



Technology and architecture of HP ProLiant AMD-based 300-series G6 (Generation 6) servers



Technology brief

Abstract.....	3
Introduction.....	3
Processor technology.....	3
Direct Connect architecture 2.0	3
HyperTransport Technology.....	4
HT Assist	5
Memory technologies	5
Double data rate memory	6
Determining memory bus speed.....	6
Memory management technologies	6
Memory bank interleaving.....	7
Memory Channel Interleaving	7
Memory node interleaving.....	7
X8 error correction	7
I/O technologies	8
PCI Express technology.....	8
HP Smart Array and SAS/SATA technology	8
SAS-2 standard	9
Battery backed write cache	9
Zero Memory RAID.....	10
HP Smart Array compatibility.....	10
Smart Array Advanced Pack.....	10
Power and thermal technologies	11
Efficient power delivery.....	11
Common-slot power supplies.....	11
Increasing power efficiency in redundant power operation.....	12
Voltage regulation.....	12
Thermal sensors and fan control.....	12
APML Remote Power Management Interface	13
Dynamic Power Capping	13
Managing processor technologies.....	14
PowerNow! Technology	14
Power Cap Manager	14
AMD Core Select and HP Core Disable	14
Security	15
Trusted Platform Module.....	15
BitLocker Drive Encryption	15
Systems management and monitoring	16
HP ProLiant Onboard Administrator powered by iLO 2.....	16
HP Insight Control Environment	16

HP Insight Dynamics – VSE.....	16
OS support	16
Summary	17
For more information.....	18
Call to action	19

Abstract

This technology brief describes the key technologies implemented in HP ProLiant 300-series G6 servers based on AMD™ processors. As of this writing, the AMD-based 300-series G6 server platforms are limited to the ProLiant DL385. A link to the QuickSpecs for this server is listed in the “For more information” section at the end of this technology brief.

Introduction

HP ProLiant 300-series G6 servers include these new technologies:

- AMD Opteron™ six-core 2400 Series processors
- The ProLiant Onboard Administrator powered by Integrated Lights-Out 2 (iLO 2) that delivers power and temperature management through multiple sensors and fan control
- Input/output (I/O) technologies such as PCI Express (PCIe) and faster Smart Array controllers that incorporate common form factor components
- “Right Size” common-slot power supplies in multiple sizes to improve power efficiency

The technologies discussed in this paper are implemented in all AMD-based ProLiant 300-series G6 servers. Exceptions are noted where different levels of technology implementation or service exist in individual ProLiant 300-series G6 platforms.

For complete specifications of all ProLiant 300-series servers, see the HP website at www.hp.com/products/servers/platforms.

Processor technology

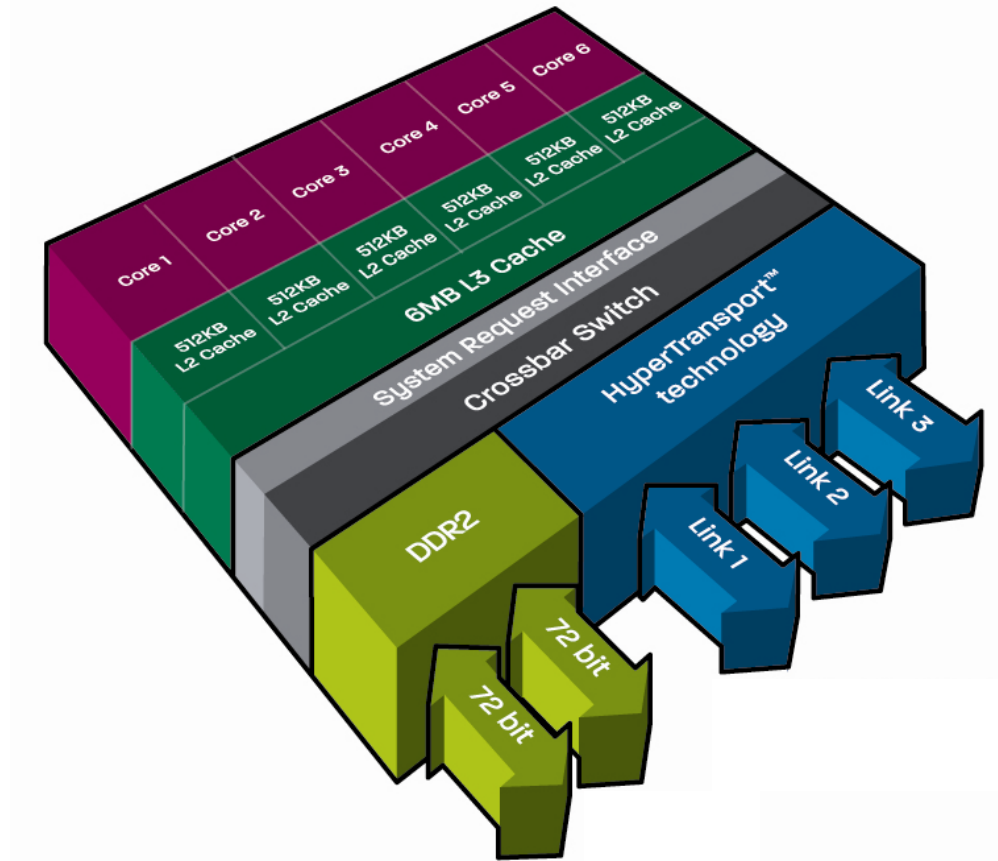
The HP G6 AMD-based 300-series servers use the AMD Opteron™ 2400 series Six-Core processor. These processors are based on AMD's 45 nanometer (nm) quad-core architecture. In addition these processors use Direct Connect Architecture 2.0, HyperTransport 3.0, HT Assist, Advanced Platform Management Link (APML) Remote Power Management Interface, and 8x ECC error correction. The Six-Core processors fit into socket 1207 architecture just as previous four-core AMD Opteron processors do. Each processor operates at speeds of up to 2.9 GHz, has 512 KB of L2 cache memory, and shares a total of 6 MB of L3 cache.

