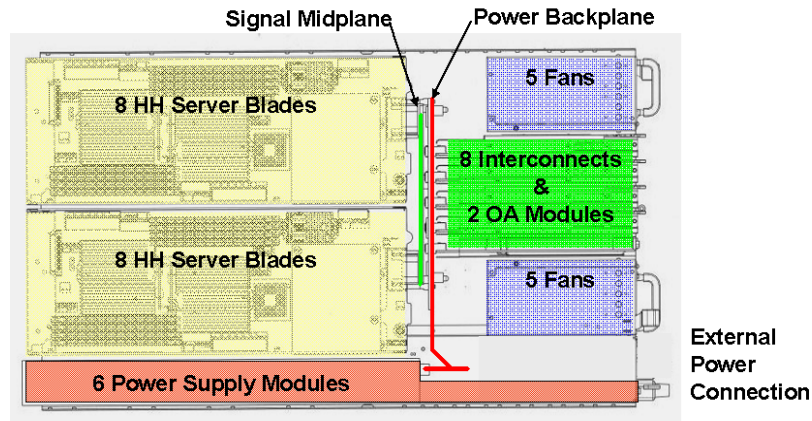
Figure 2. HP BladeSystem c7000 Enclosure – rear view



The c7000 enclosure is 10U high and includes a shared, 5-terabit-per-second, high-speed NonStop midplane for wire-once connection of server blades to network and shared storage. A pooled-power backplane delivers power and ensures that the full capacity of the power supplies is available to all server blades and interconnects. The enclosure is available with either a single-phase AC, three-phase AC, or -48V DC power subsystem to meet the needs of the data center power infrastructure.

The c7000 Enclosure has redundant signal paths between servers and interconnect modules (Figure 3). The NonStop signal midplane and separate power backplane in the c7000 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.

Figure 3. HP BladeSystem c7000 Enclosure – side view



Enclosure management

The HP BladeSystem c7000 Enclosure has extensive embedded management capabilities based on three management elements:

- BladeSystem Onboard Administrator or BladeSystem Onboard Administrator with KVM
- ProLiant Onboard Administrator powered by Integrated Lights-Out 2 (iLO 2) management processors that are integrated on the server blades
- Interconnect module management processors such as the HP Virtual Connect Manager

These integrated management elements provide powerful hardware management for remote administration and local diagnostics, as well as component and enclosure troubleshooting.

For detailed information about c-Class management technologies and capabilities, refer to the HP white paper titled "Managing HP BladeSystem c-Class systems":

<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00814176/c00814176.pdf>.

BladeSystem Onboard Administrator

The heart of c-Class enclosure management is the BladeSystem Onboard Administrator module located in the enclosure. It performs four management functions for the entire enclosure:

- Detecting component insertion and removal
- Identifying components and required connectivity
- Managing power and cooling
- Controlling components

An optional second BladeSystem Onboard Administrator in the c7000 enclosure provides complete redundancy for these functions.

IT administrators can access the BladeSystem Onboard Administrator in three different ways: remotely through the web browser graphical user interface (GUI), through the scriptable command line interface (CLI), or on site through the built-in Insight Display diagnostic LCD panel included in the front of every c-Class enclosure.

The BladeSystem Onboard Administrator with KVM module adds the ability to directly connect the BladeSystem c7000 enclosure to a keyboard, mouse, monitor, or KVM switch through a VGA port. It provides approximately a 1.7x performance increase over the basic BladeSystem Onboard Administrator module and contains a PowerPC 440EPx 400 MHz processor, 1 Gb Ethernet support, and 512 MB DDR2 memory. With a firmware upgrade on both BladeSystem Onboard Administrator modules to OA version 2.41 or later, the BladeSystem Onboard Administrator with KVM is interoperable with the basic BladeSystem Onboard Administrator.

Detecting component insertion and removal

BladeSystem Onboard Administrator provides component control in c-Class enclosures. When a component is inserted into a bay, the BladeSystem Onboard Administrator immediately recognizes and identifies the component through presence signals on each bay. If a component is removed from a bay, the BladeSystem Onboard Administrator deletes the information about that component.

Identifying components

To identify a component, the BladeSystem Onboard Administrator reads a Field-Replaceable Unit (FRU) Electrically Erasable Programmable Read-Only Memory (EEPROM) that contains specific factory information about the component such as product name, part number, and serial number. All FRU EEPROMs in c-Class enclosures are always powered, even if the component is turned off, so the BladeSystem Onboard Administrator can identify the component prior to granting power. For devices such as fans, power supplies, and Insight Display, the BladeSystem Onboard Administrator reads the FRU EEPROMs directly. The BladeSystem Onboard Administrator accesses server blade FRU EEPROMs through their ProLiant Onboard Administrator (iLO 2) management processors.

Server blades contain several FRU EEPROMs: one on the server board that contains server information and embedded NIC information, and one on each of the installed mezzanine option cards. Server blade control options include auto login to the ProLiant Onboard Administrator (iLO 2) web interface and remote server consoles, virtual power control, and boot order control. Server blade control options also include extensive server hardware information including BIOS and ProLiant Onboard Administrator (iLO 2) firmware versions, server name, NIC and option card port IDs, and port mapping. The BladeSystem Onboard Administrator provides easy-to-understand port mapping information for each of the server blades and interconnect modules in the enclosure.

From the NIC and mezzanine option FRU information, the BladeSystem Onboard Administrator determines the type of interconnects each server requires. Before granting power to a server blade, the BladeSystem Onboard Administrator compares this information with the FRU EEPROMs on installed interconnect modules to check for electronic keying errors. For interconnect modules, the BladeSystem Onboard Administrator provides virtual power control, dedicated serial consoles, and management Ethernet connections, based on the specific interconnect features that are included.

