

HP ProLiant DL980 G7 server with HP PREMA Architecture

Technical overview

Table of contents

Executive Summary	2
Introduction	2
HP PREMA Architecture overview	2
A scalable 8-socket x86 server with HP PREMA Architecture	3
Balanced scaling with smart CPU caching	3
Addressing cache coherency and memory latency	4
Memory and I/O	6
Performance results	7
Self-healing resiliency	8
Redundant system fabric	8
Additional redundant components	9
High-speed link resiliency	10
Correction mechanisms	10
Mechanisms to contain failures	11
Breakthrough efficiency from ProLiant innovations	11
Consolidation	11
Integrated Lights-Out 3	12
HP Insight Software	12
High efficiency power supplies	12
Common components	13
Thermal Logic technologies	13
Benefits of ProLiant DL980 G7 with HP PREMA Architecture for applications	13
Conclusion	14
Appendix: Benchmark results	15
Glossary	16
For more information	17



Executive Summary

The HP ProLiant DL980 G7 is the newest member of HP's industry-leading scale-up x86 ProLiant family. HP designed this server to take full advantage of Intel's latest and most capable Xeon® processor, the Intel® Xeon® processor 7500/6500 series with Quick Path Interconnect (QPI). Building on the advanced performance and reliability capabilities of the Intel processor technology and alignment with industry standard operating systems such as Microsoft® Windows® Server and Linux, the DL980 G7 delivers an enhanced level of scale-up x86 performance, availability, and reliability. The ProLiant DL980 G7 is the first server to use the HP PREMA Architecture, incorporating a node controller design with smart CPU caching and redundant system fabric. HP, Intel, and server OS technologies combine in the DL980 G7 server to deliver the balanced scaling and self-healing resiliency needed to support your most demanding and data-intensive workloads. ProLiant innovations add breakthrough efficiency to enable you to control costs and focus more of your resources on service delivery.

This paper focuses on the ProLiant DL980 G7 server architecture and key technologies that contribute to its scalability, resiliency, and efficiency. It should help IT professionals understand the capabilities of the DL980 G7 to support mission-critical applications. For a summary of technical specifications for this server, see the *ProLiant DL980 G7 Server Data Sheet* at http://h18004.www1.hp.com/products/quickspecs/DS_00190/DS_00190.pdf. For detailed technical specifications see the server *QuickSpecs* at http://h18000.www1.hp.com/products/quickspecs/13708_div/13708_div.pdf.

Introduction

Data-intensive workloads, including business processing and business intelligence, and large scale consolidation and virtualization are placing today's enterprise application environments under tremendous stress. Challenges include:

- Inadequate scaling of compute capabilities. You need scalable solutions to effectively handle increasingly complex and demanding workloads and support exponential data growth over the life of your database infrastructure.
- Insufficient reliability. You need resilient solutions, particularly to handle large single-system databases and high-density virtualization. The ability to ensure availability can have a clear economic benefit. Downtime can result in revenue loss, damage to company reputation, and lower employee productivity.
- Increased management complexity and operating cost. Underutilized IT resources consume too much space, and cost too much to power, cool, maintain, administer, and service. You need efficient solutions to enable you to free up and redirect operational dollars into business innovations that will strengthen your competitive advantage.

HP PREMA Architecture overview

HP PREMA Architecture represents HP's technology direction for scale-up x86 servers. The PREMA Architecture is the design foundation for x86 servers that need to deliver more scalability, resiliency, and efficiency to meet the requirements of the most demanding, data-intensive workload environments as well as large scale consolidation and virtualization.

With PREMA Architecture HP addresses customer requirements that have exceeded the performance and capacities offered by 4-socket x86 systems. The architecture supports an appropriately balanced system with more processors, more memory, and more I/O than previous generation x86 systems

have provided. However, simply adding processors, memory, and I/O slots is not sufficient to achieve the needed scalability and resiliency.

When a system scales to a larger number of interconnected processors, the communication and coordination between the processors grow at an exponential rate, creating a system bottleneck. To solve this issue in our 8-socket x86 server, HP looked to the design of our higher-end mission-critical servers. At the core of the HP PREMA Architecture is a node controller ASIC, derived from technology powering the HP Integrity Superdome 2. The node controller enables two key functionalities: smart CPU caching and the redundant system fabric. These serve to reduce communication and coordination overhead as well as to enhance system resiliency.

HP has developed the HP PREMA Architecture to address emerging and future requirements for x86 systems in scale-up environments. Future systems using this architecture have the potential to scale up to 32 processor sockets, with proportionate increases in the amount of memory supported—accommodating larger DIMMs—and I/O expansion potential. Today, the ProLiant DL980 G7 based on PREMA Architecture scales to eight processors supporting up to 64 processor cores and up to 128 logical processors with hyper-threading enabled. The server’s processing capacity is balanced with 128 DDR3 DIMM slots for a maximum of 2TB of memory using 16GB DIMMs, and up to 16 I/O slots. This paper discusses the design and capabilities of the PREMA Architecture as implemented in the ProLiant DL980 G7.

