



# Technology and architecture of HP ProLiant AMD-based 100-series G6 (Generation 6) servers



Technology brief

Abstract.....	2
Introduction.....	2
Processor technology.....	2
Direct Connect architecture 2.0 .....	2
HyperTransport technology .....	3
HT Assist .....	4
Memory technologies .....	4
Double data rate memory .....	5
Determining memory bus speed .....	5
Memory management technologies .....	5
Memory bank interleaving .....	5
Memory channel interleaving .....	6
Memory node interleaving .....	6
8X error correction .....	6
I/O technologies .....	6
PCI Express technology .....	7
HP Smart Array and SAS/SATA technology .....	7
SAS-2 standard .....	8
Battery backed write cache .....	8
Zero Memory RAID .....	8
HP Smart Array controller compatibility .....	9
Smart Array Advanced Pack .....	9
Power and thermal technologies .....	10
Efficient power delivery .....	10
Industry standard power supply .....	10
Thermal sensors and fan control .....	10
PowerNow! technology .....	10
Server management .....	11
Remote management .....	11
HP Lights-Out 100i .....	11
LO100i shared and dedicated networks .....	12
IPMI 2.0 and DCMI 1.0 .....	13
OS support .....	13
Summary .....	14
For more information .....	15
Call to action .....	15

# Abstract

This technology brief describes the key technologies implemented in HP ProLiant 100-series G6 servers based on AMD™ processors. As of this writing, the AMD-based 100-series G6 server platforms are limited to the ProLiant DL165. 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 AMD-based 100-series G6 server platforms are high performance, low cost, ultra-dense rack server nodes for compute and I/O intensive environments. ProLiant AMD 100-series G6 servers include these new technologies:

- AMD Opteron™ six-core 2400 Series processors, including high efficiency (HE) processor options
- The ProLiant Onboard Administrator Powered by Lights-Out 100i (LO 100i)
- Input/output (I/O) technologies such as PCI Express (PCIe) and faster Smart Array controllers that incorporate common form factor components
- High efficiency power supplies to provide the required amount of power and improve power efficiency

The technologies discussed in this paper are implemented in the ProLiant DL165 G6 server.

## Processor technology

The HP G6 AMD-based 100-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™, and 8x ECC error correction. The Six-Core processors fit into the same socket 1207 architecture as previous four-core AMD Opteron processors. 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.