Managing power and cooling

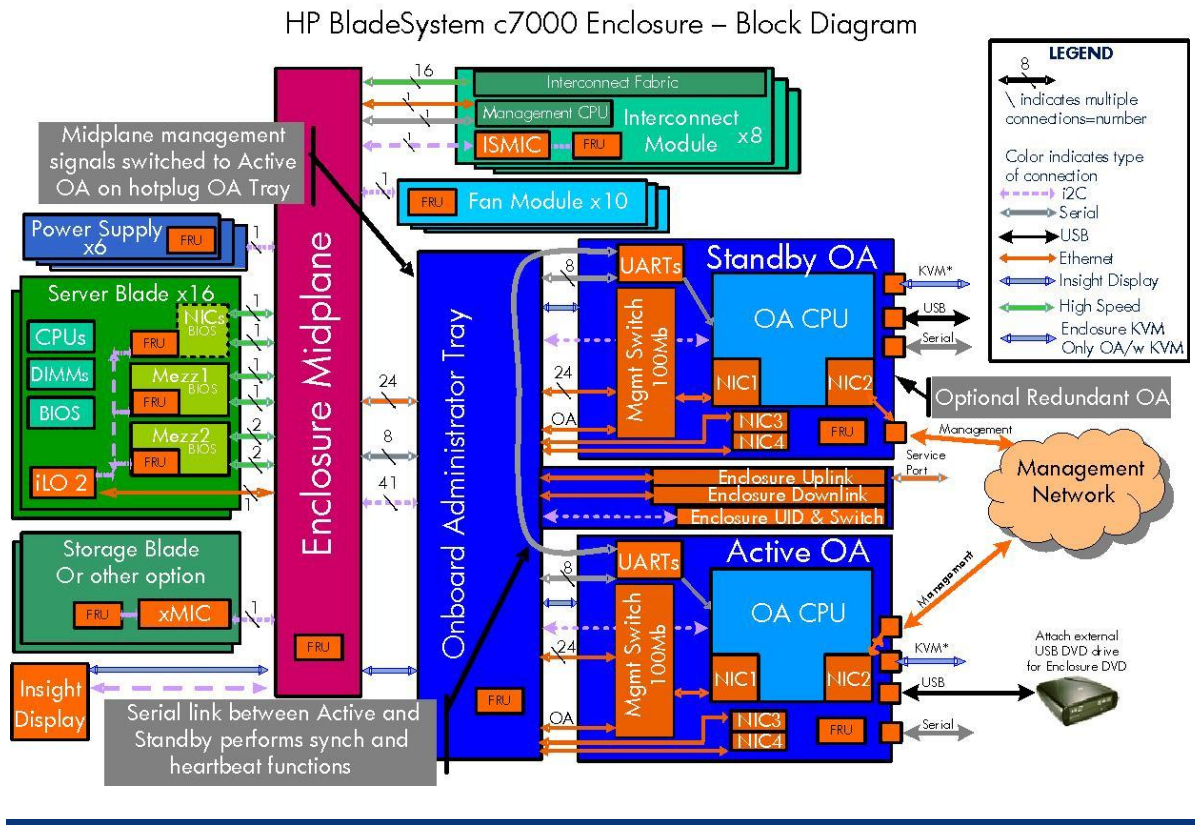
The most important Onboard Administrator tasks are power control and thermal management. The BladeSystem Onboard Administrator can remotely control the power state of all components in c-Class enclosures. For components in device bays in the front of each enclosure, the BladeSystem Onboard Administrator communicates with the ProLiant Onboard Administrator (iLO 2) to control servers and communicates with a microcontroller to control options such as storage blades. A separate microcontroller controls power to interconnect modules.

Once components are granted power, the BladeSystem Onboard Administrator begins thermal management with Thermal Logic. The Thermal Logic feature in the c7000 enclosure minimizes fan subsystem power consumption by reading “a sea of sensors” located throughout the enclosure. Thermal Logic adjusts fan speed in the four different cooling zones within the enclosure to minimize power consumption and maximize cooling efficiency.

Controlling components

The BladeSystem Onboard Administrator uses embedded management interfaces to provide detailed information and health status for all bays in the enclosure (Figure 4). The BladeSystem Onboard Administrator also reports firmware versions for most components in the enclosure and can be used to update those components.

Figure 4. Management communications between the BladeSystem Onboard Administrator and other components in an HP BladeSystem c7000 Enclosure



c7000 internal management interfaces

The BladeSystem Onboard Administrator monitors and communicates with several hardware interfaces to each bay in the c7000 enclosure. The management hardware interfaces include unique presence pins, Inter-Integrated Circuit (I2C), serial, and Ethernet connections. These management interface connections are completely isolated from the server blade connections to interconnect modules.

c7000 external management interfaces

Each c7000 enclosure has several external management interfaces connected to the BladeSystem Onboard Administrator as shown in Figure 4. The primary external management interface is the management port for each BladeSystem Onboard Administrator, which is an RJ-45 jack providing Ethernet communications not only to each BladeSystem Onboard Administrator, but also to every device or interconnect bay with a management processor. This includes ProLiant Onboard Administrator (iLO 2) communication for the server blades and any interconnect module using the c-Class embedded Ethernet management network, such as Virtual Connect Manager. For redundant BladeSystem Onboard Administrators, both BladeSystem Onboard Administrator management ports

are connected to the management network, providing redundant management network connections to each enclosure.

A serial port on each BladeSystem Onboard Administrator module provides full out-of-band CLI access to the BladeSystem Onboard Administrator and is used for BladeSystem Onboard Administrator firmware flash recovery. USB ports on the BladeSystem Onboard Administrator are used to connect DVD drives to support the DVD feature in the c7000 enclosure. All c-Class enclosures support two enclosure link connectors that provide private communications between enclosures linked with CAT5 cable. In addition, the enclosure link-up connector provides an enclosure service port for temporarily connecting a laptop PC to any of the linked enclosure BladeSystem Onboard Administrators for local diagnostics and debugging.

Updating firmware

The BladeSystem Onboard Administrator manages firmware updates for the enclosure's management devices. Updating server blade firmware including server BIOS, NIC and mezzanine BIOS, and ProLiant Onboard Administrator (iLO 2) is possible using HP System Update Manager or the blade firmware update maintenance CD. These utilities can be connected to all the server blades in the enclosure using the BladeSystem Onboard Administrator enclosure DVD feature. When the active BladeSystem Onboard Administrator detects an external USB DVD drive plugged into the USB port, it scans the DVD drive for a CD or DVD disk. This disk can then be connected to one or more server blades at the same time using the BladeSystem Onboard Administrator GUI, CLI, or Insight Display.

Redundant enclosure management

Redundant enclosure management is an optional feature of the c7000. It requires installing a second BladeSystem Onboard Administrator module in the c7000 enclosure to act as a completely redundant controller in an active-standby mode. Using redundant modules in the c7000 enclosure provides complete fault tolerance. The redundancy logic is based on a continuous heartbeat between the two modules over a dedicated serial connection (Figure 4). If the period between heartbeats exceeds a timeout, the standby module automatically takes control of the enclosure and becomes the active BladeSystem Onboard Administrator.

User interfaces for the BladeSystem Onboard Administrator

Three user interfaces to the BladeSystem Onboard Administrator allow control and provide information about the enclosure and installed components:

- Web browser GUI
- Scriptable CLI
- Insight Display diagnostic LCD panel

Additionally, the BladeSystem Onboard Administrator with KVM adds the ability to directly connect a KVM switch through the VGA port. Remote network access to the GUI and CLI is available through the management Ethernet port. A serial port is available for local CLI access and flash recovery. The c-Class enclosure link-up port is also available as the service port for temporary local Ethernet access to the BladeSystem Onboard Administrators and devices in linked enclosures. Insight Display is accessed directly through the buttons on the display or remotely through the GUI.

Security

User authentication maintains security for all user interfaces. User accounts created in the BladeSystem Onboard Administrator define three user privilege levels and the component bays to which each level is granted access. The BladeSystem Onboard Administrator stores the passwords for local user accounts and can be configured to use Lightweight Directory Access Protocol (LDAP) authentication for user group accounts.

Role-based user accounts

The BladeSystem Onboard Administrator provides configurable user accounts that can provide complete isolation of multiple administrative roles such as server, LAN, and SAN. User accounts are configured with specific device bay or interconnect bay permissions and one of three privilege levels: administrator, operator, or user. An account with administrator privileges including BladeSystem Onboard Administrator bay permission can create or edit all user accounts on an enclosure. Operator privileges allow full information access and control of permitted bays. User privileges allow information access but no control capability.

The BladeSystem Onboard Administrator requires user login to the web GUI or CLI with an account ID and password. For a local account, the password is stored on the BladeSystem Onboard Administrator. For an LDAP account, the BladeSystem Onboard Administrator contacts the defined LDAP server to check the user credentials. Two-factor authentication allows even tighter security for the user management session to the BladeSystem Onboard Administrator.

Enclosure linking

Rather than requiring separate logins to access multiple management processors within each enclosure, the BladeSystem Onboard Administrator allows single point access. Thus, an administrator can use single sign-on to log in to a single BladeSystem Onboard Administrator and use the web GUI to graphically view and manage the HP BladeSystem c-Class components in up to seven linked enclosures. For example, an IT administrator could automatically propagate management commands—such as changing the enclosure power mode—throughout the linked enclosures.

