



# HP ProLiant DL585 G5/G6 server technology

technology brief

Abstract.....	2
Introduction.....	2
Operating system support.....	2
DL585 G5 processor architecture .....	3
Quad-core AMD Opteron processors .....	3
DL585 G6 processor architecture .....	4
HyperTransport™ technology Assist .....	4
AMD Dual Dynamic Power Management™.....	5
AMD Virtualization™ technology .....	5
AMD Core Select.....	5
Processor socket technology .....	5
NVIDIA NForce Professional 2200 and 2050 chipsets.....	6
Memory technologies .....	7
I/O technologies .....	8
Storage technologies.....	9
Smart Array P400 controller.....	9
Smart Array P410 controller.....	10
Serial Attached SCSI .....	11
SAS and SATA Small Form Factor hard drives .....	11
Networking technologies .....	12
TCP Offload Engine.....	12
Accelerated iSCSI .....	12
Server management technologies .....	12
HP Insight Control – Essential server management for ProLiant .....	12
HP Integrated Lights-Out 2.....	13
Power management technology .....	13
Conclusion.....	14
For more information.....	15
Call to action .....	15

# Abstract

This technology brief describes the architecture and the implementation of major technologies in HP ProLiant DL585 G5/G6 servers based on AMD Opteron® processors. These technologies include processors, memory, connections, power, management, and the latest serial input/output (I/O) technologies.

The fifth generation (G5) and the sixth generation (G6) HP ProLiant DL585 servers offer the performance and efficiency of quad-core AMD Opteron® processors and six-core AMD Opteron® processors respectively, enhanced by improvements to all major subsystems in the server.

This paper explains the technologies implemented in the HP ProLiant DL585 G5/G6 server that makes it an excellent choice for enterprise customers needing increased performance and investment protection with a 64-bit migration path. This paper is intended for IT professionals familiar with system administration and HP ProLiant industry-standard servers. It discusses only key technologies of the ProLiant DL585 server.

## Introduction

For the HP ProLiant DL585 G5/G6 servers, HP improved upon the design of previous DL585 generations by integrating quad-core and six-core processors and by enhancing key subsystems. The DL585 G5 and the DL585 G6 provide excellent system performance by using powerful processors, on-chip memory controllers, a large memory footprint, and improved disk subsystems. These servers are particularly well suited for the following computing environments:

- Data center virtualization and consolidation projects
- High-performance technical computing (HPTC)
- Electronic design automation (EDA) semiconductor design
- Financial applications
- Large database applications
- Enterprise resource planning/Customer relationship management
- Petrochemical applications
- Life science and material science applications
- Video rendering applications

## Operating system support

The ProLiant DL585 supports the following operating systems and virtualization software:

- Microsoft® Windows® Server 2003 R2 Standard (32- and 64-bit), Enterprise (32- and 64-bit), Compute Cluster Server, and Datacenter editions
- Microsoft Windows Server 2008 (when available)
- Red Hat Enterprise Linux 4 and 5
- SuSE Linux Enterprise Server 9 and 10
- VMWare ESX virtualization software
- Microsoft Virtual Server
- Xen virtualization software

For a complete and up-to-date listing of supported OSs and versions, visit the ProLiant OS Support Matrix: <http://www.hp.com/go/supportos>.

## DL585 G5 processor architecture

The DL585 G5 is a powerful, 4U enterprise server incorporating technologies that extend the capabilities of industry-standard x86 computing. This high performance server can run both 32-bit and 64-bit applications simultaneously with no performance penalty (when using an operating system that supports 64-bit extensions). This is possible because the AMD64 instruction set is a super-set containing the x86 instruction set architecture. The DL585 G5 architecture includes the following essential features:

- Support for up to four quad-core AMD Opteron 8300-series processors
- NVIDIA nForce Professional 2200 and 2050 chipsets
- Dual Dynamic Power Management support
- Two 100MHz PCI-X slots
- Four PCI Express (PCIe) x4 slots
- Three PCIe x8 slots
- Two embedded 10/100/1000 Ethernet ports with iSCSI and Receive Side Scaling (RSS)
- HP Integrated Lights-Out (iLO) 2 management

### Quad-core AMD Opteron processors

The AMD 8300-series quad-core processor includes four separate microprocessor cores on one physical die (Figure 1). AMD quad-core technology delivers high performance and reduced latency for multi-threaded and multi-tasking environments. Quad-core processors provide two primary benefits:

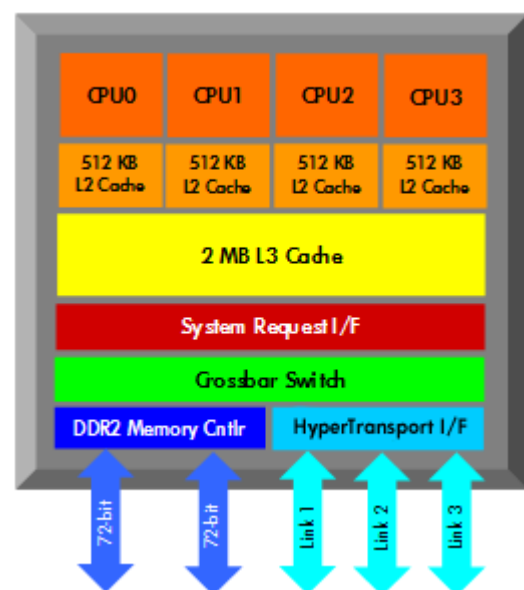
- Enhanced performance per watt—significant performance increase while operating within the same power and thermal envelopes as dual-core processors
- Increased virtualization—on-die memory controller providing optimum performance in virtual machine (VM) environments

Increasing the number of processors available to the operating system makes it possible to execute multiple threads more efficiently. In addition, a quad-core processor uses less power than two (equivalent) dual-core processors which reduces power use and heat production.

The DL585 G5 server supports the AMD Opteron 8300 processor's Dual Dynamic Power Management (DDPM) capabilities. DDPM architecture separates the core and memory controller voltage planes. The HP ProLiant DL585 G5 features a split-plane motherboard that separates the voltage feeds to the CPU core and the DDR2 memory controller. This split-plane design provides peak operating efficiency and increased performance.

As an alternative to traditional front-side bus technology, AMD Opteron processors use Direct Connect Architecture and

**Figure 1.** AMD Opteron 8300 series quad-core architecture



HyperTransport® technology. Direct Connect architecture is the AMD designation for the coherent HyperTransport connection between processors. It eliminates the bottlenecks inherent in front-side bus technology by integrating the memory controller into the processor and directly connecting CPUs to the I/O subsystem and other processors. HyperTransport is a parallel, point-to-point interconnect that replaces parallel front-side bus technology.

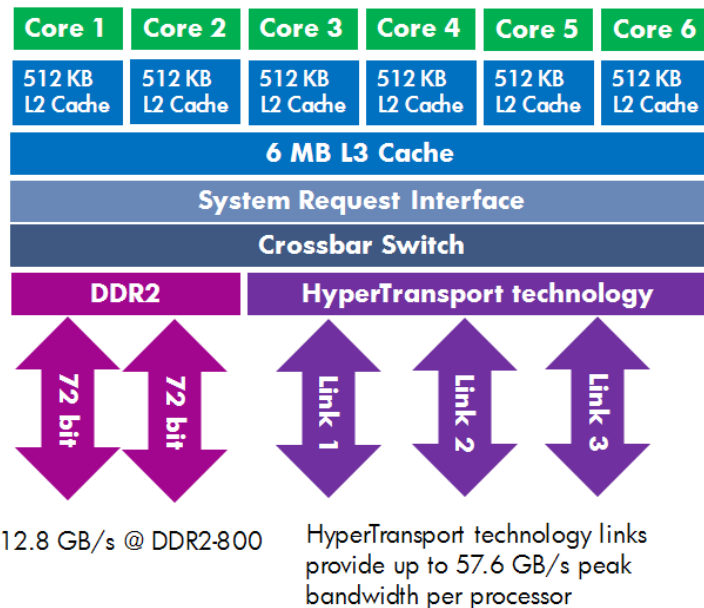
A 64-bit architecture has much more directly addressable (flat) memory space than a 32-bit processor. The AMD64 instruction set allows the OS to access a flat memory address space greater than 4 GB without incurring the overhead of Physical Address Extensions (PAE). This can result in performance advantages, particularly in the ability to use large amounts of memory with, for instance, intensive floating-point calculations used in scientific and engineering modeling programs.