Direct Connect architecture 2.0

Instead of a front-side bus architecture, Direct Connect integrates the memory controller into the processor and directly connects CPUs to the I/O subsystem and other processors (Figure 1). Direct Connect Architecture uses direct communication links (HyperTransport links) between each CPU, between CPU and I/O, and between CPU and memory. Direct Connect Architecture currently scales up to 12 cores, provides superior memory and I/O capability, near native virtualization performance, and a range of power bands¹ that place a priority on low power consumption.

¹ Power bands refer to a new metric developed by AMD to reflect power consumed by the processor and its integrated memory controller during peak workloads. This metric is based on AMD's measurement of Average CPU Power (ACP). For more information on ACP, see the whitepaper at www.amd.com/us-en/assets/content_type/white_papers_and_tech_docs/43761C_ACP_WP.pdf

Figure 1. Block diagram of Direct Connect 2.0 architecture in the AMD 2400-series processors



HyperTransport Technology

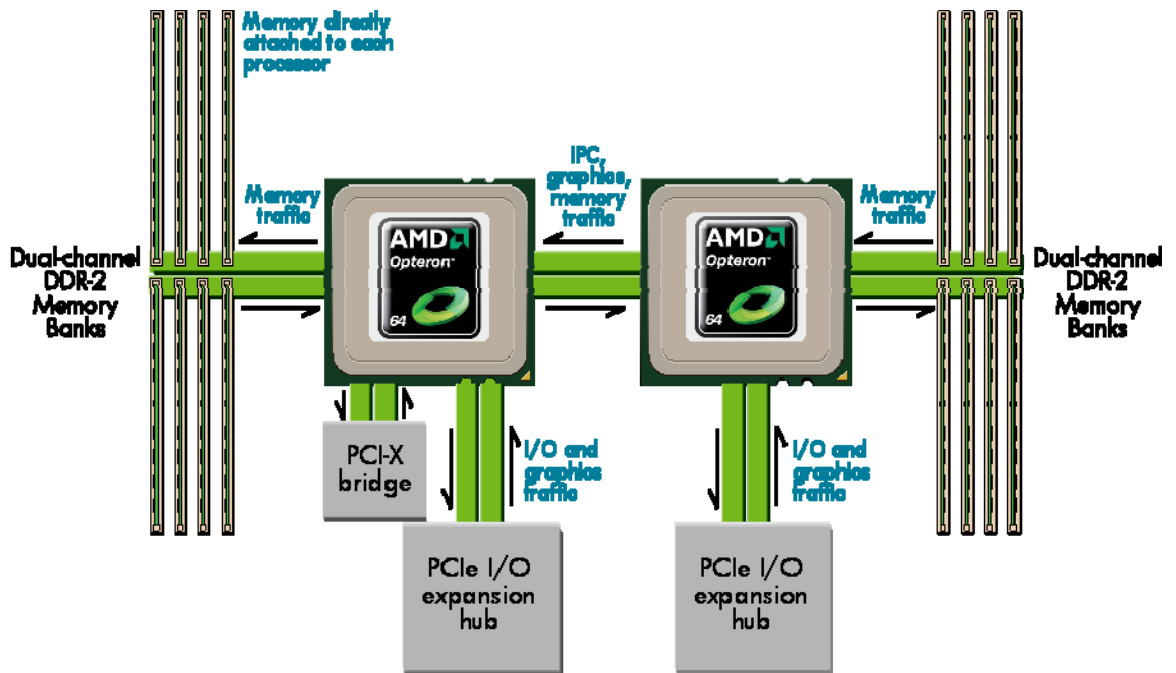
HyperTransport is a point-to-point interconnect with two unidirectional links (Figure 2). It directly connects the processors to each other and connects each processor to its dedicated memory banks, as well as to other I/O chipsets.² Compared to a shared, parallel front-side bus, HyperTransport has the advantages of no overhead for bus arbitration and easier signal integrity maintenance which results in a scalable, high-bandwidth architecture.