The c7000 enclosure contains two enclosure link ports to allow any active BladeSystem Onboard Administrator module to access linked enclosures. On a standalone enclosure or upper enclosure in a series of linked enclosures, the upper enclosure link-up port functions as a service port for temporary connection to a PC with a CAT5 patch cable. It provides quick access to any BladeSystem Onboard Administrator module, ProLiant Onboard Administrator (iLO 2), or interconnect module with Ethernet management ability.

The enclosure link-down port connects to the enclosure link-up port on the enclosure below it. Linking the enclosures enables the rack technician to access all the enclosures through the open link-up/service port. If more c-Class enclosures are added to the rack, they can be linked through the open enclosure link-up port on the upper enclosure or the link-down port on the bottom enclosure.

NOTE

The linked enclosures will enforce a common rack name on all the linked enclosures unless different rack names are established prior to linking the enclosures and not changed after linking the enclosures.

IMPORTANT

The HP BladeSystem c-Class Enclosure link ports are not compatible with the HP BladeSystem p-Class Enclosure link ports.

ProLiant Onboard Administrator (Integrated Lights-Out 2) for ProLiant server blades

HP BladeSystem c-Class employs ProLiant Onboard Administrator (iLO 2) to configure, update, and operate individual server blades remotely. The c7000 enclosure includes an Ethernet management network to aggregate all ProLiant Onboard Administrator (iLO 2) management communications across the entire enclosure. This management network connects ProLiant Onboard Administrator (iLO 2) processors to the BladeSystem Onboard Administrator through the BladeSystem Onboard Administrator tray, as illustrated in Figure 4. The BladeSystem Onboard Administrator provides direct access to each ProLiant Onboard Administrator (iLO 2) through the enclosure management network. The BladeSystem Onboard Administrator uses this network to manage pooled enclosure power and cooling, which results in substantial energy savings over the same number of individual rack-mounted servers.

Insight Display

The Insight Display, Figure 5, is an ever-ready, enclosure-mounted information exchange device with access to all BladeSystem Onboard Administrator setup, management, and troubleshooting features. It is a quick and easy-to-use device that allows technicians to initially configure the enclosure. It also provides information about the health and operation of the enclosure. The Insight Display is effective mechanically because it is large enough to see ample information, and it can slide back and forth to allow access to the power supplies.

Figure 5. Insight Display on the c7000 enclosure




When the c7000 enclosure is initially powered on, the enclosure UID LED and the Insight Display are illuminated blue to identify the enclosure being configured. The Insight Display automatically launches an installation wizard to guide the user through the configuration process. After the enclosure is configured, the Insight Display verifies that there are no installation or configuration errors. The Installation Wizard turns off the enclosure UID when the installation is complete.

When an error or alert condition is detected, the Insight Display Health Summary screen displays the total number of error conditions and their locations in the order of error severity (Figure 6). Failure alerts (if any exist) display first and then caution alerts. Providing this level of diagnostic information for each enclosure dramatically shortens setup, repair, and troubleshooting time.

For example, in Figure 6 the Insight Display diagnostic screen reports an error in power supply bay 5. The Health Summary screen shows the power supply in bay 5 as red. When the technician selects View Alert, the Device Error Summary screen indicates the same condition. The Device Error detail in the third screen shows that the power supply in bay 5 has failed. When the technician selects fix on

the Device Error screen, suggestions for corrective action appear (in this case an animation indicating the proper positions when replacing power supplies in the c7000 enclosure).

Figure 6. BladeSystem c-Class Insight Display diagnostic screens indicating an error and suggested corrective action



More information about the Insight Display is available in the technology brief titled “Managing the HP BladeSystem c-Class”:

<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00814176/c00814176.pdf>.

HP Insight Control suite

HP Insight Control suite delivers essential management for HP BladeSystem lifecycles, including proactive health management, lights out remote control, optimization of power usage, rapid server deployment, performance analysis, vulnerability scanning and patch management, and virtual machine management.

The software is delivered on DVD media. It includes an integrated installer to rapidly and consistently deploy and configure HP Systems Insight Manager (HP SIM) and essential infrastructure management software. The integrated installer includes a wizard-based interface that presents a series of configuration questions. When the user has answered the configuration questions, each of the selected components will be deployed in a single process. The HP Insight Software Advisor checks to ensure that the host central management server meets all installation pre-requisites before installation. When installation is complete, the Insight Software Update Utility will check for available software updates.

HP Insight Control suite installs and licenses the following components for use:

- HP Systems Insight Manager
- HP iLO Advanced for BladeSystem
- HP Insight Power Manager
- HP Insight Rapid Deployment Pack
- HP Insight Performance Manager module
- HP Insight Vulnerability and Patch Manager module
- HP Insight Virtual Machine Manager software

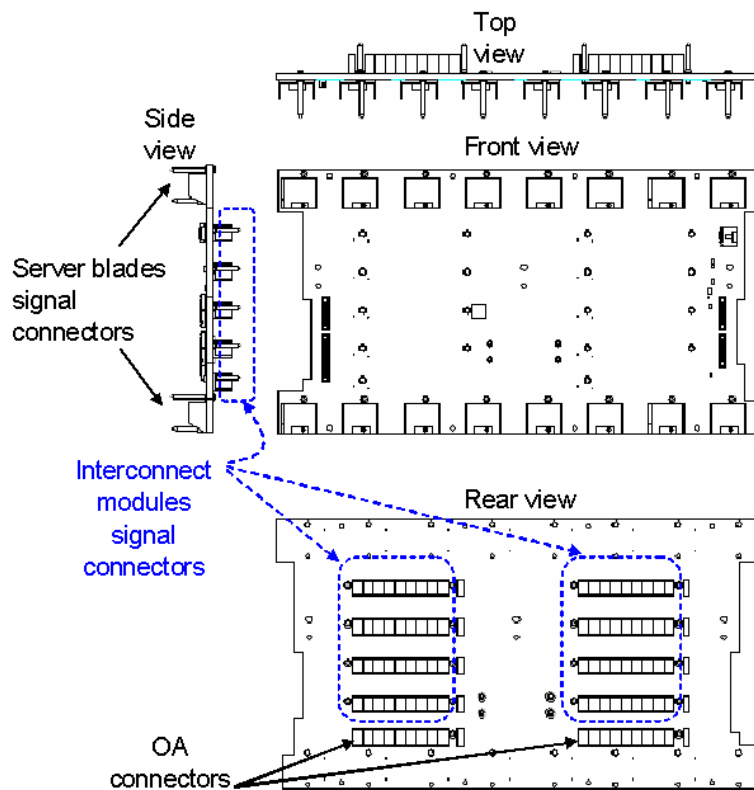
HP Insight Control suite integrates with leading enterprise management platforms through industry standards. It includes one year of 24 x 7 HP Software Technical Support and Update Service.

Interconnect options and infrastructure

A key component of the c7000 enclosure is the I/O infrastructure—essentially, 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 takes advantage of serializer/deserializer (SerDes) technology to support 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 sixteen connectors for the server/storage blades. Internal traces link them to the eight connectors on the other side of the board for the interconnect modules (Figure 7).

The signal midplane also includes the management signals from each bay to the BladeSystem Onboard Administrator modules. These management signals are completely isolated from the high-speed server-to-interconnect signals. The BladeSystem Onboard Administrator is the terminating point for all interconnect bays. An interconnect module cannot use the connection to the BladeSystem Onboard Administrator to communicate with another interconnect module.