A scalable 8-socket x86 server with HP PREMA Architecture

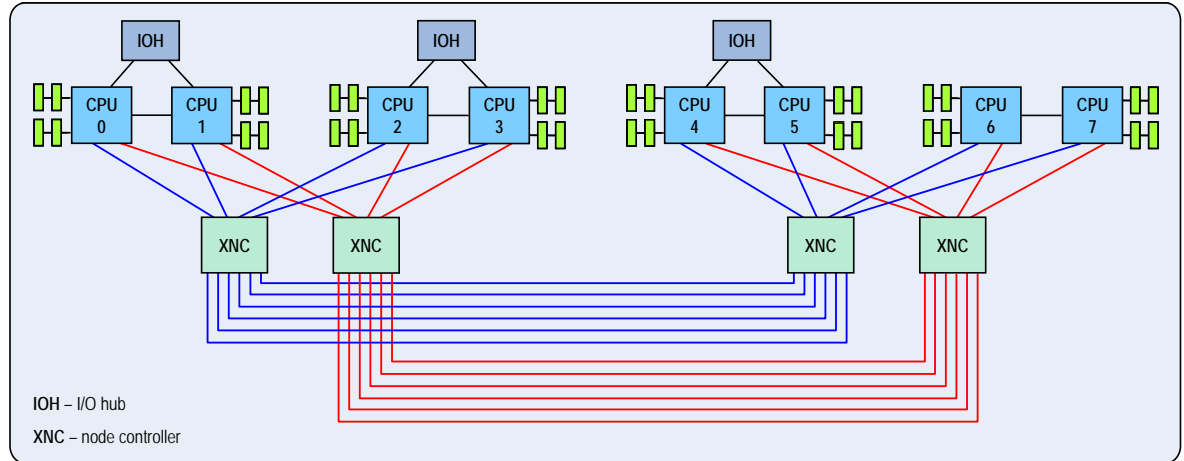
The HP PREMA Architecture in the ProLiant DL980 G7 server leverages HP’s years of experience designing mission-critical servers in RISC, EPIC, and UNIX environments and applies that knowledge and expertise to our design of the DL980 G7. As a result, the DL980 G7 provides significant benefits beyond basic 4-socket system scaling. PREMA Architecture makes use of a node controller with a directory cache to reduce memory latency and provide more efficient performance scaling and improved isolation of errors. PREMA Architecture also employs a high-speed HP fabric, leveraged from the HP Integrity Superdome 2 Crossbar Fabric, to interconnect the system, adding resiliency features used in mission-critical servers. These technology advances, along with the industry-leading innovations that ProLiant servers are known for, combine in the DL980 G7 to deliver balanced scaling, self-healing resiliency, and breakthrough efficiency. The ProLiant DL980 G7 provides capabilities more commonly associated with “big iron” systems while offering the economics associated with industry-standard x86 servers.

Balanced scaling with smart CPU caching

Server manufacturers can take one of two approaches in building an 8-socket system based on the Intel Xeon processor 7500/6500 series: using a “glueless” architecture in which the processor sockets are directly connected or using “glue” in the form of node controllers to connect sockets. When scaling to eight processors, the “glued,” node controller implementation provides performance benefits beyond those offered by a glueless implementation. (We’ll discuss the reasons for those differences later in this section.) The HP developed node controller ASIC, which is integral to the HP PREMA Architecture, plays a critical role in enabling the ProLiant DL980 G7 system to scale to eight processors while adding availability features. (For definitions of many of the terms used in this section, see the [Glossary](#).)

As Figure 1 shows, the ProLiant DL980 G7 contains up to eight processor sockets; each socket accepts a processor with up to eight cores. The HP PREMA Architecture groups the processor sockets into multiple “QPI islands” of two directly connected sockets. This direct connection provides the lowest latencies. Each QPI island connects to two node controllers (labeled “XNC” in the diagram). The system contains a total of four node controllers. Multiple links connect the node controllers in pairs.

Figure 1. Architectural block diagram of the HP PREMA Architecture in the ProLiant DL980 G7



Each node controller stores information about all data located in the processor caches. This functionality is called “smart CPU caching.” Smart CPU caching provides significant benefits for system performance:

- Minimized inter-processor coherency communication and reduced latency to local memory—Processors in each QPI island have access to the smart CPU cache state stored in the node controller, thus eliminating the overhead of requesting and receiving updates from all other processors.
- Dynamic routing of traffic—When an inter-node-controller link is overused, HP’s dynamic routing avoids performance bottlenecks by routing traffic through the least-used path. In this way, the system uses all available lanes and maintains full bandwidth.

Addressing cache coherency and memory latency

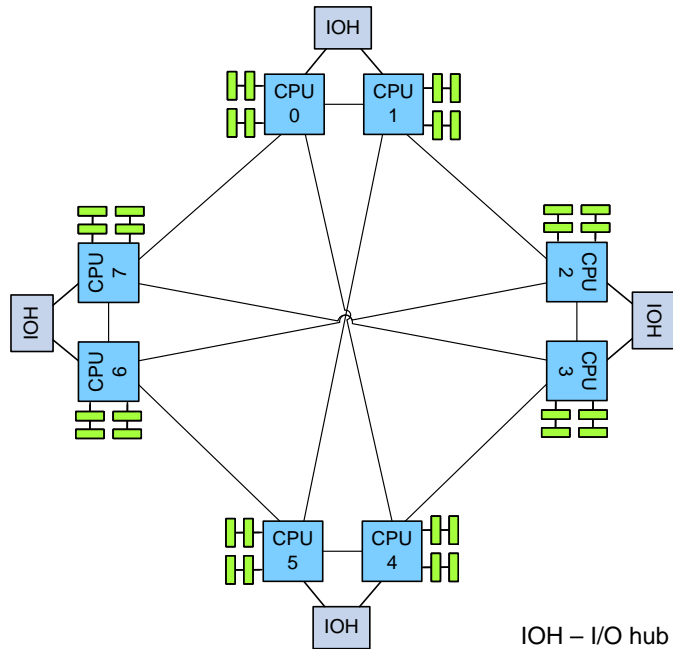
A scale-up x86 system using the Intel Xeon processor 7500/6500 with an embedded memory controller implies a Cache Coherent Non-Uniform Memory Access (ccNUMA) system. In a ccNUMA system, the hardware ensures cache coherency by tracking where the most up-to-date data is for every cache line held in a processor cache. Latencies between processor and memory in a ccNUMA system vary depending on the location of these two components in relation to each other. The Intel Xeon processor 7500/6500 series uses the source snoop variant of the Intel QuickPath Interconnect (QPI). HP’s goal in designing the PREMA Architecture was to reduce average memory latency and minimize bandwidth consumption resulting from coherency snoops. The HP node controller works with the processor’s coherency algorithms to provide system-wide cache coherency. At the same time, it minimizes processor latency to local memory and maximizes usable link bandwidth for all links in the system. The following discussion of the differences between a glueless and a glued implementation helps to elucidate how the HP architecture achieves its goals and results in balanced scaling.