## DL585 G6 processor architecture

HP ProLiant DL585 G6 servers support four AMD Opteron six-core 8400 Series processors. AMD six-core technology delivers high performance and reduced latency for multi-threaded and multi-tasking environments.

DL585 G6 servers feature HyperTransport™ 3.0 (HT3) technology which provides a direct, scalable bandwidth interconnect between the processor, the I/O subsystem, and the chipset (Figure 2). HT3 is an enhancement of HT1, and increases the interconnect rate from 2 giga-transfers per second (GT/s) available on previous AMD processors to a maximum of 4.8 GT/s. Each processor operates at speeds of up to 2.8 GHz, maintains 512 KB of L2 cache memory, and shares a total of 6 MB of L3 cache. An integrated memory controller supports PC2-6400 (DDR2-800) DIMMs.

**Figure 2.** Major components of AMD Opteron six-core processors



## HyperTransport™ technology Assist

HyperTransport™ technology Assist (HT Assist) reduces cache probe traffic between processors, resulting in faster queries that increase performance for cache sensitive applications such as database, virtualization, and compute intensive applications. HT Assist™ resides in the first 1MB of L3

cache and reduces coherence traffic overload on the HT links, resulting in faster queries in both cache- and compute-intensive applications.

## AMD Dual Dynamic Power Management™

AMD Opteron processors feature AMD Dual Dynamic Power Management. This feature powers the processor core and memory controller voltage planes independently and delivers increased performance and improved power management.

## AMD Virtualization™ technology

AMD Virtualization (AMD-V) technology is a hardware-based feature that translates virtual to physical memory addresses using Rapid Virtualization Indexing to provide increased performance and reduced latency.

## AMD Core Select

AMD Core Select uses the BIOS to select the number of software-visible cores per processor. Compliant operating systems and applications can recognize the reduced core count which decreases software licensing costs. This can improve performance of applications not written to make use of the full number of processor cores in the system by giving them access to more cache and greater memory bandwidth per thread. The DL585 G6 server has a ROM-Based Setup Utility (RBSU) option (Advanced->Processor Options) that allows two different configurations:

- All cores enabled [default]
- Enable only half of the cores [half disabled] per socket


In a multi-processor configuration the number of active cores per processor must be equal.

## Processor socket technology

The latest AMD Opteron processor packages use a processor socket technology called Land Grid Array (LGA) to support higher CPU bus speeds. 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.

The processors must be carefully installed to prevent 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 3).

