



Technologies in HP ProLiant G6 c-Class server blades with Intel® Xeon® processors

technology brief

Abstract.....	2
ProLiant c-Class server blade architecture.....	2
Processor technologies	3
Multi-level caches	3
QuickPath Interconnect controller	4
Hyper-Threading	4
Turbo Boost technology.....	5
Thermal Logic technologies	5
Processor socket technology.....	6
Memory technologies	6
I/O technologies	7
PCI Express technology.....	8
HP Smart Array controllers	9
Serial Attached SCSI technology.....	10
SAS and SATA Small Form Factor hard drives	10
Solid state drives.....	10
Optional mezzanine cards.....	10
Networking technologies	11
TCP/IP Offload Engine.....	11
Receive-side Scaling (RSS)	11
iSCSI Acceleration.....	12
iSCSI boot for Linux	12
Virtual Connect	12
Configuration and management technologies	13
BladeSystem Onboard Administrator	13
ProLiant Onboard Administrator (Integrated Lights-Out 2) for ProLiant server blades	14
HP Insight Control suite.....	14
Power management technologies.....	15
Power meter	15
HP Power Regulator for ProLiant.....	15
HP Dynamic Power Capping and HP Power Capping	15
Data security technology with the Trusted Platform Module	16
For more information.....	17
Call to action	17

Abstract

This technology brief describes the architecture and the implementation of major technologies in HP ProLiant G6 c-Class server blades based on Intel® Xeon® processors. Much of the content of this paper is similar or identical to the content of a companion paper about the technologies in HP ProLiant G6 server blades based AMD Opteron™ processors because many of the same technologies have been implemented in both.

It is assumed that the reader is familiar with HP ProLiant server technology and has some knowledge of BladeSystem architecture. For more information about the infrastructure components and complete specifications of each server blade, see the HP website: www.hp.com/go/bladefsystem/.

ProLiant c-Class server blade architecture

An HP ProLiant c-Class server blade is a complete server that slides into an HP BladeSystem c-Class Enclosure. Three different c-Class enclosures are available to meet the needs of large or small IT environments:

- The HP BladeSystem c7000 rack enclosure is 10U high and holds up to 16 ProLiant c-Class server blades.
- The HP BladeSystem c3000 rack enclosure is 6U high and holds up to 8 ProLiant c-Class server blades.
- The HP BladeSystem c3000 tower enclosure is designed with casters for sites without racks. It holds up to 8 ProLiant c-Class server blades inserted vertically.

The rack enclosures fit in HP 10000 series racks and can operate with as few as one server blade installed. The greatest advantage of blade architecture, however, is the ease of adding more server blades. ProLiant c-Class server blades are built in standard form-factors, referred to as *half-height* (4U) and *full-height* (8U). Both half-height and full-height server blades fit into any device bay in a BladeSystem c-Class enclosure.¹

ProLiant G6 c-Class server blades based on Intel Xeon processors include enterprise-class technologies:

- Two Intel Xeon 5500 Series processors
- Thermal Logic technologies
- Advanced memory technologies
- Multiple slots for I/O cards
- Integrated multifunction Ethernet network adapters that support TCP/IP offload engine (TOE)² and iSCSI acceleration
- Hot-plug internal disk drives
- Power management

Half-height server blades support up to 2 processors, 6 to 18 DIMM slots, up to 2 hot-plug or non hot-plug drives, 1 or 2 integrated Ethernet adapters, and up to 2 mezzanine slots. Optional mezzanine cards provide a variety of different I/O fabric connectivity to the interconnect bays.

¹ More information about BladeSystem c-Class enclosure configuration options can be found at <http://www.hp.com/go/bladefsystem/>.

² TOE technology moves processing transactions from the main processor on a server blade to a processor embedded on the network interconnect card. This frees the main processor for other work.

Because of their larger size, full-height server blades support up to 4 processors, up to 32 DIMM slots, 2 or 4 hot-plug drives, 4 integrated Ethernet adapters, and 3 mezzanine slots for optional I/O cards. Full-height server blades also provide twice as much I/O bandwidth to the interconnect modules as the half-height server blades provide. As of this writing, there are no full-height ProLiant G6 server blades based on Intel Xeon processors.

Processor technologies

ProLiant G6 server blades that support two Intel Xeon processors are equipped with Xeon 5500 Series processors. They include an integrated memory controller and Intel QuickPath Technology to boost bandwidth between processors, memory, and I/O subsystems.

