



Hewlett Packard
Enterprise

Technical white paper

Architecture and technologies in the HPE BladeSystem c7000 Enclosure

Contents

Introduction.....	3
HPE BladeSystem design.....	3
Scalable blade form factors and device bays.....	5
Scalable interconnect form factors.....	6
Star topology.....	7
Signal midplane design and function.....	7
Power backplane scalability and reliability.....	10
Enclosure connectivity and interconnect fabrics.....	11
HPE Virtual Connect.....	11
Interconnect modules.....	12
Server blades.....	12
Server blade products.....	12
Storage and tape options.....	13
Mezzanine cards.....	14
Fabric connectivity and port mapping.....	14
BladeSystem c7000 Enclosure bay-to-bay crosslinks.....	16
BladeSystem c7000 Enclosure management links.....	17
HPE Thermal Logic technologies.....	17
HPE Active Cool 200 fans.....	17
HPE PARSEC architecture.....	18
Power and thermal monitoring.....	20
3D sea of sensors.....	20
Enclosure power subsystem and capabilities.....	20
HPE Location Discovery Services.....	23
Conclusion.....	24
Appendix: Fan and server population guidelines.....	24
Fan bay numbering.....	24
Server blade bay numbering.....	25
Enclosure blade zones.....	25
Additional information.....	27



Introduction

The HPE BladeSystem c7000 Enclosure is a wire-once, Hybrid IT and cloud-ready system architecture designed to deploy into traditional enterprise solutions and build on top of existing infrastructure with applications and virtualization. When you're ready, you can move into cloud-based deployments with HPE CloudSystem. HPE CloudSystem is at the core of our cloud-ready converged infrastructure and part of the HPE Converged Cloud architecture. HPE CloudSystem offers a complete system to build and manage services across public, private, and hybrid clouds. You can do this with the [HPE BladeSystem](#) infrastructure as the foundation with upgrades along the way. We realize that budgets are constrained and the underlying technologies engineered in HPE BladeSystem help leverage the greatest value from every hour, watt, and dollar.

As a part of HPE BladeSystem, c7000 Enclosures use an Intelligent Infrastructure managed by Onboard Administrator (OA) and Integrated Lights-Out (iLO) Management Engine to monitor health and dynamically allocate power and cooling resources for greater efficiency. HPE Thermal Logic technologies such as HPE Active Cool fans, HPE OneView power management capabilities, and Enclosure Dynamic Power Capping help 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 demanding applications and their availability requirements. With the local Insight Display or HPE OneView management software, administrators can perform many functions locally or with a remote console station.

The enclosure features support for [HPE OneView](#) to automate the delivery and operations of IT services, transforming everyday management of server, storage, and network resources in physical and virtual environments. Its unique architecture features a software-defined approach to converged management that shifts the focus from "how devices run" to "how people work." With a single integrated view of your IT infrastructure, you can experience 66X faster workload deployment of servers, networking, and storage resources at data center scale. HPE OneView is an intelligent automation hub that provides out-of-the-box integration with HPE, VMware®, Microsoft®, and Red Hat® solutions, as well as easy integration with other management products and with your own custom solutions. HPE OneView architecture combines management of servers, virtual connect, and storage with open integration to deliver unprecedented ease of use that lets you deploy and manage your infrastructure faster, at a lower cost, and at virtually any scale.

Find more information on HPE CloudSystem at: hpe.com/info/cloudsystem.

Find more information on HPE OneView management at: hpe.com/info/oneview.

HPE BladeSystem design

The HPE BladeSystem c7000 Enclosure represents an evolution of the entire rack-mounted infrastructure consolidating and repackaging featured infrastructure elements—computing, storage, networking, and power—into a single infrastructure-in-a-box that accelerates data center integration and optimization.

The BladeSystem enclosure infrastructure is adaptive and scalable. It transitions with your IT environment and includes modular server, interconnect, and storage components. The enclosure is 10U high and holds full-height and/or half-height server blades that may be mixed with storage blades, plus redundant network and storage interconnect modules. The enclosure includes a shared high-speed NonStop passive midplane with aggregate bandwidth of up to 7 Tbps¹ for wire-once connectivity of server blades to network and shared storage. Power is delivered through a passive pooled-power backplane that enables the full capacity of the power supplies to be available to the server blades for improved flexibility and redundancy. Power input is provided with a very wide selection of AC and DC power subsystems for flexibility in connecting to data center power.

You can populate a BladeSystem c7000 Enclosure with these components:

- Up to eight full-height (FH) or 16 half-height (HH) server, storage, or other option blades
- Up to eight interconnect modules (four redundant fabrics) featuring a variety of industry standards including:
 - Ethernet
 - Fibre Channel
 - Fibre Channel over Ethernet (FCoE)
 - InfiniBand
 - iSCSI
 - Serial Attached SCSI (SAS)

¹ 7 Tbps calculated as follows: 16 half-height device bays x 16 lanes per device bay x 14 Gbps device bay to interconnect bandwidth per lane x 2 for duplex = 7,168 Gbps = 7 Tbps. The NonStop midplane also provides additional bandwidth through device bay-to-bay crosslinks, interconnect crosslinks, and management signals.



- Up to six hot-plug power supplies supporting N+1 and N+N redundancy
- Choice of a single-phase AC, three-phase AC, universal AC-DC high voltage, or a -48 VDC power subsystem
- Up to 10 hot-plug HPE Active Cool 200 fans
- Redundant BladeSystem OA management module (optional active-standby design)

A BladeSystem c7000 Enclosure provides these benefits:

- Local and remote hardware management is integrated across the solution so that one full enclosure can be managed as easily as one server.
- Management and network interconnects extend scalability beyond a single enclosure, allowing resources to be pooled and shared across multiple enclosures.
- Multiple server and network designs are accommodated in one enclosure.
- Costs, power consumption, and airflow requirements are lower when compared to rack-mounted servers.
- Cabling is reduced and simplified when compared to rack-mounted servers.

Note

In addition to the BladeSystem c7000 Enclosure, HPE offers the BladeSystem c3000 Enclosure for remote sites or small businesses. Find more information about the BladeSystem c3000 Enclosure at: hpe.com/servers/bladeSystem/enclosures.

Figures 1 and 2 show front and rear views of the BladeSystem c7000 Platinum Enclosure. The BladeSystem c7000 Platinum and the BladeSystem c3000 Platinum Enclosures are compatible with many of the same critical components such as servers, interconnects, mezzanine cards, and storage blades.




Figure 1. Front view of HPE BladeSystem c7000 Platinum Enclosure




Figure 2. Rear view of HPE BladeSystem c7000 Platinum Enclosure



The BladeSystem c7000 Enclosure delivers power through the power subsystem. It includes:

- A wide choice of power input options including single-phase high-line AC, three-phase high-line AC, universal AC/DC high voltage, or -48 VDC.
- Up to six hot-plug power supplies to 2,650 watts each with support for N+1 and N+N redundancy.
- A passive pooled power backplane enabling the full capacity of the power supplies is available to many components for improved flexibility and redundancy.