**Figure 3.** Processor installation tool

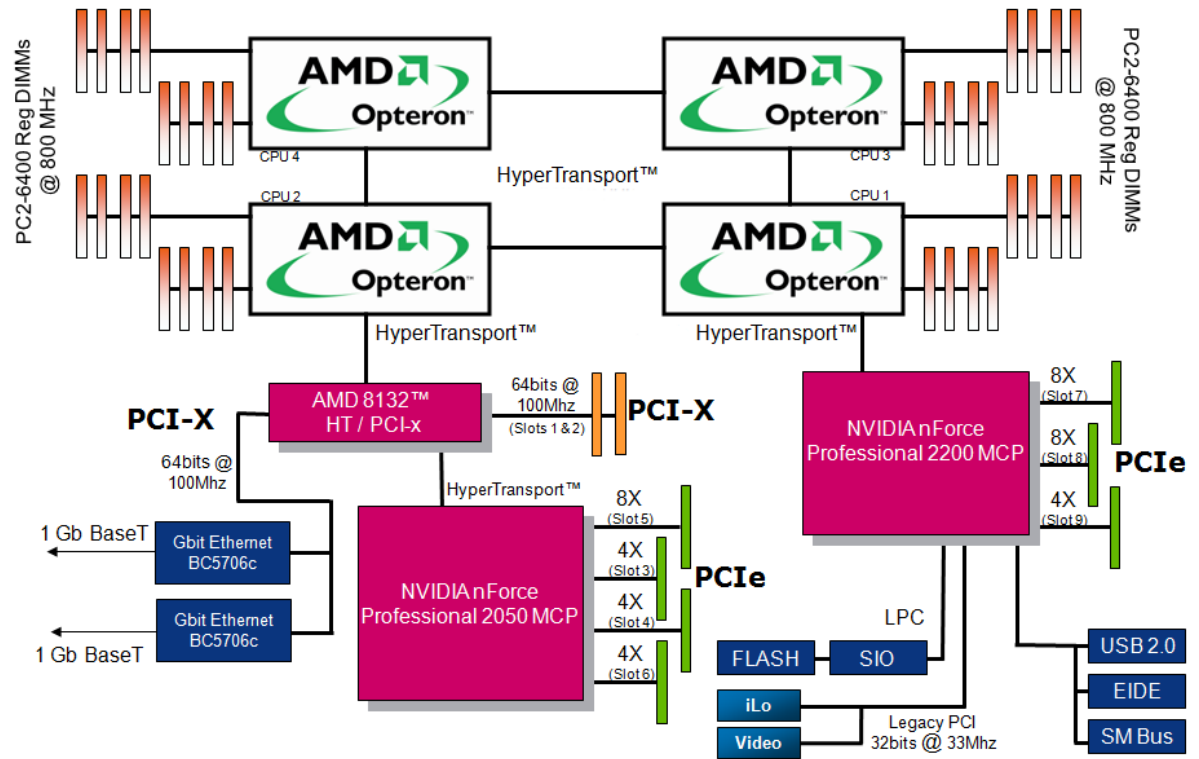


## NVIDIA NForce Professional 2200 and 2050 chipsets

The ProLiant DL585 makes use of both the NVIDIA nForce Professional 2200 and 2050 chipsets to offer a large number of expansion slots for enhanced flexibility. The 2050 chipset connects to the AMD-8132 HyperTransport tunnel with a 16-bit, 1-GHz HyperTransport interface. It provides 20 lanes of PCIe support divided into three x4 slots and one x8 slot. The 2200 chipset connects directly to the processors with an additional 16-bit, 1-GHz HyperTransport connection. In addition to supporting system devices and USB ports, the 2200 offers an additional PCIe x4 slot and two PCIe x8 slots.

The AMD-8132 HyperTransport PCI-X Tunnel provides two PCI-X bridges (Figure 3). One bridge (16-bit, 1-GHz HyperTransport interface) connects to the CPUs. The other bridge (also 16-bit, 1-GHz HyperTransport interface) connects to the downstream NVIDIA nForce Professional 2050 PCIe Bridge. The AMD-8132 supports two PCI-X bridges operating at 100 MHz. One PCI-X bridge supports two 100-MHz expansion slots (one full length and the other half length). The other PCI-X bridge supports two NC371i gigabit Ethernet interfaces.

**Figure 4.** Block diagram of DL585 G5/G6 architecture



## Memory technologies

Typical multiprocessor PC server architecture connects IA-32 processors to memory DIMMs by means of a north bridge chip, which provides the memory controller function. However, in AMD Opteron architecture, the memory controller is integrated into the processor chip for enhanced performance. The on-die memory controller reduces memory latency by eliminating the bus contention between memory and I/O cycles on a typical north bridge. In addition, with each processor containing its own memory controller, the aggregate bandwidth for system-accessible memory is scalable in multi-processor systems such as the DL585.

The Opteron processor supports dual-channel memory. Two 64-bit-wide memory channels operate in parallel to provide a 128-bit interface. Since the DL585 supports dual-width memory channels, DIMMs must be installed in pairs.

By default, ProLiant DL585 memory operates in a linear configuration. This provides optimum performance for Microsoft Windows operating systems and for many applications, such as Microsoft SQL Server. Linear memory is allocated and de-allocated at the thread level, and access is defined on all nodes sequentially. Sequential addresses are assigned to all memory locations starting on node 0, then to all locations on node 1, and so on, until memory locations on all nodes have been assigned.

For those applications that cannot take advantage of linear memory configuration, ProLiant DL585 performance may be improved by activating node interleaving. System administrators can activate node interleaving using the HP RBSU. Node-interleaving breaks memory into 4-KB addressable entities. Applications that use a common allocation thread will benefit from node interleaving.

The ProLiant DL585 G5 and G6 support PC2-6400 and PC2-5300 DDR2 memory. As more DIMMs are installed on a single memory controller (that is, per processor), the memory bus for that processor

will clock down in order to compensate for the additional load. Table 1 lists available memory configurations for the DL585.