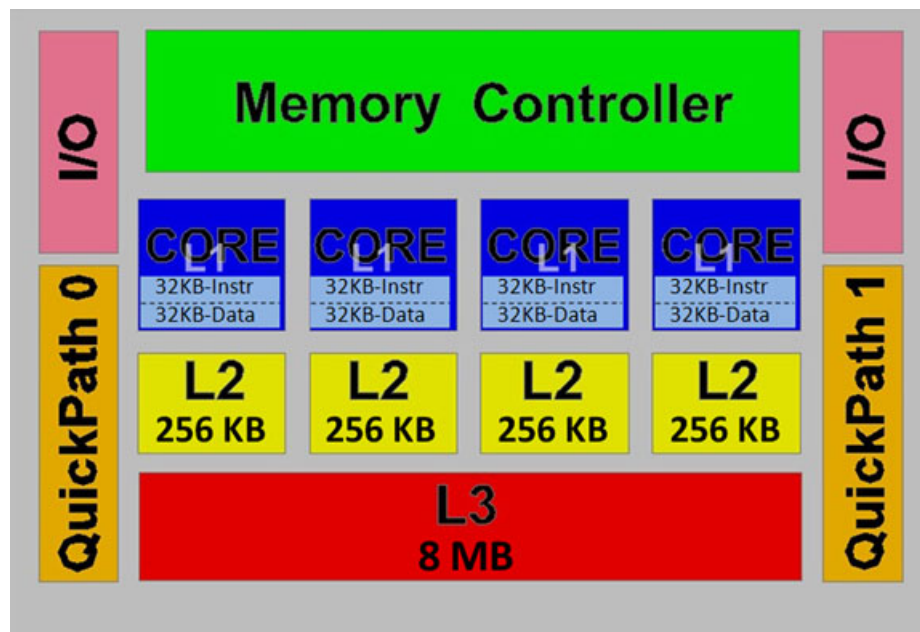
Intel Xeon 5500 Series microarchitecture is built on hafnium-based, 45-nanometer, high-k metal gate, silicon technology to reduce electrical leakage. These small, energy-efficient, high-performance processors support distributed shared memory, Intel Hyper-Threading technology, and Intel Turbo Boost Technology with Intelligent Power Technology.³

Multi-level caches

Intel Xeon 5500 Series processors have a three-level cache hierarchy (Figure 1):

- An on-core 64-kilobyte Level 1 cache, split into two 32-kilobyte caches: one for data and one for instructions
- 256-kilobyte, Level 2 cache for each core to reduce latency
- A Level 3 cache of up to 8 megabytes that is shared by all cores

Figure 1. Block diagram of three-level cache hierarchy for Intel Xeon 5500 Series processors



³ For additional information about Intel processors, see the HP technology brief titled "The Intel processor roadmap for industry-standard servers":


<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00164255/c00164255.pdf>.

The Level 3 cache is shared and inclusive, which means that it duplicates the data stored in each core's Level 1 and 2 caches. This duplication minimizes latency by eliminating unnecessary core snoops to the Level 1 and 2 caches. Flags in the Level 3 cache track the cache source of data. If one core modifies another core's data in Level 3 cache, the Level 1 and 2 caches of those cores are updated as well. This eliminates excessive inter-core traffic and ensures multi-level cache coherence.

QuickPath Interconnect controller

Intel Xeon 5500 Series processors attain their performance potential through Intel QuickPath architecture (Figure 2). High-speed, point-to-point interconnects directly connect the processors with one another. Intel QuickPath architecture also connects each processor to distributed shared memory and to the I/O chipset.

Figure 2. Block diagram of QuickPath architecture



Each QuickPath Interconnect (QPI) consists of two unidirectional links that operate simultaneously in opposite directions using differential signaling. Unlike a typical serial bus, the QPIs transmit data packets in parallel across multiple lanes, and packets are broken into multiple parallel transfers. Each link is comprised of twenty 1-bit lanes. Data transfer uses a maximum of 16 bits, while the protocol and error correction use the remaining 4 bits. The interconnect performs a maximum of 6.4 gigatransfers per second and has a bandwidth of 12.8 gigabytes per second in each direction, for a total bandwidth of 25.6 gigabytes per second. If an application requests data from the memory of another processor, the QPI uses high-bandwidth inter-processor communication to retrieve the data.

To reduce power requirements, administrators can use the ROM-Based Setup Utility (RBSU)⁴ to set QPI links to enter a low power state when the QPI links are not active. Once this feature is enabled, the Intel processor determines when to put the QPI links into a low power state. This reduces power use with minimal performance impact.

Hyper-Threading

With Intel Hyper-Threading Technology, also called Simultaneous Multi-threading Technology (SMT), each core can execute two computational threads at the same time. This means that a single processor can simultaneously execute up to eight threads. In addition, the high-bandwidth memory

⁴ Additional information is provided in the "HP ROM-Based Setup Utility User Guide":
<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00191707/c00191707.pdf>.

subsystem supplies data faster to the two computational processes than traditional front side buses, and the low-latency cache hierarchy allows more instructions to be processed simultaneously. Hyper-Threading improves performance per watt, allowing Intel-based ProLiant G6 servers to accomplish more using the same or less power than servers based on previous-generation Intel processors.