The BladeSystem c7000 Enclosure has redundant signal paths between servers and interconnect modules. The enclosure NonStop signal midplane and separate power backplane (Figure 3) are passive (have no active components). Separating the power delivery in the backplane from the high-speed interconnect signals in the midplane results in reduced thermal stress to the signal midplane.

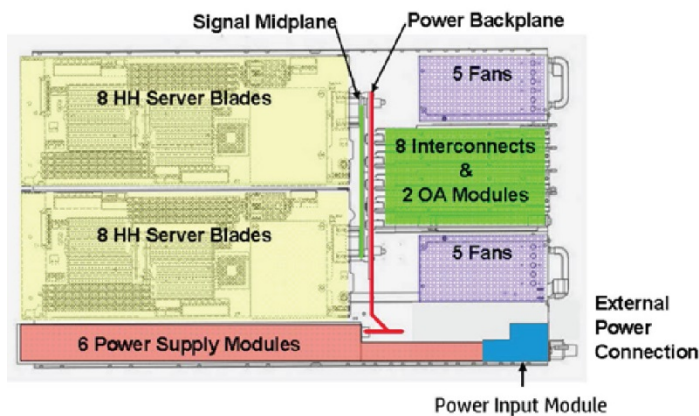


Figure 3. Enclosure NonStop signal midplane and power backplane

The dimensions of the c7000 Enclosure device bays provide several advantages over both narrower and wider alternatives:

- Holds components that support industry-standard protocols
- Holds enough blades to amortize the cost of the shared enclosure infrastructure
- Uses cost-effective, standard-height DIMMs in the server blades
- Uses vertical, rather than angled, DIMM connectors for:
 - Enhanced signal integrity
 - Additional space for heat sinks
 - Improved airflow across the DIMMs

Scalable blade form factors and device bays

The half-height and full-height blade form factors, which scale vertically in the c7000 Enclosure, provide several benefits, including reduced cost, increased reliability, improved ease-of-use, and configuration flexibility. Placing full-height form factors, half-height form factors, or both, in the same enclosure lets you exploit this flexibility. You can use two half-height bays together (one above and one below) to accommodate a full-height, single-wide server blade. The ProLiant BL660c Gen9 is one example of a full-height, single-wide server blades. Even HPE Integrity server blades are supported that combine multiple full-height, single-wide blades to create 2-, 4-, or 8-socket systems. These scalable Integrity systems use a Blade Link connector. The Blade Link connector joins selected QuickPath Interconnect (QPI) ports among processors; the required clock signals and side band signals for the system to operate are combined as a scale-up multiprocessor system. Common server blade configurations include 16 two-processor (2P) half-height server blades, eight four-processor (4P) higher performing full-height, single-wide server blades, or a mixture of the two form factors. Please see the [Appendix](#) at the end of this document for form factor population guidelines.



Figure 4 shows BladeSystem c-Class form factors scaling vertically with half-height and full-height blades.




Figure 4. HPE BladeSystem c-Class form factors with vertical scaling

Figure 5 shows horizontal scaling with Integrity blades and Blade Link connections.

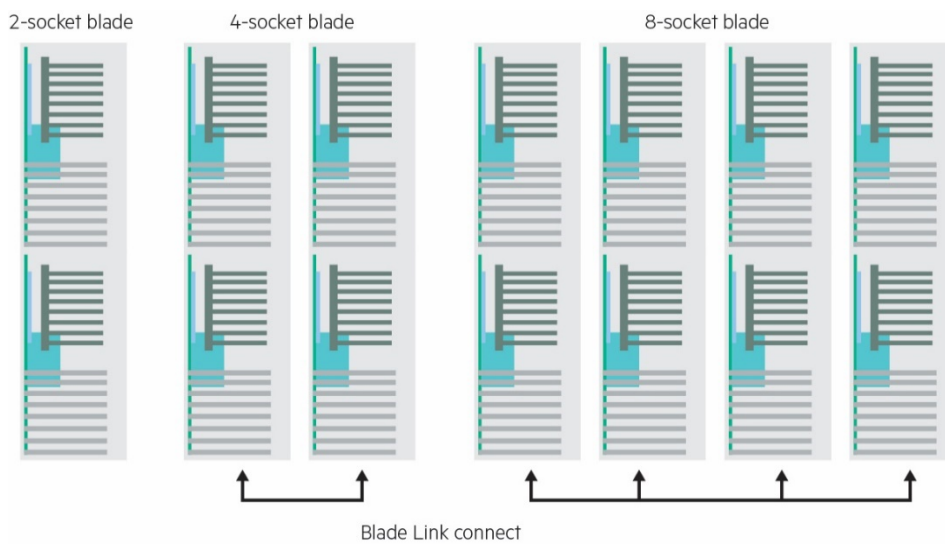


Figure 5. HPE BladeSystem c-Class form factors with horizontal scaling

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, double-wide interconnect module for higher performance, greater bandwidth interconnect modules, as shown in Figure 6.






Figure 6. Scalable interconnect bays

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

- It supports the greatest number of interconnect modules.
- It allows enough space in a double-wide module to include two signal connectors on the same printed circuit board (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 as shown in the left side of Figure 7, 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. The right side of Figure 7 shows the redundant configuration using double-wide interconnect modules.




Figure 7. Redundant configuration using double-wide interconnect modules

Signal midplane design and function

In addition to the scalable device bay and interconnect bay form factors, the c7000 Enclosure includes a high-bandwidth NonStop signal midplane and a separate power backplane.

NonStop signal midplane scalability and reliability

A key component of the BladeSystem c7000 Enclosure is the I/O infrastructure NonStop signal midplane that provides the internal wiring between the server or storage blades and the interconnect modules.

The NonStop signal midplane is a passive board that uses serializer/deserializer (SerDes) technology to incorporate multiple protocols and provide point-to-point connectivity between device bays and interconnect bays. To provide high reliability, HPE designed the NonStop signal midplane as a completely passive board. It has no active components along the high-speed signal paths. The midplane consists primarily of traces and connectors. While there are a few components on the midplane, they are limited to passive devices that are very unlikely to fail. The only active device is an Electrically Erasable Programmable Read-Only Memory (EEPROM) utility, which the OA uses to get information such as the midplane serial number. If this device fails, it does not affect the NonStop signal midplane.



On one side of the board are the 16 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 as shown in Figure 8.




Figure 8. NonStop signal midplane provides scalability and reliability

In the BladeSystem c7000 Platinum Enclosure, the aggregate bandwidth between device bays and interconnect bays is 7 Tbps across the NonStop signal midplane. The NonStop signal midplane in c7000 Platinum Enclosures features signal rates of up to 14 Gbps per lane (with each lane consisting of four SerDes transmit/receive traces). Each half-height server bay has 16 lanes to the interconnect bays delivering a cross-sectional bandwidth of up to 224 Gbps per direction or 448 Gbps full duplex. This provides the 7 Tbps total aggregate bandwidth for the 16 half-height server bays (448 Gbps per bay x 16 bays = 7,168 Gbps = 7 Tbps).

