

TECHNOLOGY BRIEF

August 1999

Compaq Computer Corporation
ISSD Technology Communications

CONTENTS

Introduction	3
High-Performance Technologies	3
<i>ProLiant</i> 8500 Architecture	4
Pentium III Xeon Processor Technology	10
High-Availability Technologies	12
PCI Hot Plug Technology	12
Redundant Network Interface Controller Technology	13
Redundant Hot-Plug Fans	13
Redundant Processor Power Modules	14
Redundant Hot-Plug Power Supplies	15
Hot-Plug Hard Disk Drives	17
Auto Processor Bus Recovery	17
Storage Technologies	17
Hard Drive Technology	17
Integrated Array Controller Technology	18
Management Technologies	19
Remote-Flash Redundant ROM	20
Auto-Default ROM Configuration	21
Integrated Management Display	21
Integrated Remote Console	21
Remote Insight Board/PCI Option	22
Serviceability Features	22
Conclusion	24

Compaq ProLiant 8500 Server Technology

This technology brief describes the high-performance, high-availability, storage, and management technologies built into the Compaq ProLiant 8500 server. The high-performance technologies include a new 8-way system architecture driven by the Profusion chipset codeveloped by Compaq, Intel, and Corollary, as well as Intel's Pentium III Xeon processors. In addition, new disk drive technology and a new integrated array controller optimize the ProLiant 8500 server's internal storage for operating system and boot support while allowing maximum external storage capacity.

The high-availability technologies include PCI Hot Plug technology; redundant hot-plug power supplies, drives, and fans; redundant processor power modules; and auto processor bus recovery. The ProLiant 8500 server uses a highly modular design that allows easy access to components, thus reducing downtime for service and upgrades. The ProLiant 8500 server also incorporates key management technologies, such as the Integrated Remote Console, Integrated Management Display, Remote-flash Redundant ROM, and auto-default ROM configuration, that further enhance availability.

The intended audience for this paper is engineers and system administrators familiar with existing Compaq technology and servers. For those less familiar with Compaq technology, please see the related technology briefs referenced in this document. For more information about the ProLiant 8500 server, see the complete list of features at <http://www.compaq.com/products/servers/ProLiant8500/quickspecs.html>.



NOTICE

The information in this publication is subject to change without notice and is provided "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES WHATSOEVER (INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal state or local requirements.

Compaq, Contura, Deskpro, Fastart, Compaq Insight Manager, LTE, PageMarq, Systempro, Systempro/LT, ProLiant, TwinTray, ROMPaq, LicensePaq, QVision, SLT, ProLinea, SmartStart, NetFlex, DirectPlus, QuickFind, RemotePaq, BackPaq, TechPaq, SpeedPaq, QuickBack, PaqFax, Presario, SilentCool, CompaqCare (design), Aero, SmartStation, MiniStation, and PaqRap, ProSignia, Concerto, Vocalyst, and MediaPilot are registered with the United States Patent and Trademark Office.

Change is Good, Compaq Capital, Colinq, Armada, SmartQ, Counselor, CarePaq, Netelligent, Smart Uplink, Extended Repeater Architecture, Scalable Clock Architecture, QuickChoice, Systempro/XL, Net1, LTE Elite, PageMate, SoftPaq, FirstPaq, SolutionPaq, EasyPoint, EZ Help, MaxLight, MultiLock, QuickBlank, QuickLock, UltraView, Innovate logo, and Compaq PC Card Solution logo are trademarks and/or service marks of Compaq Computer Corporation.

Microsoft, Windows, Windows NT, Windows NT Advanced Server, SQL Server for Windows NT are trademarks and/or registered trademarks of Microsoft Corporation.

NetWare and Novell are registered trademarks and IntranetWare, NDS, and Novell Directory Services are trademarks of Novell, Inc.

Pentium is a registered trademark and Xeon is a trademark of Intel Corporation.

Other product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

©1999 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.

Compaq ProLiant 8500 Server Technology

First Edition (August 1999)
Document Number 0091-0899-A

INTRODUCTION

In today's business environment, enterprise customers require powerful servers with fault-tolerant features to keep their businesses running 24 hours a day, 7 days a week. They also need servers that are easy to manage and will scale to meet future requirements.

The *Compaq ProLiant* 8500 server delivers outstanding 8-way scalable performance for 24x7 multiserver rack environments. Performance is maximized with up to eight Intel Pentium III Xeon processors and 16 GB of synchronous dynamic random access memory (SDRAM). Input/output (I/O) performance is further enhanced by eleven 64-bit PCI slots—two at 66 MHz—an integrated array controller, and Wide Ultra-2 SCSI hard drives.

High-availability features include PCI Hot Plug slots, redundant hot-pluggable power supplies, redundant hot-pluggable fans, redundant processor power modules (PPMs), and redundant network interface controllers (NICs). The *ProLiant* 8500 server includes server management and configuration tools that further enhance availability, such as Remote-flash Redundant ROM and Integrated Remote Console (IRC). Additionally, a highly modular design and System Interconnect Status Indicators improve serviceability and further reduce downtime associated with upgrades and field maintenance.

This technology brief explains the features of the *ProLiant* 8500 server and describes the management and storage technologies that enhance performance and availability. Many of these technologies are covered in more detail in separate technology briefs, which are referenced in this document.

HIGH-PERFORMANCE TECHNOLOGIES

The *Compaq ProLiant* 8500 server is a powerful, industry-leading technology solution for distributed enterprise computing. A revolutionary new 8-way system architecture allows the *ProLiant* 8500 server to deliver unsurpassed performance. This new architecture is the first implementation in the industry of the Profusion chipset and features Pentium III Xeon processors. Compaq servers using this architecture provide new levels of performance without requiring modifications or special releases of operating systems (OSs) or applications. The Compaq 8-way architecture provides nearly linear scalability for up to eight processors running under Microsoft Windows NT 4.0, Microsoft Windows 2000, Novell NetWare, or SCO UnixWare. The 8-way architecture delivers superior performance and price:performance for processor-intensive applications such as Microsoft Terminal Server and Oracle databases; for memory-intensive applications such as Lotus Notes, Microsoft Exchange, and SAP solutions; and for I/O-intensive applications such as Microsoft SQL Server.