Each 16-bit (2-byte) HyperTransport link is double-pumped, performing two data transfers per clock cycle. From HyperTransport 1.0 (HT1) in 2001 to HyperTransport 3.0 in 2008, the maximum clock speed and transfer rate increased from 800 MHz (1.6 MT/s³) to a maximum of 2.6 GHz (4.8 GT/s) in each direction. This gives each HyperTransport 3.0 (HT3) link a maximum data rate of 4.8 GT/s \times 2 bytes per transfer, or 9.6 GB/s.

² HyperTransport Technology was invented at AMD with contributions from industry partners and is managed and licensed by the HyperTransport Technology Consortium, a Texas non-profit corporation.

³ MT/s, or megatransfers per second, equals the speed of the link in millions of cycles per second times the number of transfers per cycle.

Figure 2. The HyperTransport interconnect is designed to separate memory and I/O traffic and directly attaches memory to each processor, allowing memory capacity to scale with the number of processors.



HyperTransport can be configured through the ROM-Based Setup Utility (RBSU). In the RBSU Advanced Options menu, administrators can use the HyperTransport Selection option to choose HT1 or HT3. HT1 is active by default.

HyperTransport is designed to provide a direct, scalable bandwidth interconnect between the processor, the I/O subsystem, and the chipset. Due to chipset architecture in ProLiant AMD G6 servers, HT3 supports processor to processor communication, but not I/O operations.

HT Assist

HT Assist™ reduces cache coherence traffic on the HT links. By tracking where data is stored in cache and guiding the processor directly to the data in the other processors' caches, HT Assist reduces cache probe traffic between processors, especially in 4-socket servers. Therefore, HT Assist results in faster queries that can increase performance for cache-sensitive applications such as database, virtualization, and compute-intensive applications.

Memory technologies

In the AMD Opteron 2400-series architecture, the memory controller is integrated 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. AMD's memory bandwidth increases as processors are added to a configuration, compared to legacy designs that scale poorly because access to main memory is limited by external northbridge chips.

The AMD Opteron 2400-series processor supports dual memory channels. Two 64-bit-wide memory channels operate in parallel to provide a 128-bit interface, so memory must be installed in pairs. The AMD architecture incorporates the following enhancements:

- Write bursting to minimize read/write transitions for greater throughput
- Optimized DRAM paging algorithm for greater throughput
- DRAM prefetcher to intelligently predict and retrieve needed data from main memory
- Core prefetchers that fetch data directly to L1 cache to decrease latency and reduce L2 bandwidth

Double data rate memory

HP ProLiant G6 300-series servers with AMD processors support PC2 Registered double data rate (DDR-2) DIMMs. These DL 300-series G6 servers use 1-GB, 2-GB, 4-GB, or 8-GB registered DDR-2 DIMMs across 16 DIMM slots. DDR-2 memory devices operate at 1.8V. They use high clock frequencies to increase data transfer rates and on-die termination control to improve signal quality. For example, at a clock frequency of 400 MHz, the data transfer rate is 800 MT/s, which translates to a memory bandwidth of 6400 MB/s per DIMM⁴.

Determining memory bus speed

Depending on the quantity and speed of installed DIMMs, memory bus speeds on the DL300-series G6 servers can be 533-MHz, 667-MHz, or 800-MHz. The memory bus speed is determined by the BIOS and RBSU using the following two-step process:

Step 1- The BIOS interrogates each DIMM that is installed in the system and sets the memory bus speed to the least common denominator of the DIMM speeds detected, regardless of the RBSU menu option selected. For example, with 667-MHz and 800-MHz DIMMs installed, this first step should complete with the memory bus speed set to 667 MHz regardless of how many DIMMs are installed on each processor.

Step 2- With this step, the following RBSU menu option settings determine the actual memory bus speed that is selected.

- Menu Option 1 (Auto default) – Memory bus speed is determined by the number of DIMMs installed for each processor:
 - Auto Case 1: Four or fewer DIMMs = 800-MHz memory bus speed provided that all DIMMs detected in step 1 are rated for 800 MHz.
 - Auto Case 2: Six DIMMs = 667-MHz memory bus speed provided that all DIMMs detected in step 1 are rated for 667 MHz or greater.
 - Auto Case 3: Eight DIMMs = 533-MHz memory bus speed provided that all DIMMs detected in step 1 are rated for 533 MHz or greater.
- Menu Option 2
 - Memory bus speed is set to 533 MHz regardless of how many DIMMs are installed provided that all DIMMs detected in step 1 are rated for 533 MHz or greater.

Memory management technologies

Memory interleaving on the AMD Opteron 2400-series processors can occur between the processor memory banks, memory channels, and between processor nodes in a multiprocessor system. All three

⁴ For additional information, refer to the HP technology brief titled "Memory technology evolution: an overview of system memory technologies": <http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00256987/c00256987.pdf>

memory interleaving technologies are supported on the ProLiant AMD 300-series G6 servers. These technologies are independent of each other and can operate simultaneously.

Memory bank interleaving

When memory bank interleaving is engaged, data is routed alternately to memory banks through the common memory channel connecting the DIMM banks and the integrated memory controller. However, memory bank interleaving does increase the probability that more DIMMs need to be kept in an active state (requiring more power) since the memory controller alternates between memory banks and therefore between DIMMs.