Glueless implementation

Each processor socket has four QPI links. A glueless 8-socket system implementation (Figure 2) uses one of these links to connect the socket to I/O and the remaining three QPI links to interconnect the processor sockets. This forces a topology in which each processor socket connects directly to three other sockets, but its connections to the other four processor sockets are indirect—requiring multiple hops for any request to those processors. For example, in the glueless block diagram, CPU0 does not

have a direct connection to CPU6. CPU0 must send any requests and responses for CPU6 through CPU7, resulting in two hops in each direction (out and back).

Figure 2. Block diagram of 8-socket glueless system



To achieve cache coherency in an implementation having no node controller, a read request must be reflected to all other processor caches as a snoop. Each processor must check for the requested memory line and provide the data if it has the most up-to-date version. If the read was for “exclusive access,” then all other caches must also invalidate their copies. In cases where the modified line is available in another cache this source snoopy protocol provides the minimum latency when the line is copied from one cache to the next. However this solution has limited scalability for some workloads.

In a source snoopy coherency protocol all reads result in snoops to all other caches. This consumes link and cache bandwidth as these snoop packets use cache cycles and link resources that could otherwise be used for data transfers. These source snoops can also impact memory latency when snoops or snoop responses require multiple hops. A source snoopy memory controller cannot return the memory data until it has collected all the snoop responses and is sure that no cache provided a more recent copy of the memory line. Accessing local memory is sufficiently fast that the multi-hop snoops and snoop responses delay the delivery of data from the memory read. In 8-socket glueless systems, snoops consume 50 to 65% of QPI bandwidth.¹

Glued implementation using HP node controller

In contrast, in the HP PREMA Architecture with smart CPU caching, coherency snoops and responses consume only 10 to 20% of QPI bandwidth and that of the HP fabric.² The HP implementation provides latency for local memory access comparable to a 4-socket glueless system and 30% lower latency when compared to an 8-socket glueless system.³ By recording when a cache in a remote QPI island has a copy of a memory line, the node controller can respond on behalf of all remote caches to

¹ Based on HP internal testing.

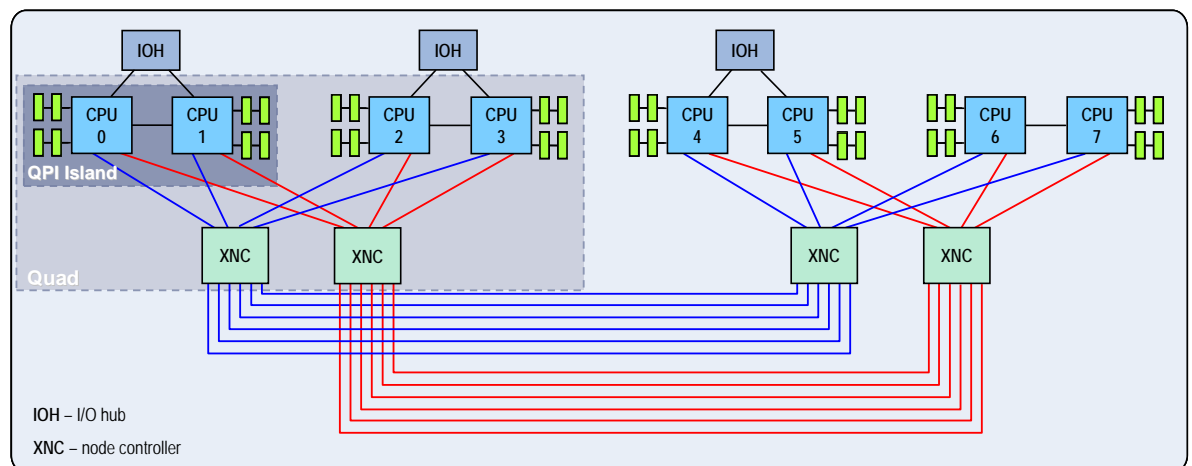
² Based on HP internal testing.

³ Based on HP internal testing.

each source snoop. This not only removes snoop traffic from consuming bandwidth on the links and remote socket caches, it also reduces the memory latency in the cases where the data is not held in any other cache.

The annotated architectural diagram (Figure 3) shows the 2-socket source snoop QPI islands. A pair of node controllers (XNCs) supports two islands in a 4-socket quad. These quads are then connected to create an 8-socket system. Within a 2-socket source snoop island all snoops have at most one QPI link hop between the requesting core, the paired socket cache, the smart CPU cache in the node controller, and the memory controller. Therefore all accesses to local memory have the bandwidth and latency of a small glueless system. Through its tagging of remote ownership of memory lines, the node controller targets any remote access to the specific location of the requested memory line. With the HP PREMA Architecture smart CPU caching technology, the HP system effectively provides more links connecting processor sockets—the equivalent of six QPI links connecting the two quads; a glueless 8-socket system has four QPI links. In addition, smart CPU caching uses the links more efficiently because it reduces the overhead of cache coherency snoops. Because of the reduction in local memory latency compared to glueless 8-processor systems, virtual environments will have higher performance on the ProLiant DL980 G7. With NUMA-aware OS support, system performance will scale nearly linearly.

Figure 3. Architectural block diagram of HP PREMA Architecture in the ProLiant DL980 G7 – annotated



NUMA-aware OS support

Microsoft Windows and Linux are both NUMA-aware operating systems that attempt to allocate memory local to or close to a requesting core or thread to minimize memory latency and link bandwidth consumption. Both operating environments support tuning parameters and other mechanisms to adapt the default NUMA behavior for various workloads to achieve improved performance scalability. Generally speaking, launching tasks on the processor or island or quad where they will run, and constraining the tasks to run on that single processor or island or quad, will maximize performance and scalability for long-running enterprise workloads.