Figure 7. Illustration showing both the server blade connectors and the interconnect module connectors




By taking advantage of the similar four-trace, differential SerDes transmit and receive signals, the NonStop signal midplane can support either network-semantic protocols (such as Ethernet, Fibre Channel, and InfiniBand) or memory-semantic protocols (PCI Express), using the same signal traces.²

² 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

Figure 8 illustrates how the physical lanes can be logically overlaid onto sets of four traces. Interfaces such as Gigabit Ethernet (1000-base-KX) or Fibre Channel need only a 1x lane, or a single set of four traces. Higher bandwidth interfaces, such as InfiniBand DDR, use up to four lanes.

Figure 8. Logically overlaying physical lanes (right) onto sets of four traces (left)



The NonStop signal midplane has eight 200-pin connectors to support 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 HP BladeSystem c-Class Enclosure requires two interconnect switches or two pass-thru modules, side-by-side, for a fully redundant configuration.

Interconnect modules

The c7000 enclosure offers a variety of interconnect options, including pass-thru modules, Ethernet and Fibre Channel switches, and high-bandwidth fabrics such as InfiniBand. The most up-to-date information about c-Class interconnect modules is available at this website:

www.hp.com/go/ bladesystem/interconnects.

Switches offer a traditional approach to administering the 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 direct one-to-one connections between servers and LAN or SAN are required. 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 c7000 are available in two widths: single- and double-wide. Single-wide modules provide sixteen internal ports, each connected to a separate device bay in the front of the enclosure. Double-wide modules provide sixteen internal ports, each double-wide, providing connectivity to DDR InfiniBand and other 4-lane, high speed interconnects.

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.

Each interconnect module also provides external connectors that vary based on its particular design. In the c7000 enclosure, pairs of single-wide interconnect modules installed in adjacent horizontal bays provide redundant connectivity for dual-port interfaces in each device bay. Adjacent interconnect modules also have high-speed cross-connect capability through the enclosure's NonStop signal midplane. For double-wide interconnects such as DDR InfiniBand, two modules are installed in c7000 interconnect bays 5 and 7 to provide redundant high bandwidth connectivity.

NOTE

The c-Class Ethernet Pass-Thru Module only supports fixed-speed gigabit Ethernet. Because the server, storage, or other option blades are connected through SerDes to the interconnect bays, and SerDes Ethernet does not have an auto-negotiation protocol, a switch is required to connect to 10/100 networks outside of the enclosure. The NICs themselves are capable of different modes of operation, but the outbound wiring to which they are connected does not auto-negotiate. Note that this is a limitation of an Ethernet Pass-Thru Module only. The Fibre Channel Pass-Thru Module ports do auto-negotiate.

Server blades

Server blades for the BladeSystem c7000 enclosure are built according to c-Class standard form-factors referred to as half-height and full-height. The enclosure can hold either full-height or half-height server blades or a combination of the two. The enclosure is preconfigured with device bay dividers to house half-height server or storage blades. Dividers must be removed to accommodate full-height server blades.

For connectivity, every server blade ships with at least two built-in Ethernet connections. To maintain flexibility, the server blades use optional mezzanine cards to provide additional interconnect fabric connections such as Gigabit Ethernet, InfiniBand, and Fibre Channel.

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 Gigabit NIC mezzanine card provides eight independent Gigabit 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 three quad-port Gigabit NIC mezzanine cards provides sixteen independent Gigabit NICs. The inherent flexibility of c-Class design allows customers to configure up to four different interconnect fabrics without sacrificing redundancy or performance.

The most up-to-date information about c-Class server blades is available at this website: www.hp.com/go/bladeSystem/.

Storage options inside the BladeSystem enclosure

Storage options inside the enclosure provide an alternative to local disk drives or SAN connectivity. The c7000 enclosure supports 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 would not fit inside the server blade.

Direct attach storage blades

The StorageWorks SB40c Storage Blade delivers direct attach storage for c-Class server blades, with support for up to six hot plug small form factor (SFF) SAS or SATA hard disk drives. The enclosure backplane provides a PCIe connection to the adjacent c-Class server blade and enables high

performance storage access without any additional cables. The SB40c Storage Blade features an onboard Smart Array P400 controller with 256 MB battery-backed write cache, for increased performance and data protection. For mechanical compatibility, storage blades use the same half-height form factor as server blades. The SB40c direct attach storage blade must be paired with an adjacent server blade in the same cooling zone. This is 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. The half-height server blade and direct attach storage blade must be paired in specific bays. Refer to the appendix for bay number details. A mezzanine card is not required to connect a half-height server blade to a direct attach storage blade.

A mezzanine card is required to connect a full-height server blade to a direct attach storage blade. The card must be in the Mezzanine 3 connector to allow full use of the interconnect bays with Type I or Type II mezzanine cards, to be consistent with the half-height blades, and to enable mixing half-height and full-height server blades in the same enclosure.

NOTE

When used with a full-height server blade, the direct attach storage blade must be in the bottom bay. To maintain essential airflow, a blank must be attached above the storage blade to block the empty upper bay, or a half-height server blade must be inserted into the upper bay. For this configuration, the storage blade should be installed before the half-height server blade is installed, and the half-height server blade should be removed before the storage blade is removed.

Shared storage

HP LeftHand P4000 Virtual SAN Appliance (VSA) software turns the SB40c into an iSCSI SAN for use by all servers in the enclosure and any server on the network. HP LeftHand VSA software is installed in a virtual machine on a VMware ESX host server adjacent to the SB40c. HP LeftHand P4000 VSA features storage clustering for scalability, network RAID for storage failover, thin provisioning, snapshots, remote replication, and cloning. Capacity can be expanded within the same enclosure or to other BladeSystem enclosures by adding additional SB40c Storage Blades and HP LeftHand VSA licenses.

External SAS connectivity with direct connect SAS storage for HP BladeSystem

With direct connect SAS storage for HP BladeSystem, administrators can build local server storage on the fly with zoned storage, or enable low cost shared storage within the rack with high performance 3 Gb/s SAS architecture. Each HP server blade requires an HP Smart Array P700m controller installed in a mezzanine slot. Redundant HP StorageWorks 3 Gb SAS BL switches are installed in the interconnect bays of the HP BladeSystem enclosure. The switches are then connected through an SAS cable to external storage depending on application requirements.

NAS/SAN Gateway

The X3800sb Network Storage Gateway Blade is a ready-to-deploy SAN gateway solution, with Microsoft Storage Server 2008 Enterprise x64 Edition pre-installed. The X3800sb Network Storage Gateway Blade is used to access FC, SAS, or iSCSI SAN storage, translating file data from the server into blocks for storage to provide consolidated file, print, and management hosting services in a cluster-able package.

Storage server

The X1800sb Network Storage Blade comes with Microsoft Windows Storage Server 2008 Standard Edition pre-installed. The X1800sb can be paired with the SB40c storage blade to create shared

storage and file serving inside the BladeSystem enclosure. The X1800sb can also be used as a low cost gateway to external FC, SAS, or iSCSI storage.

For the most up-to-date information about HP StorageWorks storage blade solutions, go to the HP website: www.hp.com/go/storageblades.

Mezzanine cards

HP offers 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 only in the amount of 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 to them and are slightly smaller. Type I mezzanine cards are compatible with all ProLiant c-Class server blades in all mezzanine connectors (Table 1). Type II mezzanine cards are compatible with Mezzanine 2 or 3 connectors in full-height c-Class server blades. Type II mezzanine cards are also compatible with Mezzanine 2 connectors in half-height c-Class server blades.