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

- DIMMs are installed in identical pairs
- Four single-rank DIMMs are populated per node (4 x 1 GB or 4 x 2 GB for example)
- Two dual-rank DIMMs are populated per node (2 x 4 GB or 2 x 8 GB for example)
- Four dual-rank DIMMs are populated per node (4 x 4 GB or 4 x 8 GB for example)

Using four single-rank DIMMs or two dual-rank DIMMs results in two-way bank interleaving. Using four dual-rank DIMMs results in 4-way bank interleaving.

DIMMs must be installed in identical pairs. If single and dual ranked DIMMs are mixed on the same node, bank interleaving will not be enabled. DIMMs must be installed in decreasing capacity with the largest DIMMs installed in the banks furthest away from each processor.

Memory Channel Interleaving

With memory channel interleaving, data is transferred by means of alternate routing through the two available memory channels. The result is that when the memory controller needs to access a block of logically contiguous memory, the requests are distributed more evenly across the two channels rather than potentially stacking up in the request queue of a single channel. This alternate routing decreases memory access latency and increases performance. As with memory bank interleaving, memory channel interleaving increases the probability that more DIMMs need to be kept in an active state.

Memory channel interleaving is always active on the AMD 2400-series processor.

Memory node interleaving

Node interleaving allows memory to be interleaved across any subset of nodes in the multiprocessor system. Node interleaving breaks memory into 4 KB addressable entities. Addressing starts with address 0 on node 0 and assigns sequential addresses through address 4095 to node 0, addresses 4096 through 8191 to node 1, addresses 8192 through 12287 to node 2, and addresses 12888 through 16383 to node 3. Address 16384 is assigned to node 0, and the process continues until all memory has been assigned in this fashion. An application that uses a common allocation thread will benefit from node interleaving.

Memory node interleaving is disabled by default. Administrators can activate node interleaving using the RBSU. Node interleaving can only be configured if the memory footprint for both processors is the same.

X8 error correction

In AMD Opteron 2400-series processors, the memory controller supports error correction circuitry (ECC) for both x4 and x8 DIMMs.

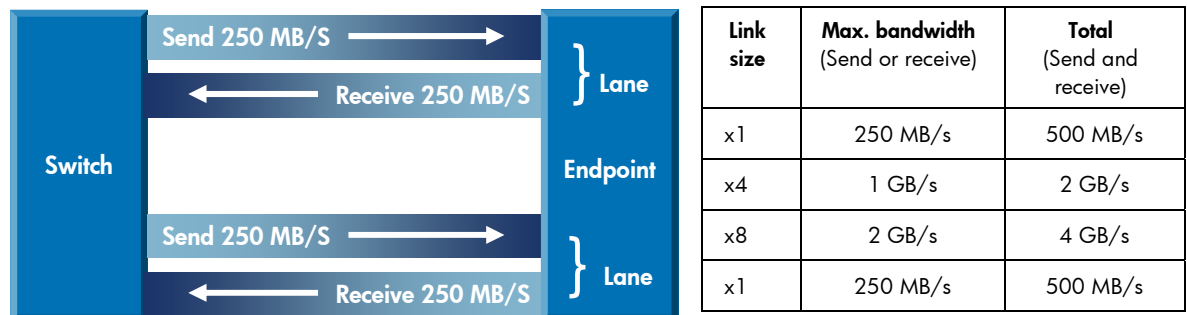
I/O technologies

ProLiant 300-series G6 servers incorporate PCI Express, Serial-Attached SCSI (SAS), and Serial ATA (SATA) I/O technologies. This server architecture lets administrators add PCI Express-compliant expansion cards to the system. SAS is a serial communication protocol for direct-attached storage devices such as SAS and SATA hard drives.

PCI Express technology

The PCI Express (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. ProLiant 300-series servers with AMD processors support PCIe 1.0 slots which have a signaling rate of 2.5 Gb/s per direction per lane. After accounting for 20% serializing/deserializing encoding overhead, the resulting effective maximum bandwidth is 2 Gb/s (250 MB/s) per direction per lane (Figure 3).

Figure 3. PCIe data transfer rates



A PCIe 2.0 device can be used in a PCIe 1.0 slot. For best performance, however, each card should be used in a slot that supports its logical link size (Table 1).

Table 1. 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 and SAS/SATA technology

Present generation Smart Array controllers use SAS technology, 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 higher signaling rates.

