Technologies in HP ProLiant G7 c-Class server blades with AMD Opteron™ processors

Technology brief, 3rd edition

Introduction 2

ProLiant server design principles 2

ProLiant c-Class server blade architecture 2

Processor technologies 5

Configuration and management technologies 6
- BladeSystem OA 6
- Integrated Lights-Out for ProLiant server blades 6

Power management technologies 7
- Power meter 7
- HP Power Regulator for ProLiant servers 7
- HP Dynamic Power Capping and HP Power Capping 7

HP Thermal Logic technologies 8

Memory technologies 9
- Unbuffered and registered DIMMs 9
- Memory interleaving 9
  - Memory bank interleaving 9
  - Memory channel interleaving 10
  - Memory node interleaving 10
  - Error correction 10

I/O technologies 10
- PCI Express technology 11
- HP Smart Array controllers and supported drive technology 11
- Optional mezzanine cards 11

Networking technologies 11
- Virtual Connect 12

Data security technology with the Trusted Platform Module 12

Conclusion 13

For more information 14

Introduction

This paper describes the architecture and major technologies implemented in HP ProLiant G7 c-Class Server Blades based on AMD Opteron™ processors. Specifically, the paper applies to the BL465 G7 and BL685c G7 Server Blades updated with the AMD Opteron 6200 series processors.

To gain the most value from this paper, you should be familiar with HP ProLiant server technology and be somewhat acquainted with BladeSystem architecture. You can find more information about the infrastructure, components, and complete specifications of each HP server blade and enclosure at www.hp.com/go/bladesystem/.

ProLiant server design principles

We design industry-standard servers with a balanced architecture to use all subsystems effectively under a broad range of applications and workloads. For example, we distribute memory capacity across processors and I/O devices instead of increasing memory capacity asymmetrically. Many applications have improved performance when the memory is closely coupled to the processor. Asymmetric memory, on the other hand, can be inefficient and contribute to increased power consumption and cost.

We develop balanced architecture through superior engineering of fundamental elements such as mechanical infrastructure, power, cooling, processor, memory, I/O devices, storage, boot, networking, and interconnecting components. A balanced architecture includes the following characteristics:

- Computing capability based on processor core count, cache size per processor, and processor socket count
- Low-latency processor-to-memory bandwidth commensurate with core count
- Memory footprint and capacity that maximize bandwidth and capacity with power efficiency and performance without compromising quality or reliability
- Application-appropriate I/O devices
- Closely-coupled and balanced processor-to-memory and processor-to-I/O ratios
- Mechanical design that ensures optimum levels of cooling, stability, and serviceability through space-efficient, modular partitioning

ProLiant c-Class server blade architecture

HP ProLiant c-Class server blades can slide into either c7000 or c3000 HP BladeSystem c-Class Enclosures to meet the needs of large or small IT environments (Table 1).

Table 1: Characteristics of c-Class BladeSystem chassis

Enclosure	Height in Data Rack Units	Server Blade Capacity	Blade position
C7000 rack enclosure	10U	8 c-Class full-height 16 c-Class half-height server blades	Vertical
C3000 rack enclosure	6U	4 c-Class full-height 8 c-Class half-height server blades	Horizontal

The rack enclosures fit in HP 10000 series racks and can operate with as few as one server blade installed. A major advantage of blade architecture is the ease of adding and changing server blades. We build ProLiant c-Class server blades in standard form-factors, referred to as half-height (1 chassis slot) and full-height (2 chassis slots). Both half-height and full-height server blades fit into any device bay in a BladeSystem c-Class enclosure.

In designing ProLiant G7 c-Class server blades based on AMD Opteron processors, we used enterprise-class technologies to provide maximum flexibility, scalability, performance, and control to meet many types of customer needs:

- Two or four AMD Opteron 6200 Series processors
- HP Thermal Logic technologies
- HP advanced memory technologies
- Multiple PCIe Gen 2 x8 slots for I/O cards
- HP NC551i dual-port FlexFabric 10 Gb converged network adapters (CNAs) that support Ethernet, iSCSI, Flex-10, and Fibre Channel over Ethernet (FCoE)
- Up to two HP hot-plug internal disk drives and the HP Smart Array P410i controller with 1 GB flash-backed write cache (FBWC)
- HP advanced power monitoring and management

Half-height server blades support up to 2 processors, 16 DIMM slots for DDR3 memory, up to 2 hot-plug drives, 1 dual-port FlexFabric 10 Gb CNA, and 2 PCIe Gen 2 x8 mezzanine expansion slots (Figure 1). Optional mezzanine cards support multiple types of I/O fabric connectivity to the interconnect bays in the BladeSystem enclosure.