Table 1. Compatibility of mezzanine cards and mezzanine connectors in ProLiant c-Class server blades

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

For the most up-to-date information about the c-Class mezzanine card options, go to the HP website: <http://h18004.www1.hp.com/products/blades/components/c-class-interconnects.html>.

Virtual Connect

Virtual Connect technology is a set of interconnect modules and embedded software for c-Class enclosures that simplifies the setup and administration of server connections. HP Virtual Connect includes the following components:

- HP 1/10 Gb Virtual Connect Ethernet Module
- HP 1/10 Gb-F Virtual Connect Ethernet Module
- HP Virtual Connect Flex-10 10 Gb Ethernet Module
- HP Virtual Connect 4 Gb Fibre Channel Module
- HP Virtual Connect 8 Gb 24-Port Fibre Channel Module
- HP Virtual Connect Manager

Virtual Connect implements server-edge virtualization so that server administrators can upgrade, replace, or move server blades within their enclosures without changes being visible to the external LAN and SAN environments. HP recommends using Virtual Connect or managed switches to reduce cabling and management overhead.

Like other Ethernet and Fibre Channel switches, Virtual Connect modules slide into the interconnect bays of c-Class enclosures. To support Fibre Channel, the enclosure must have at least one Virtual

Connect Ethernet module, because the Virtual Connect Manager software runs on a processor that resides on the Ethernet module.

When the LAN and SAN connections are made to the pool of servers, the server administrator uses Virtual Connect Manager to define a server connection profile for each server. Instead of using the default media access control (MAC) addresses for all network interface controllers (NICs) and default World Wide Names (WWNs) for all host bus adapters (HBAs), the Virtual Connect Manager creates bay-specific profiles, assigns unique MAC addresses and WWNs to these profiles, and administers them locally. Network and storage administrators can establish all LAN and SAN connections once during deployment. If servers are later deployed, added, or changed, no connection changes are needed because Virtual Connect keeps the profile for that LAN and SAN connection constant.

To help administrators fully utilize 10 GbE connection bandwidth, HP introduced Flex-10 technology in the BladeSystem c-Class architecture. Using Flex-10, administrators can partition the bandwidth of a single 10 Gb pipeline into multiple *FlexNICs*. In addition, administrators can regulate the bandwidth for each partition by setting it to a user-defined portion of the total 10 Gb connection. Speed can be set from 100 Mb per second to 10 Gb per second in 100 Mb increments.

There are advantages to partitioning a 10 GbE pipeline:

- More NIC 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

Flex-10 technology uses two hardware components:

- The HP Virtual Connect Flex-10 10Gb Ethernet Module
- Either the 10 Gb Flex-10 LAN-on-motherboard (LOM) or the HP NC532m Flex-10 10 GbE Network Adapter mezzanine card

The HP Virtual Connect Flex-10 10 Gb Ethernet Module is required to manage the 10 GbE (Flex-10) server connections to the data center network. The 10 Gb Flex-10 LOM and mezzanine cards are NICs, 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 as an individual NIC.

Full details about Flex-10 technology are available in the technology brief titled “HP Flex-10 technology” on the HP technology website:

<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01608922/c01608922.pdf>.

Full details about Virtual Connect technology are available in the technology brief titled “HP Virtual Connect technology implementation for the HP BladeSystem c-Class” on the HP technology website:

<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00814156/c00814156.pdf>.

Fabric connectivity and port mapping

Because the connections between the device bays and the interconnect bays are hard-wired through the NonStop signal midplane, the 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.

Interconnect bays 1 and 2 are reserved for Ethernet switches or pass-thru modules supporting server LAN on Motherboard (LOM) NIC connections to ports on the Ethernet switch or pass-thru module. Supported bays for additional Ethernet switch modules include unpopulated interconnect bays 3/4, 5/6, or 7/8. Redundant switches must be adjacent to one another in interconnect bays 3/4, 5/6, or 7/8.

Connecting the ports of embedded devices to the interconnect bays in the c7000 enclosure is relatively simple. For port mapping purposes, it does not matter in which bay a server blade is installed. The mezzanine connectors always connect to the same interconnect bays. Figure 9 shows BladeSystem c7000 port mapping of half-height server blades to interconnect bays. Figure 10 shows BladeSystem c7000 port mapping of full-height server blades to interconnect bays.

Figure 9. HP BladeSystem c7000 port mapping of half-height server blades to interconnect bays

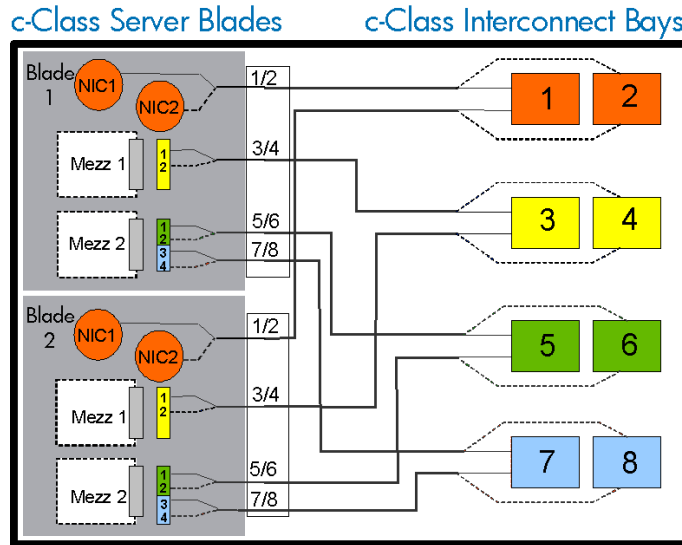
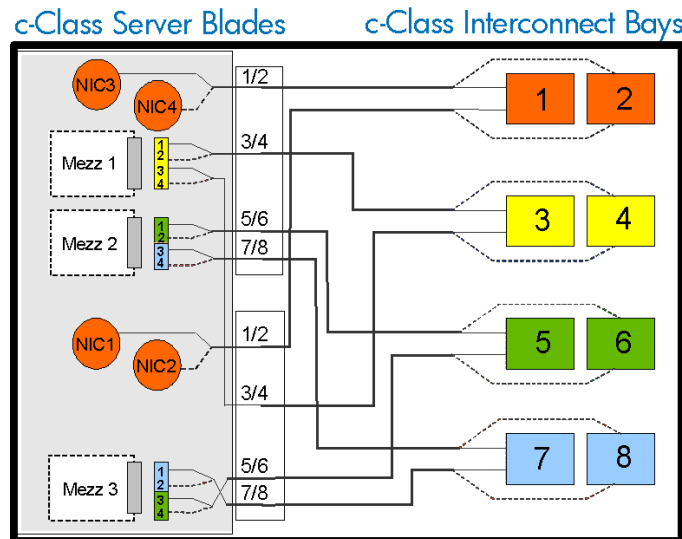


Figure 10. HP BladeSystem c7000 port mapping of full-height server blades to interconnect bays



Port mapping differs slightly between full-height and half-height server blades due to the support for additional mezzanine cards on the full-height version. HP simplified the process of mapping mezzanine ports to switch ports by providing management tools through the BladeSystem Onboard Administrator and HP Systems Insight Manager software. Specific port mapping details can be found

in the HP BladeSystem Onboard Administrator User Guide available at this website:
<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00705292/c00705292.pdf>.


c7000 bay-to-bay crosslinks

Four-trace SerDes signals between adjacent bays in the c7000 midplane permit horizontal bay-to-bay communications.