The latest Smart Array controllers are compatible with SATA technology and include the following features to enhance performance and maintain data availability and reliability:

- SAS and SATA compatibility — The ability to use either SAS or SATA hard drives lets administrators deploy drive technology that fits each computing environment. HP Smart Array controllers can manage both SAS arrays and SATA arrays. Smart Array configuration utilities help administrators configure arrays correctly so that data remains available and reliable.
- SAS wide port operations — Wide ports contain four single lane (1x) SAS connectors and the cabling has all four lanes bundled together. SAS wide ports allow balanced SAS traffic distribution across the links for enhanced performance. In addition, wide ports provide redundancy by tolerating up to three physical link failures while maintaining the ability to communicate with the disk drives. The most common use of these wide links is to a JBOD or to an internal server expander connecting to large numbers of drives. No special configuration is required for this functionality.
- SAS expanders — Low-cost, high-speed switches called expanders can combine multiple single links to create wide ports and increase available bandwidth. SAS expander devices also offer higher system performance by expanding the number of hard drives that can be attached to an HP Smart Array controller. SAS expanders are an aggregation point for large numbers of drives or servers providing a common connection. By cascading expanders, administrators can chain multiple storage boxes together. For more information on the HP SAS Expander Card, go to <http://h18004.www1.hp.com/products/servers/proliantstorage/arraycontrollers/sas-expander/index.html>.

SAS-2 standard

The second-generation SAS (SAS-2) link speed⁵ of 6 Gb/s is double the SAS-1 transfer rate. SAS-2 link speeds require SAS-2 compliant hard drives. SAS-2 eliminates the distinction between fanout and edge expanders by replacing them with self-configuring expanders. SAS-2 enables zoning for enhanced resource deployment, flexibility, security, and data traffic management.

SAS-2 connections have the potential to deliver peak raw data bandwidth of up to 600 MB/s per physical link in each direction. SAS-2 devices are capable of sending and receiving data simultaneously across each physical link, which is known as full duplex. When effectively implemented, full duplex, 6 Gb/s SAS connections can deliver peak raw data bandwidth of up to 1200 MB/s between the controller and storage device. It is important to note that the SAS-2 data bandwidths described here are theoretical speeds identified by the SAS-2 standard. Real world performance will be affected by the performance of storage devices attached to the SAS-2 connection.

Smart Array controllers, with releases beginning in the first quarter of 2009, incorporate SAS-2 connections. The SAS-2 standard is compatible with both Serial SCSI and Serial ATA protocols for communicating commands to SAS and SATA devices. SAS-2 compliant controllers are fully compatible with 1.5 Gb/s and 3 Gb/s SATA technology.

For the most up-to-date listing of HP Smart Array controllers that support the SAS-2 specification, see the Smart Array controller matrix: <http://www.hp.com/products/smarrays>.

Battery backed write cache

The BBWC is required for capacity expansion (adding one or more physical disks to an existing array). The controller recalculates parity and balances the data across all the disks. During the

⁵ Serial Attached SCSI-2 (SAS-2) is an American National Standards Institute (ANSI) standard from the INCITS T10 Technical Committee on SCSI Storage Interfaces. SAS-2 is the successor to SAS-1.1 and SAS-1.

expansion, the BBWC preserves data and logical structures on the array. The HP 650 mAh P-Series battery extends battery life up to 48 hours before recharging is necessary.

Zero Memory RAID

Using Zero Memory RAID (ZMR), administrators can create a RAID 0-1 configuration without additional memory. ZMR uses memory embedded in the controller, approximately 1K in size, and supports limited configurations. ZMR supports up to eight drives in Zero Memory Mode, or seven drives and one tape drive. ZMR does not include caching; however, all systems can be upgraded to a BBWC memory module that can significantly increase performance.

Present generation Smart Array controllers support ZMR for internal, direct connections only. This includes the Smart Array P410 supported in the ProLiant DL385 G6 server. Modular Smart Array (MSA) products are not supported in ZMR mode.

HP Smart Array compatibility

HP Smart Array controllers are modular solutions with a common form factor, hardware, and firmware. The present generation of SAS/SATA based Smart Array controllers are compatible with ProLiant 300-series G6 servers. Zero Memory RAID (ZMR) is available as an entry level hardware based RAID, but users have the option to choose the cache size and to include the battery back write cache (BBWC). These options provide upgradeability from ZMR to 512 BBWC.

Smart Array Advanced Pack

HP Smart Array Advanced Pack (SAAP) firmware provides advanced functionality within Smart Array controllers. This firmware further enhances performance, reliability, and data availability. SAAP is hosted on the Smart Array controller hardware firmware stack. It can be enabled beginning with the new generation of Smart Array controllers released in the first half of 2009.

SAAP requires a license key for activation. After activation, administrators can use several standard capabilities:

- RAID 6 with Advanced Data Guarding (ADG) protects against failure of any two drives. It requires a minimum of four drives, but only two will be available for data. ADG can tolerate multiple simultaneous drive failures without downtime or data loss and is ideal for applications requiring large logical volumes, because it can safely protect a single volume of up to 56 disk drives. RAID ADG also offers lower implementation costs and greater usable capacity per U than RAID 1.
- RAID 60 allows administrators to split the RAID storage across multiple external boxes. It requires a minimum of eight drives, but only four will be available for data.
- Advanced Capacity Expansion (ACE) automates higher capacity migration using capacity transformation to remove logical drives by shrinking and then expanding them online. Standard drive migration and expansion remain unchanged.
- Mirror Splitting and Recombining in Offline Mode breaks a RAID 1 configuration into two RAID 0 configurations. This is similar to a scaled down rollback functionality that requires two disk drives.
- Drive Erase completely erases physical disks or logical volumes. This capability is useful when decommissioning, redeploying, or returning hard drives.
- Video On Demand Performance Optimization decreases latency and improves video streaming.