Figure 1: This HP BL465c G7 server is an example of a half-height HP c-Class BladeSystem server.

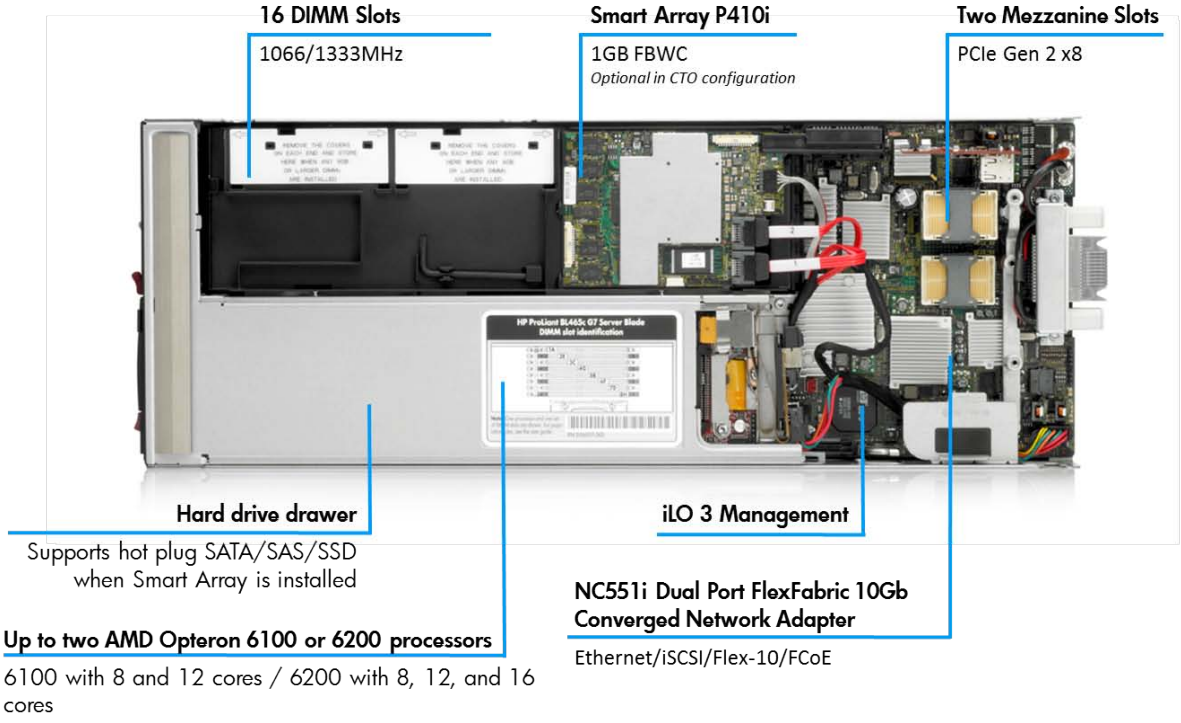

The unique drive drawer design on the BL465 G7 enables hot-plug storage. You can service each drive independently. The drawer slides out to reveal two disk drives, one behind the other. The front drive slides over to provide access to the rear drive (Figure 2).

Figure 2: The HP BL465c G7 hot-plug drive drawer slides out for access to both disk drives.