**Table 1.** Memory configuration options in the DL585 G5 & G6

<b>DIMMs per processor</b>	<b>Maximum memory speed</b>
2, 4	800 MHz
6	667 MHz
8	533 MHz

## I/O technologies

The DL585 supports two types of I/O expansion slots: PCI-X and PCI Express. The server includes two 64-bit, 100-MHz, PCI-X slots (one full-length, one half-length) that provide a 64-bit, 100-MHz interface for expansion devices.

PCI Express (PCIe) is an expansion bus technology that builds upon earlier PCI technology, but it uses a serial interface rather than the traditional parallel PCI bus. A PCIe interface can support one or more lanes, as determined by the PCIe slot. Multiple-lane connections are described as x4 for a 4-lane connection, x8 for an 8-lane connection, and so forth. Each lane is a separate serial connection of four wires and has a bandwidth of 250 MB/s in each direction (for a total bandwidth of 500 MB/s per single PCIe lane link). PCIe links automatically negotiate the highest number of lanes that the slot and the card support. Therefore, a PCIe expansion card can be installed into any slot in which it will physically fit and it will work correctly.

The DL585 has seven standard PCIe slots: three full-length x8 slots, three full-length x4 slots, and one half-length x4 slot. Since PCIe interface handling is split 4/3 between the two processor pairs and their associated chipsets, optimal performance can be achieved by balancing the load of multiple PCIe expansion cards across the HyperTransport links. Table 2 identifies the preferred installation order.

**Table 2.** Preferred order for populating PCIe expansion slots in the ProLiant DL585 server

Installation sequence	Slot	Speed
First	7	X8
Second	5	X8
Third	8	X8
Fourth	3	X4
Fifth	4	X4
Sixth	9	X4
Seventh	6	X4

## Storage technologies

### Smart Array P400 controller

The DL585 G5 includes a Smart Array P400 controller that supports Serial Attached SCSI (SAS) storage. Some configurations of the DL585 G5 include a battery-backed write cache (BBWC) on the P400 controller. For models that do not include the BBWC, it is available as an option. The battery will last two days without receiving any power from the computer.

The BBWC buffers disk writes so that disk I/O can be handled efficiently. If there is an unexpected system shutdown, the battery maintains the information in the buffer so that data will not be lost. In the event of a complete system failure, the controller and disks can be moved to a different server, where the controller will flush out the cache to the disks once power has been restored. If a controller fails, the cache module and disks can be removed from the failed controller and installed on a working controller, where the cache will be flushed out to the disks.

The Smart Array P400 supports RAID levels 0, 1, 1+0 and 5. RAID 6 (double parity) is available; it requires the BBWC. Mirror splitting is available for RAID 1 arrays. This allows the user to split a RAID 1 mirror into two separate RAID 0 arrays (breaking the mirror). Mirror recombining is the opposite — combining two RAID 0 arrays into a RAID 1 mirror. BBWC is not required for mirror recombining.

The BBWC is required for capacity expansion functionality, which allows the user to add a physical disk to an existing array. The controller then recalculates parity and balances the data across the disks. During the expansion, data and logical structures on the array are preserved.

The P400 controller supports a recovery ROM. In the event of a failed attempt to flash the controller's ROM it reverts to the previous valid ROM. Online Drive Flashing is also supported. With Online Drive Flashing, disk drive firmware updates can be pre-loaded onto the controller, and then at the next reboot, the controller will flash the firmware on the hard drives.

## Smart Array P410 controller

The DL585 G6 includes the Smart Array P410 controller. The HP Smart Array P410 is the HP PCI Express 2.0 (PCIe) Serial Attached SCSI RAID controller. The low profile, half height card has 8 ports and utilizes DDR2-800 memory. The P410 is ideal for RAID 0/1, 1+0, and 5+0. It can be upgraded with the 512MB BBWC module and Smart Array Advanced through the license key for RAID 6 and 6+0. The P410 supports up to 12 drives and up to 2TB of total storage (6 x 500GB HDD). It also supports tape storage.

On the Smart Array P410, a battery module is available with the 512MB BBWC upgrade module. In the event of a controller/server failure, the BBWC equipped Smart Array cache can be removed from the SA-P410 controller and placed on another SA-P410 controller board; the cached data will be flushed to the disk drives.