Turbo Boost technology


Intel Turbo Boost Technology complements Hyper-Threading by increasing the performance of both multi-threaded and single-threaded workloads. For workloads and applications that do not benefit from multi-threading, Turbo Boost Technology can increase performance. Turbo Boost is engaged by default and automatically increases the clock frequency of active cores operating below power and thermal design points determined by the processor. Turbo Boost Technology is operating system independent, which means that Advanced Configuration and Power Interface-aware (ACPI) operating systems require no changes to support it.

Thermal Logic technologies

For ProLiant c-Class server blades, HP designed smaller heat sinks than rack-mount servers use. The server blade heat sinks have vapor chamber bases, thinner fins, and tighter fin pitch than previous designs to achieve the largest possible heat transfer surface in the smallest possible package (Figure 3). The smaller heat sink allows more space on the server blades for DIMM slots and hard drives.

Precise ducting on ProLiant c-Class server blades manages airflow and temperature based on the unique thermal requirements of all the critical components. The airflow is tightly ducted to ensure that no air bypasses the server blade and to obtain the most thermal work from the least amount of air moved. Ducting produces high pressure that reduces the amount of required airflow, which in turn reduces the power draw of fans. The lower airflow requirement has the added benefit of optimizing the available data center cooling capacity.

Figure 3. Processor heat sink using fully ducted design (left) and traditional heat sink used in a 1U rack-mount server (right)




More information about HP Thermal Logic technologies is available on the HP website:
www.hp.com/go/thermallogic.

Processor socket technology

The latest Intel 5000 Sequence processors use a processor socket technology called Land Grid Array (LGA). The processor package designs no longer have pins. Instead, the processor package has pads of gold-plated copper that touch processor socket pins on the motherboard.

Technicians must be careful when installing processors to avoid damaging the delicate processor socket pins. Because pin damage could require replacing the motherboard, HP engineers developed a special installation tool to simplify processor installation and reduce the possibility of damaging the socket pins (Figure 4).

Figure 4. Diagram showing how the processor installation tool simplifies installation



Memory technologies

At this writing, HP uses DDR3 memory only on ProLiant G6 server blades with Intel 5500 Series processors. DDR3, the third generation of DDR memory, delivers higher speeds, lower power consumption, and more effective heat dissipation than DDR1 and DDR2 memory. It is an ideal memory solution for bandwidth-hungry systems equipped with multicore processors. Depending on the memory configuration and processor model, the memory may run at 1333, 1066, or 800 MT/s.⁵ DDR3 currently operates at 1.5V, and the voltage is expected to drop after introduction. DDR3 DIMMs have 240 pins, the same number as DDR2, and are the same size. However, they are electrically incompatible and have a different key notch location.

DDR3 memory is available in two forms: Unbuffered (UDIMM) and Registered (RDIMMs). Both UDIMMs and RDIMMs support error correcting code (ECC). There are three types of DDR3 memory:

- PC3-8500R (RDIMM) — 1066 or 800 megatransfers per second (MT/s) data rate, depending on memory configuration and processor installed
- PC3-10600E (UDIMM) — 1333, 1066, or 800 MT/s data rate, depending on memory configuration and processor installed
- PC3-10600R (RDIMM) — 1333, 1066, or 800 MT/s data rate, depending on memory configuration and processor installed

⁵ For additional information about HP server memory technology, go to the HP server memory site: <http://h18004.www1.hp.com/products/servers/options/memory-description.html>.

Administrators can use either RDIMMs or UDIMMs, but RDIMM and UDIMM memory cannot be mixed within a single server. When choosing memory configurations using DDR3 memory, the following guidelines are useful:

- UDIMM configurations are limited to a maximum of two UDIMMs per memory channel because the memory controller must drive the address and command signals to each DRAM chip on a channel. This results in a 24-GB maximum configuration. Because they require fewer components, UDIMMs are typically less expensive than RDIMMs.
- RDIMM configurations can provide larger memory capacity because the memory controller only drives the address and command signals to a single register chip, thereby reducing the electrical load on the memory controller. Users requiring large memory footprints can install up to three 8-GB RDIMMs per channel, for a total of 144 GB. BL460c G6 server blades have only two DIMM slots per channel, for a maximum memory configuration of 96 GB.
- For smaller memory configurations, installing only one or two DIMMs per memory channel can potentially increase memory performance. In many instances this allows administrators to clock the memory channel at a higher data rate. Table 1 summarizes the choices available for DDR3 memory.