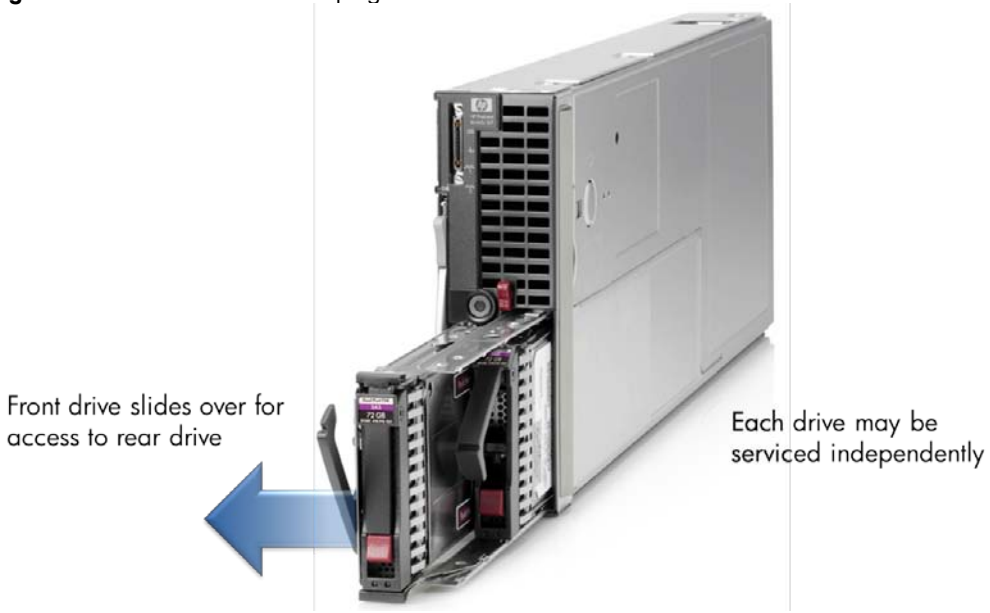

Because of their larger size, full-height server blades support up to 4 processors, up to 32 DIMM slots for DDR3 memory, 2 hot-plug drives, 2 dual-port FlexFabric 10 Gb CNAs (4 ports total), and 3 PCIe x8 Gen 2 mezzanine slots for optional I/O cards. Full-height server blades also provide twice as much I/O bandwidth to the interconnect modules as half-height server blades (Figure 3).

Figure 3: A full-height HP BL685c G7 server supports more processors, memory, and I/O than the half-height server blade.

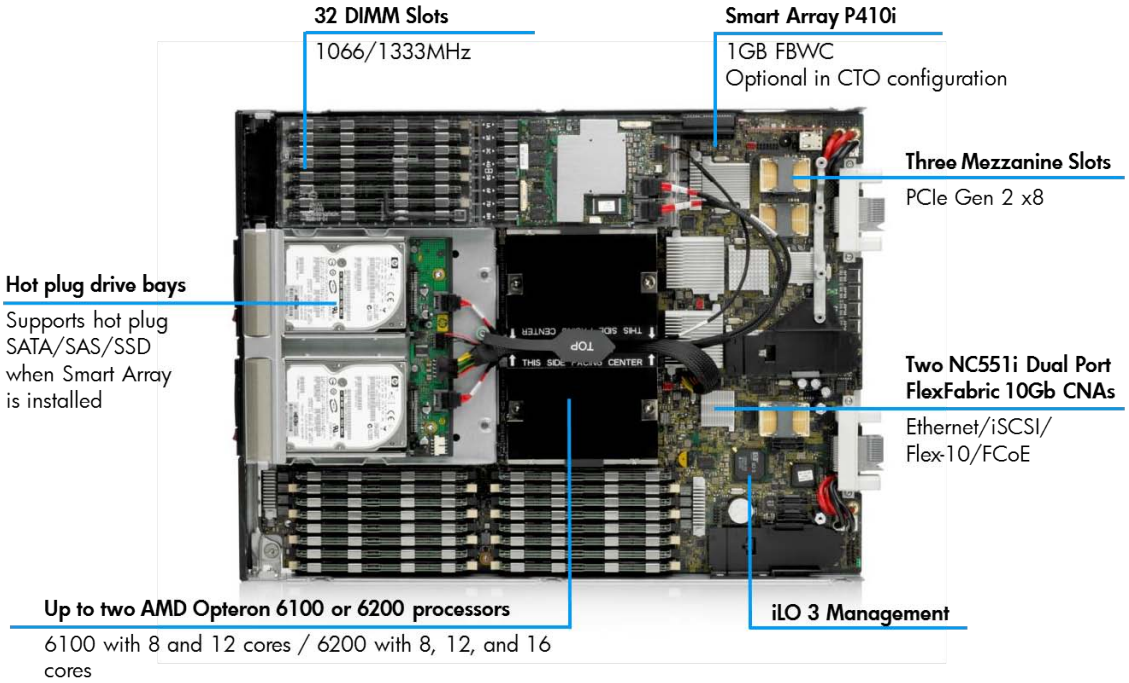

Processor technologies

ProLiant G7 server blades with AMD processors contain 8-core, 12-core, or 16-core AMD Opteron 6200 Series processors. All ProLiant G7 server blades with AMD Opteron processors use AMD SR5690/SP5100 chipsets to manage data throughput.

AMD Opteron processors use Direct Connect architecture: An integrated memory controller replaces the front side bus. The architecture includes dedicated memory banks for each processor and point-to-point HyperTransport™ 3.0 technology (HT3) links between the processors, memory, and I/O chipsets. The integrated memory controller supports only DDR3 memory. HT3 operates at a maximum interconnect rate of 6.4 gigatransfers per second (GT/s). HT Assist™ reduces coherence traffic overload on the HT links, resulting in faster queries in both cache and compute-intensive applications.