More information about SAAP is available at www.hp.com/go/SAAP.

NOTE:

At a minimum, a 256 MB cache and battery kit is required to enable the SAAP license key. SAAP is not available on Zero Memory Configurations.

Power and thermal technologies

HP engineers have developed a robust set of power and thermal technologies and components to manage power within ProLiant 300-series G6 servers. The following technologies improve power efficiency throughout the power delivery chain:

- Efficient power delivery
- Thermal sensors and fan control
- Dynamic Power Capping
- Processor Management technologies

Administrators can disable certain components and capabilities within ProLiant 300-series G6 servers or reduce capabilities to bring the components to a lower power state.

Efficient power delivery

Oversized and lightly loaded power supplies do not run as efficiently as those that are heavily loaded. ProLiant G6 servers use common-slot, “right-size” power supplies and highly efficient DC power regulators to deliver significantly higher power efficiencies than previous ProLiant servers.

Common-slot power supplies

The HP G6 common slot power strategy provides power supply commonality across ProLiant G6 server lines. HP reduced the number of power supply designs, which reduces the number of spares that customers must keep in the data center.


By using “right-sized” power supplies, customers can closely meet power needs and avoid oversubscribing power requirements. HP “right-sized” power supplies have achieved efficiency ratings of up to 92%. With the exception of the 1200W power supplies, they meet Climate Savers Gold requirements. The 1200W and 1200W 48VDC power supplies meet Climate Savers Silver requirements.

- 460W AC up to 92% efficiency
- 750W AC up to 92% efficiency
- 1200W AC up to 90% efficiency
- 48Vdc 1200W up to 90% efficiency

The power loading efficiency curve for the 750 W power supply shown in Figure 4 illustrates the high levels of power efficiency achieved by power supplies used in ProLiant G6 servers.

The HP Power Advisor is available at the following link to help users define power supplies that will best meet their needs: <http://h71028.www7.hp.com/enterprise/cache/600307-0-0-121.html>.

Figure 4. Power/Efficiency curve for the 750 W HP power supply



Increasing power efficiency in redundant power operation

With ProLiant G6 servers, customers can select a power supply operation mode for redundant power systems. Through the RBSU, administrators can select 'High Efficiency' mode or 'Balanced' mode. In Balanced mode, both power supplies provide power equally. This mode ensures full redundancy but can result in higher power consumption when power supplies are operating with reduced loads and lower power efficiency. The High Efficiency mode means the system will only use one power supply until system load exceeds a certain threshold. The second power supply stays online maintaining redundancy but does not supply power until needed. Either selection still provides full power redundancy.

Voltage regulation

Voltage regulators convert the 12V DC supplied from the server power supply into a variety of lower voltages used by the different system components. HP has developed voltage regulators that maintain greater than 90% efficiency over a broad range of power requirements. The net result is near an 8% gain in DC power efficiency, which results in almost 10% efficiency gain in AC input power. These efficiency gains come with no loss in performance and require no configuration by the user.

Thermal sensors and fan control

Thermal sensors and fan control in the ProLiant AMD-based 300-series G6 servers remain unchanged from the previous generation G5p servers⁶. Sensor data from the processor, hard drive enclosure, and mother board provide thermal information to firmware that controls fan speed. Fan speed is

⁶ The 'G5p' indicates performance and capacity enhancements over the G5 platforms. Those additions include higher performance Quad-Core AMD Opteron processors, increased memory DIMM slot count by 2x and high efficiency options such as processors, memory and power supplies.

altered to increase cooling based on the sensor data and location. If one fan fails, all the other fans default to high speed in order to assure the server remains within thermal specifications.

APML Remote Power Management Interface

Multiple processor-based thermal sensors are new for the AMD Opteron six-core 2400-series processors. The APML Remote Power Management Interface from AMD allows users to remotely monitor and control P-state limits from a systems management device such as iLO 2. Thermal readings indicating the hottest part of the processor provide more accurate information about the operational limits of the processor. These processor-based sensors are included in the data processed by the fan controller.

Dynamic Power Capping

All ProLiant 300-series G6 servers incorporate Dynamic Power Capping, which allows administrators to limit maximum power use and more accurately allocate power within a server. Dynamic Power Capping can bring a server experiencing a sudden increase in workload back under its power cap in less than one-half second, preventing any surge in power demand that could trip a typical data center circuit breaker. Dynamic Power Capping has been designed and tested (at 50 degrees C and 150% overload) to ensure that it can prevent tripping circuit breakers that have a specified trip time of three seconds or longer.

This ability to keep server power consumption below the power cap in real time means that Dynamic Power Capping can be used to effectively plan and manage both electrical provisioning and cooling requirements in the data center. An administrator can electrically provision a PDU or a rack to less than the indicated full faceplate power rating.