Table 1. DDR3 memory options

	Type	DIMM capacity (GB)	Rank	Max. data rate in MT/s (1 or 2 DIMMs/channel)	Max. data rate in MT/s (3 DIMMs/channel)
PC3-8500R	RDIMM	4	quad	1066	800
PC3-8500R	RDIMM	8	dual	1066	800
PC3-10600R	RDIMM	2	dual	1333	800
PC3-10600R	RDIMM	4	dual	1333	800
PC3-10600R	RDIMM	8	dual	1333	800
PC3-10600E	UDIMM	1	single	1333	N/A
PC3-10600E	UDIMM	2	dual	1333	N/A

An RBSU setting allows PC3-10600R memory modules to run at 1333 MT/s with two DIMMs per channel and without performance degradation. Only HP branded DIMMs have been fully validated to operate at this speed. When using third-party DIMMs, HP does not recommend configuring this setting for 1333 MT/s data rates with two DIMMs per channel. This is because third-party DIMMs may not meet stringent design requirements. At this writing, operating the memory at 1333 MT/s is supported only on select ProLiant G6 server blades⁶ using the Xeon 5500 Series processors operating at 95W. For detailed memory configuration guidelines, use the Online DDR3 Memory Configuration Tool available on the HP website: www.hp.com/go/ddr3memoryconfigurator.

I/O technologies

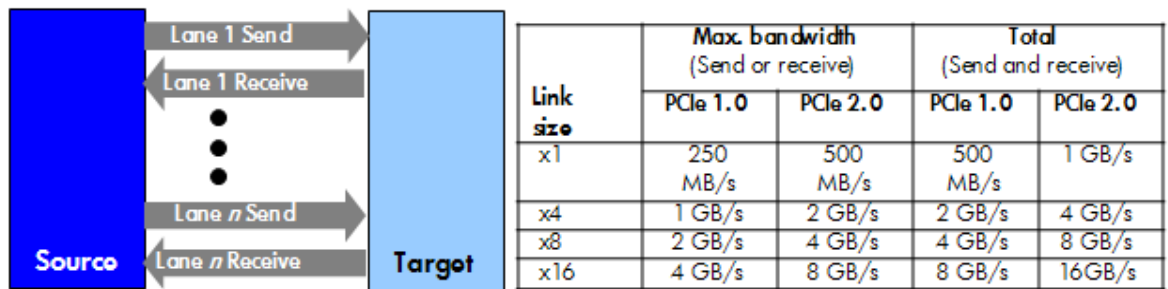
HP ProLiant c-Class server blades support PCI Express (PCIe), serial attached SCSI (SAS), serial ATA (SATA) I/O technologies, Multifunction 1 Gb or 10 Gb Ethernet, 4 Gb Fibre Channel, and 4X DDR (20 Gb) InfiniBand.

⁶ This speed increase for the PC3 10600R memory modules also applies to select BladeSystem 200-series servers such as the BL280c.

PCI Express technology

PCIe allows administrators to add internal expansion cards that support various system capabilities and connection to external storage blades. The PCIe serial interface provides point-to-point connections between the chipset I/O controller hub and I/O devices. Each PCIe serial link consists of one or more dual-simplex lanes. Each lane contains a send pair and a receive pair to transmit data at the signaling rate in both directions simultaneously (Figure 5). PCIe 1.0 has a signaling rate of 2.5 Gb/s per direction per lane. PCIe 2.0 doubles the per-lane signaling rate of PCIe 1.0 from 2.5 Gb/s to 5 Gb/s. This flexibility allows slower devices to transmit on a single lane with a relatively small number of pins while faster devices can transmit on more lanes as required.⁷

Figure 5. PCIe bandwidth



PCIe 2.0 is backward compatible with PCIe 1.0: A PCIe 2.0 device can be used in a PCIe 1.0 slot and a PCIe 1.0 device can be used in a PCIe 2.0 slot. For best performance, however, each card should be used in a slot that supports its logical link size. A ProLiant G6 option allows all expansion slots to run at PCIe 1.0 rather than PCIe 2.0 speed. Enabling this option saves power and provides backward compatibility with cards that may not correctly operate in PCIe 2.0 slots. Administrators can control expansion slot speed through the RBSU under Advanced Power Management Options (Figure 6).

⁷ For additional information about PCI Express technology, see the technology brief titled “HP local I/O technology for ProLiant and BladeSystem servers”:
<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00231623/c00231623.pdf>.