PCI – peripheral component interconnect

SCSI – small computer system interface

ROM – read only memory

ProLiant 8500 Architecture

The system architecture is the key to the high level of performance offered by the *ProLiant 8500* server. Compaq worked closely with Intel and Corollary for three years to perfect the 8-way symmetric multiprocessing (SMP) architecture so that the processors, memory, and I/O subsystem work in harmony to deliver breakthrough levels of performance. Figure 1 illustrates the architecture used in the *ProLiant 8500* server. The essential features of the architecture include:

- Profusion five-point crossbar switch
- Dual 100-MHz processor buses
- Dual 100-MHz memory buses, each with its own memory controller
- Dedicated 100-MHz I/O bus
- Support for up to eight Pentium III Xeon processors
- Support for up to 16 GB of two-way, cache-line interleaved SDRAM
- Dual cache accelerators
- I/O filter
- Three Compaq designed host-to-PCI bridges

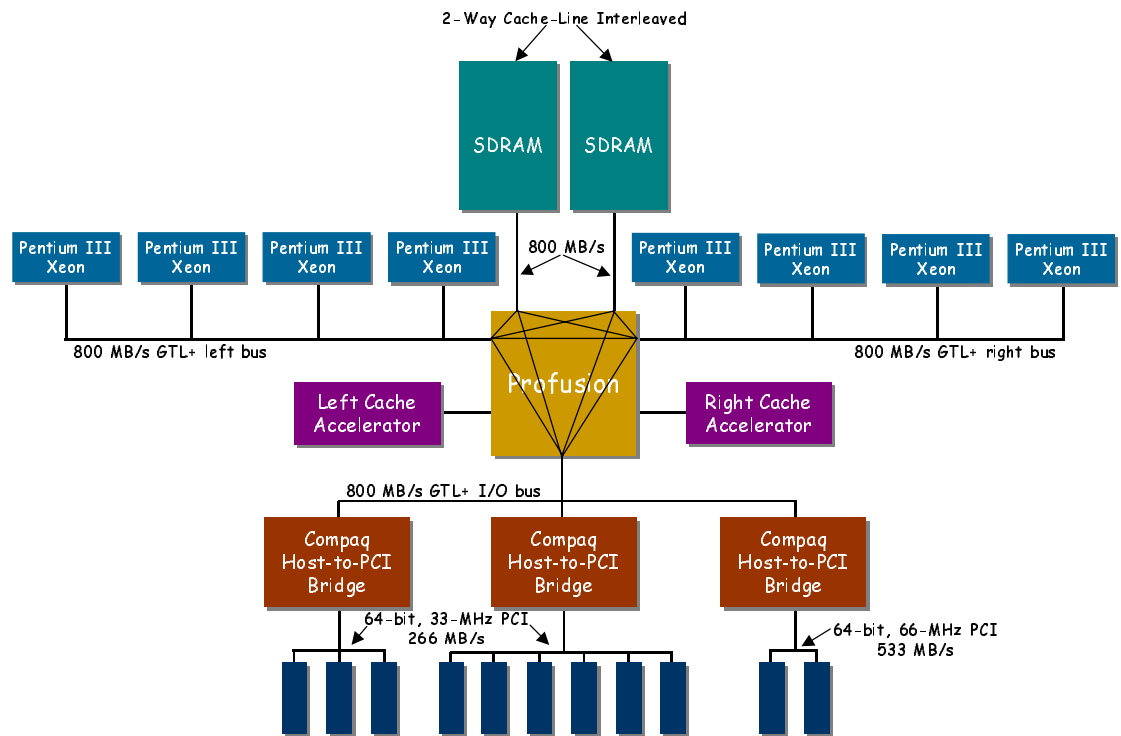


Figure 1: Block diagram of the ProLiant 8500 system architecture.

Profusion Chipset

At the heart of the 8-way architecture is the Profusion chipset. The chipset uses a five-point crossbar switch (Figure 2) to connect processor buses, memory ports, and the I/O bus. The crossbar switch contains static random access memory (SRAM) with ten ports—five read and five write—that appear as five bidirectional ports, one for each of the processor, memory, and I/O buses. The switch may connect two ports directly or may store data from the originating bus in the SRAM before it is transferred to the destination bus. This nonblocking design allows simultaneous read and write accesses from all five buses, which results in better system performance.




Figure 2: Block diagram of the Profusion crossbar switch.

The Profusion crossbar switch provides direct paths from each memory bus to each processor bus and to the I/O bus. Depending on the status of the system, the direct paths can be used to bypass the SRAM, thus reducing latency and improving performance. The direct paths are used only to read data. To improve processor utilization and performance, write data is always posted to the SRAM and written to main memory later.

The Profusion crossbar switch consists of two physical chips—the memory address controller (MAC) and the data interface buffer (DIB). This functional partitioning of the application-specific integrated circuit (ASIC), as shown in Figure 3, improves system performance. For every transaction from a processor or an I/O controller, the address and command portions are routed through the MAC and the data is routed through the DIB. The MAC manages the external cache accelerators and tracks the information stored in the DIB. The DIB allows simultaneous data transfer on all five ports, has 64-cache-line buffers, and uses error-correcting code to maintain data integrity. The cache-line buffers can be used by any transaction for any device on any bus; and since there are no dedicated queues between buses, the efficiency of the buffers is high. This improves system performance.

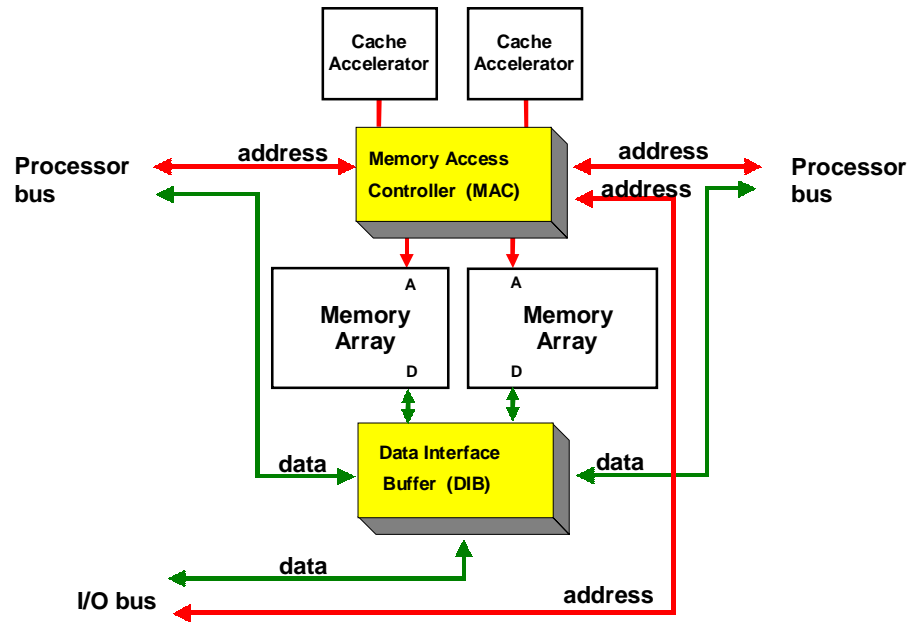


Figure 3: ASIC partitioning in the Profusion chipset.

Processor and I/O Buses

The *ProLiant* 8500 server includes two 64-bit processor buses and a dedicated 64-bit I/O bus. All three are 100-MHz advanced Gunning transceiver logic plus (AGTL+) buses with a theoretical maximum throughput of 800 MB/s. AGTL+ buses facilitate higher clock speeds without significantly reducing the bus length or number of electrical loads. Each 100-MHz AGTL+ bus can support a maximum of five loads per bus. This allows four processors and one connection to the memory controller on each processor bus. The use of two processor buses enables the *ProLiant* 8500 server to support up to eight Pentium III Xeon processors.

The Profusion chipset joins the two processor buses, the I/O bus, and the two memory ports. The otherwise independent processor and I/O buses are joined by a logical connection that is made only when required to transfer data. Each of the three AGTL+ buses has independent access to the two memory ports. This architecture prevents I/O traffic from consuming bandwidth on the processor bus. In addition, the use of 100-MHz buses and five independent paths allows the crossbar switch to deliver an aggregate instantaneous peak throughput of 4 GB/s—unprecedented high performance for customers.