To implement Dynamic Power Capping, the iLO 2 management processor works in conjunction with a power microcontroller to measure and control power use. When enforcing the Dynamic Power Cap, the power microcontroller keeps the processor's performance and power use under the set cap (see Figure 5). Administrators can set Dynamic Power Caps for individual servers from the iLO 2 Advanced user interface and for multiple rack-mount servers from the power management module within HP Insight Control Environment (ICE).

HP Dynamic Power Capping is OS-independent and functions with all operating systems and software applications. HP Dynamic Power Capping will continue to function even if the software fails because it uses OS-independent hardware.

Since Dynamic Power Capping can affect server performance if set too aggressively, HP recommends that Dynamic Power Caps be set at values that match or exceed the highest observed power consumption over a representative server workload sample.

For a more detailed explanation of HP Dynamic Power Capping, see the technology brief "HP Power Capping and HP Dynamic Power Capping for ProLiant servers" available at <http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01549455/c01549455.pdf>.

Figure 5. Rapid response of Dynamic Power Capping avoids circuit breaker trips



Managing processor technologies

Built into AMD Opteron processors is the AMD-P suite of power management technologies, including AMD PowerNow™ Technology, and AMD Power Cap Manager™.

PowerNow! Technology

AMD PowerNow! Technology with Independent Dynamic Core technology and Dual Dynamic Power Management™ is technology that allows the processors to run dynamically at different frequencies and voltages depending upon CPU computing demand. As a result, PowerNow! can lower server power consumption without compromising performance.

AMD PowerNow! can be enabled on ProLiant AMD-based 300-series G6 servers through the BIOS-controlled Dynamic Mode of Power Regulator for ProLiant, which does not require an OS driver.

Power Cap Manager

Power Cap Manager allows administrators to set a fixed limit on a server's processor power consumption. In a multi-core environment, administrators control the P-state of individual cores. The caveat is that a single voltage is supplied to all cores in the processor even if different P-states are requested by the user. If an administrator requests different P-states for different cores in the processor, the actual P-state will equal that of the highest voltage required among the selected P-states. Administrators can control this functionality through HP Dynamic Power Capping, accessed from iLO 2 advanced menus or the HP ICE management suite.

AMD Core Select and HP Core Disable

AMD Core Select lets administrators use the RBSU to select the number of software-visible cores per CPU (minimum of one, up to full number supported). Operating systems and applications can

recognize the reduced core count so that fewer software licenses may be required. By reducing core count recognition, Core Select has the potential to reduce software licensing costs. Also, applications can benefit from increased memory bandwidth and cache per core. This feature is offered on HP ProLiant platforms that are configured with an AMD 2400 or 8400 series processors and the latest System ROM supporting this feature. For HP ProLiant platforms using earlier models of the AMD processor, the RBSU offers a similar core disable option based on HP Core Disable technology. With this option, the user is able to configure the server such that all cores are enabled, half the cores are enabled, or one core is enabled per processor socket. Please consult the HP ROM-Based Setup Utility User Guide for more details on this option.

<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00191707/c00191707.pdf>

Security

The Trusted Platform Module™ (TPM) and Microsoft® BitLocker® technology are supported in all ProLiant 300-series G6 servers by means of the Trusted Platform Module option kit.

Trusted Platform Module

The Trusted Platform Module v1.2, supported on ProLiant G6 servers, is a microcontroller chip that can create, securely store, and manage artifacts such as passwords, certificates, and encryption keys that are used to authenticate the server platform. The TPM 1.2 chip provides a unique Endorsement Key (EK) and a unique Storage Root Key (SRK). It provides data encryption and uses RSA, SHA-1, RNG cryptographic functions to provide access protection, OS level protection, and stolen disk protection.

The TPM 1.2 chip can also store platform measurements to help ensure that the platform remains trustworthy. TPM enables Microsoft BitLocker, part of Windows Server 2008.

For more information about TPM, go to <http://www.hp.com/go/TPM>

BitLocker Drive Encryption

Microsoft BitLocker Drive Encryption (BitLocker) is a data protection feature available in Windows Server 2008. BitLocker uses the enhanced security capabilities of TPM version 1.2 to protect data and ensure that a server running Windows Server 2008 has not been compromised while the system was offline.

Implementing BitLocker requires the following:

- The Master Boot Record (MBR), a small, encrypted system partition of approximately 50 MB to contain boot utilities
- TPM version 1.2
- Trusted Computing Group (TCG) compliant firmware including support of “Static Root of Trust”
- Two NTFS partitions on the boot drive

During the boot process, the TPM will not release the encryption key until completing a comparison of OS configuration information, or hash, with an earlier snapshot of the same data. If any part of the hash is compromised (for example by introduction of malicious code), the TPM ensures that the volume encryption key is never released.

Systems management and monitoring

HP offers management tools to program and control all aspects of the dynamic server environment. HP ProLiant Onboard Administrator powered by the iLO 2 management processor provides remote management with other core-embedded management functions to simplify setup, health monitoring, power and thermal control, and remote administration. The HP ICE suite provides a foundation for deploying, managing, optimizing, and controlling the entire server environment from any location. HP Insight Dynamics – VSE suite for ProLiant delivers comprehensive functions for optimizing and balancing resources and workloads in real time.