Figure 6. Example of RBSU PCI Express Generation 2.0 support menu

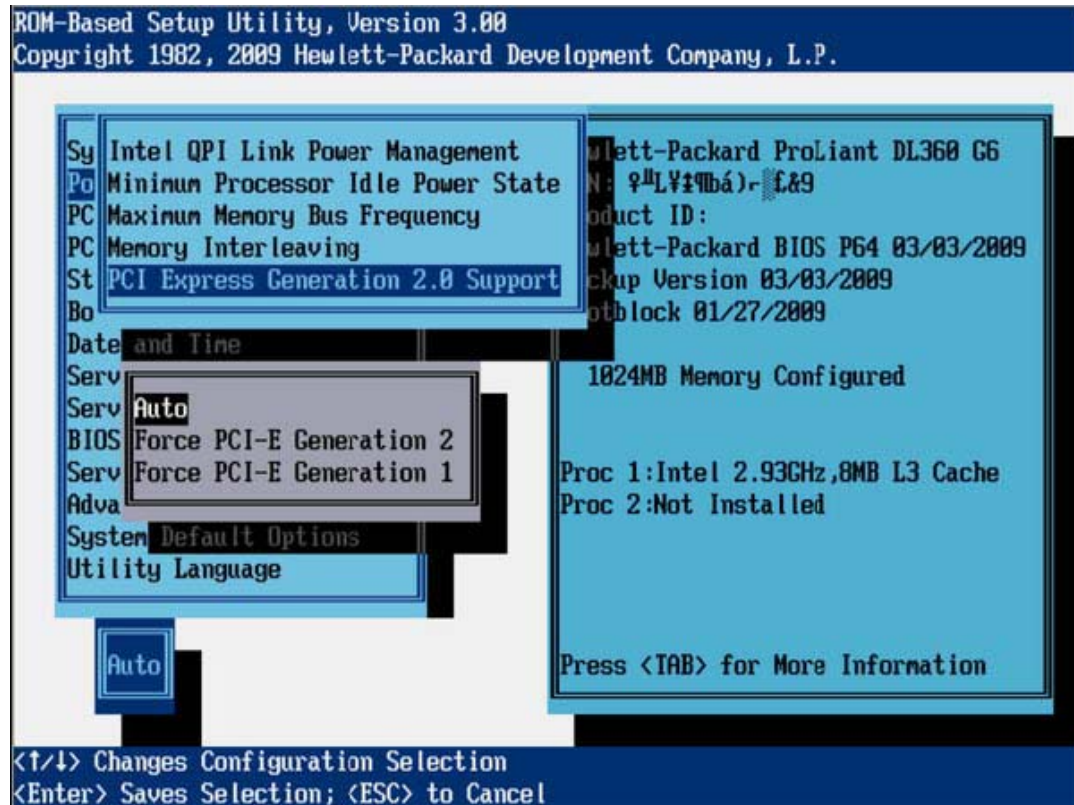


Table 2 shows the level of interoperability between PCIe cards and PCIe slots.

Table 2. PCIe device interoperability

PCIe device type	x4 Connector x4 Link	x8 Connector x4 Link	x8 Connector x8 Link	x16 Connector x8 Link	x16 Connector x16 Link
x4 card	x4 operation	x4 operation	x4 operation	x4 operation	x4 operation
x8 card	Not allowed	x4 operation	x8 operation	x8 operation	x8 operation
x16 card	Not allowed	Not allowed	Not allowed	x8 operation	x16 operation

HP Smart Array controllers

ProLiant c-Class server blades with Intel Xeon 5500 Series processors support internal hard drives through integrated or optional HP Smart Array controllers. The ProLiant BL460c G6 server blade uses the Smart Array P410i controller and supports SAS and SATA SFF drives. The BL280c G6 and BL490c G6 have an embedded SATA controller. An optional Smart Array controller is available for the BL280c G6 to support SAS SFF drives. The BL490c G6 supports only SATA SFF drives internally.

A battery-backed write cache (BBWC) is available as an option for the Smart Array controllers. The battery prevents information in the buffer from being lost in case of an unexpected system shutdown. In the case of a complete system failure, IT administrators can move the controller and disks to a different server, where the controller will flush out the cache to the disks after power is restored. In the case of a controller failure, administrators can move the cache module and disks to a working

controller, where the cache will be flushed out to the disks. The battery will last up to two days without receiving any power from the computer.

Serial Attached SCSI technology

SAS is a serial communication protocol for direct-attached storage devices such as SAS and SATA Small Form Factor (SFF) disk drives. It is a point-to-point architecture in which each device connects directly to a SAS port rather than sharing a common bus, as parallel SCSI devices do. Point-to-point links increase data throughput and improve the ability to locate and fix disk failures. More importantly, SAS architecture solves the parallel SCSI problems of clock skew and signal degradation at high signaling rates.⁸

SAS and SATA Small Form Factor hard drives

The SAS architecture enables system designs that deploy high-performance SAS and high-capacity SATA⁹ SFF drives. This capability provides a broad range of storage solutions that give IT managers the flexibility to choose storage devices based on reliability, performance, and cost.

SFF drives provide higher performance than large form factor drives. The smaller SFF platters reduce seek times because the heads have a shorter distance to travel. At this writing, the peak data transfer rate for SAS drives supported with the P410i firmware is 6 Gb/s in full duplex mode. RAID performance improves by increasing the numbers of spindles.