The aggregate bandwidth is measured between the device bays and interconnect bays. It does not include additional traffic capacity between interconnect modules (interconnect crosslinks), device bays (device bay crosslinks), and server blades and OA management modules. Please see the Fabric connectivity and port mapping section within this document for information on these midplane crosslinks and management signals.

For the c7000 Platinum Enclosure, one of the areas of improved design focus was on high-speed signal integrity. Obtaining the 14 Gbps level of bandwidth per lane required special attention to high-speed signal integrity.

- Using general proven practices for signal integrity to reduce end-to-end signal losses across the signal midplane
- Moving the power into an entirely separate backplane to independently enhance 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

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) and memory-semantic protocols (PCI Express), using the same signal traces. Compared to the previous enclosure midplane, the midplane in the c7000 Platinum Enclosure includes the following server blade to interconnect bandwidth improvements:

- Hundred percent networking bandwidth increase with 20 Gb Ethernet (10 Gb to 20 Gb Ethernet support)
- Hundred percent storage bandwidth with 16 Gb Fibre Channel (8 Gb to 16 Gb Fibre Channel support)
- Forty percent increase in I/O bandwidth with 4x FDR InfiniBand (40 Gb 4x QDR to 56 Gb 4x FDR InfiniBand support)



Figure 9 illustrates how you can logically overlay the physical lanes onto sets of four traces. Interfaces such as Gigabit Ethernet and Fibre Channel need one lane (1X), or a single set of four traces. Higher bandwidth interfaces, such as InfiniBand, use up to four lanes (4X).




Figure 9. Logical overlay of physical lanes onto sets of four traces

Note

Network-semantic interconnect protocols use network addresses in the packet headers to exchange data between two nodes, such as media access control (MAC) addresses and IP addresses for Ethernet, worldwide port name for Fibre Channel, or globally unique identifier (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 that connect to the interconnect bays. The rear of the enclosure includes eight interconnect bays that can accommodate eight single or four redundant interconnect module. The interconnect modules plug directly into these interconnect bays. Each c-Class Enclosure requires at least two interconnect switches or two pass-thru modules, side-by-side, for a fully redundant configuration.

Proven practices

To enable BladeSystem high-speed connectivity among the blades and interconnect modules, HPE leveraged its many years of experience in designing [HPE Superdome](#) computers. Specifically, HPE paid special attention to:

- Controlling the differential signal impedance along each end-to-end signal trace across PCBs and through the connector stages
- Using a ground plane to isolate, receive, and transmit signal pins as shown in Figure 10
- Keeping signal traces short to reduce losses
- Routing signals in groups to reduce signal skew




Figure 10. Ground plane to isolate, receive, and transmit signal pins



Channel topology and equalization settings

Even when using proven 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, the shape of the transmitter’s signal waveform must be enhanced by selecting the signal equalization settings. 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 embedded network adapters or other embedded I/O devices. As shown in Figure 11, the channel topology for device 1 on server blade 1 (a-b-c) is different from the channel topology for device 1 on server blade 4 (a-d-e). A link configuration mechanism in the OA (assisted by iLO on each server blade) identifies the channel topology for each device and configures the proper equalization settings for that device.




Figure 11. Server blade channel topologies

Power backplane scalability and reliability

The PCB power backplane is separate from the NonStop signal midplane. This design improves the signal midplane by reducing its PCB thickness, reducing electrical noise from the power components that would affect high-speed signals, and improving the thermal characteristics. These design choices result in reduced cost, improved performance, and improved reliability. To reduce losses while distributing power, the power backplane has solid copper plates and integrated power delivery pins (Figure 12). Solid copper plates reduce voltage drops, produce high current density, and provide high reliability.




Figure 12. Power backplane with solid copper plates and integrated power delivery pins



Enclosure connectivity and interconnect fabrics

Whether you need basic connectivity for a remote office or a high-bandwidth, low-latency link for a high-performance computing cluster, the HPE BladeSystem c-Class offers an enclosure connectivity option to meet your needs. The c7000 Enclosure supports the industry's broadest portfolio of networking solutions spanning Ethernet, Fibre Channel, InfiniBand, SAS, iSCSI, and FCoE. The c7000 Enclosure includes up to four redundant interconnect fabrics supported simultaneously within the enclosure. BladeSystem c-Class interconnect technology features HPE Virtual Connect providing wire-once connectivity to simplify and converge your server edge connections.

HPE Virtual Connect

HPE Virtual Connect lets you virtualize your network, storage, and power infrastructure so that resources can be dynamically allocated based on server workloads. The technology removes traditional infrastructure constraints by abstracting servers from their uplinks, simplifying management between networks and servers. HPE Virtual Connect virtualizes server I/O connections. With Virtual Connect wire-once technology you can add, move, or change servers in minutes. The HPE Virtual Connect family includes a wide variety of Flex-10/Flex-20, FlexFabric, and Fibre Channel modules.

Virtual Connect implements server-edge virtualization so you can add, replace, or move server blades without having to make changes to LAN and SAN connections. HPE recommends 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 Virtual Connect FlexFabric module. The Virtual Connect Manager software runs on these modules. You can also use HPE 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 MAC addresses for the network adapters (NICs) and default Worldwide Names (WWNs) for the 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 the LAN and SAN connections once during initial deployment. If you later add, deploy, or change a server, you do not need to change a LAN and SAN connection because Virtual Connect keeps the profile for that server bay constant.

HPE Virtual Connect includes Flex-10/20 technology, which allows you to fully utilize 10/20GbE connection bandwidth. Using Flex-10/20, you can allocate the bandwidth of a single 10/20 Gb pipeline into multiple network connections called FlexNICs. With Virtual Connect FlexFabric modules and FlexFabric Adapters, each 10/20 Gb port can be allocated across eight² network connections, or FlexNICs. Or one of those eight³ 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 10/20 Gb connection. You can set speed from 10 Mb per second to 10/20 Gb per second in 100 Mb increments.

Advantages to partitioning a 10/20GbE pipeline include:

- 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 HPE Virtual Connect Flex-10/10D and FlexFabric-20/40 F8 modules to manage the 10/20GbE (Flex-10/20) server connections to the data center network. The Flex-10/20 and FlexFabric embedded adapters and mezzanine cards are converged network adapters, each with two 10 Gb or 20 Gb ports. Each port can be configured from one up to four individual FlexNICs. The server ROM and the operating system or hypervisor recognize each FlexNIC and FlexHBA as an individual NIC or HBA, respectively.

Find up-to-date information about the c-Class Virtual Connect options at: hpe.com/info/virtualconnect.

^{2,3} Up to eight FlexNICs with HPE Virtual Connect FlexFabric-20/40 F8 with VC v4.50 firmware, HPE FlexFabric 556FLB or HPE FlexFabric 650FLB or 650M and managed with HPE OneView 3.0. Up to four FlexNICs otherwise.