The following features offer failure detection. They keep the server running and keep data available while a failed drive is being replaced, which allows preventive action to be taken:

- Self Monitoring Analysis and Reporting Technology was developed at HP and detects possible hard disk failure before it occurs, allowing replacement of the component before failure occurs.
- Drive Parameter Tracking monitors drive operational parameters, predicting failure and notifying the administrator.
- Dynamic Sector Repairing continually performs background surface scans on the hard disk drives during inactive periods and automatically remaps bad sectors, ensuring data integrity.
- RAID 6 with ADG (Advanced Data Guarding) allocates two sets of parity data across drives and allows simultaneous write operations. This level of fault tolerance can withstand two simultaneous drive failures without downtime or data loss.
- RAID 5 (Distributed Data Guarding) allocates one set of parity data across drives and allows simultaneous write operations. This level of fault tolerance can withstand a single drive failure without downtime or data loss.
- RAID 5+0 is a RAID 0 array striped across 5 elements. It combines the straight block-level striping of RAID 0 with the distributed parity of RAID 5.
- RAID 1, 1+0 (Drive Mirroring) allocates half of the drive array to data and the other half to mirrored data, providing two copies of every file. It is a high-performance RAID.
- Smart Array Cache Tracking monitors integrity of controller cache, allowing pre-failure preventative maintenance.
- Recovery ROM protects from firmware image corruption by storing a redundant copy of the image. If the active image becomes corrupt, the controller will use the redundant image and continue operating.
- DRAM ECC detects and corrects data bit errors.
- Battery-backed write cache upgrade provides up to two days of battery power for data cache retention. The data backup duration can be extended any time the server's auxiliary power is available during system power down.
- On-Line spares minimize downtime, reconstruct data, and facilitate a quick recovery from drive failure. Up to two spare drives can be installed prior to drive failure. If a failure occurs, recovery begins with an On-Line Spare and data is reconstructed automatically.

Online Drive Flashing is available on the SA-P410. Refer to the Smart Array P400 controller section for a description of this feature.

Mirror Splitting and recombining with HP Smart Array Advanced Pack (SAAP) allows the user to divide a RAID 1 array into two RAID 0 arrays and later combine the two RAID 0 arrays into a single RAID 1 array. The Array Configuration Utility (ACU) lets the user combine any two RAID 0 arrays of the same size and select which drive contains the data to keep. Currently, this feature is only

available offline. The user must boot to the Smart Start CD and run ACU from there. Typically this feature is used when testing out a software patch. The user could split the mirror as a means to save the current data and then perform any type of destructive software update necessary, keeping the resulting data set or reverting back to the old data. BBWC is not required for this feature to be enabled.

Capacity expansion is the process of adding physical drives to the array that has already been configured. The logical drives (or volumes) that exist in an array before the expansion takes place are unchanged, only the amount of free space in the array changes. BBWC is required for this feature.

All Smart Array controllers use the same configuration utility and diagnostic software (Array Configuration Utility (ACU)) and management software (HP Insight Manager). In addition, the SA-P410 also provides Option ROM Configuration for Arrays (ORCA) that allows a simplified configuration tool at the time of controller boot.

To view the HP QuickSpecs for this product, visit the HP web site:

[http://h18000.www1.hp.com/products/quickspecs/13201\\_div/13201\\_div.html](http://h18000.www1.hp.com/products/quickspecs/13201_div/13201_div.html)

## Serial Attached SCSI

Serial Attached SCSI (SAS) leverages a common electrical and physical connection interface with Serial ATA (SATA), while offering logical SCSI compatibility along with SCSI reliability, performance and manageability. SAS provides investment protection in compatible SCSI software and middleware and the choice of direct-attach storage devices (SAS or SATA). In addition, greater performance, longer cabling distances, smaller form factors and greater addressability will all lead to a new level of flexibility when deploying mainstream data center servers and subsystems. This compatibility provides users with many choices for server and storage subsystem deployment, by leveraging the SATA development effort on smaller cable connectors, providing customers a downstream compatibility with desktop class ATA technologies.

## SAS and SATA Small Form Factor hard drives

The SAS architecture enables system designs that deploy high-performance SAS and high-capacity SATA Small Form Factor (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. RAID performance improves by increasing the numbers of spindles.

HP ships SATA drives with Drive Write Cache (DWC) disabled. The preset configuration was selected to provide greater safety for drive data in case of sudden power loss when there is no battery on the controller to protect the cache. Enabling DWC may result in data loss if power is lost to the server and there is no power protection configured for the server.