Device bay crosslinks

Device bay crosslinks are wired between adjacent device bay pairs as indicated by the arrows in the c7000 enclosure front view (Figure 11). 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. The two blades always occupy a pair of slots. If the server blade is installed in an odd-numbered bay, the partner blade is installed in the adjacent even-numbered bay to the right. If the server blade is installed in an even-numbered bay, the partner blade is installed in the adjacent odd-numbered bay to the left. 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 server blade is installed in an odd-numbered bay. The partner blade is installed in the adjacent even-numbered bay to the right and is always located on the lower row of the enclosure in slot 10, 12, 14, or 16. The BladeSystem Onboard Administrator disables the device bay crosslinks in instances where they cannot be used, such as two server blades residing in adjacent device bays.

Figure 11. HP BladeSystem c7000 device bay crosslinks as indicated by the arrows



Interconnect bay crosslinks

Interconnect bay crosslinks are wired between adjacent interconnect bay pairs as indicated by the arrows on rear view of the c7000 enclosure in Figure 12. Crosslink signals can be enabled to 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 12. HP BladeSystem c7000 interconnect bay crosslinks as indicated by the arrows



HP Thermal Logic technologies

The HP BladeSystem c7000 Enclosure incorporates a variety of HP Thermal Logic technologies, including mechanical design, built-in power and thermal monitoring, and control capabilities. Thermal Logic technologies can provide significant power and cooling savings compared to traditional rack and tower based servers.³ 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 and cooling use while maintaining adequate cooling for all devices and ensuring high availability.

HP Thermal Logic technologies include the following elements and capabilities:

- Active Cool 200 fans
- Parallel Redundant Scalable Enclosure Cooling (PARSEC) design
- 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 200 fans

Quite often, dense, full-featured, small form-factor servers use very small fans designed to provide localized cooling in the specific areas needed by the server blade. Because such fans generate fairly low airflow (in cubic feet per minute, or CFM) at medium back pressure, 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 that can provide cooling across an entire enclosure. Such fans generate high airflow, but they typically require higher power input and more space. They are loud and must be 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.


³ For more information see the IDC technical brief titled "Next-Generation Power and Cooling for Blade Environments": <http://h20195.www2.hp.com/v2/GetPDF.aspx/4AA2-5177ENW.pdf>.

To overcome these issues in the c7000 enclosure, HP 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. HP has 20 patents pending for its Active Cool fan technology and implementation.

HP Active Cool 200 fans can cool 16 server blades using as little as 150 watts of power. Active Cool 200 fans use ducted fan technology with a high-performance motor and impeller to deliver high CFM at high pressure (Figure 13). 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. This design provides cooling capacity to support blade products beyond current roadmaps. 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 13. Ducted fan cross-section and ducted fan blade compared to traditional server fan

Ducted fan cross-section



Active Cool 200 fans are controlled by the BladeSystem Onboard Administrator, which can ramp cooling capacity 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 c7000 enclosure can accommodate full-featured servers that are 60% more dense than traditional rack-mount servers, while the Active Cool 200 fans consume only 50% of the power typically required and use 30% less airflow. As new Active Cool fans are introduced, they may not interoperate with Active Cool 200 fans. The BladeSystem Onboard Administrator will indicate a fan compatibility problem if mixed fans that are not interoperable are detected in an enclosure.

HP PARSEC architecture

The c7000 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). The enclosure is divided into four cooling zones with fans in each. The Active Cool 200 fans provide cooling for their own zone and redundant cooling for the rest of the enclosure. To ensure scalability, HP 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, HP developed a relatively airtight center air plenum, or air chamber. In the c7000 enclosure, all device bays include a shutoff door that is normally closed to prevent air leakage through that device bay. When a server blade is inserted, it seals into the center air plenum


docking collar, and the server shut-off door opens to allow airflow across that server blade. Similarly, the fan seals into the center air plenum docking collar. Each fan bay includes louvers that automatically open when a fan is installed. If a fan is not installed or is not functional, the pressure distribution around the fan changes. This pressure change causes the louvers to close, ensuring that cooling air is not diverted through the inoperative fan (Figure 14).

Figure 14. HP BladeSystem c7000 self-sealing enclosure



The enclosure and the components within it optimize the cooling capacity through unique mechanical designs. Managed airflow through the enclosure ensures that every device gets cool air, that no device is in the path of hot exhaust air from another device, and that air goes only where it is needed for cooling. Fresh air is pulled into the interconnect bays through a slot in the front of the enclosure. Ducts allow the air to move from the front to the rear of the enclosure, where it is then pulled into the interconnects and the central plenum. The air is then exhausted out the rear of the system (Figure 15).


Figure 15. Airflow through the HP BladeSystem c7000 Enclosure – side view



Redundant

As illustrated in Figure 16, fans located in each of four cooling zones provide direct cooling for server blades in their respective zones and redundant cooling for adjacent zones. Each zone can contain four server blades. Full-height server blades couple the two vertical zones into a single zone, as in the example in Figure 16 where a full-height server blade in bay 1 has resulted in zones 1 and 3 operating as a single zone. Fans 3 and 8, shown in the Figure 16 rear view, support two adjacent horizontal zones and will operate at the higher speed of either of the two zones. If any fan fails, the result is only a 10 to 25% loss in cooling capacity.

Figure 16. Four cooling zones in the c7000 enclosure




Zone cooling minimizes the power consumption of the fan subsystem by increasing fan efficiency in a single zone if one of the server blades requires more cooling. This saves operating costs and

minimizes fan noise. HP recommends using at least eight fans. Using ten fans optimizes power and cooling. See the appendix for more detailed fan and server population guidelines.

Scalable

Operating c-Class server blades requires installing a minimum of four fans at the rear of the c7000 enclosure. Up to ten fans can be installed so that cooling capacity can scale as needs change. Using more fans allows the fans to spin slower to move the same volume of air, so each fan uses less power. Eight fans are almost always more power-efficient than four fans. As the air flow rate increases, ten fans are even more efficient (Figure 17). Slower spinning fans also create less noise.

Figure 17. General relationship between the number of fans in a c7000 enclosure and the associated power draw



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 thermal work from the least amount of air. This concept allows much 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, HP was 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 18). The smaller heat sink allows more space on the server blades for DIMM slots and hot-plug hard drives.

Figure 18. Processor heat sink using fully ducted design (left) versus traditional heat sink in a 1U rack-mount server (right)



Instant thermal monitoring provides a real-time view of heat, power, and cooling data. The BladeSystem Onboard Administrator 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 BladeSystem Onboard Administrator keeps fan and system power at the lowest level possible. However, if the thermal load within 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 ProLiant Onboard Administrator (iLO 2) and BladeSystem Onboard Administrator 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 Onboard Administrator (iLO 2), and BladeSystem Onboard Administrator. In combination, these algorithms minimize the power and cooling required to maintain the proper HP BladeSystem environment.

Power supplies and enclosure power subsystem

Expecting that power supply technology will advance during the lifecycle of the c7000 enclosure, HP designed the c7000 power subsystem to be flexible and upgradeable. Customers can choose from four types of HP BladeSystem c7000 Enclosures to select the one that will work best with their data center power infrastructure:

- Single-phase enclosure that accepts IEC C19-C20 power cords (available worldwide for use with in-rack power distribution units)
- Three-phase enclosure with a pair of US/Japan power cords with NEMA L15-30P power connectors
- 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)

Figure 19 illustrates the c7000 enclosure power options.