Interconnect modules

The BladeSystem c7000 Enclosure offers a variety of interconnect module options including:

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

Switches offer a traditional approach to administering a network. The primary value in switches is cable consolidation through high-speed uplinks. Each interconnect module also provides external connectors that vary based on its particular design. Interconnect modules in the BladeSystem c7000 Enclosure are available in two widths:

- Single-wide modules provide 16 internal ports with each connecting to a separate device bay in the front of the enclosure.
- Double-wide modules provide 16 double-wide internal ports enabling connectivity to InfiniBand and other four-lane, high-speed interconnects.

In the BladeSystem 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 InfiniBand, two modules are installed in bays 5/6 and 7/8 to provide redundant high-bandwidth connectivity.

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. HPE Ethernet and Fibre Channel pass-thru modules provide 16-port, transparent, one-to-one port connectivity between the server and an external switch.

Find up-to-date information about the c-Class interconnect options at: hpe.com/servers/blades/interconnects.

Note

The c-Class Ethernet pass-thru module is only compatible with fixed-speed Gigabit Ethernet. Because the 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. This is restricted to the Gigabit Ethernet pass-thru modules only. The 10 Gb Ethernet pass-thru modules is compatible with 1 Gb or 10 Gb connections to the server. The Fibre Channel pass-thru modules ports have the auto-negotiate protocol.

Server blades

The BladeSystem c-Class Enclosure is compatible with many different server blade options including HPE ProLiant server blades Intel® Xeon® x86 processors, HPE ProLiant workstation blades, and HPE Integrity server blades. These server blade options are available in multiple form factors, performance levels, and scalability needs supported simultaneously within the enclosure, allowing nearly endless configuration possibilities. HPE server blades for the BladeSystem Enclosures are designed according to the c-Class standard form factors referred to as half-height and full-height. The flexible design of the enclosure allows the combination of full-height server blades, half-height server blades, or a combination of the two. Additionally, storage, tape, and other option blades may be mixed with full-height and half-height server blades extending even further the very wide variety of possible configurations. Please see the [Appendix](#) at the end of this document for form factor population guidelines. Currently available ProLiant server blades including the BL460c Gen10, BL460c Gen9, the WS460c Gen9, and the BL660c Gen9.

Server blade products

HPE ProLiant BL460c Gen10: As the next step up in scaling business performance with secure workload delivery, the new server blade delivers innovation and new levels of performance. The half-height blade features the new Intel Xeon Scalable processor family with up to 26 cores per processor, up to 2 TB DDR4-2666 MT/s Memory, 16 GB Persistent Memory NVDIMM module support, NVMe/SFF/uFF drives, and a new iLO USB port, among other key innovations. Gen10 also enables secure operations and firmware root of trust, reducing the risk of firmware hijacking, ensuring a more secure foundation for the blades infrastructure.

HPE ProLiant BL460c Gen9: The half-height blade features the new Intel Xeon E5-2600 v4 processors with up to 2 TB of DDR4 memory, NVMe drive and M.2 support, among other key features.



HPE ProLiant WS460c Gen9: This half-height graphics server blade enables high-performance graphics for applications ranging from Virtual Desktop Infrastructures (VDIs) to workstation graphics for high performance compute and other demanding sectors. Supporting a wide range of sustaining workstation class graphics adapters from NVIDIA® and AMD, the WS460c accelerates graphics performance coupled with the new Intel Xeon E5-2600 v4 processors with up to 1 TB of DDR4 memory support, NVMe and M.2 drive options to boost application performance.

HPE ProLiant BL660c Gen9: The full-height server blade delivers up to four-socket processor performance for in-memory database and demanding virtualization applications. It is powered by up to four Intel Xeon E5-4600 v4 processors with up to 4 TB of DDR4 memory support, NVMe and M.2 drive options.

For connectivity, every server blade ships with built-in Ethernet connections. With G7 and prior generations, ProLiant server blades included Ethernet ports embedded on the system board. Starting with ProLiant Gen8 through Gen10 server blades, these Ethernet ports were moved onto a user-selectable card known as a FlexibleLOM adapter. Half-height server blades support up to two Ethernet built-in ports and full-height, single-width server blades support up to four. Server blades larger than full-height, single-width are available that have even more built-in Ethernet ports. Optional mezzanine cards can be installed for additional interconnect fabric connections such as Ethernet, InfiniBand, Fibre Channel, and SAS.

There is a very extensive array of possibilities for configuring server blade I/O connectivity. Here are just two options:

- Half-height server blades, including the ProLiant BL460c Gen10, ProLiant BL460c, and ProLiant WS460c Gen9 blades, typically have two built-in 10 Gb or 20 Gb Ethernet ports and two PCIe mezzanine option connectors. A half-height server can be configured with two built-in Ethernet ports, one dual-port Ethernet mezzanine card, and one quad-port Ethernet mezzanine card for a total of eight independent Ethernet ports.
- Full-height, single-width server blades, including the ProLiant BL660c Gen9, typically have four built-in 10 Gb or 20 Gb Ethernet ports and three PCIe mezzanine option connectors. A full-height, single-width server can be configured with four built-in Ethernet ports, one dual-port Ethernet mezzanine card, and two quad-port Ethernet mezzanine cards to form a total of 14 independent Ethernet ports.

Find up-to-date information about the c-Class server blades at: hpe.com/info/bladessystem.

Storage and tape options

Storage and tape options may be attached externally to the enclosure or installed inside the enclosure. For external options, the c-Class Enclosure interoperates and connects with other HPE infrastructure pieces. This includes external storage components such as DAS, NAS, and SAN solutions. External tape backup solutions are also available such as HPE StoreOnce Backup and HPE StoreEver Tape Automation.

Storage options inside the enclosure provide an alternative to local disk drives or external SAN and tape backup connectivity. They allow you to quickly add storage expansion and data protection without adding a single cable. The BladeSystem c7000 Enclosure is compatible with several types of storage solutions. Tape and PCI Expansion option blades are also available for c-Class. Storage blades, tape blades, and the PCI Expansion Blade (collectively referred to as expansion blades) expand the capability of the server blades and increase the overall HPE BladeSystem configuration flexibility.

The expansion blade is paired with an adjacent server blade. The physical connection between the expansion blade and its adjacent server blade is a dedicated PCIe x4 connection across the NonStop midplane that connects the adjacent bays. The design enables the expansion capability, such as high-performance storage access, without any additional cables. For mechanical compatibility, expansion blades use the same half-height form factor as server blades. A half-height server blade can be paired with one expansion blade. Some full-height server blades support up to two expansion blades that may be the same or different options. Please see the [Appendix](#) at the end of this document for more information.

Storage blades

Storage blade options for c-Class server blades are available per the following:

- **HPE D2220sb Storage Blade**—The HPE D2220sb Storage Blade is the newest direct attached storage blade designed for c-Class server blades starting with Gen8. It supports up to 12 SFF SAS hard disk drives, or SAS/SATA SSDs or SATA Midline hard drives. It includes an onboard Smart Array P420i controller with 2 GB flash-backed write cache for increased performance and data protection. Up to eight D2220sb storage devices can be supported in a single c7000 Enclosure for up to 115.2 TB of storage.
- **HPE StoreEasy 3840 Gateway Storage Blade**—The HPE StoreEasy 3840 Gateway Storage Blade is a converged storage platform from HPE that supports both file and application workloads over ubiquitous Ethernet. HPE StoreEasy delivers file and application storage made easy. It gives you the capability to add file services from a c-Class Enclosure to your array or Fibre Channel SAN.