Memory and I/O

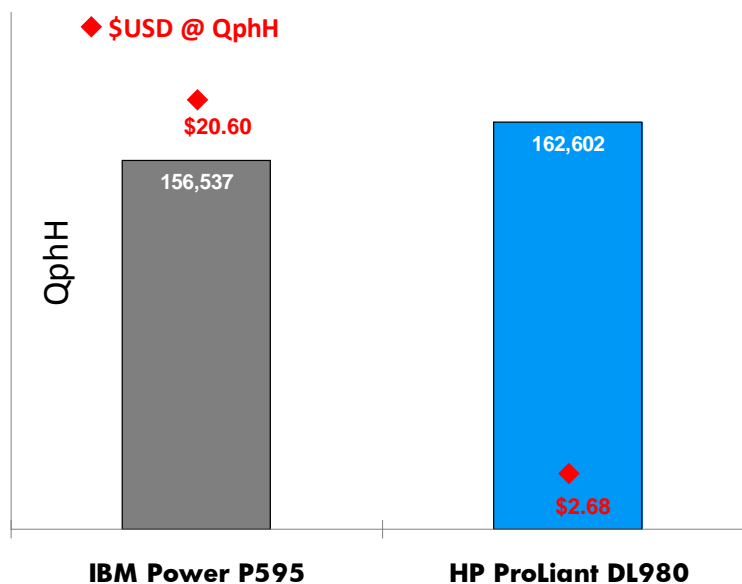
The well-balanced architecture of the ProLiant DL980 G7 encompasses memory and I/O as well as processors. As a scale up system the DL980 G7 will handle applications that require large amounts of shared memory. An 8-socket DL980 G7 using 16GB DIMMs will provide up to 2TB of main memory.

High performance processors and large memories may not be valuable unless a system has enough I/O to feed all of the components. I/O bottlenecks are one of main issues in large enterprise deployments, and limited I/O capacity can constrain virtual machines (VMs). The ProLiant DL980 G7 provides up to 16 PCI Express (PCIe) expansion slots (an alternate expansion option permits a combination of slots, supporting up to 13 PCIe slots and up to two PCI-X slots). This enables you, for example, to provide dedicated I/O for VMs when consolidating and boost data warehousing and decision support application performance with more parallel access to storage. In addition, smart CPU caching benefits I/O bandwidth when you have a combination of I/O and processor traffic. By reducing coherency snoops and responses required to support both types of memory access, more bandwidth is available for transferring requested data.

Performance results

Performance testing using industry standard benchmarks confirms the advantages of the HP PREMA Architecture. By reducing processor overhead, The ProLiant DL980 G7 server delivers better system performance for a diverse set of workloads than competitive 8-socket—and larger—systems. With up to 30% faster local memory access⁴ provided by the PREMA Architecture, workloads with affinity to local memory and the applications isolated to a processor or a QPI island can gain additional performance advantages. For example, VMs, SAP, Microsoft SQL Server, and High Performance Computing applications will experience performance gains on a DL980 G7 when compared with glueless 8-processor systems.

Figure 4. TPC-H @ 3000GB benchmark⁵



TPC-H performance

On the TPC-H @ 3TB benchmark, the HP ProLiant DL980 G7 achieved 162,601.7 QphH@3000GB⁶ and the top price performance for non-clustered systems at \$2.68 USD/QphH@3000GB (Figure 4).⁷ The TPC-H benchmark measures performance for complex data warehouse transactions for business intelligence solutions. This result showed that the ProLiant DL980 G7 server delivers better

⁴ Based on HP internal testing.

⁵ For compliance details, see Table 1 in [Appendix](#).

⁶ TPC-H Composite Query-per-Hour Performance Metric

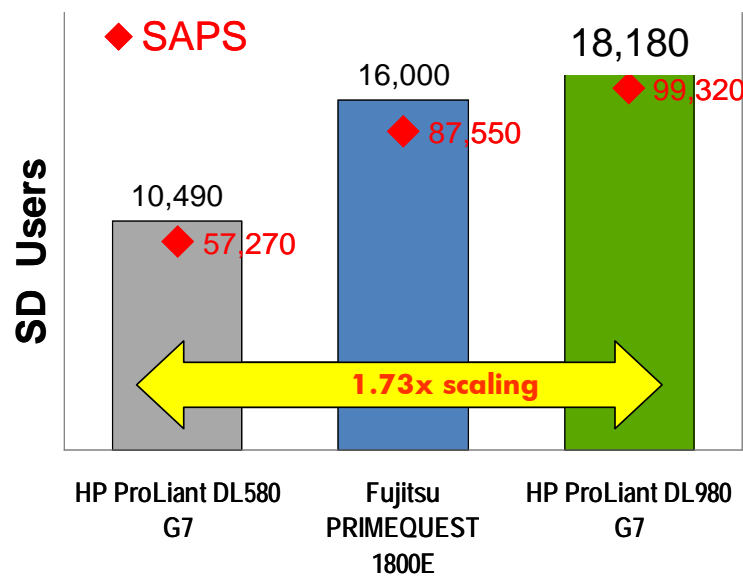
⁷ For compliance details, see Table 1 in [Appendix](#).

performance and much better price performance in one-fifth the rack space compared to the IBM POWER6-based POWER 595.

SAP SD performance

On the two-tier SAP Sales and Distribution (SD) standard application benchmark, the ProLiant DL980 G7 server achieved 18,180 SAP SD Benchmark users, equivalent to a throughput of 1,986,330 fully processed order line items per hour or 99,320 SAPs (Figure 5).⁸ The ProLiant DL980 G7 achieved leading Windows performance results on the benchmark, exceeding the Fujitsu 8-processor result by 13.6%. In addition, the DL980 G7 achieved excellent 8-processor performance scalability results compared to 4-processor results with the same processors (1.73x scaling over the 4-processor ProLiant DL580 G7) and compared to 8-processor results with a previous generation HP ProLiant platform (2.2x scaling over the 8-processor ProLiant DL785 G6). The exceptional 4-processor to 8-processor scaling is a direct result of the smart CPU caching capability of the HP PREMA Architecture.