Memory Subsystem

The *ProLiant* 8500 memory subsystem includes dual 100-MHz buses, each with its own memory controller. The use of two memory buses increases memory bandwidth, reduces access conflicts, and increases the quantity of memory supported. Memory attached to each of the buses is cache-line interleaved, which means the buses share a common address range. One memory bus responds to even-numbered cache lines, and the other bus responds to odd-numbered cache lines. This configuration allows simultaneous use of both memory buses, which theoretically doubles throughput. It is especially advantageous for applications that access memory randomly. In random accesses, roughly half the requests at any one time are even-numbered lines, while the other half are odd-numbered lines.

The memory subsystem uses uniform memory access, which reduces latency and gives all processors equal access times to either memory bus. In systems using nonuniform memory access architectures, a processor has quick access to one memory bus but incurs a lag time (or latency) when accessing a second memory bus.

The *ProLiant* 8500 server supports up to 16 GB of error checking and correcting SDRAM that corrects all single-bit errors and detects double-bit errors. Memory is divided into eight banks, each consisting of two dual inline memory modules.

Although the Profusion chipset supports up to 32 GB of memory, industry-standard OSs provide only minimal support and scalability for this memory capacity, and enhanced support will not be available for an extended period. Also, in discussions with customers, Compaq learned that very few server implementations are fully configured with memory. With this in mind, Compaq used the internal server space to provide additional drive capacity in the *ProLiant* 8500 server. As customer requirements and OS capabilities increase in the future, Compaq will continue to modify servers to match these requirements.

Cache Accelerators

One of the main challenges of designing an efficient SMP architecture is maintaining cache coherency. To allow faster access to memory, most processors write data to cache memory rather than main memory. When a processor writes data to its cache, that cache has a newer copy of the data than main memory. Cache coherency ensures that the most recent copy of the data is read by any device that requests it. The cache coherency protocol essentially makes the cache look like main memory. Cache coherency is critical for the proper operation of an SMP architecture, and the performance and scalability of the architecture is affected by how efficiently it maintains cache coherency.

With multiple processor buses and a separate I/O bus, it is extremely challenging to maintain cache coherency in the 8-way architecture. Each memory access must look at, or snoop, the caches on its local processor bus and snoop all caches on the remote processor bus and the I/O bus. The amount of snoop traffic can significantly impact the scalability of the system.

The *ProLiant* 8500 architecture uses cache accelerators to minimize snoop traffic to the remote processor bus and I/O bus. The cache accelerators store the address and state of the data for all caches on their respective buses. The Profusion crossbar switch uses this information to determine whether to snoop the remote processor and I/O buses. Depending on how often a software application shares data, the reduction in snoop traffic can significantly improve overall system performance and scalability.

I/O Filter

The *ProLiant* 8500 server also includes three Compaq host-to-PCI bridges with prefetch buffers, so they act as caching bridges. The Profusion chipset contains a built-in I/O filter for the caching bridges on the I/O bus. The I/O filter enhances performance by reducing snoop traffic on the I/O bus. This I/O filter is designed to work with all three of the Compaq host-to-PCI bridges. When a processor requests a cache line with the intent to modify it, the MAC performs a lookup into the I/O filter to determine if that line resides in one of the caching bridges. If it does reside there, the MAC initiates a transaction on the I/O bus to invalidate that cache line. If the cache line is not present in one of the bridges, then no transaction is run on the bus.

Compaq Host-to-PCI Bridges

The Compaq host-to-PCI bridges included in the *ProLiant* 8500 server were engineered to enhance performance of the Profusion chipset. They include the following features:

- 64-bit, 66-MHz PCI support with a theoretical peak throughput of 533 MB/s
- Delayed transaction support
- Asynchronous design to accommodate multiple bus frequencies
- Multiple prefetch buffers
- PCI Hot Plug technology
- Peer-to-peer transaction support

64-Bit, 66-MHz PCI Support

Compaq licensed the host-to-PCI bridges to Intel, and they will ship with a Compaq logo on top in servers produced by other original equipment manufacturers. They are compliant with the PCI 2.2 Specification¹ and support 64-bit, 66-MHz PCI transactions.

The *ProLiant* 8500 supports three bridges with up to eleven I/O expansion slots. PCI Bus 1 operates at 33 MHz and includes three 64-bit PCI Hot Plug slots. PCI Bus 2 is also a 33-MHz bus and includes six 64-bit PCI Hot Plug slots. PCI Bus 3 operates at 66 MHz with two 64-bit PCI Hot Plug slots. Each of the 33-MHz buses has a maximum throughput of 267 MB/s. The 66-MHz bus has a maximum throughput of 533 MB/s. This combination of 11 slots provides adequate I/O bandwidth for the most demanding applications.

The Compaq 8-way architecture will support up to four Compaq host-to-PCI bridges. However, by incorporating only three host-to-PCI bridges, the *ProLiant* 8500 server has room on the I/O board to deliver an additional PCI slot, a meaningful advantage for enterprise customers who consistently request additional capacity.

Delayed Transactions

One of the most important features of Compaq's host-to-PCI bridge is that it supports delayed PCI transactions, which improve bus performance. In a delayed PCI transaction, the device that initiates the transaction polls the host-to-PCI bridge to determine if the requested data is there, rather than holding the bus while waiting for the data. This allows other devices to use the bus while the transaction is completed. To further improve bus efficiency, Compaq built in additional features to reduce the amount of polling.

Asynchronous Design

The bridge is split into two sections: upstream (host I/O side) and downstream (PCI side). All the functions on the upstream side are in the host processor clock domain at 100 MHz. All functions on the downstream side are in the PCI clock domain at up to 66 MHz. This asynchronous design is a flexible solution for supporting multiple clock speeds.

Multiple Prefetch Buffers

The Compaq host-to-PCI bridge is designed with multiple prefetch buffers to ensure optimum I/O-to-processor performance. Each buffer can hold multiple cache lines. These buffers are sized to provide optimal performance at a reasonable and cost-effective die size. Because of the delayed transaction support, the bridge can get data for multiple PCI devices concurrently. Other controllers can hold only a single cache line and a single request at a time.

¹ *PCI Local Bus Specification*, Revision 2.2, announced January 25, 1999, PCI Special Interest Group.

PCI Hot Plug Technology

The host-to-PCI bridges include electronics to support PCI Hot Plug technology, which was pioneered by Compaq. PCI Hot Plug technology provides the ability to remove, replace, upgrade, and add PCI adapters without shutting down the server. Any PCI adapter can be placed in a PCI Hot Plug slot; however, PCI Hot Plug device drivers and OS support are required to enable PCI Hot Plug functionality.

Peer-to-Peer Transactions

The host-to-PCI bridges also support PCI peer-to-peer transactions. They allow communications between two devices on the same PCI bus segment. They also allow communication across the I/O bus between PCI devices on different PCI bus segments (Figure 4).