In ProLiant G7 server blades with the standard feature set (see the product QuickSpecs for standard features), each 16-core AMD processor operates at speeds of up to 2.3 GHz, has access to 512 KB of L2 cache memory per core, and shares a total of 16 MB of L3 cache. The 12-core processor operates at speeds of up to 2.6 GHz, and the 8-core at speeds up to 3.0 GHz. AMD Dual Dynamic Power Management™ powers the processor core and the voltage planes of the memory controller independently for increased performance and improved power management.

The top processor speeds cited here occur under full core-count operation. AMD Core Select lets you choose the number of cores per CPU that are visible to software. You can access AMD Core Select through the ROM Based Setup Utility (RBSU). Compliant operating systems and applications recognize a reduced core count, so a lower core count may decrease your software licensing costs. This option also reduces processor power use and improves overall performance for applications that benefit from higher performing cores rather than from additional processing cores. In multi-processor configurations, AMD Core Select enables the same number of cores for each CPU. AMD Core Select function can also depend on support by user application software. Check with software vendors to determine application support.

The latest AMD Opteron processor packages use a processor socket technology called Land Grid Array (LGA). AMD calls this a G34 socket. The processor package designs have no pins. Instead, gold-plated pads touch the processor socket pins on the motherboard. Because pin damage could require replacing the motherboard, we designed a tool to simplify processor installation and reduce the possibility of damaging the socket pins (Figure 4). We recommend that you use HP Services for field processor upgrades.

Figure 4: The HP processor installation tool protects processor pins from damage and simplifies processor installation.

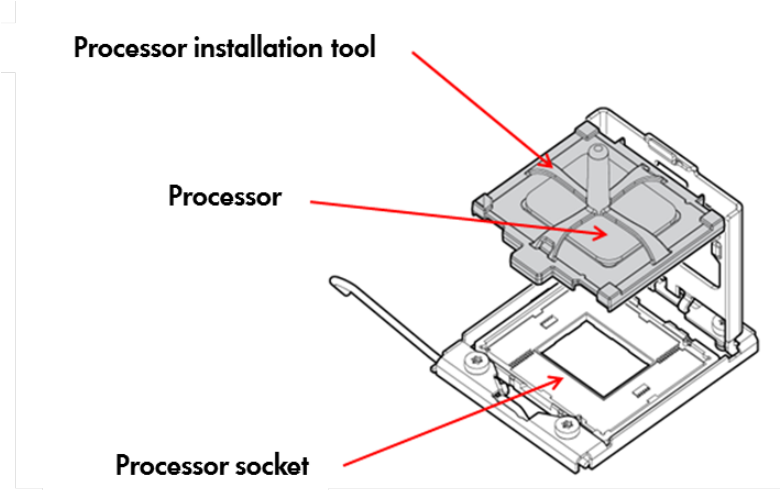

Configuration and management technologies

HP BladeSystem G7 c-Class servers provide an intelligent infrastructure through extensive integration of three management components:

- BladeSystem Onboard Administrator (OA) module located in the enclosure
- Integrated Lights-Out 3 (iLO 3) processor on the server blade
- Interconnect module management firmware such as the HP Virtual Connect Manager

Integrating these management capabilities provides powerful hardware management for remote administration, local diagnostics, and troubleshooting.

For detailed information about c-Class management technologies and capabilities, refer to the HP paper titled "Managing HP BladeSystem c-Class systems" at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00814176/c00814176.pdf.

BladeSystem OA

The heart of c-Class enclosure management is the BladeSystem OA. It works with the iLO 3 management processor on each G7 server blade to form the core of the management architecture for the HP BladeSystem c-Class Enclosure. The OA performs four management functions for the entire enclosure:

- Detects component insertion and removal
- Identifies components and required connectivity
- Manages power and cooling
- Controls components

To identify a component, the OA reads a FRU EEPROM that contains factory information about the component such as product name, part number, and serial number. The OA accesses server blade FRU EEPROMs through iLO 3. The 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 iLO 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 iLO firmware versions, server name, network adapter and option card port IDs, and port mapping.

The OA provides easy-to-understand port mapping information for each of the server blades and interconnect modules in the enclosure. To simplify installation of mezzanine cards and interconnect modules, the OA uses an electronic keying process to detect mismatches.

Integrated Lights-Out for ProLiant server blades

HP BladeSystem G7 c-Class Server Blades include the iLO 3 management processor to configure, update, and operate individual server blades remotely. The c-Class enclosure includes an Ethernet management network to aggregate all iLO management communications across the enclosure. The management network connects iLO processors to the OA through the OA tray. The OA provides direct access to each iLO through the enclosure management network and uses this network to manage pooled enclosure power and cooling.