Figure 5. SAP SD Standard Application Benchmark⁹



For details supporting these benchmark results, see the tables in the [Appendix](#). For more information on the latest ProLiant benchmark results, see

<http://h18004.www1.hp.com/products/servers/benchmarks/products.html#8p>.

Self-healing resiliency

The ProLiant DL980 G7 takes advantage of the enhancements that Intel has provided in its processor technology and builds on these to provide the resiliency demanded by mission-critical enterprise environments. Through tight collaboration with OS vendors HP ensures that the OS has awareness of the resiliency features that are enabled in the system.

Redundant system fabric

HP PREMA Architecture extends the advanced reliability of the Intel Xeon processor 7500/6500 series in the ProLiant DL980 G7 with a redundant system fabric. This interconnect fabric provide

⁸ For compliance details, see Table 2 in [Appendix](#).

⁹ For compliance details, see Table 2 in [Appendix](#).

higher interconnect bandwidth and lower data error rates for improved resiliency. The redundant system fabric enables:

- Redundant data paths—The fabric’s provision of 50% more interconnect links (six here versus four in most competitive 8-socket systems with no node controller) not only improves system performance, it minimizes downtime in the event of an interconnect link failure.
- Rapid recovery—Improved error logging and diagnostics information means that OS and virtual machines can attempt error recovery or administrators can easily take corrective actions. If a fatal error occurs, the DL980 G7 does a warm reset and captures the error log to assist in diagnosis. With the system running, the administrator can then use the log to diagnose the error and take corrective action based on a single failure.

Additional redundant components

The ProLiant DL980 G7 server provides full redundancy for components that are most susceptible to failure. This includes the following components:

- Fully redundant (dual grid) hot-swap bulk power supplies
- Fully redundant hot-swap system fans
- Spare wires, clocks, and lanes in the fabric

For an extra level of redundancy the ProLiant DL980 G7 offers options for:

- Dual path I/O cards.

Customers may optionally use two or more identical I/O cards on different I/O backplanes (IOHs) to provide I/O card and I/O path redundancy. Some OS support I/O card “teaming” to get the best performance by spreading the load among multiple I/O cards.

- Optional memory spare rank.

Customers concerned about memory failures can optionally set aside a spare rank. With this option in place, the system tracks corrected memory found by memory scrubbing (see Correction mechanisms). If a section of memory exceeds the threshold for corrected errors, the system copies its contents to the spare rank. Activating the spare and taking the questionable memory out of use before an uncorrectable memory error occurs avoids a system crash.

- Optional mirrored memory.

For customers wanting the highest level of memory resilience, the DL980 G7 offers optional memory mirroring functionality. Memory mirroring allocates half of the system memory as a “mirror.” To ensure that mirrored memory maintains consistency with primary memory, all writes update both the primary memory and the mirror. If an uncorrectable memory error occurs, the system automatically switches from the primary memory to the mirror. Mirrored memory offers a higher level of protection than memory spare rank, since the system is able to recover even if an uncorrectable memory error occurs. This high level of memory error protection comes at a price: mirrored memory consumes half of the available memory capacity and half of the overall memory bandwidth.

- RAID storage and hot plug drives.

The HP Smart Array P410i Controller in the DL980 G7 supports up to eight hot plug SFF Serial Attached SCSI (SAS) drives. The P410i enables RAID configurations 0, 1, 1+0, 5, and 5+0. The controller also supports RAID 6 with Advanced Data Guarding, as well as RAID 6+0, with the addition of optional 512MB or 1GB flash-backed write cache (FBWC) and the optional HP Smart Array Advanced Pack (SAAP). (The 512MB FBWC is standard on some server models.) The FBWC uses flash devices and super-capacitors to retain cache data indefinitely in the event of a power loss. The SAAP provides a license key to activate firmware for advanced functionality on the Smart Array P410i controller.

For a detailed discussion of the controller technology and RAID levels supported, see the HP technology brief: *HP Smart Array Controller* technology at <http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00687518/c00687518.pdf>. For details on the features of the disk drives used, see the HP technology brief: *Drive technology overview* at <http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01071496/c01071496.pdf>.

High-speed link resiliency

The point-to-point links that connect the chips in the system have many reliability, availability, and serviceability (RAS) features in common. RAS features, including Cyclic Redundancy Checksum (CRC), link level retry, link width reduction (LWR), and link retrain, work together to keep the system operational in the presence of link errors. These features demonstrate the level of RAS sophistication in the ProLiant DL980 G7.

The HP PREMA Architecture of the ProLiant DL980 G7 uses multiple high-speed links:

- Intel QuickPath Interconnect (QPI)—Connects processor to processor, processor to I/O hub, and processor to the HP node controller
- Intel Scalable Memory Interconnect (SMI)—Connects processor to DDR3 memory buffer
- HP redundant system fabric—Connects HP node controllers to each other

Normally, transaction requests such as memory reads and responses such as data returned from reads are sent to their destinations through these narrow (typically 10 to 20 bits wide), but high-speed (4.8 to 6.4GHz) links. The system uses error correction codes (ECC) similar to those used to protect memory DRAMs to detect and correct transient errors. However ECC can miss errors if too many of the bits are corrupted. To ensure detection of any error on the link, a CRC detects transient errors that affect many bits in each transaction request and response.

Once a CRC detects an error that cannot be corrected by simple ECC, then the transaction is retried on the link (link level retry). If the link continues to experience uncorrectable errors detected by CRC, then the link is reset and retrained (link retrain), and the transaction is retried again. If the link reset and retrain are unsuccessful in making the link operational, the system analyzes the failure and, if possible, invokes a spare wire (for either data or clock), and retries the transaction once again. If using the spare wire does not make the link functional, then the system attempts to find a part of the link that is functional—the upper half or lower half—and again retries the transaction (LWR).