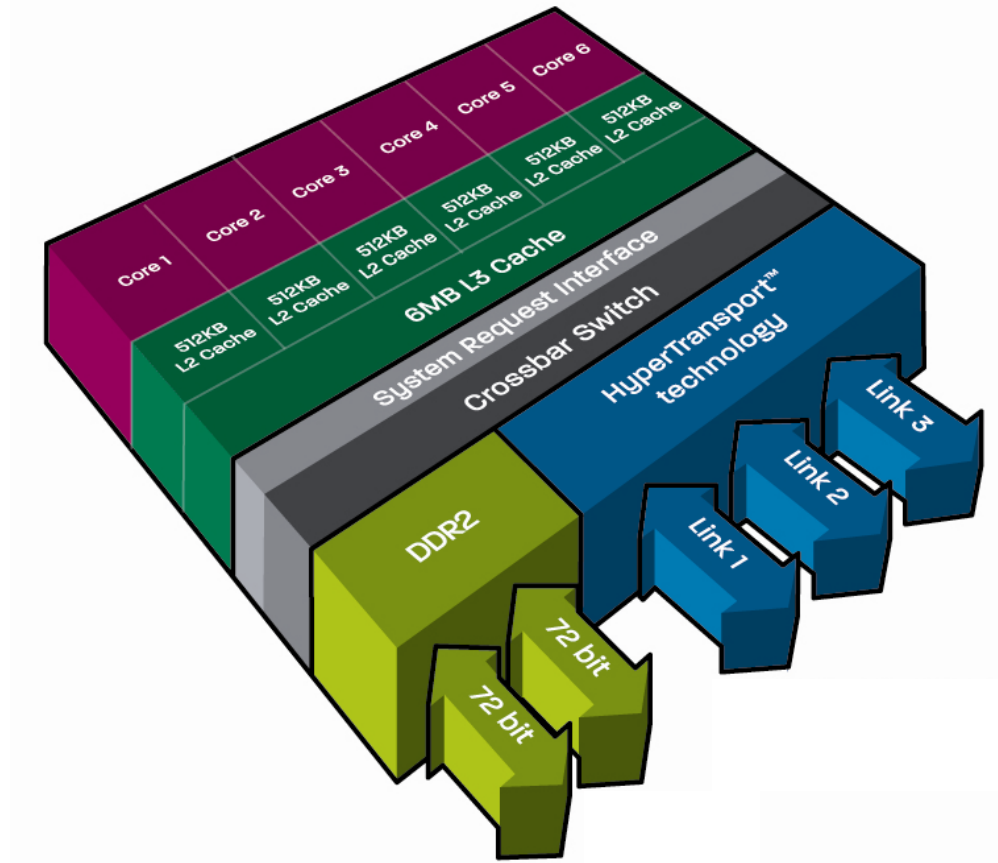
### Direct Connect architecture 2.0

Instead of 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 which provide superior memory and I/O capability, near native virtualization performance, and a range of power bands<sup>1</sup> that place a priority on low power consumption.

---

<sup>1</sup> 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](http://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 (see Figure 2) that is designed to connect the processors directly and to connect each processor to its dedicated memory banks, as well as to other I/O chipsets.<sup>2</sup> Compared to a shared, parallel front-side bus, HyperTransport has the advantages of having 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/s<sup>3</sup>) to a maximum of 2.6 GHz (5.2 GT/s) in each direction.

<sup>2</sup> 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.

<sup>3</sup> 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



HyperTransport can be configured through the Basic Input/Output System (BIOS) and BIOS Setup Utility (BSU). In the BSU 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 uses:

- 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 100-series servers with AMD processors support PC2 Registered double data rate (DDR-2) DIMMs. These 100-series G6 servers use 1-GB, 2-GB, 4-GB, or 8-GB registered DDR-2 DIMMs across 8 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 megatransfers per second (MT/s), which translates to a memory bandwidth of 6400 MB/s per DIMM<sup>4</sup>.

## Determining memory bus speed

Depending on the quantity and speed of installed DIMMs, memory bus speeds on the 100-series G6 servers can be 533-MHz, 667-MHz, or 800-MHz. The memory bus speed is determined by the BIOS and BSU 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 BSU 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- The following BSU 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 memory interleaving technologies are supported on the ProLiant AMD 100-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.

---

<sup>4</sup> 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>

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 four-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 for 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 BSU. Node interleaving can only be configured if the memory footprint for both processors is the same.

## **8X error correction**

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

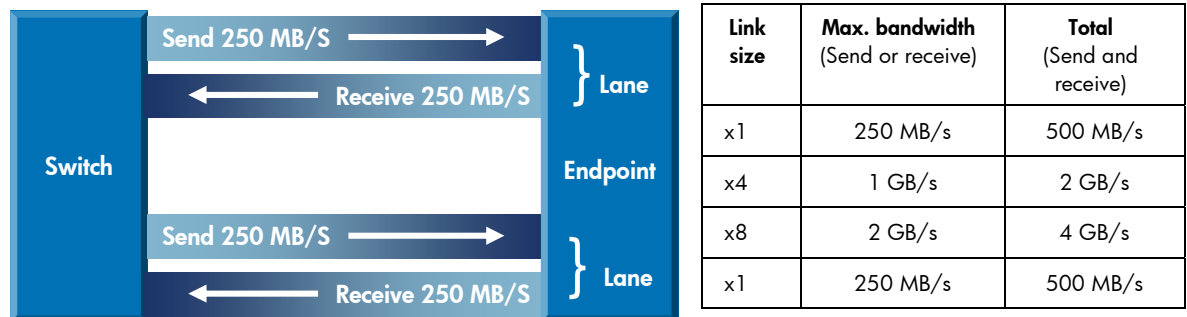
## **I/O technologies**

ProLiant 100-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 100-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 speed<sup>5</sup> 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. Actual 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/smartarray>

### **Battery backed write cache**

The battery backed write cache (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 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.

---

#### **NOTE:**

In the ProLiant DL165 G6 server, the Smart Array P212 controller cache size is limited to 256 MB

---

### **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

---

<sup>5</sup> 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.

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.

ZMR is supported on present generation Smart Array controllers for internal, direct connections only. This includes the Smart Array P212 supported in the ProLiant DL165 G6 server. Modular Smart Array (MSA) products are not supported in ZMR mode.