Shared storage

With the HPE P4000 Virtual Storage Appliance (VSA) Software, you can deploy the D2220sb as an iSCSI SAN for use by the servers in the enclosure and any server on the network. HPE StoreVirtual VSA software is installed in a virtual machine on a VMware ESX® host server adjacent to the storage blades. HPE StoreVirtual 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 storage blades and HPE StoreVirtual VSA software licenses.

External SAS connectivity with direct connect SAS storage for HPE BladeSystem

With direct connect SAS storage for HPE BladeSystem, you can build local server storage with zoned storage. Alternatively, you can enable low-cost shared storage within the rack with high performance 6 Gbps SAS architecture. Keep in mind the following considerations when using external SAS storage:

- Each HPE server blade requires an HPE Smart Array SAS controller installed in a mezzanine slot for access to external storage.
- A single or redundant HPE 6 Gb SAS switch is installed in the interconnect bays of the enclosure.
- Depending on application requirements, the switches in the enclosure may need to be connected by a SAS cable to the external storage.

Mezzanine cards

HPE BladeSystem c-Class offers a variety of mezzanine card options to provide connectivity to networks and storage. HPE ProLiant c-Class server blades use two types of mezzanine cards (Type I and Type II) to connect to the various interconnect fabrics such as Fibre Channel, Ethernet, SAS, 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 ProLiant c-Class server blades in all mezzanine connectors. You can also use Type II mezzanine cards in all mezzanine connectors except mezzanine connector 1 in half-height c-Class server blades in Table 1.

Table 1. Demonstrating use of Type I and Type II mezzanine cards in c-Class server blades

	c-Class server blades	Type I mezzanine card	Type II mezzanine card
Mezzanine connector 1	Full-height	Yes	Yes
	Half-height	Yes	No
Mezzanine connector 2	Full-height	Yes	Yes
	Half-height	Yes	Yes
Mezzanine connector 3	Full-height	Yes	Yes

Find up-to-date information about the c-Class mezzanine card options at: hpe.com/servers/proliantnics.

Fabric connectivity and port mapping

BladeSystem c7000 Enclosure device bay to interconnect bay links

The I/O signals from the on-board network adapters and mezzanine cards in the server blades in the front of the enclosure are sent to the interconnect modules in the rear of the enclosure via the NonStop passive midplane. Since the NonStop midplane is passive with no active components, these connections between the device bays and the interconnect bays are hard-wired. As a result, the server blade on-board network adapter ports and mezzanine cards are matched to the appropriate type of interconnect module. For example, an Ethernet switch or Ethernet pass-thru module is used in an interconnect bay that connects to the on-board Ethernet network adapter ports. Likewise, a Fibre Channel switch would be installed in an enclosure interconnect bay that corresponds to a Fibre Channel mezzanine card in a server blade mezzanine slot.



Interconnect bays 1 and 2 are reserved for Ethernet switches or Ethernet pass-thru modules that connect to the server blade on-board network adapter (NIC) ports. The on-board NIC ports may be either embedded Ethernet ports (ProLiant G7 servers and earlier) or ports supplied on a FlexibleLOM adapter (starting with ProLiant Gen8 servers). Mezzanine cards are added for additional I/O connectivity and connect to interconnect bays 3 through 8. Redundant interconnect modules are placed adjacent to one another in interconnect bays 3/4, 5/6, and 7/8. For port mapping purposes, it does not matter which device bay houses a server blade. The built-in NIC ports and mezzanine card ports connect to the same interconnect bays. Figure 13 shows the c7000 Enclosure port mapping of half-height server blades to interconnect bays.

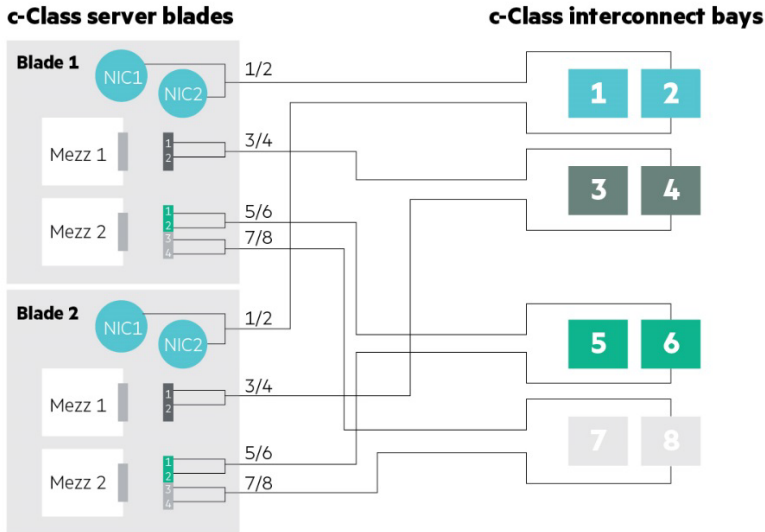


Figure 13. c7000 Enclosure port mapping of half-height server blades to interconnect bays

Figure 14 shows the c7000 Enclosure port mapping of full-height, single-wide server blades to interconnect bays.

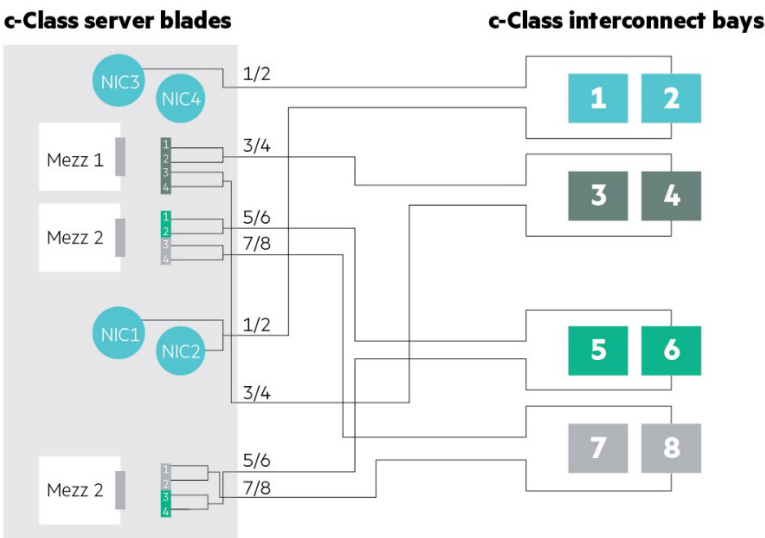


Figure 14. c7000 Enclosure port mapping of full-height, single-wide server blades to interconnect bays

As shown in Figures 13 and 14, 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 HPE OneView software include tools that simplify the process of mapping mezzanine ports to interconnect module ports. For port-mapping details, see the HPE BladeSystem Onboard Administrator User Guide at: h20566.www2.hp.com/hpsc/doc/public/display?sp4ts.oid=3188465&docLocale=en_US&docId=emr_na-c00705292.