If reset and retraining are successful, the error logs record a warning and an on-screen message alerts the system administrator that the system is running with less than full protection (for example, no additional spare wires are available). The administrator can then schedule a service call for a convenient future time.

If the system invokes a spare wire, but this does not resolve the link errors, then the system again re-initializes the link, but at reduced width (LWR). For example, if the failed link was 20 bits wide, it would be reduced to 10 bits wide. The system selects upper or lower 10 bits based on which trains as properly functional. As in the previous case, the system error logs record a warning entry and inform the system administrator of the reduced protection as well as the reduced performance due to the lower bandwidth on the affected link.

In all of these instances, the system remains operational without requiring any special user intervention.

Correction mechanisms

The HP PREMA Architecture in the ProLiant DL980 G7 enables key features of the Intel Xeon processor 7500/6500 series designed to detect and correct errors before they result in a system crash and to keep a record of errors that caused a crash to assist in diagnosing and servicing the system later:

- Memory demand and patrol scrubbing —The system monitors and removes transient single-bit errors before they accumulate into uncorrectable memory errors. Exceeding the threshold for error correction rate triggers predictive memory de-allocation, or optionally invokes the use of the memory spare rank in place of the suspect memory rank.

Demand scrubbing corrects errors detected during a normal data read transaction. Patrol scrubbing preemptively walks through all memory periodically to correct errors before the OS or an application accesses the memory and crashes because of an uncorrectable multi-bit memory error. Demand scrubbing is enabled in the initial server release; patrol scrubbing will be added in a subsequent firmware release. The functionality is independent of which OS you are running.

- Memory Correction Codes (ECC)—This traditional memory error correction coding can detect all double-bit errors and correct any single-bit error (Single Error Correct, Double Error Detect) or any single DRAM device failure (Single Device Data Correction).

Mechanisms to contain failures

The ProLiant DL980 G7 enables features to isolate and contain failures, and log errors leading to failures, in the rare event that these do occur.

Machine Check Recovery

The HP ProLiant DL980 G7 will provide run-time error logging to enable the OS or VM host to recover from errors. This logging information will enable the server to isolate and contain memory errors to an application or OS VM guest, preventing the errors from bringing down the whole system.

In the case of an uncorrectable memory error, the memory controller will mark the corrupt data line to identify it as bad or “poison.” Patrol scrub (see [Correction mechanisms](#)), and Last-Level-Cache-Write-Back error data poisoning will be supported in a subsequent firmware release. Machine Check Recovery requires that the OS or VM host be notified and be able to take corrective action before the application or VM guest OS consumes the poisoned data. MCA (Machine Check Architecture) logging currently provides failure records for improved diagnostics and serviceability in cases where the error is unrecoverable.

Machine Check Recovery is a complex software recovery process, and the system firmware, OS, and VM hosts are expected to add additional recovery cases over time, and to improve their ability to isolate these fatal errors to applications or VMs. The ProLiant DL980 G7 will provide the key error information to support additional recovery cases as the OS and VM hosts improve over time.

I/O fire-door error containment

HP PREMA Architecture in the ProLiant DL980 G7 server works with features of the Intel Xeon 7500/6500 processor series to protect data integrity. If a hardware component (processor, I/O hub or HP node controller) detects uncorrectable errors it immediately signals the rest of the system. The DL980 G7 reacts to this quick propagation of the error condition by shutting down the I/O subsystem; in this way the error is contained within the server. This prevents any data integrity issues from contaminating external storage or the network.

Breakthrough efficiency from ProLiant innovations

In addition to the balanced scaling and self-healing resiliency enabled by the HP PREMA Architecture, the HP ProLiant DL980 G7 incorporates key ProLiant innovations to deliver breakthrough efficiency.

Consolidation

The performance gains enabled by HP PREMA Architecture make it possible for a single ProLiant DL980 G7 server to support workloads previously running on many legacy servers. As a result you lower the space, power and cooling requirements needed to run those workloads. As an example of

what this means in your datacenter, it is possible to consolidate 197 legacy servers to a single HP DL980 G7, to achieve a payback in as little as 60 days.¹⁰

Integrated Lights-Out 3

The Integrated Lights-Out 3 (iLO 3) management processor integrated into every ProLiant DL980 G7 server provides secure remote management capabilities, regardless of server status or location. The iLO 3 provides remote console performance up to eight times faster than the previous generation iLO 2 processor—and equal to the performance of KVM and software-based remote management solutions. iLO 3 also offers three times faster virtual media performance compared to iLO 2.¹¹ Booting a remote system from an iLO virtual media device enables administrators to perform many tasks without visiting the server, for example, to perform upgrades, deploy an OS from network drives, and perform disaster recovery of failed operating systems. The iLO Advanced option upgrades the iLO firmware to enable virtual media and other advanced functionality.

For more information on iLO and iLO Advanced, visit www.hp.com/go/ilo.

HP Insight Software

HP Insight Software is an integrated portfolio of infrastructure management software for HP ProLiant and Integrity servers.

HP Insight Control is essential infrastructure management that can help save time and money by making it easy to deploy, migrate, monitor, control, and optimize your IT infrastructure through a single, simple management console. Insight Control supports both Windows and Linux-based central management servers. Insight Control enables you to deploy a ProLiant server 12 times faster than manual deployment using vendor provided CDs or DVDs.¹²

HP Insight Dynamics is advanced infrastructure lifecycle management software. Insight Dynamics includes:

- Integrated infrastructure design with automated activation of servers, storage and networking
- Built-in capacity planning and rebalancing tools
- Automated disaster recovery and failover capabilities

Insight Control gives you the potential to reduce operational expenses by \$48,300 per 100 users over 3 years.¹³ In addition, managing health proactively has enabled HP ProLiant customers to reduce unplanned downtime by up to 77%.¹⁴

HP Insight Control and Insight Dynamics are options for the ProLiant DL980 G7. For more information on HP Insight Software, visit: www.hp.com/go/insightsoftware.