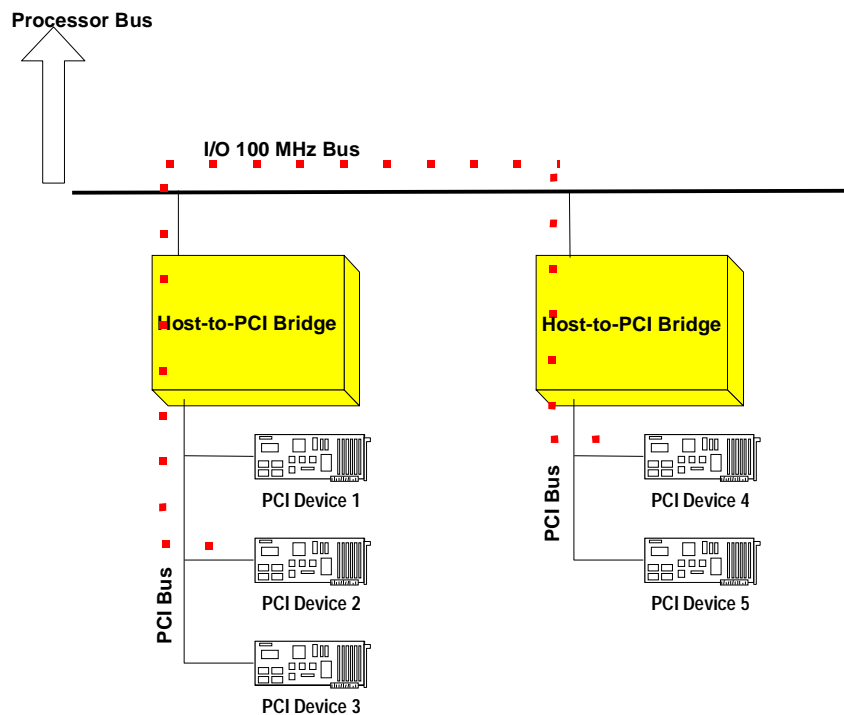


Figure 4: Schematic of peer-to-peer communications between PCI devices on different I/O bus segments.

Microprocessors embedded in the PCI devices make these devices “intelligent” and perform the communication tasks otherwise handled by the processors. Thus, the host-to-PCI bridges enhance performance by freeing the system processors to do other work.

Pentium III Xeon Processor Technology

The Pentium III Xeon processor is based on the Pentium II Xeon processor cartridge form factor. The Pentium III Xeon processor provides a higher internal core frequency than the Pentium II Xeon processor and includes 70 new streaming instructions that permit more efficient use of memory above 4 GB. The Pentium III Xeon processor includes the following features:

- 550-MHz core frequency
- 330-contact slot connector form factor (formerly called Slot 2)
- Full-speed cache on module
- 512-KB, 1-MB, or 2-MB Level 2 (L2) cache
- Katmai core logic

Level 2 Cache

Each Pentium III Xeon processor includes an L2 cache sized at 512 KB, 1 MB, or 2 MB that operates on a full-speed backside bus. The large cache size and full-speed backside bus enable very efficient access to the most frequently used data. This reduces traffic on the host bus and notably enhances performance.

Katmai Core Logic

Aside from speed and cache size increases, the main difference between the Pentium II Xeon and Pentium III Xeon processors is the migration to the Katmai core logic. The Katmai core logic adds the following enhancements to the existing Intel core:

- Memory streaming
- Concurrent Single Instruction Multiple Data – Floating Point (SIMD-FP)
- Katmai New Instructions (KNI)

Memory Streaming

Memory streaming is a set of techniques that increase the utilization of the memory bus during memory reads and writes. Processors must frequently wait for data not stored in cache memory to be retrieved from system memory. Because the system bus connecting the memory to the processor is slower than the processor core speed, this results in a performance bottleneck. Katmai's memory streaming enables the processor to prefetch data from system memory to cache memory.

During writes to memory, other new instructions can tell the processor to bypass the L2 cache and write data directly back to memory. Bypassing the L2 cache frees room in the cache for more data. This technique is useful when the software programmer knows in advance that the data will not be requested again in the near future.

Concurrent SIMD-FP

Concurrent SIMD-FP applies the same architecture used by MMX for integer operations to floating-point numbers. Floating-point numbers are used extensively in 3D graphics and scientific computations. Katmai processors have eight new registers that can each hold four single-precision floating-point numbers. Once the registers are filled, a single instruction can be used to add, subtract, or do other operations between the registers. For example, eight floating-point numbers could be added together using one instruction call.

Katmai New Instructions

KNI refers to the addition of 70 new instructions to the MMX instruction set. The new instructions primarily refer to the new SIMD-FP registers, but some new integer instructions also have been added. The new instructions are particularly useful for multimedia applications such as 3D rendering and video encoding and decoding.

Thermal Design

Compaq designed an innovative thermal cooling solution for the Pentium III Xeon processor that contributes to increased system reliability and performance scalability. Heat pipes transfer heat away from the Slot 2 cartridge thermal plate to cooling fins mounted atop the cartridge (Figure 5).





Figure 5: Heat pipes and fins for the Pentium III Xeon processor cartridge.

This top-mounted position and the direction of the cooling fins are advantageous in two important ways. First, this design cools efficiently regardless of the direction of airflow (vertical or horizontal). Therefore, the *ProLiant 8500* server can use the same Slot 2 cartridge and heat sink combination used in other servers, such as the *ProLiant 7000* and the *ProLiant 8000* servers. Second, because the heat fins are top mounted, the redundant PPM can be placed in its optimal position, next to the processor, to maintain short signal trace lengths across the 100-MHz system bus. The resulting improvements to power supply regulation and system bus signal quality increase system reliability.

To further reduce any cooling issues that may arise from packing so much power into such a small form factor, Compaq also improved the design of the *ProLiant 8500* chassis. As shown in Figure 6, a spring-loaded trap door has been added to each power supply bay. When a power supply is removed, the trap door will close to block the opening. This door preserves the air path required to cool the internal components of the server. It also allows the server to maintain availability with only one power supply and gives the administrator more time to replace the power supply.



Spring-loaded trap doors

Figure 6: Spring-loaded trap doors in the power supply bays of the ProLiant 8500 server.

The *ProLiant* 8500 server includes two 200-cubic-feet-per-minute hot-plug fans. These heavy-duty fans and thermal design modifications allow the *ProLiant* 8500 server to maintain proper airflow and temperature with only one hot-plug fan. The second hot-plug fan provides redundancy for the server.

For rack configurations of multiple *ProLiant* 8500 servers, Compaq designed a new rack series. The series includes racks with perforated doors and door inserts that provide better airflow and help these configurations maintain the proper temperature.

HIGH-AVAILABILITY TECHNOLOGIES

Tolerance of server downtime continues to decrease, as companies become more global, decentralized, and aware of downtime costs. The *ProLiant* 8500 server answers the demand for high availability by incorporating the following features into the server hardware:

- PCI Hot Plug technology
- Redundant NIC technology
- Redundant, hot-plug fans
- Redundant PPMs
- Redundant, hot-plug power supplies
- Hot-plug drives
- Auto processor bus recovery

Redundant components help ensure that the server is highly fault tolerant. In redundant configurations, if a power supply, fan, or PPM fails, the redundant standby component can take over operation with no downtime to the server. With the addition of hot-plug capabilities, server downtime can be eliminated while components are being replaced, added, or upgraded.