You can take advantage of the full functionality of iLO 3 by purchasing an iLO Advanced license or by purchasing Insight Control. The iLO Advanced license includes several enhancements:

- Remote console performance equal to that of KVM and software-based remote management solutions, and over 8 times the performance of iLO 2

- Streamlined Web 2.0 interface with JavaScript Object Notation (JSON) technology to create a more dynamic interface experience. Content updates do not require a web page refresh.
- Streamlined user experience with remote console direct access. It provides immediate access to the iLO Advanced remote console without having to navigate through the web interface.
- Enhanced standards support with AES encryption for customers with elevated security requirements
- IPMI LAN/DCMI support for performing simple monitoring and control operations

Power management technologies

To accurately monitor and provision server power use, ProLiant G7 server blades include three HP power management tools, enabled through iLO 3:

- HP power meter
- HP Power Regulator for ProLiant
- HP Dynamic Power Capping and HP Power Capping

Power meter

An HP integrated power meter in ProLiant c-Class server blades analyzes server power use. The BladeSystem OA can access the power meter through iLO or through external power management software such as HP Insight Power Manager (IPM). The BladeSystem OA provides instant and time-averaged views of the power consumption of individual servers or of all servers within the c-Class BladeSystem enclosure. IPM also consolidates power data for multiple servers to a central location. IT departments can use this information to charge business units or third parties for the actual energy costs associated with workload processing.

HP Power Regulator for ProLiant servers

HP Power Regulator for ProLiant servers provides firmware-based speed stepping for recent AMD processors. Power Regulator improves server energy efficiency by giving processors full power when they need it and reducing power when they don't. This power management feature allows ProLiant servers with policy-based power management to control processor power states. IT administrators can configure Power Regulator for HP Static High Performance mode, HP Static Low Power mode, or HP Dynamic Power Savings mode. The latter automatically adjusts available power to match processor demand. You can find additional information on the HP Power Regulator in the paper titled "Power Regulator for ProLiant servers" at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00300430/c00300430.pdf.

HP Dynamic Power Capping and HP Power Capping

All BladeSystem G7 c-Class server blades with iLO 3 support Dynamic Power Capping and HP Power capping (basic Power Capping). Both technologies achieve the same general goal of maintaining server power consumption at or below the cap value set by an administrator. But HP Dynamic Power Capping monitors power consumption and responds to changes much more rapidly than basic Power Capping.

Dynamic Power Capping can bring a server experiencing a sudden increase in workload back under its power cap in less than one-half second. This fast response prevents any surge in power demand that could trip a typical data center circuit breaker. Dynamic Power Capping prevents tripping circuit breakers that have a specified trip time of 3 seconds or longer at 50 degrees C and 150 percent overload.

Dynamic Power Capping can be an effective tool in planning and managing both electrical provisioning and cooling requirements in the data center. An administrator can electrically provision a PDU or a rack to something less than the full faceplate power rating of all the servers supported. That's because Dynamic Power Capping guards against any unexpected change in power consumption that might trip a circuit breaker.

For detailed information about HP Power Capping technologies and capabilities, refer to the HP technology brief titled "HP Power Capping and HP Dynamic Power Capping for ProLiant servers" at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01549455/c01549455.pdf.

HP Thermal Logic technologies

We use the term *Thermal Logic* to refer to the unique HP mechanical design and control technologies of BladeSystem c-Class systems that let you optimize your power and thermal environment.

Thermal Logic encompasses technologies at every level of the c-Class architecture: processors, memory, server blades, Active Cool fans, and the c-Class enclosure. Through the OA controller, you can access real-time power and temperature data to understand your current power and cooling environments. OA allocates power to the device bays based on the specific configuration of each blade in the enclosure. As you insert blades into the enclosure, the OA automatically discovers each blade and allocates power, based on actual measured power requirements.

HP DDR3 DIMM modules have integrated thermal sensors that signal the chipset to limit memory traffic to the DIMM if the DIMM temperature exceeds a programmable trip point. The OA in the c-Class enclosure uses data from the thermal sensors to reduce fan speed when memory is idle. That reduces power consumption. The BIOS in ProLiant G7 server blades verifies the presence of the thermal DIMM sensor during POST.

Note that some third-party DIMMs may not include this DIMM thermal sensor. If you use third-party DIMMs without thermal sensors, a POST message will warn that the DIMMs don't have thermal sensors. The enclosure fans will run at higher speeds, using more power.