Figure 19. c7000 enclosure power options



- Single-phase AC enclosure available worldwide for use with in-rack PDUs which accept C19 – C20 power cords



C-19 16A



- Three-Phase AC enclosure with a pair of NA/JPN power cords with NEMA L15-30P power connectors



NA/JPN L15-30p



- Three-Phase AC enclosure with a pair of International power cords with IEC 309 16A power connectors



Intl IEC309 5-Pin, 6h, 16A

- -48V DC enclosure with DC lug connectors



DC lugs

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, HP 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.

NOTE

The c7000 power module in the bottom rear of the enclosure determines whether the enclosure operates on single-phase AC, three-phase AC, or DC power. This choice is available when the enclosure is ordered.

The enclosure can contain up to six 2250W self-cooled or 2400W self-cooled high efficiency power supplies, for a maximum of 14400W output capacity per enclosure. The 2400W power supplies provide higher efficiency to save energy—92% peak and over 88% efficient with as little as a 10%

load.⁴ They also lower standby power for reduced power consumption when servers are idle. The HP 2400W high efficiency power supplies operate at 200-240V high line only and require BladeSystem Onboard Administrator firmware version 2.40 or later. The high efficiency 2400W power supplies do not interoperate with existing 2250W power supplies. A pair of power distribution units (PDUs) is required for AC line redundancy. PDU options are available in the c-Class infrastructure QuickSpecs: http://h18000.www1.hp.com/products/quickspecs/12517_div/12517_div.html.

HP recommends single-phase enclosure models for most situations. At the time of this writing, most typical c-Class customer configurations that HP builds do not exceed 6 kVA for a worst case load. For example, under normal operating conditions, a system configuration of 16 server blades with 2 CPUs each, 4 to 8 GB of RAM, Ethernet, Fibre Channel, and all components at 100% load uses approximately 4 to 4.5 kVA. In customer environments, most systems will probably consume less than 4 kVA because all components will not be at 100% load simultaneously. Actual measurements of real customer applications show usage numbers between 3 and 3.5 kVA. A rack with four three-phase enclosures would require 8 x 30A three-phase power whips and would supply roughly 70 kVA to the rack, of which no more than 18 to 24 kVA would be used.

The three-phase c-Class enclosure is generally useful in only a few situations:

- When there is only one enclosure and there will never be more than one enclosure
- When the configuration is extreme (for example, 16 server blades with 2 high-power CPUs, 32 GB of RAM, all mezzanine connectors and switch bays filled, running high-performance technical computing loads)
- When the enclosure is a drop-in replacement for a p-Class solution where three-phase power already exists
- When the customer plans to maintain the enclosure for a long period of time and to update the internal components with newer ones that could take advantage of full three-phase power capacity

Greater power efficiency results from delivering three-phase power to the rack and using single-phase power within the rack. This can be achieved by using the S332 or S348 PDU to run 2 x 60A three-phase power to the rack, connected to three or four single-phase enclosures in the rack. If power capping technology is used in this scenario, enclosure power consumption will never exceed 4.3 kVA, meaning it will not exceed the capacity of the PDU, even if an AC line were to fail.

NOTE

US/Japan three-phase enclosures use a different type of three-phase power than International enclosures. For more information, see the "Power basics for IT professionals" technology brief available on the HP website: <http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01234421/c01234421.pdf>.

Pooled power

All the power in the c7000 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 customers can choose the required level of redundancy. Because this power design has no zones, it facilitates both N+N and N+1 power modes, which future-proofs the enclosure for higher power requirements, if needed. Therefore, looking forward at least five years, HP believes there is sufficient power capacity to handle future power-hungry devices.

⁴ For more information about power supply efficiency, see the 80 PLUS® website: http://www.80plus.org/manu/psu/psu_detail.aspx?id=41&type=1.


The c7000 enclosure has three redundancy modes: AC or DC redundant, power supply redundant, and no redundancy mode. The BladeSystem Onboard Administrator or the Insight Display can be used to select the power redundancy mode. For more information, the HP BladeSystem Onboard Administrator User Guide is available at this website:

<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00705292/c00705292.pdf>.

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

If the user changes the power mode of the enclosure to AC redundant, the BladeSystem Onboard Administrator ensures that equal numbers of power supplies are available on the A and B sides as indicated in Figure 20. This logical grouping ensures that when the enclosure is configured with the three-phase AC input module, all the power supplies on one side will maintain enclosure power if AC input power is lost on the other side. If the user configures the enclosure power mode to N+1 or not redundant, the power supplies can be installed in any bay.

Figure 20. HP BladeSystem c7000 power supplies



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. Up to six power supplies can be installed 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 will cause the system to power off if the remaining power supplies are unable to handle the full load.

The BladeSystem Onboard Administrator manages power allocation rules of various components and can limit overall power capacity for the enclosure. More information on power management is

available in the technology brief entitled “Managing the HP BladeSystem c-Class”:
<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00814176/c00814176.pdf>.

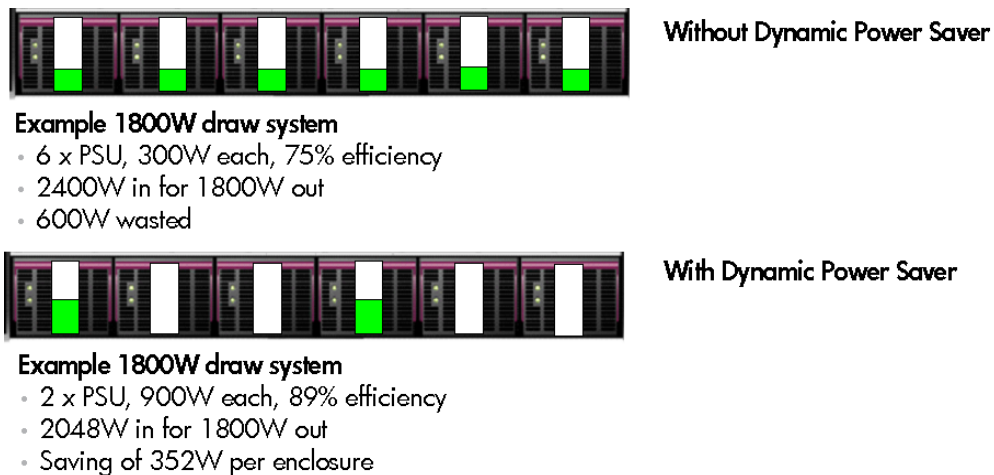
Dynamic Power Saver mode

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 since it saves power in the majority of 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%. However, at 50% load, the efficiency rating could be up to 92%, providing a significant savings in power consumption.

In the first example in Figure 21, without Dynamic Power Saver, power demand is low and spread inefficiently across six power supplies. In the second example, with Dynamic Power Saver active, the power load is shifted to two power supplies for more efficient operation. The remaining power supplies are placed in a standby condition. When power demand increases, the standby power supplies instantaneously deliver the required extra power. This enables the power supplies to operate at optimum efficiency, with no effect on redundancy.

Figure 21. Example of power efficiency with and without Dynamic Power Saver



Dynamic Power Saver is enabled by the BladeSystem Onboard Administrator module. 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 that can be reached on any power supply is 50%. Once the 50% load is reached, another two power supplies are activated to ensure that redundancy is maintained at all times.