PCI Hot Plug Technology

The *Compaq ProLiant* 8500 server incorporates industry-standard PCI Hot Plug technology. This technology, pioneered by Compaq, adds hot-plug capabilities to existing PCI adapters for increased system availability and serviceability. Compaq led the industry by licensing its PCI Hot Plug implementation to Intel, thus broadening customer access to this industry-standard technology.

Compaq's implementation of PCI Hot Plug incorporates electronics into the server to control each PCI Hot Plug slot individually, so that a single slot can be powered down without affecting the operation of the other slots. Compaq's implementation of the PCI Hot Plug standard allows the following capabilities while the system is running:

- *Hot replacement* – replacing a single PCI adapter with a similar adapter
- *Hot upgrade* – replacing a PCI adapter with an upgraded adapter
- *Hot expansion* – adding a PCI adapter to an empty slot

The *ProLiant* 8500 server supports all three hot-plug capabilities for all 64-bit PCI Hot Plug slots. These capabilities can be added incrementally, depending on the OS implementation.

For more information about major OS vendors that support PCI Hot Plug, refer to the following Compaq technology briefs:

- *PCI Hot Plug Technology*, document number [ECG080/0698](#)
- *PCI Hot Plug Technology with Novell Architecture*, document number [ECG081/0698](#)
- *PCI Hot Plug Technology with SCO Software Architecture*, document number [ECG082/0698](#)
- *PCI Hot Plug Technology with Microsoft Windows Architecture*, document number [ECG071/0399](#)

Redundant Network Interface Controller Technology

Compaq's redundant Netelligent NIC technology allows two similar NICs to share a single instance of device driver code. One NIC becomes the active network controller and the other NIC acts as a standby controller. If the active NIC fails, the network traffic is automatically switched to the standby NIC. This redundancy eliminates the NIC or cable as a single point of failure. With PCI Hot Plug technology, the failed NIC can be replaced without rebooting the system. Thus, the end user can have continuous service and the administrator can greatly reduce planned and unplanned downtime.

The *ProLiant* 8500 server supports a Compaq NC3131 Fast Ethernet controller on the PCI local bus. The NIC has two ports for 10BaseT or 100TX Ethernet. The ports are seen as two separate local area networks, both supporting full-duplex Ethernet (20 Mb/s) and Fast Ethernet (200 Mb/s). Redundancy is available across the two ports with one NIC. In addition, the NIC is upgradeable with a daughter card to an additional two ports of 100BaseTX Ethernet, a single port of 100 BaseFX, or single-port Gigabit Ethernet SX or LX Fiber support. Redundancy is also available here; the NIC can be configured to failover from the Gigabit daughter card to the dual 10/100 ports in the event of a failure, thus allowing continued access to the network.

Redundant Hot-Plug Fans

The *ProLiant* 8500 server includes two hot-plug system fans in the I/O module of the server (Figure 7). One fan is required to cool the system, and the other fan is redundant. If the primary fan fails, a system alert triggers the secondary (or redundant) fan to take over automatically. The redundant hot-plug system fan protects the various server components from overheating and helps prevent possible system interruption.




Figure 7: Hot-plug fans in the *ProLiant* 8500 server.

The fan control logic operates independently from the system processor and OS. Regardless of which OS is running or its condition, the fan control logic and the server temperature monitoring circuitry work together to cool the server effectively.

The fans have two speeds: normal and high. Under typical operating conditions, the normal fan speed is sufficient to cool up to 8 processors, 4 hard disk drives, and 11 I/O expansion boards. As a precaution against increases in ambient air or localized internal temperatures, Compaq established three thermal trip points, as shown in Figure 8. If the first thermal trip point is reached, the fans adjust automatically to high speed to increase cooling.

Users can select whether to enable a server response to the second thermal trip point. When enabled, if the internal air temperature exceeds the second trip point, the OS initiates a controlled shutdown of the server. Simultaneously, a warning message is sent to the Integrated Management Display (IMD) and Compaq Insight Manager. The server restarts automatically when it has cooled sufficiently.

Finally, in the unlikely event server temperature continues to rise, a third trip point shuts down the power supplies to protect critical components from overheating. The server restarts automatically when it has cooled sufficiently.

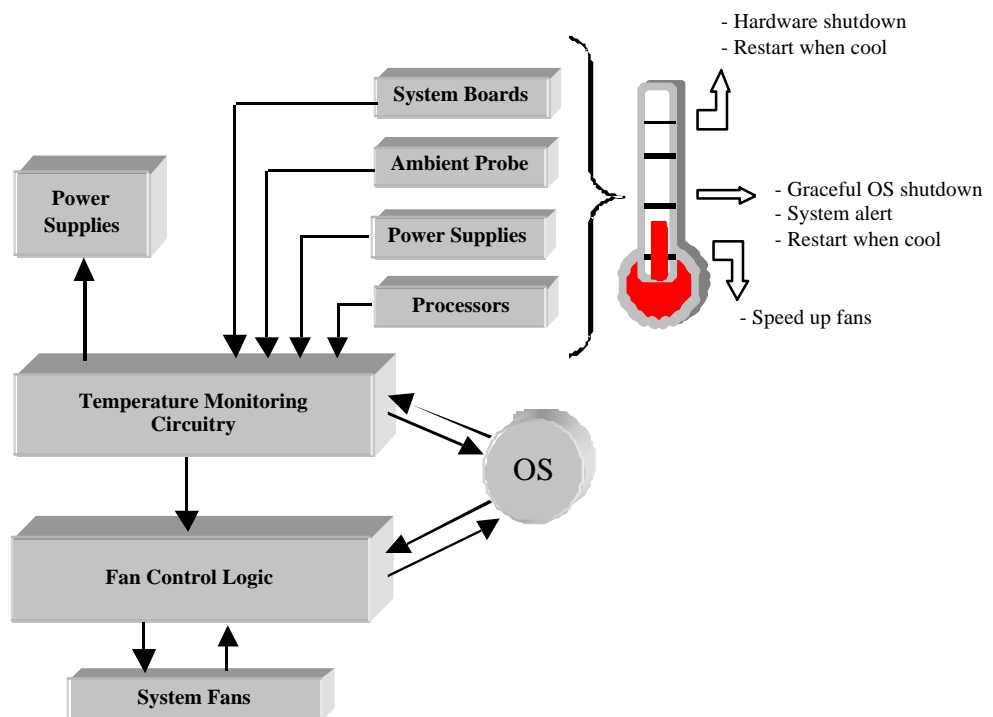


Figure 8: Schematic diagram of fan control logic.

Redundant Processor Power Modules

Each processor in the *ProLiant* 8500 server has its own redundant PPM. PPMs are also known as voltage regulator modules. An additional PPM is included in the server to power the system board. The PPMs deliver the precise voltage required by the processors.

Compaq PPMs are programmable and support all current and future versions of Xeon processors. Each PPM has redundant circuitry. If a failure occurs in one circuit of the module, the other circuit automatically takes over the task of regulating power to the processor.