For ProLiant c-Class server blades, we designed smaller heat sinks than those used in rack-mount servers. Our server blade heat sinks have vapor chamber bases, thinner fins, and tighter fin pitch than previous designs. The new design achieves the largest possible heat transfer surface in the smallest possible package (Figure 5). Using smaller heat sinks also creates more space on the server blades for DIMM slots and disk drives.

Figure 5: The left processor heat sink uses fully ducted design. The right heat sink is a traditional heat sink used in a 1U rack-mount server.

Precise ducting on ProLiant c-Class server blades manages airflow and temperature based on the unique thermal requirements of all the critical components. Tightly ducted airflow ensures that no air bypasses the server blade. Ducting produces high pressure that reduces the amount of required airflow, which in turn reduces the power draw of fans. The lower airflow requirement optimizes the available cooling capacity. More information about HP Thermal Logic technologies is available at www.hp.com/go/thermallogic.

Memory technologies

In AMD Opteron 6200 Series architecture, the design integrates the memory controller into the processor chip to optimize memory performance and bandwidth per CPU. The memory controller reduces latency inherent in front side bus architectures by eliminating the bus contention between memory and I/O cycles.

Memory bandwidth increases as you add processors to a configuration. One reason we chose to use AMD Opteron 6200 Series processors in ProLiant G7 c-Class server blades was to use HP DDR3 DIMMs for energy efficiency, high performance, and improved signal quality. We designed specific BIOS and RBSU functions to manage memory configurations, so you can optimize configurations for maximum performance while reducing power consumption and cooling requirements. These management options also include memory protection and latency reduction.

The increased reliability of DDR3 on-DIMM thermal sensors allowed us to incorporate DIMM thermal data into the advanced algorithms controlling thermal and power states within the server.

Unbuffered and registered DIMMs

ProLiant G7 server blades with AMD Opteron processors support Registered DDR3 DIMMs (RDIMMs). Depending on memory size, rank, and configuration, RDIMMs support data rates of 1600, 1333, 1066, or 800 MT/s.

The BL465c G7 server blade with AMD Opteron processors supports Unbuffered DDR3 DIMMs (UDIMMs). UDIMM configurations support two UDIMMs per memory channel. Depending on memory size, rank, and configuration, UDIMMs support data rates of 1333, 1066, or 800 MT/s.

You cannot mix RDIMM and UDIMM memory within a single server. Both UDIMMs and RDIMMs support error correcting code (ECC).

If you require large memory footprints, you can install two 32 GB RDIMMs per channel for a total of 512 GB for the BL465c G7 and 1 TB for the BL685c G7. For smaller memory configurations, you should spread memory as evenly as possible across the memory channels on a given processor and as evenly as possible across processors in a given system.

You can take advantage of the HP low voltage (LV) DDR3 memory option. LV memory operates at 1.35-1.5 V with no reduction in performance in some configurations and with reduced cooling requirements. If you mix LV and standard voltage DDR3 DIMMs in a processor memory bank, all DIMMs will run at the higher voltage.

See the QuickSpecs to identify exactly what memory each server blade supports. For detailed memory configuration guidelines, use the Online DDR3 Memory Configuration Tool available on the HP website at www.hp.com/go/ddr3memory-configurator.

Memory interleaving

Memory interleaving on the AMD Opteron 6200 Series processors can occur between the processor memory banks, between memory channels, and between processor nodes in a multi-processor system. ProLiant G7 c-Class server blades with AMD Opteron 6200 Series processors support all three memory interleaving technologies. These technologies are independent of each other and can operate simultaneously.

Memory bank interleaving

When you use memory bank interleaving, data goes alternately to memory banks through the common memory channel connecting the DIMM banks and the integrated memory controller. Memory bank interleaving increases the probability that more DIMMs will remain in an active state (requiring more power) because the memory controller alternates between memory banks and between DIMMs.

Memory bank interleaving is automatically enabled on a processor node under the following conditions:

- Two single-rank DIMMs per channel result in two-way bank interleaving.

- Two dual-rank DIMMs per channel result in four–way bank interleaving.
- Two quad-rank DIMMs per channel result in eight-way bank interleaving.
- Two dual-rank DIMMs and one quad-rank DIMM result in eight-way bank interleaving, in servers using three DIMMs per channel.

Memory channel interleaving

Memory channel interleaving transfers data by alternate routing through the two available memory channels. As a result, when the memory controller must access a block of logically contiguous memory, the requests don't stack up in the queue of a single channel. Alternate routing decreases memory access latency and increases performance. However, memory channel interleaving increases the probability that more DIMMs must remain in an active state.

Memory channel interleaving is always active on AMD Opteron 6200 Series processors.