Native Command Queuing (NCQ) increases SATA HDD performance by internally prioritizing read and write command execution. This reduces unnecessary drive head movement and results in increased performance especially in server or storage-type applications with outstanding multiple simultaneous read/write requests. Without NCQ, the drive can process and complete only one command at a time. NCQ must be supported and turned on in both the controller and the drive. Please see the SATA Hard Drive QuickSpecs for specific SATA hard drive capacities that support NCQ: [http://h18000.www1.hp.com/products/quickspecs/11940\\_div/11940\\_div.html](http://h18000.www1.hp.com/products/quickspecs/11940_div/11940_div.html).

## Networking technologies

The DL585 includes two BC5706i multifunction gigabit Ethernet (GbE) network adapters. Technologies supported by the BC5706i include TCP Offload Engine, Receive Side Scaling (RSS), and accelerated iSCSI.

### TCP Offload Engine

Network bandwidth has improved steadily since the early days of TCP/IP networking. Along with this improvement in speed has come increased demand for CPU cycles to manage the network protocol stack. A busy, full-duplex gigabit Ethernet connection can consume the available computing power of a 1-GHz Pentium 4 processor. Unfortunately, this means that even a modern, high-powered processor will show degraded performance in processing application instructions while data is being transferred onto or off of the network. Computers most susceptible to this problem are application, web, and file servers that have a high number of concurrent connections.

To reduce this burden on the processor, the embedded network controller in the DL585 is designed with TCP Offload Engine (TOE) capabilities. TOE NICs (network interface cards) are designed with on-board logic to process common and repetitive tasks of TCP network traffic. Because the processor does not have to devote cycles to processing these tasks, it can be used more efficiently to significantly increase application performance on servers attached to gigabit Ethernet networks.

### Accelerated iSCSI

iSCSI (internet SCSI) is a standard that implements the SCSI protocol for interacting with storage devices over a TCP/IP network. While iSCSI can be implemented over any TCP/IP network, the most common implementation is over gigabit Ethernet. iSCSI serves the same purpose as Fibre Channel in building Storage Area Networks (SANs), but without the cost, complexity, and compatibility issues associated with Fibre Channel SANs. A more efficient approach is to offload the management of the protocol to the network adapter. Adapters that implement the iSCSI protocol are known as iSCSI Host Bus Adapters (HBAs). The embedded NIC of the DL585 is an iSCSI HBA and appears to the operating system as a SCSI HBA.

## Server management technologies

### HP Insight Control – Essential server management for ProLiant

The ProLiant DL585 server infrastructure costs less to operate because Insight Control lets server administrators deploy the ProLiant DL585 quickly and increase uptime by proactively managing health. Additionally, the DL585 server reduces overall power requirements and extends data center capacity by optimizing power consumption. The DL585 server reduces travel costs and increases the number of ProLiant servers managed per administrator by providing complete remote control for ProLiant servers from any location.

With Insight Control, administrators can manage ProLiant servers running Windows, Linux, VMware ESX, and Citrix XenServer environments, or even integrate Insight Control with leading third party enterprise management consoles, including HP Operations Center and Microsoft System Center.

HP Insight Control taps into the ProLiant Onboard Administrator – a set of embedded management capabilities designed into every ProLiant server. These embedded management capabilities exposed through Insight Control help make ProLiant the best managed server in the industry.

HP Insight Control is a full suite of essential ProLiant management software that integrates into the HP System Insight Manager console, which helps system administrators in several ways:

- Deploy ProLiant servers quickly – brings reliability and consistency to each ProLiant deployment with a fast, drag-and-drop tool that can turn a manual, resource-intensive server deployment into an unattended, repeatable, and highly automated ProLiant server deployment.
- Proactively monitor ProLiant server health – system administrators can monitor an entire ProLiant infrastructure with one simple, integrated interface and receive proactive notification of any actual or impending component or server failure.
- Control ProLiant servers from any location – administrators can take control of ProLiant servers at any time, regardless of location, and realize reduced travel time and costs and improved system recovery times.
- Optimize ProLiant server power confidently – data center capacity can be extended by fitting more ProLiant servers within an existing power envelope.

Insight Control is available bundled with the ProLiant DL585 G6 and is backed by HP Service and Support.