Power Regulator

HP ProLiant Power Regulator provides iLO-controlled speed stepping for Intel x86 and recent AMD 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. Power Regulator can be configured for continuous, static low power mode or for Dynamic Power Savings mode in which power is automatically adjusted to match processor demand. Additional information on the HP Power Regulator is provided in the paper titled "Power Regulator for ProLiant servers":

<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00593374/c00593374.pdf>.

Basic Power Capping for each server blade

Using ProLiant Onboard Administrator (iLO 2) firmware version 1.30 and System ROM/BIOS dated May 1, 2007, or later, ProLiant server blades can limit the amount of power consumed. Customers can set a limit in watts or BTUs per hour. The purpose of this limit is to constrain the amount of power consumed, which reduces heat output into the data center. The ProLiant Onboard Administrator (iLO 2) firmware monitors server power consumption, checks it against the power cap goal, and, if necessary, 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.

Using the Insight Power Manager (IPM) v1.10 plug-in to HP Systems Insight Manager v5.1, customers can set power caps on groups of supported servers. The IPM software statically allocates the group power cap among the servers in the group. The group cap is allocated equitably among all servers in the group based on a calculation using the idle and maximum measured power consumption of each server. In addition, IPM can track and graph over time the actual power use of groups of servers and enclosures. This provides data center facilities managers with measured power consumption for various time 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 with ProLiant Onboard Administrator (iLO 2) version 1.70. Supported servers contain an internal hardware circuit that monitors server power consumption on a sub-second basis. If consumption approaches the power cap set in ProLiant Onboard Administrator (iLO 2), 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 administrators set an enclosure level power cap. The BladeSystem Onboard Administrator manages the power distribution to the server blades and will change server caps over time as workloads change. Enclosure Dynamic Power Capping requires redundant power wired to separate A and B feeds, redundant BladeSystem Onboard Administrator modules using firmware version 2.30 or later, and supported servers using ProLiant Onboard Administrator (iLO 2) version 1.70 or later.

HP BladeSystem Power Sizer

The HP BladeSystem Power Sizer is a tool that assists facilities teams and IT staff in sizing their power and cooling infrastructures to meet the needs of an HP BladeSystem solution. The BladeSystem Power Sizer is based on actual component-level power measurements of a system stressed to maximum capability. It allows a customer to select the type and number of components within each server blade and enclosure and to see the effect of changes on power consumption and heat loading.

Values obtained from the BladeSystem Power Sizer tool are based on heavy enterprise application loads and are intended as guidelines only. Actual measured power consumption may be significantly lower or higher than the numbers provided by the Power Sizer. Power consumption will vary with application type, application utilization, and ambient temperature. The BladeSystem Power Sizer is available at this website: <http://www.hp.com/go/bladeSystem/powercalculator>.

Conclusion

The HP BladeSystem c7000 Enclosure is the foundation of a modular computing architecture that consolidates and simplifies infrastructure, reduces operational cost, and delivers IT services more effectively. Thermal Logic technologies provide the mechanical design features, built-in monitoring, and control capabilities that enable IT administrators to optimize their power and thermal environments. The shared, high-speed, NonStop midplane and pooled-power backplane in the enclosure accommodate new bandwidths and new technologies. The BladeSystem Onboard Administrator supplies the infrastructure to provide essential power and cooling information and help to automate infrastructure management. The HP BladeSystem c7000 Enclosure provides all the power, cooling, and I/O infrastructure to support c-Class modular servers, interconnects, and storage components, today and for the next several years.

Appendix: Fan and server population guidelines

For correct operation, fans and server blades must be installed in the correct bays of the HP BladeSystem c7000 Enclosure. The BladeSystem Onboard Administrator ensures that fans and blades are correctly placed before allowing systems to power on.

Fan bay numbering

Figure A-1 indicates the location of the ten fan bays in the back of the enclosure. Table A-1 identifies the correct bays in which up to ten fans must be located.

Figure A-1. Required order for populating fans in the HP BladeSystem c7000 Enclosure





Table A-1. Fan bay population locations

Number of fans	Fan bays populated (cooling zones)	Device bays supported
4	4, 5 (Zone 1) and 9, 10 (Zone 3)	Two devices located in bays 1, 2, 9 or 10
6	3, 4, 5 (Zone 1) and 8, 9, 10 (Zone 3)	1-4, 9-12
8	1, 2, 4, 5, 6, 7, 9, 10 (all four zones)	1-4, 9-12, 5-8, 13-16
10	1, 2, 3, 4, 5, 6, 7, 8, 9, 10 (all four zones)	1-4, 9-12, 5-8, 13-16

Server blade bay numbering

Half-height server blades should be populated in the front of the enclosure from top to bottom and from left to right. The first two half-height server blades should be placed in bays 1 and 9; the second two half-height server blades should be placed in bays 2 and 10 and so on, until the enclosure is full. Full-height server blades should be populated from left to right.


Figure A-2. Numbering of half-height server bays in the HP BladeSystem c7000 Enclosure



Enclosure blade zones

The c7000 enclosure is divided into four zones by vertical support metalwork. These are not the same as the cooling zones referred to above. Within each zone a removable divider is used to support half-height devices. To install a full-height server blade in any zone, this divider must be removed. Storage blades and tape blades can be installed in the same zone as either full-height or half-height server blades. Each SB40c or tape blade comes with a bracket that allows a half-height blade to be mounted on top of the storage blade.

Figure A-3. Mixed configuration – full-height and half-height population rules



A half-height SB40c storage blade must be installed within the same zone as its partner server blade. If the SB40c is partnered with a full-height server blade, the SB40c must be installed in the bottom

position adjacent to the full-height server blade. The storage blade requires an HP SB40c PCIe pass-thru mezzanine card in the mezzanine 3 connector.

To maintain proper airflow, either a half-height server blade or a blade blank must be installed in the half-height position above the SB40c. Installing a half-height server blade above the SB40c requires first installing a half-height bracket.

For more information

For additional information, refer to the resources listed below.

Resource description	Web address
General HP BladeSystem information	http://www.hp.com/go/bladefsystem/
HP BladeSystem c7000 Site Planning Guide	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01038153/c01038153.pdf
HP BladeSystem c-Class documentation	http://h71028.www7.hp.com/enterprise/cache/316735-0-0-0-121.html
HP BladeSystem c-Class Enclosure Setup and Installation Guide	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00698286/c00698286.pdf
HP BladeSystem Onboard Administrator User Guide	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00705292/c00705292.pdf
HP BladeSystem c-Class interconnects	www.hp.com/go/bladefsystem/interconnects
Technical white papers and podcasts about HP BladeSystem	http://h18004.www1.hp.com/products/servers/technology/whitepapers/proliant-servers.html
HP BladeSystem Power Sizer	http://www.hp.com/go/bladefsystem/powercalculator
Managing HP BladeSystem c-Class Systems technology brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00814176/c00814176.pdf
HP ProLiant c-Class Server Blades technology brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01136096/c01136096.pdf
HP BladeSystem c3000 Enclosure technology brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01508406/c01508406.pdf
HP Power Capping and Dynamic Power Capping for ProLiant Servers technology brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01549455/c01549455.pdf?jumpid=reg_R1002_USEN
HP Flex-10 technology	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01608922/c01608922.pdf

Call to action

Send comments about this paper to TechCom@HP.com.

© 2006, 2008, 2009 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

Intel and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

AMD is a trademark of Advanced Micro Devices, Inc.

Linux is a U.S. registered trademark of Linus Torvalds.

TC090606TB, June 2009