Solid state drives

In late 2008, HP introduced solid state drives (SSDs) for use in specific BladeSystem environments. 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 hard 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.

Solid state memory (NAND) provides higher capacity, reliability, and performance for local, low-power boot drives than USB keys provide. HP server SSD interfaces are compatible with traditional disk drives connected to a SATA controller. This allows benchmarking and direct comparison of their external performance with that of disk drives to determine their suitability in various application environments.¹⁰

Optional mezzanine cards

HP ProLiant c-Class server blades use two types of mezzanine cards to connect to the various interconnect fabrics such as Fibre Channel, Ethernet, serial-attached SCSI, or InfiniBand. Type I and Type II mezzanine cards differ only in the amount of power allocated to them by the server and in the physical space they occupy on the server blade. Type I mezzanine cards have slightly less power available to them and are slightly smaller. Type I mezzanine cards are compatible with all mezzanine connectors in ProLiant c-Class server blades. Type II mezzanine cards are compatible with Mezzanine

⁸ For more information about SAS technology, refer to the HP technology brief titled "Serial Attached SCSI storage technology":

<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01613420/c01613420.pdf>.

⁹ For more information about these features, refer to the technology brief "Serial ATA technology":

<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00301688/c00301688.pdf>.

¹⁰ For more information about Solid state drive technology, refer to the HP technology brief titled "Solid state drive technology for ProLiant servers":

<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01580706/c01580706.pdf>.

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. Because the connections between the device bays and the interconnect bays are hard-wired through the signal midplane, the mezzanine cards must be matched to the appropriate type of interconnect module. For example, a Fibre Channel mezzanine card must be placed in the mezzanine connector that connects to an interconnect bay holding a Fibre Channel switch. For the most up-to-date information about the c-Class mezzanine card options, go to the HP website: <http://h18004.www1.hp.com/products/blades/components/c-class-interconnects.html>.

Networking technologies

Multifunction 1 Gb or 10 Gb Ethernet network adapters integrated on all c-Class server blades provide several advantages:

- TCP/IP Offload engine (TOE) for Microsoft® Windows® operating systems improves CPU efficiency.
- Receive-side Scaling (RSS) for Windows dynamically load balances incoming traffic across all processors in a server.
- iSCSI Acceleration (available on some integrated network adapters) offloads some of the work in creating iSCSI packets from the processor onto the network controller, freeing up the processor for other work.
- iSCSI boot for Linux® operating systems make it possible to boot the server from a storage area network (SAN) and eliminates the need for disk drives in a server.
- HP Virtual Connect (VC) and Flex 10 provide up to 16 FlexNICs across 4 ports to simplify server connection setup and administration.

For complete specifications about HP network adapter products, go to www.hp.com/go/ProLiantNICs.

TCP/IP Offload Engine

The increased bandwidth of Gigabit Ethernet networks increases demand for CPU cycles to manage the network protocol stack. This means that performance of even a fast CPU will degrade while simultaneously processing application instructions and transferring data to or from the network. Computers most susceptible to this problem are application servers, web servers, and file servers that have many concurrent connections.

The ProLiant TCP/IP Offload Engine for Windows speeds up network-intensive applications by offloading TCP/IP-related tasks from the processors onto the network adapter. TOE network adapters have on-board logic to process common and repetitive tasks of TCP/IP network traffic. This effectively eliminates the need for the CPU to segment and reassemble network data packets. Eliminating this work significantly increases the application performance of servers attached to gigabit Ethernet networks.

TOE is included on integrated Multifunction Gigabit Ethernet adapters and optional multifunction mezzanine cards. It is supported on Microsoft Windows Server 2003 when the Scalable Networking Pack is installed. With the delivery of Windows Server 2008, the TCP/IP Offload Chimney that shipped in the Scalable Networking Pack is included as part of the latest Windows operating system.

Receive-side Scaling (RSS)

RSS balances incoming short-lived traffic across multiple processors while preserving ordered packet delivery. Additionally, RSS dynamically adjusts incoming traffic as the system load varies. As a result, any application with heavy network traffic running on a multi-processor server will benefit. RSS is

independent of the number of connections, so it scales well. This makes RSS particularly valuable to web servers and file servers handling heavy loads of short-lived traffic.

For RSS support on servers running Windows Server 2003, Scalable Networking Pack must be installed. Windows Server 2008 supports RSS as part of the operating system.

iSCSI Acceleration

Accelerated iSCSI offloads the iSCSI function to the NIC rather than taxing the server CPU.