Memory node interleaving

Node interleaving can interleave memory across any subset of nodes in the multi-processor system. Node interleaving breaks memory into 4 KB addressable entities and assigns blocks of addresses to the nodes in the sequence indicated in Table 2.

Table 2: Sequencing of memory node interleaving across multiprocessor systems

Node	Assigned Addresses
0	0 – 4095
1	4096 – 8191
2	8192 – 12287
3	12888 – 16383

Address 16384 is assigned to node 0, and the process continues to assign all memory in this fashion. An application that uses a common allocation thread will benefit from node interleaving.

The default configuration disables memory node interleaving. You can activate node interleaving using the RBSU. You can only configure memory node interleaving if the memory footprint for all processors is identical.

Error correction

In AMD Opteron 6200 Series processors, the memory controller supports ECC for both x4 and x8 DIMMs.

In 1996, we introduced Advanced Memory Protection (AMP). It used enhanced ECC algorithms to detect and correct up to four bad bits of data. The HP intelligent correctible memory error threshold algorithm is a development outgrowth of AMP. The algorithm analyzes the probability of system failure when two independently correctable errors on the same rank combine to form an uncorrectable error.

The HP intelligent correctable memory error threshold algorithm replaced existing AMP protection in mid-2011 for AMD G7 servers.

I/O technologies

HP ProLiant c-Class Server Blades support PCI Express (PCIe), serial attached SCSI (SAS), and serial ATA (SATA) I/O technologies. They support multi-function 1 Gb or 10 Gb Ethernet, 4 Gb Fibre Channel, and 4X DDR (20 Gb) or QDR InfiniBand.

PCI Express technology

All ProLiant G7 server blades support the PCIe 2.0 specification. PCIe 2.0 has a per-lane signaling rate of 5 Gb/s, which is double the per-lane signaling rate of PCIe 1.0.

A ProLiant G7 option allows all expansion slots to run at PCIe 1.0 speed rather than at PCIe 2.0 speed. Enabling this option saves power. You can control expansion slot speed through the RBSU under the Advanced Power Configuration submenu.

HP Smart Array controllers and supported drive technology

ProLiant c-Class server blades with AMD processors support hot-plug small form factor (SFF) SAS, SATA, and solid state (SSD) drives through integrated HP Smart Array controllers. ProLiant G7 server blades with AMD Opteron processors use the Smart Array P410i controller with 1 GB FBWC.

A 1 GB FBWC is standard for the Smart Array P410i. It prevents losing information in the cache during an unexpected system shutdown. In a controller failure, you can move the cache module and disks to a working controller that will flush out the cache to the disks.

HP Smart Array technology enables system designs that deploy high-performance SAS and high-capacity SATA SFF drives. SFF drives provide higher performance than large form factor (LFF) drives. The smaller platters of SFF drives reduce seek times because the heads have a shorter distance to travel. RAID performance improves by increasing the number of spindles.

SSDs connect to the host system using the same protocols as disk drives, but they store and retrieve file data in flash memory arrays rather than on spinning media. SSDs eliminate the latency of traditional disk drives by eliminating seek times and by powering up quickly. They also achieve high random-read performance. HP SSDs provide a level of reliability equivalent to or slightly greater than current HP Midline disk drives for servers.

Optional mezzanine cards

HP ProLiant c-Class server blades use two types of mezzanine cards to connect to interconnect fabrics such as Fibre Channel, Ethernet, SAS, or InfiniBand. Type I and Type II mezzanine cards differ only in the amount of power the server allocates to them and in the physical space they occupy on the server blade. Type I mezzanine cards require less power and are slightly smaller than Type II mezzanine cards. Type I mezzanine cards are compatible with all mezzanine connectors in ProLiant c-Class server blades. 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.

Both types of mezzanine cards use a 450-pin connector, enabling up to eight lanes of differential transmit and receive signals. The connections between the device bays and the interconnect bays are hard-wired through the signal midplane. For that reason, you must match the mezzanine cards with the appropriate type of interconnect module. For example, you must place a Fibre Channel mezzanine card in the mezzanine connector that connects to an interconnect bay holding a Fibre Channel switch.

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

Networking technologies

ProLiant G7 c-Class AMD-based server blades use integrated HP NC551i dual-port FlexFabric 10 Gb CNAs. The NC551i delivers the performance benefits and cost savings of converged network connectivity without requiring an add-on mezzanine card. The dual-port NC551i LOM optimizes network and storage traffic with hardware acceleration and offloads for stateless TCP/IP, TCP Offload Engine (TOE), Fibre Channel over Ethernet (FCoE), and iSCSI.