HP ProLiant Onboard Administrator powered by iLO 2

HP ProLiant Onboard Administrator powered by iLO 2 provides core management functionality including server power control, ROM-based setup, embedded health and management logs, power consumption level, SSH and SSL support, and flexible access through a dedicated or shared network port. With ProLiant G6 servers, HP introduces support for viewing CPU configuration and status, along with fan and power supply status. HP plans to introduce the ability to view configuration and status of memory, and array controllers. The iLO 2 management processor is incorporated into all ProLiant 200-series and above servers.

ProLiant Onboard Administrator management capabilities can be extended with HP iLO 2 Advanced. Additional capabilities include high performance graphical and text console, diagnostic capabilities (including automatic video footage of the server's last boot and last fault), global team collaboration for up to four users, 24-hour power measurement, and single server Dynamic Power Capping. HP Dynamic Power Capping, enabled through iLO 2 Advanced, is available for all ProLiant ML and DL 300-series servers. For additional details, see the HP product support matrix at <http://www.hp.com/servers/ilo/supportedservers>.

HP Insight Control Environment

HP recommends the Insight Control Environment (ICE) management suite to manage the infrastructure in every HP G6 server. ICE provides infrastructure management functions, including a complete set of lifecycle management tools.

For more information about ICE, go to www.hp.com/go/ice

HP Insight Dynamics – VSE

HP Insight Dynamics – VSE suite for ProLiant servers helps users continuously analyze and optimize adaptive server infrastructure. HP Insight Dynamics – VSE delivers four key capabilities for HP ProLiant 300-series G6 servers:

- Continuous capacity and power planning
- Balanced physical and virtual resources
- Cost-effective availability
- Consistent infrastructure provisioning

For more information about HP Insight Dynamics – VSE suite, go to <http://h18013.www1.hp.com/products/servers/management/insightdynamics.html>

OS support

HP performs extensive testing, qualification, and certification of the latest server operating systems on ProLiant servers to ensure maximum performance and reliability. HP resells and provides full service

and support for Microsoft Windows operating systems, Red Hat Linux® subscriptions, Novell SUSE Linux subscriptions, Sun Solaris subscriptions, Citrix XenServer, and VMware hypervisors. The latest information regarding support and deployment can be found online at www.hp.com/go/ossupport.

Summary

The HP ProLiant 300-series G6 servers equipped with AMD processors help administrators increase business performance, lower power costs, and manage their server hardware more easily. To improve performance, these servers use the AMD Opteron 2400-series processor technologies with integrated memory controllers and DDR-2 memory. The latest Smart Array controllers and firmware improve RAID performance compared to the previous generation of controllers. Using HP common power supplies is another means for customers to refine and constrain server power based on their data center requirements. ProLiant Onboard Administrator powered by iLO 2, HP Insight Control Environment management suite, and HP Insight Dynamics – VSE all facilitate server management. These servers can be easily deployed with SmartStart, Rapid Deployment Pack, and RBSU.

For more information

For additional information, refer to the resources listed below.

Resource description	Web address
Dynamic Power Capping TCO and Best Practices White Paper	http://h71028.www7.hp.com/ERC/downloads/4AA2-3107ENW.pdf
HP Insight Control Environment	www.hp.com/go/ice
HP Network Adapters for ProLiant DL and ML Servers	http://media.hpvtc.veplatform.com/content/HP_Network_Adapters_for_ProLiant_DL_Family_data_sheet_1237839147.pdf
HP ProLiant DL385 G6 Server QuickSpecs	http://h18004.www1.hp.com/products/quickspecs/13360_na/13360_na.html
HP ProLiant DL385 G6 server user guide	http://bizsupport2.austin.hp.com/bc/docs/support/SupportManual/c01759017/c01759017.pdf
HP ROM-Based Setup Utility User Guide	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00191707/c00191707.pdf
HP iLO 2 product information	www.hp.com/go/ilo
HP SAS and SATA technology	www.hp.com/go/serial
HP Smart Array Advanced Pack	http://h18004.www1.hp.com/products/servers/proliantstorage/arraycontrollers/smartarray-advanced/index.html
HP Smart Array controllers	www.hp.com/products/smartarray
ISS Technology Communications briefs: "HP Power Capping and HP Dynamic Power Capping for ProLiant servers"	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01549455/c01549455.pdf
ISS Technology Communications briefs: "Memory technology evolution: an overview of system memory technologies"	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00256987/c00256987.pdf
ISS Technology Communications briefs: "Serial Attached SCSI storage technology"	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01613420/c01613420.pdf

Call to action

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

© 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, Windows, and BitLocker are U.S. registered trademarks of Microsoft Corporation.

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

AMD, Opteron, HT Assist, and Direct Connect are trademarks or registered trademarks of Advanced Micro Device Corporation or its subsidiaries in the United States and other countries and is used under license.

TC090909TB, September 2009