### **HP Smart Array controller 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 AMD 100-series G6 servers. The ProLiant DL165 G6 server is available as a non-hot plug SATA hard drive model, or as a hot plug SATA/SAS hard drive model. The non-hot plug SATA model uses an HP embedded SATA controller and does not support RAID configurations. The hot plug SATA/SAS model uses an HP Smart Array P212 controller. Zero Memory RAID (ZMR) is available as an entry level, hardware based RAID solution, but users have the option to choose the cache size and to include the BBWC. These options facilitate upgrades 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, redeploing, 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](http://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 100-series G6 servers. The following technologies improve power efficiency throughout the power delivery chain:

- Efficient power delivery
- Thermal sensors and fan control
- Processor management technologies

Administrators can disable certain components and capabilities within the ProLiant 100-series G6 servers or reduce capabilities to bring the components to a lower power state.

## Efficient power delivery

Power supplies for the ProLiant AMD DL165 G6 server are not hot-pluggable. Common slot power supplies are only an option in those G6 platforms supporting the common slot architecture. All 100-series G6 servers use highly efficient power supplies and DC power regulators to deliver significantly higher power efficiencies.

### Industry standard power supply

The ProLiant AMD-based 100-series G6 servers are not configured for common slot power supplies. Instead, ProLiant DL165 G6 servers are equipped with the High-efficiency 500W multi-output power supply that improves performance-watt ratio. This power supply is neither hot pluggable nor auto-switching.

## Thermal sensors and fan control

Thermal sensors and fan control in the ProLiant AMD based 100-series G6 servers remain unchanged from the previous generation G5p servers.<sup>6</sup> Sensor data from the processor, hard drive enclosure, and mother board provide thermal information to firmware that controls fan speed. Fan speed is 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.

## 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 100-series G6 servers through the BIOS-controlled Dynamic Mode of Power Regulator for ProLiant, which does not require an OS driver.

---

<sup>6</sup> The G5p platforms included 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.

# Server management

ProLiant ML and DL 100-series G6 server users each have different computing requirements. Consequently, the way in which customers manage and control servers can differ. With these requirements in mind, this section examines the following management topics:

- HP ProLiant Onboard Administrator Powered by Lights-Out 100i remote management and control (LO100i)
- Intelligent Platform Management Interface (IPMI) 2.0 and Data Center Management Interface (DCMI) 1.0 Standards

Some of these technologies are new tools for the ProLiant 100-series G6 servers, while others have been available with previous generations of ProLiant 100-series servers.

## Remote management

LO100i is included in all ProLiant 100-series G6 servers. LO100i represents the core embedded management functions in ProLiant ML, DL, and SL100-series G6 servers. LO100i works in concert with HP Systems Insight Manager (SIM), BSU, and Option ROM Configuration for Arrays (ORCA) to provide remote management, deployment, and control functions without additional software. This functionality can be accessed locally through the BSU, or remotely with a web browser through HP SIM. Additional software functions can be added with the HP Lights-Out 100i Advanced Licenses which include Virtual KVM (remote graphical console) and virtual media capabilities.

### HP Lights-Out 100i

HP LO100i Remote Management is hardware and firmware that provides remote server access and control capabilities through an Ethernet interface. The HP LO100i management interface is active even when the OS is not operating. The LO100i management processor obtains its power from the auxiliary power plane of the server, so it is available as long as the server is plugged into an active power source. HP LO100i Remote Management is compatible with industry standards including IPMI 2.0 for hardware health, DCMI 1.0, and Secure Sockets Layer (SSL) and Secure Shell (SSH) technology for secure communications over public and private local area networks. HP LO100i is fully accessible using popular web browsers including Microsoft Internet Explorer 6.0 (Microsoft Windows clients), Firefox 1.5 (Windows and Linux clients), and Mozilla 1.7.12 (Windows and Linux clients). HP LO100i is also accessible using System Management Architecture for Server Hardware Command line Protocol (SMASH CLP) for Telnet and SSH sessions.

---

#### **NOTE:**

ProLiant 100-series G6 servers do not support LO100i Select and LO100c

---

LO100i Advanced Pack contains all the features of LO100i Standard. Additional LO100i Advanced Pack features can be obtained through an optional license key. Table 3 shows the differences in functionality between LO100i Standard, which comes with all 100-series G6 servers, and LO100i Advanced Pack.