When connected to a converged enhanced Ethernet (CEE) capable switch, Fibre Channel and Ethernet I/O are separated and routed to the corresponding network.

For iSCSI storage, the NC551i LOM supports full protocol offload that provides better CPU efficiency when compared to software initiators. Because of this, the server can handle increased virtualization workloads and compute-intensive applications. Accelerated iSCSI offloads the iSCSI function from the server CPU to the NIC. You can use the HP ProLiant Essentials Accelerated iSCSI Pack with certain multi-function NICs in Windows and Linux environments. To see the supported servers, NICs, and operating systems, go to the QuickSpecs at http://h18002.www1.hp.com/products/servers/proliantessentials/aip/index.html.

The NC551i LOM also supports Virtual Connect Flex-10 that allows dividing each 10 Gb port into four physical NICs to optimize bandwidth management. This combination of high-performance network and storage connectivity reduces cost and complexity. It provides the flexibility and scalability to optimize the data center.

Virtual Connect

HP Virtual Connect (VC) is a portfolio of interconnect modules and embedded software for c-Class enclosures. VC management capability includes control of networks, SAN fabrics, server profiles, and user accounts that make it simpler to set up and manage server connections. VC implements server-edge virtualization so that you can upgrade, replace, or move server blades within their enclosures without visible changes to the external LAN and SAN environments.

VC Flex-10 technology helps administrators fully use 10 GbE connection bandwidth. Using Flex-10, you can partition the bandwidth of a single 10 Gb pipeline into multiple FlexNICs. You can also regulate the bandwidth for each from 100 Mb/s to 10 Gb/s in 100 Mb increments.

HP VC FlexFabric extends VC Flex-10 technology to include FCoE and iSCSI converged storage traffic. The result is a CEE-based FCoE network with Flex-10 variable bandwidth and 10 Gb partitioning abilities that can handle Ethernet, iSCSI, and Fibre Channel traffic. This means that you can connect BladeSystem servers to network and storage infrastructure with a single server connection and a single VC interconnect module. You can connect not only to Fibre Channel SANs but also to iSCSI storage via accelerated iSCSI technology (hardware initiated iSCSI) over a Flex-10 connection. A single pair of VC modules will aggregate an entire enclosure's traffic and connect to existing 1 Gb or 10 GbE and native Fibre Channel external networks.

Full details about VC technology are available in the technology brief titled "HP Virtual Connect technology implementation for the HP BladeSystem c-Class" at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00814156/c00814156.pdf.

Data security technology with the Trusted Platform Module

The Trusted Platform Module (TPM) is a hardware-based system that can securely store information such as passwords and encryption keys to authenticate the platform. You can also use TPM to store platform measurements that help ensure that the platform remains trustworthy. HP ProLiant G7 AMD-based server blades support an optional TPM v1.2. A rivet supplied with the optional TPM v1.2 module attaches and secures the module to the TPM riser board. To prevent possible damage to the TPM module or to the TPM riser board, you cannot remove the TPM from the board once you have installed it.

For additional information about the TPM, see the HP paper titled "Data security in HP ProLiant servers using the Trusted Platform Module and Microsoft® Windows® BitLocker™ Drive Encryption" at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01681891/c01681891.pdf.

Conclusion

The HP ProLiant G7 c-Class server blades with AMD Opteron 6200 series processors offer improved performance and give you more DIMM choices, including higher density DDR3 memory. HP BladeSystem c-Class architecture uses full-featured server blades in a highly dense form factor to accommodate new technologies such as serialized I/O technologies, AMD Opteron multi-core architectures, virtualization, higher power loads, and memory formats requiring more power. This architecture makes network, computing, and storage resources modular, enabling a truly adaptive infrastructure that can accommodate continually changing business needs. HP ProLiant c-Class Server Blades address the growing concern of balancing computing performance with the power and cooling capacity of the data center.

For more information

Visit the URLs listed below if you need additional information.

Resource description	Web address
HP ProLiant BL465c G7 QuickSpecs	http://h18000.www1.hp.com/products/quickspecs/13654_na/13654_na.html
HP ProLiant BL685c G7 QuickSpecs	http://h18004.www1.hp.com/products/quickspecs/13655_na/13655_na.html
General HP BladeSystem information	http://www.hp.com/go/bladesystem/
Technology briefs about HP BladeSystem	http://h18004.www1.hp.com/products/servers/technology/whitepapers/proliant-servers.html

Send comments about this paper to TechCom@HP.com

Follow us on Twitter: http://twitter.com/ISSGeekatHP

© Copyright 2011 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.

TC1111838, January 2011