Redundant Hot-Plug Power Supplies

The *ProLiant* 8500 server uses one of the newest Compaq technologies to reduce downtime: intelligent, redundant, hot-plug power supplies. The server supports two dual-rated 1150-watt/500-watt power supplies with new handles for easier insertion and removal. Even if one of these power supplies fails, its internal fans will continue to operate and help cool the server. A single power supply running at 220 volts can support a fully loaded *ProLiant* 8500 server, but a second power supply adds redundancy and fault tolerance.

The power supplies have an embedded microcontroller and an inter-integrated circuit bus that communicates the status back to the system health drivers.² The health drivers then send status information to the IMD and to Compaq Insight Manager.

The microcontroller performs the following functions:

- **Self-test.** The intelligent power supplies perform their own power-on self-test (POST) that checks temperature, RAM integrity, ROM version, analog-to-digital conversion accuracy, and nonvolatile memory integrity of the power supply.
- **Auto Line Sensing.** Because the power supplies are rated for both 1150 watts (at 220 volts) and 500 watts (at 110 volts), a line-sensing feature automatically recognizes which line voltage is connected. The customer does not need to configure the supply for voltage.
- **LED Standby Controls.** The microcontroller also drives two status light-emitting diodes (LEDs) on the back of the power supply that indicate power status, as shown in Figure 9. This gives the user vital status information at a glance.




Figure 9: Power supply LED status indicators.

² For additional details about Compaq's intelligent power supplies, see the technology brief *Intelligent Power Supply Technology*, document number [ECG016/0198](#).

Power-Down Controls

As shown in Figure 10, the server power switch connects to the Advanced Configuration and Power Interface (ACPI) controller rather than directly to the power line. This means that turning off the power switch does not turn off power instantly. Instead, the ACPI controller communicates with the system board and the OS. The power-down controls can be configured in one of three ways:

- Secure mode – Power supply will not shut down if power switch is pressed.
- Delay mode – Power supply will shut down after a 10-second delay. This allows time to reactivate the switch if it was mistakenly pressed.
- Graceful shut down – A 10-second delay counter begins as a message warning of an imminent power-down goes to Compaq Insight Manager and the IMD. Because of the communication between the ACPI controller and the OS, the OS can delay the power down until all running applications have properly closed. Should the OS not be running, the power would go down when the delay counter expired.

The power-down control feature is supported under Microsoft Windows NT and Novell NetWare. SCO is planning to implement support for the power-down control in the future.

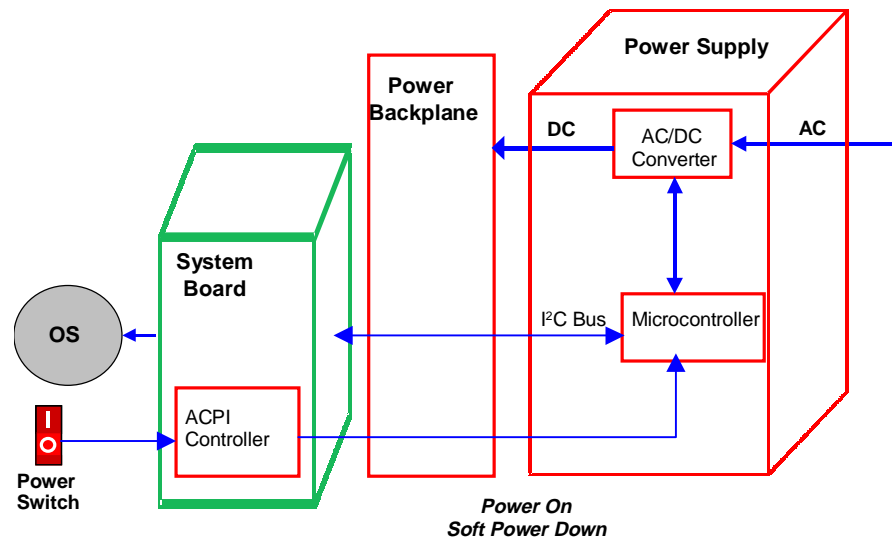


Figure 10: Power supply schematic.

Load Balancing

The power supplies automatically load balance to within 10 percent of the average current of all supplies. Before the power comes up, the master power supply matches its output load rails to those of the other power supplies so that the load is automatically balanced across all power supplies. If one supply fails, the load balancing feature allows the other supplies to pick up the remaining power load quickly.

Power Factor Correction

Compaq's intelligent power supplies have built-in power factor correction to synchronize the voltage and current phases. This helps ensure that maximum power is available. It also reduces the return current in the ground line, which reduces the overall power consumption.

Hot-Plug Capability

When a power supply is hot swapped, the microcontroller governs the power-up and power-down sequence, enables the health drivers to calculate redundancy “on the fly,” and informs the health drivers of changes in load. Compaq hot-plug power supplies are readily identifiable by their port-colored bezel, as shown in Figure 11.




Figure 11: Front view of hot-plug power supply with port-colored bezel.

Hot-Plug Hard Disk Drives

The *ProLiant* 8500 server supports four 1-inch, Wide Ultra-2 SCSI, internal hot-plug hard disk drives. The internal drives can be controlled by an Integrated Smart Array Controller (see “[Storage Technologies](#)”) to optimize them for OS and boot support. When the server is configured for maximum capacity, it can support 73 GB of internal storage. This number will increase in the future as new drive technologies emerge.

Auto Processor Bus Recovery

Auto processor bus recovery is a feature of the *ProLiant* 8500 server that maintains server availability in the unlikely event that one of the cache accelerators fails. If a cache accelerator fails, the server will go through the Automatic Server Recovery (ASR) process and reboot. Typically, a failed cache accelerator would prevent the server from rebooting. Auto processor bus recovery, however, enables the server to recover and maintain availability of the processor bus with the good cache accelerator. With the server still available, the administrator can replace the failed accelerator at the most convenient time.

STORAGE TECHNOLOGIES

As corporations commit more critical information to computers, the requirements for more storage capacity and faster data movement are increasing at a startling rate. The *Compaq ProLiant* 8500 server includes new storage technologies that help meet these requirements. New hard drive technology and a specially designed integrated array controller allow the *ProLiant* 8500 server to deliver the highest levels of external storage capacity faster and with more reliability.

Hard Drive Technology

The *Compaq ProLiant* 8500 server uses low-voltage differential (LVD) Ultra-2 SCSI hard drives. Ultra-2 SCSI hard drives use a 16-bit bus and a 40-MHz clock rate to achieve maximum transfer rates of 80 MB/s. This increase in bus speed provides the performance demanded by today’s data-intensive applications. Compared with other alternatives available today, LVD Ultra-2 SCSI technology prevents system performance from becoming I/O limited as drive throughputs

increase. In addition, LVD Ultra-2 SCSI technology integration allows faster drive-to-host transfers, enabling the I/O bus to accommodate more devices.

LVD signaling allows use of longer cables (up to 40 feet or 12 meters) than earlier forms of SCSI, which used a single wire that ended in a terminator with a ground. Ultra-2 SCSI sends the signal over two wires; the data is represented as the difference in voltage between the two wires. With this approach, even if the signal attenuates while traveling additional distance, the relative differential between the two wires remains the same. This allows longer cables. In addition, the improved noise immunity of LVD leads to greater signal integrity and data reliability. A low-voltage differential also reduces power requirements and manufacturing costs.