BladeSystem c7000 Enclosure bay-to-bay crosslinks

Four lanes of four-trace SerDes signals between adjacent bays in the BladeSystem c7000 Enclosure midplane permit horizontal bay-to-bay communications.

Device bay crosslinks

Device bay crosslinks provide a connection for adjacent device bay pairs as indicated by the arrows in Figure 15.




Figure 15. c7000 front view with device bay crosslinks

Device bays crosslinks allow connection of optional expansion blades (storage blade, tape blade, or PCI Expansion Blade) to a server blade through the midplane eliminating the need for cables. The OA disables these side-to-side device bay crosslinks when they are not needed, such as when two server blades reside in adjacent device bays.

For half-height server blades, the device bay crosslinks provide a communication pathway between the four-lane PCIe connection and an expansion blade. The two blades occupy a pair of adjacent device bays that may be in a left-right or right-left arrangement. Please see the [Appendix](#) at the end of this document for more information.

For full-height, single-wide server blades, the device bay crosslinks provide a communication pathway for a PCIe connection to an expansion blade in the lower adjacent bay. Some server blades also support a second expansion blade in the upper adjacent bay. Specific to the BL620c G7 Server Blade, the device bay crosslinks are not active; a PCIe pass-thru mezzanine card is installed in a ProLiant BL620c G7 Server Blade mezzanine connector for each expansion blade attached to the server. Please see the [Appendix](#) at the end of this document for more information.

Some examples of expansion blade options using these dedicated device bay crosslinks to a server blade include:

- HPE Storage D2220sb Storage Blade that consist of a RAID controller and additional drives
- BladeSystem PCI Expansion Blade that holds two off-the-shelf PCI-X or PCIe cards

Interconnect bay crosslinks

Interconnect bay crosslinks are wired between adjacent horizontal interconnect bay pairs as indicated by the arrows in Figure 16. You can enable interconnect crosslink signals to provide module-to-module connections (such as Ethernet crosslink ports between matching switches) or Virtual Connect modules can use them as stacking links. The OA disables the interconnect bay crosslinks when they are not needed when, for example, two different modules reside in adjacent horizontal interconnect bays.






Figure 16. Interconnect crosslinks wired between adjacent horizontal interconnect bays

BladeSystem c7000 Enclosure management links

The signal midplane also includes the management signals from each device bay and interconnect bay to the OA modules. These management signals are isolated from the high-speed device bay to interconnect bay signals. The OA is the terminating point for the interconnect bays. An interconnect module cannot use the connection to the OA to communicate with another interconnect module; this is accomplished through the interconnect bay crosslinks described earlier.

HPE Thermal Logic technologies

The BladeSystem c7000 Enclosure uses several HPE 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 reduce power usage while maintaining adequate cooling for the devices to enable high availability.

HPE Thermal Logic c7000 Enclosure technologies include the following elements and capabilities:

- HPE Active Cool 200 fans
- HPE Parallel Redundant Scalable Enclosure Cooling (PARSEC) architecture
- Power and thermal monitoring
- 3D sea of sensors

HPE Active Cool 200 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 back pressure, a single server often requires multiple fans to provide 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 largest 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, allowing for adequate airflow to the servers without leakage, over provisioning, or bypass is a challenge.

To overcome these issues in the BladeSystem c7000 Enclosure, HPE 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. HPE has 20 patents for the Active Cool fan technology and implementation.



HPE 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 as shown in Figure 17. The fan includes a bell mouth inlet with a specially designed impeller and a stator section that also provides cooling fans for the motor and acoustic treatments at the rear of the fan. The unique shape of the fan generates high-volume, high-pressure airflow at even the slowest fan speeds with low noise levels and reduced power consumption.




Figure 17. Active Cool 200 fans using ducted fan technology

The OA controls Active Cool 200 fans, ramping cooling capacity either up or down based on system needs. Along with optimizing the airflow, the control algorithm improves the acoustic levels and power consumption. As a result, the BladeSystem c7000 Enclosure can accommodate full-featured servers that are 60% more dense than traditional rack-mount servers. And, the Active Cool 200 fans consume just 50% of the power typically required and use 30% less airflow.

HPE PARSEC architecture

The BladeSystem c7000 Enclosure uses HPE Parallel Redundant Scalable Enterprise Cooling architecture. In this context, “parallel” means that fresh, cool air flows over the blades in the front of the enclosure and all the interconnect modules in the back of the enclosure.

The c7000 Enclosure is divided into four cooling zones with fans in each zone, as shown in Figure 18. The Active Cool 200 fans supply cooling for their own zone and redundant cooling for the rest of the enclosure. Zone cooling reduces 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 reduces fan noise.

Each zone can contain four half-height server blades. Full-height, single-wide server blades couple the two vertical zones into a single zone. A full-height, single-wide server blade in bay 1 results in zones 1 and 3 operating as a single zone.




Figure 18. c7000 Enclosure divided into four cooling zones with fans in each zone

To improve thermal design, the enclosure uses a relatively airtight center air plenum or air chamber. In the BladeSystem c7000 Enclosure, the 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 so that cooling air does not flow through the inoperative fan as shown in Figure 19.




Figure 19. Fan louvers in action

The enclosure and the components within it improve the cooling capacity through unique mechanical designs. With managed airflow through the enclosure:

- Every device receives cool air.
- No device stands in the path of hot exhaust air from another device.
- Air is used only where necessary 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 flows into the interconnects and the central plenum. The air then passes out the rear of the system as shown in Figure 20. Each of the modular power supplies has its own internal fan for improved cooling.




Figure 20. Air ventilation paths

Operating c-Class server blades requires installing at least four fans at the rear of the BladeSystem c7000 Enclosure. Up to 10 fans can be installed allowing cooling capacity to scale as needs change. Additionally, using more fans allows the fans to spin slower to move the same volume of air so each fan uses less power. HPE recommends using at least eight fans as they are usually more power-efficient than four fans. As the airflow rate increases, 10 fans become even more efficient as shown in Figure 21. Slower spinning fans also create less noise. See the [Appendix](#) for additional fan information.






Figure 21. Increasing airflow rates make 10 fans more efficient

Power and thermal monitoring

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

The HPE 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 HPE OneView. 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.

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

3D sea of sensors

The 3D sea of sensors technology in HPE ProLiant servers uses an array of sensors to map the server’s temperature profile more accurately. 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. Instead of using a fan curve, the iLO processor monitors multiple temperature sensors throughout the server and then employs a sophisticated proportional-integral-derivative (PID) control feedback algorithm. This adjusts fan speed in the enclosure cooling zones to reduce power consumption and increase cooling efficiency. Sensors have also been added to many HPE mezzanine option cards to obtain a three-dimensional temperature profile in the server. This additional data enables even more precise and efficient cooling of the servers.

Enclosure power subsystem and capabilities

The flexible design of the BladeSystem c7000 Enclosure power subsystem offers a choice of five types of power input options for the ideal fit with your data center power infrastructure as shown in Figure 22.