**Table 3.** LO100i functionality

Features	LO100i Standard with every ProLiant 100-series server	LO100i Advanced (license upgrade options)
<b>Technical Support and Upgrade Licensing</b>		Yes - G6 only
<b>Flatpack and electronic key delivery</b>		Yes
<b>Host access to IPMI environment HW status</b>	Yes	Yes
<b>Unencrypted browser for power, SEL, health, and key activation</b>	Yes	Yes
<b>Shared and dedicated network port *</b>	Yes	Yes
<b>DCMI 1.0</b>	Yes - G6 only	Yes - G6 only
<b>SSL &amp; SSH security (setup in factory)</b>	Yes - Included with G6	Yes
<b>DNS registered names **</b>	Yes	Yes
<b>License manager support</b>	Yes - G6 only	Yes - G6 only
<b>IPv6 support (coexistence at launch)</b>	Yes - G6 only	Yes - G6 only
<b>Virtual KVM</b>		License Upgrade
<b>Virtual media (CD/DVD, floppy, ISO Image files)</b>		License Upgrade
<b>Power Capping</b>		Yes - DL1000 only

\* Shared and dedicated network port is currently an optional feature on some ProLiant 100-series G6 servers and may be purchased as an option.

\*\* The DNS registered names require high speed network ports to accommodate the large packet sizes associated with this feature. High speed network ports are included on ProLiant DL series servers. The shared low speed network ports on the ProLiant ML series servers do not support DNS registered names.

LO100i includes the following feature support for the AMD 100-series G6 server platforms:

- Host access to in-band IPMI 2.0 features supported by IPMI-aware operating systems
- DNS Registration —The LO100i device comes with default host names and will automatically register with the DNS if DHCP is enabled

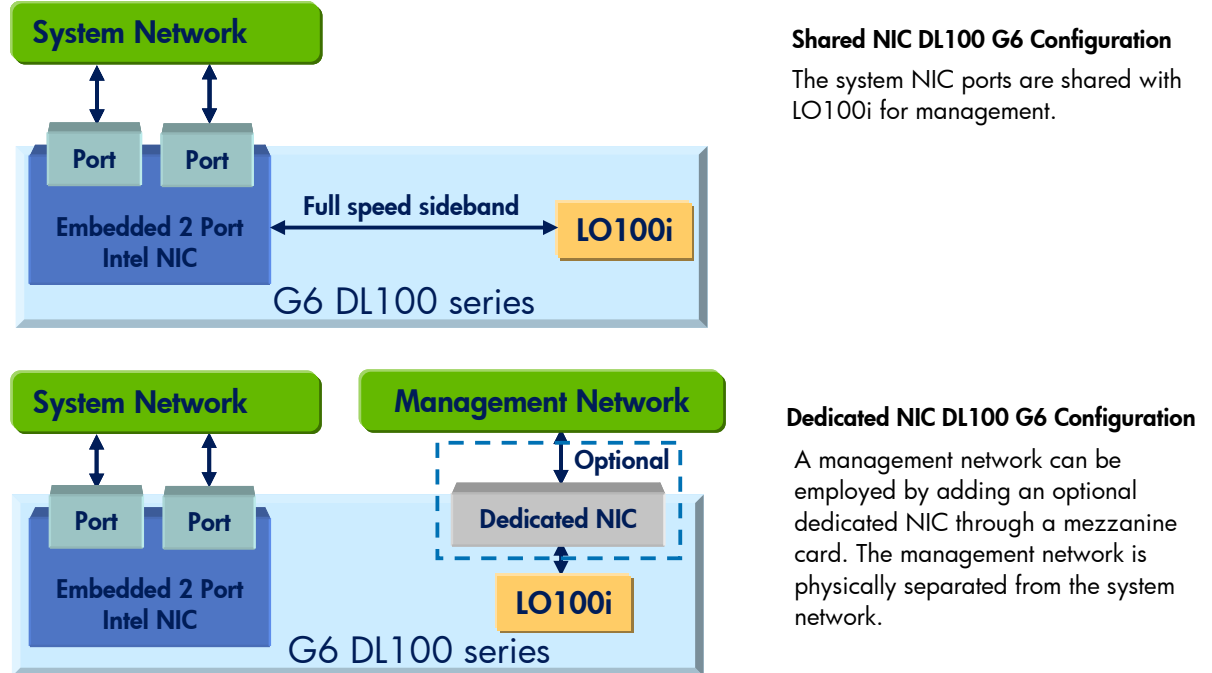
### **LO100i shared and dedicated networks**

In the ProLiant DL100-series G6 servers, LO100i supports a full-speed shared Ethernet port and a dedicated Ethernet port. LO100i and the server share the full-speed Ethernet port, utilizing the system network for both. Since the connection is full speed, it supports Graphic Remote Console and virtual media.