Accelerated iSCSI is enabled by the HP ProLiant Essentials Accelerated iSCSI Pack that is used with certain embedded Multifunction NICs in Windows and Linux environments.

iSCSI boot for Linux

iSCSI boot for Linux is available on BladeSystem NC370i, NC373i, and NC737m Gigabit Ethernet adapters. iSCSI boot allows the host server to boot from a remote OS image located on a SAN within a Red Hat or SUSE Linux environment. The host server uses an iSCSI firmware image (iSCSI boot option ROM), making the remote disk drive appear to be a local, bootable C: drive. Administrators can configure the server to connect to and boot from the iSCSI target disk on the network. It then downloads the OS image from the iSCSI target disk. The HP iSCSI boot solution also includes scripts to significantly simplify the installation process. Adding an iSCSI HBA card is not required.

Virtual Connect

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

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

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

Like other Ethernet and Fibre Channel switches, Virtual Connect modules slide into the interconnect bays of c-Class enclosures. To support the Virtual Connect Fibre Channel module, the enclosure must have at least one Virtual Connect Ethernet module, because the Virtual Connect Manager software runs on a processor that resides on the Ethernet module.

When the LAN and SAN connect to the pool of servers, the server administrator uses Virtual Connect Manager to define a server connection profile for each server. The Virtual Connect Manager creates bay-specific profiles, assigns unique MAC addresses and WWNs to these profiles, and administers them locally. Network and storage administrators can establish all LAN and SAN connections once during deployment. If servers are later deployed, added, or changed, no connection changes are needed because Virtual Connect keeps the profile for that LAN and SAN connection constant.

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

There are advantages to partitioning a 10 GbE pipeline:

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

Flex-10 technology uses two hardware components:

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

The Virtual Connect Flex-10 Ethernet Module is required to manage the 10 GbE server connections to the data center network. The 10 Gb Flex-10 LOM and mezzanine cards are NICs, each with two 10Gb ports. Each 10 Gb port can be configured from one to a maximum of four individual FlexNICs. The server ROM and the operating system or hypervisor recognize each FlexNIC as an individual NIC.

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

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

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

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

Configuration and management technologies

The HP BladeSystem c-Class provides an intelligent infrastructure through extensive integrated management capabilities. These capabilities are based on three components:

- BladeSystem Onboard Administrator
- ProLiant Onboard Administrator powered by Integrated Lights-Out 2 (iLO 2) on the server blade
- Interconnect module management processors such as the HP Virtual Connect Manager

Integrating all 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 white paper titled "Managing HP BladeSystem c-Class systems":

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

BladeSystem Onboard Administrator

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

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

The BladeSystem Onboard Administrator works with the ProLiant Onboard Administrator (iLO 2) management processor on each server blade to form the core of the management architecture for HP BladeSystem c-Class.

To identify a component, the BladeSystem Onboard Administrator reads a Field-Replaceable Unit (FRU) Electrically Erasable Programmable Read-Only Memory (EEPROM) that contains specific factory information about the component, such as product name, part number, and serial number. The BladeSystem Onboard Administrator accesses server blade FRU EEPROMs through their ProLiant Onboard Administrator (iLO 2) management processors.

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 ProLiant Onboard Administrator (iLO 2) web interface and remote server consoles, virtual power control, and boot order control. Server blade control options also include extensive server hardware information including BIOS and ProLiant Onboard Administrator (iLO 2) firmware versions, server name, network adapter and option card port IDs, and port mapping. The BladeSystem Onboard Administrator provides easy-to-understand port mapping information for each of the server blades and interconnect modules in the enclosure. To simplify the installation of the various mezzanine cards and interconnect modules, the BladeSystem Onboard Administrator uses an electronic keying process to detect mismatches.

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

HP BladeSystem c-Class Server Blades include the ProLiant Onboard Administrator powered by iLO 2 to configure, update, and operate individual server blades remotely. The c-Class enclosure includes an Ethernet management network to aggregate all ProLiant Onboard Administrator (iLO 2) management communications across the entire enclosure. This management network connects ProLiant Onboard Administrator (iLO 2) processors to the BladeSystem Onboard Administrator through the BladeSystem Onboard Administrator tray. The BladeSystem Onboard Administrator provides direct access to each ProLiant Onboard Administrator (iLO 2) through the enclosure management network. The BladeSystem Onboard Administrator uses this network to manage pooled enclosure power and cooling.

HP Insight Control suite

HP Insight Control suite provides essential management for HP BladeSystem lifecycles, including proactive health management, lights-out remote control, optimized power use, rapid server deployment, performance analysis, vulnerability scanning and patch management, and virtual machine management.