- Single-phase intelligent power module that accepts HPE Intelligent Power Distribution Unit (iPDU) IEC C19 power cables designed for single-phase high line 200–240 VAC Intelligent Power applications
- Single-phase power module that accepts IEC C19 power cables designed for single-phase high line 200–240 VAC applications
- Three-phase power module with a pair of U.S./Japan power cords with NEMA L15-30P connectors designed for three-phase high line 200–208 VAC North America and Japan applications
- Three-phase power module with a pair of international power cords with IEC 309, 5-pin, 16A connectors designed for three-phase high line 220/380 to 240/415 VAC applications
- DC power module using a screw down terminal lug (45DG 4AWG 1/4 2H) designed for –36 VDC to –72 VDC applications
- High-voltage power module that accepts power cables with APP Saf-D-Grid connectors designed for high voltage 380 VDC and single-phase high voltage 277 VAC applications





Figure 22. Five types of power input options

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, several power-related components are no longer needed and the power efficiency on the server blades and in the infrastructure is improved.

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 ordering the enclosure.

Power management

The OA manages power allocation rules of various components and can prevent a device from powering up if overall power capacity of installed power supplies output, redundancy settings, and power cap setting would be exceeded. Power is allocated by the OA based on installed device power requests. Installed devices include server blades, interconnect modules, fans, and OA modules. To identify a device in the enclosure, the OA reads its EEPROM containing the factory information (product name, part number, and serial number). The EEPROM operates on auxiliary power so the OA can identify the component before granting power requests. The OA does not power-up a server blade until it verifies the configuration. If there is a configuration issue, the Insight Display shows the issue and possible remedies.

Status and information for each power supply, as well as the power enclosure itself is available including power redundancy status, total available power, and total power consumed. An Enclosure Power Meter screen displays a variety of detailed and enclosure-level information including peak power use, average power use, and allocated power. Available in a graphical or table format, this data enables fast and easy interpretation of the power situation for the enclosure and for identifying trends in power consumption. For additional information, see the [HPE BladeSystem Onboard Administrator User Guide](#).

Power supplies and configuration

The c7000 Enclosure can contain up to six hot-plug, self-cooled power supplies with up to 15370W output capacity per enclosure. The highest amount of output capacity is determined by using six 2650W power supplies (15900W total) and adjusting for the power sharing tolerance. Power supplies are configured with an appropriate power input module. For example, the -48 VDC power supply is used with the -48 VDC power input module. Power supplies must be of the same type within an enclosure. The OA manages power allocation rules of various components. It can prevent a device from powering up if the overall power capacity of installed power supply output, redundancy settings, or power cap setting, is exceeded.



The available power supply and power input module options include:

- HPE 2650W and 2400W Platinum (80 PLUS certified) high line 200–240 VAC power supplies supported with the high line single-phase power modules
- HPE 2650W and 2400W Platinum (80 PLUS certified) high line 200–240 VAC power supplies supported with the high line three-phase power modules
- HPE 2400W Gold (80 PLUS certified) high line 200–240 VAC power supply supported with the high line single-phase power modules
- HPE 2400W Gold (80 PLUS certified) high line 200–240 VAC power supply supported with the high line three-phase power modules
- HPE 2250W –48 VDC power supply supported with the –48 VDC power module
- HPE 2650W universal high voltage 277 VAC/380 VDC power supply supported with the high voltage power module

Enclosures configured with the single-phase power modules and 2650W Platinum power supplies are recommended for many situations. The Platinum power supplies provide higher efficiency to save energy—94% peak and more than 90% efficient with as little as a 20% load. They also lower standby power for reduced power consumption when servers idle.

Enclosures configured with a three-phase power module are generally useful in a few situations when:

- Only one enclosure is present and more than one enclosure is not needed.
- The configuration is extreme with, for example, 16 server blades, each with two high-power processors, 256 GB to 512 GB of memory, and all mezzanine connectors and interconnect bays filled.
- The enclosure is a drop-in replacement for an HPE p-Class solution where three-phase power already exists.
- Plans call for the enclosure to be maintained for a long period of time and for the internal components to be updated with newer ones that can 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. For example, this can be achieved by using a pair of HPE Power Distribution Units (PDUs) to distribute 2 x 60A (17.3 kVA rated) three-phase power to the rack connected to multiple single-phase enclosures in the rack. If the PDUs are connected to three enclosures, power capping can be set to 5.7 kVA per enclosure (17.1 kVA total). Therefore, the enclosure power draw cannot exceed the capacity of one PDU even if an AC line were to fail.

Pooled power and redundancy

The power in the BladeSystem c7000 Enclosure is part of a single power pool that any server blade within the enclosure can access. This design has no zones allowing for improved flexibility in choosing the level of redundancy. Since power supplies share in delivering pooled power to the entire enclosure, the OA report shows that the power supply outputs track each other closely. This distributes the load evenly, which is particularly important when using three-phase AC power.

Configured using the OA or the Insight Display, the BladeSystem c7000 Enclosure has three redundancy modes:

- AC or DC redundant (N+N redundancy)
- Power supply redundant (N+1 redundancy)
- No redundancy mode

If the power mode of the enclosure is set to N+N redundancy, the OA enables equal numbers of power supplies to be available on the A and B sides as indicated in Figure 23. This logical grouping means that when the enclosure is configured with a three-phase AC power module, the power supplies on one side maintain enclosure power, even if AC input power is lost on the other side. If the enclosure power mode is configured to N+1 or no redundancy, the power supplies can be installed in any bay.



Figure 23. Enclosure set power modes



Connecting to PDUs with N+N redundancy

In an N+N 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 is 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 power line feeds, this configuration provides for a single power line feed that cannot trigger the powering off the enclosure.

Connecting to a UPS with N+1 redundancy

In a configuration with N+1 power supply redundancy connecting to an uninterruptible power supply (UPS), the total power available equals the total power available less one power supply. In this configuration, there can be up to six power supplies and one of them is available to provide redundancy. In the event of a single power supply failure, the redundant power supply takes 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 the installed power supplies. Any power supply or power line failure causes the system to power off if the remaining power supplies are unable to handle the full load.

HPE Power Discovery Services

The c7000 single-phase intelligent power module, HPE Platinum power supplies, and HPE iPDUs are key components of HPE Power Discovery Services (PDS). These components allow the enclosure to communicate information to iPDUs that automatically track enclosure power connections to the specific iPDU outlets to provide redundancy and prevent downtime. Together they communicate through power line communication (PLC) technology embedded into the c7000 single-phase intelligent power module. PLC allows the Platinum power supplies to share data with the c7000 single-phase intelligent power module, iPDUs, and with the HPE OneView software. PDS gives you accurate, automated, error-free power mapping, monitoring, and control on ProLiant G7 and later servers. Find additional information about HPE Power Discovery Services at: hpe.com/us/en/integrated-systems/rack-power-cooling.html.