High efficiency power supplies

The ProLiant DL980 G7 uses HP 1200W common slot power supplies. HP power supplies are the first to achieve Platinum certification. With 94% efficiency these power supplies provide \$20-\$80 savings per power supply per year compared to competitive products.

¹⁰ Calculated using HP internal ROI calculators using latest SPECpower estimates; 197 ProLiant DL360 G4 servers can be consolidated into a single ProLiant DL980 G7.

¹¹ 2010: HP iLO engineering team benchmark

¹² 20 minutes versus four hours; TTX Company Customer Success Story (June 2009), <http://h20195.www2.hp.com/v2/GetPDF.aspx/4AA2-6629ENW.pdf>.

¹³ IDC White Paper sponsored by HP, *Gaining Business Value and ROI with HP Insight Control*, Doc # #218069 (May 2009), <http://h18013.www1.hp.com/products/servers/management/IDC-ICE-ROIWhitePaperFinal.pdf>.

¹⁴ IDC: *Gaining Business Value and ROI with HP Insight Control* (May 2009). Includes the functionality provided by Systems Insight Manager and Insight Remote Support.

For more see the HP High Efficiency Power Supplies web page at http://h18004.www1.hp.com/products/servers/rackandpower/powersupplies/index.html?jumpid=reg_R1002_USEN.

Common components

Common power supplies, hard drives and memory across the ProLiant server product line allow self-service customers to reduce spare parts inventories.

Thermal Logic technologies

Thermal Logic is the HP portfolio of embedded server technologies designed to help you achieve an energy efficient data center. Thermal Logic includes the following key technologies:

- Sea of Sensors—An array of thermal sensors that adjust fan speeds and powers only slots in use to provide optimal system cooling at the lowest power.
- Power Regulator—An OS-independent power management feature that enables dynamic or static changes in processor performance and power states to adapt to changing workloads.
- Dynamic Power Capping—A feature that enables you to reclaim trapped power and cooling capacity by safely limiting server power consumption. This feature gives you the potential to expand data center capacity up to triple existing capacity.¹⁵

For more detail on these and other HP Thermal Logic technologies, visit www.hp.com/go/thermallogic.

Benefits of ProLiant DL980 G7 with HP PREMA Architecture for applications

The ProLiant DL980 G7 together with industry leading OS and application software offers a lower cost, standards-based solution for online transaction processing (OLTP), data warehouse, and business intelligence applications. The balanced scaling and self-healing resiliency of the DL980 G7 make it an ideal platform for these applications.

Business intelligence applications typically stream large amounts of data from disk into memory, while using processing power to analyze the data. Hardware requirements to support the data warehouse focus on providing sufficient amounts of storage. Customers implementing business intelligence/data warehouse solutions must ensure that they can support an ever increasing query volume, query complexity, database growth, increased sophistication of user queries and availability.

Online transaction processing (OLTP) applications generally consist of a database server running a DBMS (database management system) package along with several application servers to process transactions. OLTP systems such as large e-commerce websites must respond to spikes in demand from large numbers of users and a high volume of transactions. These require large amounts of memory to maintain connection context for every database object opened by a user.

With the scalability and expansion capabilities of the ProLiant DL980 G7 and the headroom offered by the software—whether Microsoft SQL Server or Oracle Database, on Microsoft Windows Server or Linux—you can readily accommodate needed growth in database capacity. By scaling up, you can provide more efficient database capacity expansion within a single system without adding the complexity and overhead of a database cluster. The DL980 G7 supports up to 2.4TB of internal storage and up to 960TB of external storage. The provision of up to 16 I/O slots in the server means

¹⁵ HP Performance Engineering Team, 2008 Benchmark. Percentage increase in server density compares the number of servers that can be deployed using Dynamic Power Capping with the number of servers that can be deployed using server face plate values. Density improvements may vary based on how customers budget power and cooling resources.

you can accelerate data warehouse and decision support application performance with more parallel access to storage. The balanced architecture of the DL980 G7 enables you to scale performance by increasing processors and memory as well—with support for up to eight 8-core processors for up to 64 processor cores and up to 128 logical processors with hyper-threading enabled, and up to 2TB of memory.

Ensuring availability of OLTP systems and protecting business data are typically critical to your business. The self-healing resiliency of the ProLiant DL980 G7 servers maximizes application uptime with a 200% boost in server availability compared to a previous generation ProLiant server,¹⁶ and assures you that your critical data is protected.

The ProLiant DL980 G7 also provides an excellent platform for consolidation, and virtualization with leading hypervisors such as Microsoft Hyper-V and VMware vSphere. With consolidation ratios on the DL980 G7 as high as 197:1¹⁷, you can reduce the number of OS and application instances, giving you the opportunity to save licensing costs and reduce the personnel effort to manage those instances.

For information on the ProLiant DL980 G7 with Microsoft SQL Server 2008 R2 for OLTP, see *HP recommended configurations for online transaction processing: ProLiant DL980 G7 and Microsoft SQL Server 2008 R2* at <http://h20195.www2.hp.com/v2/GetPDF.aspx/4AA2-1137ENW.pdf>.

Conclusion

HP is building the data center of the future, bringing to bear our expertise in both industry standard and mission-critical system design and our close collaboration with Intel and industry-leading software partners. The ProLiant DL980 G7 combines advanced HP PREMA Architecture and the wealth of ProLiant innovations to give you scale-up performance and availability with x86 economics. The result is a server on which you can deploy your most demanding applications with confidence.

¹⁶ Based on HP internal research comparing HP ProLiant DL785 G5 to DL980 G7 with configurations normalized to address similar performance; defined as MTBCF (mean time between critical failures); availability varies according to specific model.

¹⁷ 197 ProLiant DL360 G4 servers can be consolidated into a single ProLiant DL980 G7.