Integrated Array Controller Technology

The *Compaq ProLiant* 8500 server is equipped with an Integrated Smart Array Controller. The integrated controller, also known as RAID on a chip, or ROC, is an embedded hardware-based RAID solution that enhances system reliability and improves host processor utilization. It is ideal for data center servers, like the *ProLiant* 8500 server, in which internal storage is optimized for OSs and swap space.

RAID implementation enhances data integrity and recovery through data striping and parity generation for the data rebuild process. The integrated array controller improves host processor utilization by performing all RAID functions, thus freeing the host processor for other tasks. The controller is fully integrated and embedded on the motherboard of the *ProLiant* 8500 server, which improves I/O efficiency and frees a PCI slot for another controller. The integrated controller can also queue tasks and data, thereby increasing performance.

The integrated controller consists of a processor, memory controller and hardware exclusive-OR (XOR) engineer, SCSI chip, PCI bridge, and an internal bus structure embedded on a single substrate. The controller has fewer parts than typical array controllers, since all of the components and connections are combined into one integrated circuit. Fewer parts leads to greater reliability and increases mean time between failures.

The integrated array controller includes a 16-MB read-only memory, with an 8-MB read-ahead cache. This memory configuration is optimal for OS and boot support. Since the memory is read only, data loss is not an issue and battery backup is not required.

This controller provides two Wide Ultra-2 SCSI channels. One channel is dedicated to the internal drive bay. The other channel is connected to the external SCSI connector. It uses a high-density connector, and only supports tape backup devices.

Configuration Utilities

The Option ROM Configuration for Arrays Utility is firmware that comes with the integrated controller in the *ProLiant* 8500 server. This firmware allows users to set up a simple configuration during POST. The support includes the ability to create a single logical drive, delete a single existing logical drive, view the current configuration, and assign an online spare as part of the logical drive creation process. The integrated controller may be configured to initialize as an intelligent I/O processor supporting either dual-channel SCSI operation or RAID.

For more advanced configuration support, Compaq provides the Array Configuration Utility. The Compaq Array Configuration Utility is an easy-to-use graphical configuration utility to help users set up and change drive array configurations. It can be used to configure the standard internal drive controller as well as any optional controller boards that are installed.

RAID Support

The integrated array controller supports RAID levels 0, 1, 0+1, and 5 to ensure data integrity and availability. When configured for RAID 5, and prior to any possible drive failure, the controller proactively generates parity data so that it can keep all data available and the server running during replacement of any failed drive. The integrated array controller also supports online spare disk drives. These spares are powered up but not active; they are held in reserve in case one or more of the active drives should fail.

The controller supports logical drive expansion, which consists of three different types of operations: capacity expansion, RAID migration, and distribution factor changing. Capacity expansion allows an existing logical drive to be grown by either adding additional physical drives or using existing free space on the current physical drives. RAID migration allows an existing logical drive to change its RAID level to a new one. Distribution factor changes the amount of data that is stored on each physical drive for a RAID stripe. All of these operations can be performed without disturbing the data on the drives.

The integrated array controller in the *ProLiant* 8500 server can be bypassed using the Integrated Array Bypass Kit and an optional controller such as the *Compaq Smart Array* 3200 controller. The bypass kit is shipped standard with all *ProLiant* 8500 servers.

MANAGEMENT TECHNOLOGIES

Fault-tolerant, redundant features are only part of what makes Compaq servers highly available. Compaq offers a range of hardware- and software-based server management tools that can be used to reduce downtime and increase productivity. The *ProLiant* 8500 server includes the following standard tools available to all Compaq servers:

- **Automatic Server Recovery.** ASR automatically resets the server after a critical hardware or software error. ASR can reset the server to the OS or to the System Partition Utilities. The ASR reset function is based on a hardware timer working in conjunction with the Server Health Drivers. If the Server Health Drivers can no longer reset the hardware timer after some user-specified amount of time, the server is automatically reset. ASR can also be configured to page the administrator when an ASR event occurs.
- **System Partition Utilities.** These utilities assist the system administrator in diagnosing server problems, configuring new hardware, and upgrading ROMs. For example, an administrator can analyze the system configuration files to verify that no interrupts or other conflicts are causing system failures. The System Partition Utilities are available when the server is reset. ASR can be configured to reset automatically to the System Partition Utilities; or the administrator can reboot to the System Partition Utilities through Compaq Insight Manager.
- **Compaq Insight Manager.** This software-based management tool monitors more than 1,000 management parameters to allow fault prediction and alerting, asset and configuration inventory, and performance monitoring. Through Compaq Insight Manager, a remote administrator can request an OS console operation, reset the server, view utilities, and view system configuration information. See the white paper *Compaq Insight Manager 4.23*, document number [0167-0699-A](#), for more details.
- **Asynchronous Insight Management.** Asynchronous Insight Management provides access to Compaq Insight Manager through a dial-up networking, or asynchronous, connection. This gives an administrator additional flexibility in troubleshooting problems. If the server OS is still functioning, an administrator can dial in remotely and access Insight Manager through a point-to-point protocol (PPP) connection.

In addition to these utilities, the *ProLiant* 8500 server also ships standard with additional management technologies that reduce downtime and increase productivity even more. Remote-flash Redundant ROM allows administrators to update the ROM remotely, and the auto-default ROM configuration provides default configuration settings for most devices. The IMD and IRC provide essential information at a glance and allow users easy remote access to their servers. For additional remote management functionality, customers can request that the optional Remote Insight Board/PCI be added to the *ProLiant* 8500 server.

Remote-Flash Redundant ROM

The *ProLiant* 8500 server is equipped with a Remote-flash Redundant ROM that improves manageability. Using the Compaq Remote ROM Flash Utility, administrators can flash the system ROMs for a wide range of Compaq servers, locally or across the network. The Remote ROM Flash Utility is a combination of components that allows administrators to upgrade the system ROMs on servers from a single point of execution. The ROM upgrades can either be flashed individually or batched together to perform multiple ROM upgrades in a single step. The following components are used in the ROM flash process:

- **User interface.** The user interface is command-line driven and allows administrators to tailor functionality to a specific environment.
- **System ROM Flash Service.** The System ROM Flash Service allows the user interface to perform essential functionality on remote servers for the ROM upgrade process, including remote system shutdown, mounting, and unmounting the system partition.
- **Compaq System Configuration Utility.** The latest version of the Compaq System Configuration Utility is integrated within the Compaq Remote ROM Flash Utility. Parts of this utility have been enhanced to allow flawless firmware upgrades.
- **ROMPAQ Utility and image files.** Generally, any version of the ROMPAQ Utility and image files can be used. The latest ROMPAQ Utility and images are available from the Compaq website at www.compaq.com/support/files/server/ROMPAQ/index.html. It is possible to maintain a common collection of ROM images in a particular directory on the local administrative client or on a shared resource.

All of these components work together to deliver an easy-to-manage solution for administrators who need to manage large-scale ROM upgrades, like those required for a year 2000 update, or a few ROM upgrades at a time.