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

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

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

- HP Insight Virtual Machine Manager software

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

Power management technologies

ProLiant G6 server blades include many more embedded thermal sensors than previous ProLiant server blades. Up to 64 thermal sensors, referred to as a “sea of sensors,” are located on DIMMs, hard drives, and elsewhere throughout the server blade. Three HP power management tools, enabled through the ProLiant Onboard Administrator (iLO 2), help to accurately monitor and provision server power use:

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

Power meter

An integrated power meter in HP ProLiant c-Class server blades analyzes actual server power use. The BladeSystem Onboard Administrator can access the power meter through the ProLiant Onboard Administrator powered by iLO 2 or through external power management software such as HP Insight Power Manager (IPM). 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. The BladeSystem Onboard Administrator provides instant and time-averaged views of the power consumption of individual servers or of all servers within a c-Class BladeSystem enclosure.

HP Power Regulator for ProLiant

HP Power Regulator for ProLiant provides firmware-based speed stepping for Intel x86 processors. Power Regulator improves server energy efficiency by giving processors full power when they need it and reducing power when they do not. This power management feature allows ProLiant servers with policy-based power management to control processor power states. 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. Additional information on the HP Power Regulator is provided in the paper titled “Power Regulator for ProLiant servers”:

<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00300430/c00300430.pdf>.

HP Dynamic Power Capping and HP Power Capping

Both HP Dynamic Power Capping and HP Power Capping (basic Power Capping) achieve the same general goal of maintaining server power consumption at or below the cap value set by an administrator. HP Dynamic Power Capping, however, monitors power consumption and responds to changes much more rapidly than basic Power Capping.

Basic Power Capping does an excellent job of maintaining average server power utilization at or below a cap value. It can help manage data center cooling requirements by limiting server power consumption in a timely enough manner to prevent excessive heat generation. However, basic Power Capping cannot respond quickly enough to limit sudden increases in server power consumption that could cause an electrical circuit breaker to trip.

Dynamic Power Capping operates more than 25 times faster than basic Power Capping, and 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 cause a typical data center circuit breaker to trip. Dynamic Power Capping prevents tripping circuit breakers that have a specified trip time of three seconds or longer at 50 degrees C and 150 percent overload.

This basic difference, the ability to keep server power consumption below the power cap in real time, means that 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 because Dynamic Power Capping guards against any unexpected change in power consumption that might cause a circuit breaker to trip.

Basic Power Capping is supported on all c-Class BladeSystem server blades with iLO 2 version 1.30 or later and system BIOS 2007.05.01 or later. Support for Dynamic Power Capping requires iLO 2 version 1.70 or later, system BIOS 2008.11.01 or later, and BladeSystem Onboard Administrator firmware version 2.32 or later for HP BladeSystem enclosures.

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": <http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01549455/c01549455.pdf>.

Data security technology with the Trusted Platform Module

The Trusted Platform Module (TPM) is a hardware-based system security feature that can securely store information such as passwords and encryption keys to authenticate the platform. Administrators can also use TPM to store platform measurements that help ensure that the platform remains trustworthy. ProLiant G6 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 system board. To prevent possible damage to the TPM module or to the system board, the TPM cannot be removed from the board once it has been installed.¹¹

¹¹For additional information about the TPM, see the HP technology brief titled "Data security in HP ProLiant servers using the Trusted Platform Module and Microsoft® Windows® BitLocker™ Drive Encryption": <http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01681891/c01681891.pdf>.

For more information

For additional information, refer to the resources listed below.

Resource description	Web address
"The Intel processor roadmap for industry-standard servers" technology brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c0164255/c00164255.pdf
"Server virtualization technologies for x86-based HP BladeSystem and HP ProLiant servers" technology brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01067846/c01067846.pdf
General HP BladeSystem information	http://www.hp.com/go/bladesystem/
HP BladeSystem c-Class documentation	http://h71028.www7.hp.com/enterprise/cache/316735-0-0-121.html
HP BladeSystem c-Class Enclosure Setup and Installation Guide	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00698286/c00698286.pdf
HP ROM-Based Setup Utility User Guide	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00191707/c00191707.pdf
HP BladeSystem c-Class interconnects	www.hp.com/go/bladesystem/interconnects
Technology briefs about HP BladeSystem	http://h18004.www1.hp.com/products/servers/technology/whitepapers/proliant-servers.html
HP BladeSystem Power Sizer	www.hp.com/go/hppoweradvisor
"Serial ATA technology" technology brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00301688/c00301688.pdf
"Serial-Attached SCSI storage technology" technology brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01613420/c01613420.pdf
"Memory technology evolution: an overview of system memory technologies" technology brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00256987/c00256987.pdf

Call to action

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

© 2007, 2008, 2009 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

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

Linux is a registered trademark of Linus Torvalds.

Intel and Xeon are registered trademarks of Intel Corporation in the U.S. and other countries and are used under license.

AMD and Opteron are trademarks of Advanced Micro Devices, Inc.

HyperTransport is a licensed trademark of the HyperTransport Technology Consortium.

TC090811TB, August 2009