Appendix: Benchmark results

Table 1. TPC- H @3000GB benchmark

Platform, Processor (chips/cores/threads), Memory	Availability	OS and Database	QphH@3TB	USD \$/QphH
HP ProLiant DL980 G7 - 8 processors, 8 cores; Intel Xeon X7560 2.27GHz, (8/64/128), 512GB RAM	10/13/2010	Microsoft Windows Server 2008 R2 Datacenter x64, SQL Server 2008 Enterprise Edition x64	162,602 QphH @ 3000 GB	\$2.68 USD/QphH @ 3000 GB
IBM POWER 595 Model 9119-FHA, 32 processors, Dual-Core IBM POWER 6 – 5.0GHz (32/64/128), 512GB RAM	11/04/2009	AIX v 6.1, Sybase IQ Single Application Server Edition v.15.1 ESD #1.2	156,537 QphH @ 3000 GB	\$20.60 USD/QphH @ 3000 GB
Unisys ES7000 Model 7600R Enterprise Server (16s), 16 processors Intel Hex-core Xeon X7460 2.66GHz (16/96/96), 1TB RAM	05/06/2010	Microsoft Windows Server 2008 R2 Datacenter Edition, Microsoft SQL Server 2008 R2 Datacenter Edition	102,778 QphH @ 3000 GB	\$21.05 USD/QphH @ 3000 GB

Note: All results noted were achieved at www.tpc.org. Results as of August 2, 2010.

Table 2. Two-tier SAP® Sales and Distribution (SD) Standard Application Benchmark with SAP enhancement package 4 for SAP ERP 6.0

Platform, Processor and Memory	Certification Number	OS, Database and SAP Software	SD Users	SAPS
HP ProLiant DL980 G7, 8 processors, 8-core 2.26 GHz Intel Xeon X7560 (8/64/128), 512 GB RAM	2010028	Windows Server 2008 R2 DC x64, SQL Server 2008 EE x64, SAP enhancement package 4 for SAP ERP 6.0	18,180	99,320
Fujitsu PRIMEQUEST 1800E, 8 processors, 8-core 2.26 GHz Intel Xeon X7560 (8/64/128), 512 GB RAM	2010010	Windows Server 2008 R2 DC, SQL Server 2008 EE, SAP enhancement package 4 for SAP ERP 6.0	16,000	87,550
HP ProLiant DL580 G7, 4 processors, 8-core, 2.26 GHz Intel Xeon X7560 (4/32/64), 128 GB RAM	2010032	Windows Server 2008 Enterprise Edition, SQL Server EE 2008, SAP enhancement package 4 for SAP ERP 6.0	10,490	57,270
HP ProLiant DL785 G6, 8 processors, 6-core AMD Opteron 8439SE 2.8 GHz (8/48/48), 128 GB RAM	2009035	Windows Server 2008 R2 DC x64, SQL Server 2008 EE x64, SAP enhancement package 4 for SAP ERP 6.0	8280	45,350

Note: All results noted were achieved on the two-tier SAP SD Standard Application Benchmark and all servers shown ran SAP enhancement package 4 for SAP ERP 6.0 (Unicode). For more details, please visit: SAP benchmark details at <http://www.sap.com/benchmark>. Results as of August 2, 2010.

Glossary

Cache coherency—The maintenance of data consistency among processor memory caches in a system.

ccNUMA system—A cache coherent NUMA system. In this design, hardware ensures cache coherency, but the system has varying memory latency depending on the location of memory being accessed. *See also* Non-uniform memory access system.

Coherency protocol—An established convention for inter-processor communication between cache controllers to maintain a consistent memory image when more than one cache stores the same memory location, or if data has been modified (written) and is only stored in a processor cache. *See also* cache coherency.

Snoop protocol—A mechanism for keeping caches consistent (coherent) by explicitly checking every cache in the system to determine if it holds a copy of a particular cache line. On any processor READ operation all other caches in the system must be “snooped” to check for a second copy of the cache line.

Directory-based—Referring to an alternative mechanism for keeping caches consistent (coherent) by having the memory controller keep track of (create a directory of) all caches that have copies of any particular memory locations. Directory information is stored with the data in the memory DRAMs.

Glued—Referring to a processor/memory subsystem topology using node controllers to provide at least some of the connections between processor sockets.

Glueless—Referring to processor/memory subsystem topology in which all processor sockets are directly connected to each other.

Memory latency—The time it takes for a memory reference from a processor to be satisfied by the memory system. This includes the time spent waiting for snoop responses. *See also* snoop protocol.

Non-uniform memory access (NUMA) system—Multiprocessor system with varying memory latency depending on the location of the memory being accessed. In a NUMA system, memory access time depends on the distance between processor and memory: a processor can access its own local memory faster than non-local (remote) memory. *See also* memory latency.

QuickPath Interconnect (QPI)—Intel high-speed point-to-point links used to connect processors to each other, and to I/O hubs and node controllers.


Scalable Memory Interconnect (SMI)—Intel links used to connect a processor’s memory controller to the memory buffers.

For more information

For more information on the HP ProLiant DL980 G7 with HP PREMA Architecture, visit the server web page at <http://h10010.www1.hp.com/wwpc/us/en/sm/WF04a/15351-15351-3328412-241644-4222584>.

For more information on the Intel Xeon processor 7500 series, visit <http://www.intel.com/itcenter/products/xeon/7500/index.htm>.

Share with colleagues    

**Get connected**
www.hp.com/go/getconnected
Current HP driver, support, and security alerts
delivered directly to your desktop

Become a fan on  »

Follow on  »

© Copyright 2010 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.

Intel and Intel Xeon are trademarks of Intel Corporation in the United States and other countries.

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

Oracle is a registered trademark of Oracle Corporation and/or its affiliates.

AMD is a trademark of Advanced Micro Devices, Inc.

4AA3-0643ENW, Created August 2010