HPE Enclosure Dynamic Power Capping allows an enclosure-level power cap. The OA maintains the overall enclosure cap by monitoring and managing the power consumption of the server blades. To accomplish this, the OA collects the overall enclosure power use and the total power used by the managed server blades. From this data, the OA constructs a blade power budget representing the amount of power the blades can consume while keeping the overall enclosure power consumption below the enclosure cap. As a final step, OA software actively adjusts the power caps of the individual servers so the total matches the enclosure power budget.

HPE Power Advisor

In the latest generation of HPE BladeSystem, power is both critical and flexible. The HPE 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. The HPE Power Advisor is available at: hpe.com/us/en/integrated-systems/rack-power-cooling.html#HPEPowerAdvisor.

HPE Location Discovery Services

HPE Location Discovery is the industry's first location-based intelligence in servers and racks allowing you to improve workload placement with servers that self-identify and inventory, automating asset tracking and reducing manual errors.

HPE Location Discovery technology works with ProLiant Gen8 servers, HPE 11000 G2 or Intelligent Series racks, and the BladeSystem c7000 Platinum enclosure to send the rack or enclosure device bay identification number and precise U location to the servers. This provides important location information to HPE Insight Control Software along with power and temperature data. This information is transmitted using non-emitting and radio free technology to meet the security requirements of our customers.



Conclusion

The HPE BladeSystem c7000 Enclosure goes beyond blade servers to extend the foundation for a converged infrastructure. It consolidates server, storage, networking, and power management into a single solution that you can manage as a unified environment. It consolidates and simplifies infrastructure, reduces operational cost, and delivers IT services more effectively.

The shared, high-speed, NonStop midplane and pooled-power backplane in the enclosure accommodate new bandwidths and new technologies. The 7 Tbps aggregate bandwidth of the enclosure midplane includes 20 Gb Ethernet, 16 Gb Fibre Channel, and FDR InfiniBand capability for demanding workloads. Intelligent Infrastructure technologies such as 2650W and 2400W 80 PLUS certified platinum hot-plug power supplies, intelligent power module, and Location Discovery Services enhance the converged infrastructure foundation.

HPE OneView combines server, storage, and networking with control of your data center environment into a single, integrated management platform architected to deliver lifecycle management for your complete converged infrastructure. The HPE Onboard Administrator together with the HPE Insight Display enclosure display provides both local and remote administration of HPE BladeSystem c-Class Enclosures plus essential power and cooling control and management.

Appendix: Fan and server population guidelines

For proper operation, fans and blades are installed in the correct bays of the HPE BladeSystem c7000 enclosure. The OA ensures the correct placement of the fans and blades before it allows the system to power on.

Fan bay numbering

The c7000 requires at least four fans and supports up to 10. Ten fans are recommended for improved cooling efficiency and redundancy, and reduced power consumption and noise. Figure 24 indicates the location of the 10 hot-plug fan bays in the back of the enclosure. Table 2 provides an overview of the correct fan bays in which the fans must be installed and number of supported blades and the supported blade locations. For a detailed description of fan population rules and guidelines, please see the c7000 Enclosure QuickSpecs at: hpe.com/servers/bladeSystem/enclosures.




Figure 24. 10 hot-plug fan bays located in the back of the enclosure



Table 2. Fan bay population locations

Number of fans	Fan bays populated (cooling zones)	Device bays supported (device bay zones)
4	4, 5 (Zone 1) and 9, 10 (Zone 3)	Two blade devices located in bays 1, 2, 9, and 10 (Zone 1)
6	3, 4, 5 (Zone 1) and 8, 9, and 10 (Zone 3)	Eight blade devices located in bays 1, 2, 9, 10 (Zone 1) and 3, 4, 11, and 12 (Zone 2)
8	1, 2, 4, 5, 6, 7, 9, and 10 (all four zones)	All device bays, all zones
10	1, 2, 3, 4, 5, 6, 7, 8, 9, and 10 (all four zones)	All device bays, all zones

Server blade numbering

Figure 25 provides the c7000 enclosure blade bay numbering. Half-height server blades are populated in the front of the enclosure from top to bottom and left to right. The first two half-height servers are placed in bays 1 and 9; the second two half-height server blades in bays 2 and 10, and so on until the enclosure is full. Full-height server blades are populated from left to right. For a detailed description of blade population rules and guidelines, please see the c7000 Enclosure QuickSpecs [here](#).




Figure 25. c7000 enclosure blade bay numbering

Enclosure blade zones

The c7000 enclosure is divided into four zones by its vertical support metalwork. Within each zone a removable divider is used to support half-height devices. These are not the same as the cooling zones referred earlier. Each zone contains a removable divider that supports half-height devices. To install a full-height blade in any zone this divider is removed. See Figure 26.




Figure 26. Removable divider for installation of full-height blade in any zone



An “expansion blade” may be installed in the same zone as half-height and full-height, single-wide server blades. Expansion blades consist of storage blades, tape blades, and the PCI Expansion Blade. Up to eight expansion blades are supported per HPE BladeSystem c-Class c7000 Enclosure. An expansion blade is installed directly adjacent to a server blade, also known as a partner server blade. One half-height partner server blade supports one expansion blade. Figure 27 shows the installation of one expansion blade next to a half-height partner server blade. One full-height, single-wide partner server blade supports one or two expansion blades; Figure 28 shows the installation of one expansion blade next to a full-height, single-wide partner server blade. Expansion blades are not supported with full-height, double-wide server blades. For a detailed description of expansion blade rules and guidelines, please see the c7000 Enclosure QuickSpecs [here](#).




Figure 27. Installation of one expansion blade next to half-height server blade




Figure 28. Installation of one expansion blade next to full-height, single-wide partner server blade



Additional information

HPE BladeSystem portfolio
hpe.com/info/bladesystem

HPE BladeSystem c7000 Enclosure QuickSpecs
hpe.com/h20195/v2/default.aspx?cc=us&lc=en&oid=1844065

HPE BladeSystem c-Class documentations
h17007.www1.hpe.com/us/en/enterprise/integrated-systems/info-library/index.aspx?cat=bladesystems#.WJC7L7HnZ

HPE BladeSystem c-Class Enclosure set-up and install guide, site planning guides, user guides, and additional support and technical documentation: hpe.com/info/ossupport

HPE Virtual Connect
hpe.com/info/virtualconnect

HPE OneView management software
hpe.com/info/oneview

HPE CloudSystem
hpe.com/info/cloudsystem


HPE Power Advisor
hpe.com/us/en/integrated-systems/rack-power-cooling.html#hpepoweradvisor

HPE power and rack options
hpe.com/us/en/integrated-systems/rack-power-cooling.html

80 PLUS power supply certification
80plus.org

BladeSystem services
hpe.com/poinnext

Learn more at
hpe.com/info/bladesystem



Sign up for updates

© Copyright 2013–2014, 2017–2018 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise 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. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

AMD is a trademark of Advanced Micro Devices, Inc. Intel Xeon is a trademark of Intel Corporation in the U.S. and other countries. Microsoft is either a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries. Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. VMware and VMware ESX are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. NVIDIA is a trademark and/or registered trademark of NVIDIA Corporation in the U.S. and other countries. All other third-party marks are property of their respective owners.