## HP Integrated Lights-Out 2

The HP ProLiant DL585 server also includes the HP Integrated Lights-Out 2 (iLO 2) management technology. iLO 2 hardware and firmware provide remote server management capabilities over Ethernet. The iLO 2 management processor obtains its power from the auxiliary power plane of the server, so it is always available when the server is plugged into a power source—even when the OS is not operating.

The DL585 provides the ability to connect to iLO 2 through a dedicated Ethernet port. This port can connect to a dedicated management network that is out of the data path for the server and that can be highly secured. Alternatively, in the DL585, iLO 2 also provides a Shared Network Port (SNP) that allows network access to both iLO and the host server using a single network port. SNP allows iLO network traffic to be routed by means of a sideband connection on one of the server NIC interfaces. Although the iLO traffic shares a port with the server OS traffic, both the iLO processor and the server NIC have their own Media Access Control (MAC) address. Having separate MAC addresses enables the iLO device and the server to have separate Internet Protocol (IP) addresses. Using the SNP simplifies hardware installation and reduces overall hardware costs because both corporate and iLO network traffic comes through the system NIC.

## Power management technology

The HP DL585 includes HP Power Regulator, an innovative OS-independent power management tool. HP Power Regulator is a ROM-based utility that can be used to set the server to one of four power modes:

- Static high power – Server runs continuously in the highest performance state.
- Static low power – Server runs continuously in the lowest power state.
- OS Control – Server uses AMD Opteron PowerNow!™ Technology that allows the operating system or drivers to control processor power states.
- Dynamic power savings – server processor power is automatically adjusted according to application activity.

HP Power Regulator improves the energy efficiency of the DL585. Opteron processors run at full power when they need to, but when application activity is reduced, they run in a power savings mode without performance degradation. HP Power Regulator modes can be selected through the RBSU during POST or through the iLO 2 remote management interface any time during runtime.

For more information about HP Power Regulator, visit the HP web site at <http://h18013.www1.hp.com/products/servers/management/ilo/power-regulator.html>.

## Conclusion

The HP ProLiant DL585 is a 4U rack-optimized, four-processor server created for large data center deployments requiring enterprise-class performance, uptime, and scalability, plus ease of management and expansion. It offers customers running both 32- and 64-bit applications increased performance and memory speed. Quad-core and six-core processor support along with improved network and I/O subsystems ensure that the DL585 is scalable in high-performance computing environments.

## For more information

For additional information, refer to the resources listed below.

Resource description	Web address
"AMD Opteron™ and Intel® Xeon® x86 processors in industry-standard servers" technology brief	<a href="http://h20000.www2.hp.com/bc/docs/support/SupportManual/C02731435/C02731435.pdf">http://h20000.www2.hp.com/bc/docs/support/SupportManual/C02731435/C02731435.pdf</a>
AMD white paper "HyperTransport™ Technology: Simplifying System Design, October 2002"	<a href="http://www.hypertransport.org/docs/26635A_HT_System_Design.pdf">http://www.hypertransport.org/docs/26635A_HT_System_Design.pdf</a>
Up-to-date information on operating systems and versions supported by the HP ProLiant DL585 server	<a href="http://www.hp.com/go/supportos">http://www.hp.com/go/supportos</a>
Latest drivers available for the HP ProLiant DL585 server	<a href="http://www.hp.com/support/files">http://www.hp.com/support/files</a>
Details about HP Insight Control	<a href="http://www.hp.com/go/insightcontrol">http://www.hp.com/go/insightcontrol</a>
Information about iLO 2	<a href="http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00257345/c00257345.pdf">http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00257345/c00257345.pdf</a>
Information about SAS technology	<a href="http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01613420/c01613420.pdf">http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01613420/c01613420.pdf</a>
Details about Serial ATA technology	<a href="http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00301688/c00301688.pdf">http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00301688/c00301688.pdf</a>
"Memory technology evolution: an overview of system memory technologies" technology brief	<a href="http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00256987/c00256987.pdf">http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00256987/c00256987.pdf</a>

## Call to action

Send comments about this paper to [TechCom@HP.com](mailto:TechCom@HP.com).

© 2006, 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 trademarks or registered trademarks of Intel Corporation in the U.S. and other countries and are used under license.

AMD and Opteron are trademarks and AMD-8000 is a mark of Advanced Micro Devices, Inc.

HyperTransport is a licensed trademark of the HyperTransport Technology Consortium.

TC090907TB, September 2009