Compaq Remote-flash Redundant ROM provides a unique redundancy feature that helps ensure system availability. During the ROM flash process, it is possible for a problem to occur. Typically, a failure during the ROM flash would shut down the server and require the administrator to replace the ROM. This could be a significant problem if, for example, the flash were scheduled for a time when the administrator was unavailable. Theoretically, the server could be down for hours until the administrator became available to manage the situation. Compaq's unique redundancy feature helps ensure system availability because it gives the system the ability to recover the last known good system ROM in the event that the current system ROM is corrupted.

When a *ProLiant* 8500 server leaves the factory, both system ROMs contain the same image. Through subsequent boots of the server, if integrity errors are detected by the boot block, the system will automatically launch the redundant image and continue the POST process. If the redundant ROM is launched, the user will see an error message identifying the faulty system ROM.

Auto-Default ROM Configuration

When a *ProLiant* 8500 server is first powered on, the system ROM detects the unconfigured state of the hardware and provides default configuration settings for most devices. By providing this initialization, the system can run *Diagnostics* and other software applications before running the normal *SmartStart* and *System Configuration* applications.

If the user inserts a *System Configuration*, *Diagnostics*, or *SmartStart* CD in the CD-ROM drive before powering on the server, the system ROM will boot using that CD. If the system ROM does not detect one of these CDs, the user will be prompted for the intended OS. The system will reboot if any OS-dependent configurations have changed with the new OS selection. If the selected OS-dependent configurations match the current configurations, the system will boot normally. The OS selection may be changed during subsequent reboots.

Integrated Management Display

The IMD (Figure 12) is a backlit liquid crystal display that allows administrators to enter server and contact information. Just as importantly, the IMD displays critical information, warning messages, and error messages in an easy-to-understand format.




Figure 12: Integrated Management Display with sample alert message.

A 5-volt auxiliary power line supplies power to the IMD whenever the server power line is plugged in. Even if the server power goes down, the auxiliary line continues to power the IMD so that critical information, such as the last error message and the administrator's name, can be obtained.

The IMD unit contains its own static memory. This memory contains the software code to control the IMD and the text messages shown on the display. All software enhancements added in the future will be available through a flashable ROMPaq. During POST, all event logs, system information, and administrator information are uploaded from the system nonvolatile RAM. During runtime, new events are stored in the nonvolatile RAM and the IMD SRAM. If the IMD unit is removed and replaced, all key event and system information is still available.

Integrated Remote Console

IRC consists of a Compaq designed ASIC and associated firmware. IRC allows out-of-band, or asynchronous, management capabilities such as remote console and remote reset. These capabilities are independent of the state of the network OS. With IRC, an administrator has full text mode video and keyboard access to the server to perform diagnostics, reset the system, watch the reset process remotely, and view reset sequences, regardless of whether the OS is functional.

For more information about IRC, refer to the Compaq technology brief *Remote Server Management with Integrated Remote Console*, document number [582A/1096](#).

TCP/IP – transmission control protocol/internet protocol

Remote Insight Board/PCI Option

The optional Remote Insight Board/PCI provides seamless remote access and control independent of both the server OS and the server hardware. The Remote Insight Board/PCI is also fully integrated with Compaq Insight Manager. An administrator can remotely access the Remote Insight menu through the Insight Manager interface or through a web browser. The web browser uses the PPP dial-in connection with its user authentication and security features or a TCP/IP connection, rather than an unsecured Internet web connection.

Like IRC, the Remote Insight Board/PCI can be used when the server OS is not functioning. Unlike IRC, Remote Insight can be used even when the server hardware is not functioning or power is lost. If a server power loss occurs, Remote Insight's battery backup provides power for at least 30 minutes—long enough to send out power-down alerts.

Remote Insight provides several important troubleshooting features that none of the other Compaq tools can. Remote Insight provides both numeric and alphanumeric paging. Other essential troubleshooting features are provided by Remote Insight through its multiple diagnostic capabilities. It captures videotext of the last two reset sequences in addition to the last failure sequence. It can also access several management logs that are unavailable through IRC. For additional information about the Remote Insight Board/PCI, see the white paper *Compaq Remote Insight Board/PCI Product Overview*, document number [ECG042/0798](#).

SERVICEABILITY FEATURES

The *ProLiant* 8500 server defines a new level of serviceability with a unique modular design that includes five components: the processor/memory module, the media module, the I/O module, the power supplies, and a system midplane. As shown in Figure 13, the removable, drawer-like modules and power supplies connect to a common system midplane. This virtually cable-free design eases serviceability and future upgrades, since a module can be removed and replaced with little effort.




Figure 13: Modularity in the ProLiant 8500 server.

TECHNOLOGY BRIEF (cont.)

The system midplane is the backbone of the *ProLiant 8500* server: all power and signals pass through it. The midplane is held in place with two thumbscrews and allows a virtually cable-free design and tool-free serviceability. High-density connectors and a positive interlocking mechanism allow the midplane to provide a highly reliable interconnect.

The design of the chassis has been enhanced with a sliding top access panel that provides easy access to the I/O module and the processor and memory module. The modularity of the *ProLiant 8500* server permits the chassis to be emptied completely to eliminate heavy lifting. Cam levers allow each module to be easily removed, and four lift assist handles make rack mounting the chassis even easier. Once the server is in the rack, the components can be quickly reinstalled. These serviceability features are shown in Figure 14.




Figure 14: Serviceability features of the *ProLiant 8500* server.

System Interconnect Status Indicators (Figure 15) improve serviceability and further reduce downtime associated with upgrades and field maintenance. The status indicators are visible from the exterior of the server and provide a closed-loop checking mechanism for verifying proper component mating and interconnections between critical server components. LEDs inside the sliding top panel provide visual assistance in determining whether a component or module is not properly installed. If an LED is on, its associated component needs to be reseated. Labels next to the LEDs identify the component causing the problem.

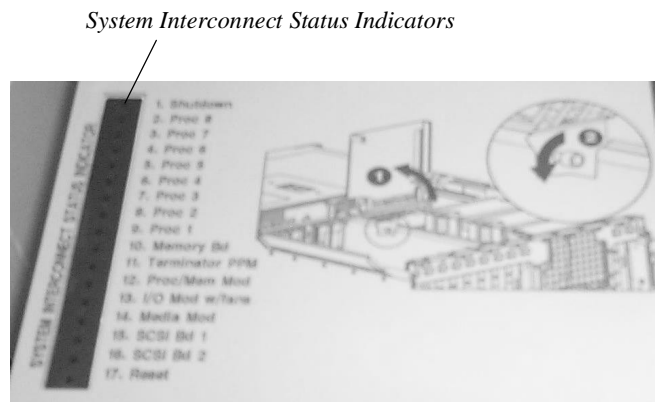


Figure 15: System Interconnect Status Indicators in the *ProLiant 8500* server.

CONCLUSION

The *Compaq ProLiant 8500* server offers customers a high-performance, highly scalable solution for their data center needs. The server also provides all the high-availability features customers have come to expect from Compaq servers. In addition, Compaq has integrated key management technologies and revolutionary serviceability features that further reduce downtime and improve manageability.

For customers needing greater performance, availability, and scalability in the data center, the *ProLiant 8500* server is an excellent choice. Development of this server demonstrates Compaq's commitment to provide industry-standard servers that exceed customer expectations. As business requirements become more demanding, Compaq will continue to deliver servers that meet those requirements and can be expanded to meet future needs.