The dedicated Ethernet port is provided for DL 100-series G6 servers by means of an optional mezzanine card and enables a separate management network. LO100i implementation varies depending on whether the 100-series G6 server is a DL, SL, or an ML platform.

Figure 4 shows the two implementations available for the DL and SL100-series G6 servers.

**Figure 4.** LO100i sideband architecture for DL and SL100-series G6 servers



## IPMI 2.0 and DCMI 1.0

ProLiant 100-series G6 servers and LO100i are architected with IPMI 2.0 and DCMI 1.0 standards so that customers in heterogeneous environments can manage these servers with either industry standard.

The following are basic compliance mandates-

- The implementation of all mandatory IPMI 2.0 and DCMI 1.0 in-band and out-of-band commands
- Reliable Local and Remote Power on/off/reset through IPMI Chassis commands
- Per IPMI 2.0, console redirection over telnet or SSH
- Server identification by device ID, globally unique identifier (GUID), asset tag, and chassis ID
- Accurate System Event Logging using IPMI
- Reliable in-band keyboard controller style (KCS) interface and out-of-band LAN interface

For more information on HP Lights-Out 100i Remote Management, go to [www.hp.com/go/LO100](http://www.hp.com/go/LO100).

## OS support

HP performs extensive testing, qualification, and certification on the latest server operating systems 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 at [www.hp.com/go/ossupport](http://www.hp.com/go/ossupport).

## Summary

The development and engineering innovations associated with the HP ProLiant AMD100-series G6 servers help administrators lower power costs and increase business performance. These innovations include the AMD Opteron 2400-series processor technologies, improved thermal control, common power supplies, DDR-2 memory, and PCIe 2.0-compliant modular Smart Array controllers. The AMD 100-series G6 servers possess superior technology differentiation and management capabilities with ProLiant OnBoard Administrator Powered by Lights-Out 100i.

## For more information

For additional information, refer to the resources listed below.

Resource description	Web address
SAS and SATA technology	<a href="http://www.hp.com/go/serial">www.hp.com/go/serial</a>
Smart Array Advance Pack	<a href="http://h18004.www1.hp.com/products/servers/proliantstorage/arraycontrollers/smartarray-advanced/index.html">http://h18004.www1.hp.com/products/servers/proliantstorage/arraycontrollers/smartarray-advanced/index.html</a>
Smart Array controllers	<a href="http://www.hp.com/products/smartarray">www.hp.com/products/smartarray</a>
HP Network Adapters for ProLiant DL and ML Servers	<a href="http://h20195.www2.hp.com/V2/GetPDF.aspx/4AA0-6064ENW.pdf">http://h20195.www2.hp.com/V2/GetPDF.aspx/4AA0-6064ENW.pdf</a>
HP ProLiant DL165 G6 Server	<a href="http://h18004.www1.hp.com/products/quickspecs/13015_na/13015_na.html">http://h18004.www1.hp.com/products/quickspecs/13015_na/13015_na.html</a>
HP ProLiant Onboard Administrator Powered by LO100 (remote management)	<a href="http://www.hp.com/go/lo100">www.hp.com/go/lo100</a>
ISS Technology Communications briefs: Memory technology evolution: an overview of system memory technologies	<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>
ISS Technology Communications briefs: Serial Attached SCSI storage technology brief	<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>
ProLiant 100 Series Core Management Software	<a href="http://www.hp.com/go/coremanagement">www.hp.com/go/coremanagement</a>

## Call to action

Send comments about this paper to [TechCom@HP.com](mailto: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, Direct Connect, AMD PowerNow!, Dual Dynamic Power Management, and Independent Dynamic Core technology are trademarks or registered trademarks of Advanced Micro Device Corporation or its subsidiaries in the United States and other countries and are used under license.

