

TECHNOLOGY BRIEF

July 1997

Compaq Computer
Corporation

CONTENTS

Introduction.....	3
ProLiant 5000	
Overview	4
Processors.....	4
I/O.....	4
Memory	5
System Architecture.....	6
Storage	6
Power.....	7
ProLiant 5000	
Architecture	7
Performance Implications of the EISA Bus.....	9
Configuration Recommendations	10
PCI Peripherals	10
Example	
Configurations.....	13
Conclusion	15
Appendix.....	16

Configuring the Compaq ProLiant 5000 Server for Peak Performance

EXECUTIVE SUMMARY

The development of the Intel Pentium Pro processor paved the way for the next technology leap in capturing, processing, and managing business-critical data with industry-standard platforms. Compaq has designed a high performance, high availability platform to exploit the full power of the Pentium Pro and the latest in high-performance options. The Compaq ProLiant 5000 Server provides midrange-class performance and reliability for running business-critical applications in the data center, work group, or department.

The ProLiant 5000 is based on the Intel 82450GX chip set. The key to realizing peak performance from any Pentium Pro server, including the ProLiant 5000, is optimal configuration of the I/O subsystems. The ProLiant 5000 contains architectural changes and enhancements, including use of the Intel 82450GX chip set, that are not available in the Compaq ProLiant 4500 Server. Because of these enhancements, achieving peak performance from the ProLiant 5000 system requires careful planning of initial system setup and configuration.

This brief explains important characteristics of Pentium Pro servers and provides technical information for configuring the ProLiant 5000 to optimize overall system performance—initially and later, as user needs expand.

Please direct comments regarding this communication to the ECG Technology Communications Group at this internet address:
tech_com@bangate.compaq.com

COMPAQ

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, Compaq Insight Manager, PageMarq, ProLiant, ROMPaq, and QVision are registered with the United States Patent and Trademark Office.

Netelligent, Systempro/XL, Net1, Vocalyst, PageMate, and SoftPaq are trademarks and/or service marks of Compaq Computer Corporation.

Microsoft, Windows NT are trademarks and/or registered trademarks of Microsoft Corporation.

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

©1997 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.

Configuring the Compaq ProLiant 5000 Server for Peak Performance

Fourth Edition (July 1997)

Document Number 679A/0697/ECG

INTRODUCTION

The Compaq ProLiant 5000 Server is a highly successful member of the ProLiant family of high availability/high performance servers. It is the midrange alternative for the distributed enterprise, delivering midrange-class performance and reliability. Its architecture is significantly different from that of its predecessor, the Compaq ProLiant 4500 server, because the ProLiant 5000 is designed to exploit performance advantages and unique characteristics of Pentium Pro processors, plus the latest in high-performance options. With proper system configurations, the quad-processor ProLiant 5000 will typically outperform an equivalent Pentium-based system by a significant margin.

The key to realizing peak performance from any Pentium Pro server, including the ProLiant 5000, is optimal configuration of the I/O subsystems. Because of architectural changes and enhancements incorporated into the ProLiant 5000, careful consideration must be given to the initial system setup and configuration. Specifically:

- The Intel Pentium Pro processor is optimized for performance using Peripheral Component Interconnect (PCI) peripherals. Therefore, the ProLiant 5000 achieves best performance with PCI peripherals. The ProLiant 5000 supports dual-peer PCI buses that provide aggregate I/O throughput capability as high as 267 MB/s. The dual-peer PCI buses increase configuration flexibility and allow higher levels of overall performance. However, attaining peak performance requires carefully balanced I/O loading across both PCI buses. This means careful planning of which expansion slot to use for each device in the initial configuration, as well as planning for future expansion.
- The ProLiant 5000 also includes an Extended Industry Standard Architecture (EISA) bus for backward compatibility with installed EISA option boards. The EISA bus is bridged off one of the PCI buses. Since the ProLiant 5000 is optimized for PCI and the EISA bus is not the primary bus architecture in the system, achieving optimal overall system performance requires care in selecting and configuring any EISA peripheral to be used.
- Compaq uses the Intel 82450GX chip set in the ProLiant 5000. To ensure that its customers do not experience the Intel-documented limitations of this chip set, Compaq has developed its own system board, and unique workarounds for many of these limitations are designed into Compaq hardware. For the rest, the only workaround is careful system configuration. Fortunately, most of the Intel chip set limitations are related to lower performance I/O adapters, so use of high-performance PCI peripherals and careful system configuration can minimize the impact of the Intel 82450GX chip set limitations on performance.

The configuration challenges cited above are not limitations specific to Compaq or the ProLiant 5000 platform. They are applicable to all 82450GX-based Pentium Pro systems. Compaq systems embody features specifically engineered to overcome most of these chip set issues.

The purpose of this brief is to explain some important characteristics of Pentium Pro servers and to provide technical information and recommendations for configuring the ProLiant 5000 to optimize overall system performance—initially and later as user needs expand.

PROLIANT 5000 OVERVIEW

Compaq continually strives to provide the highest levels of performance, scalability, and reliability for our customers. For example, the Pentium-based ProLiant 4500 is so highly optimized for EISA that it outperforms most competitors' PCI-based Pentium servers in file and print benchmarks, as well as in database benchmarks. Providing further significant performance improvements for customers deploying business-critical database applications required a new approach for the next generation Compaq server—the ProLiant 5000.

This section provides a brief overview of key subsystems in the ProLiant 5000.

Processors

The ProLiant 5000 uses the Intel 82450GX chip set, which supports up to four Pentium Pro processors. The ProLiant 5000 accommodates up to four Pentium Pro processors, two per processor board.

Customers may choose between these Pentium Pro processor options:

- 200-MHz processors, each with an integrated 16-KB L1 cache (8 KB for data and 8 KB for instructions) and an integrated 256-KB L2 cache.
- 166-MHz processors, each with the same L1 cache and an integrated 512-KB L2 cache.
- 200-MHz processors, each with the same L1 cache and an integrated 512-KB L2 cache.

The L2 cache on each of these processors runs at processor speed and is protected by Error Correcting Code (ECC). The L2 caches are 4-way set associative for significant improvement in cache hit rate as compared to 2-way set associative caches in Pentium systems, including the ProLiant 4500. In typical applications, the higher hit rate allows the Pentium Pro L2 cache to outperform 2-way caches that are twice as large.

With its faster L2 cache, the 166-MHz processor provides a data throughput rate of more than 1 GB/s. That is 2.5 times the throughput rate of a 166-MHz Pentium processor.

To select a ProLiant 5000 Server with the right Pentium Pro processor option, two key factors must be understood:

- The demands of the computing environment.
- The relationship between processor speed and the size and speed of the cache.

In a uniprocessor ProLiant 5000 Server, the 200-MHz Pentium Pro processor with 256-KB L2 cache is the better choice, particularly in a file and print environment where I/O capability and the total memory available have the greatest effect on performance. However, for scaling, particularly in applications or database environments, cache size is more important than CPU clock speed. Therefore, the 166-MHz Pentium Pro processor with 512-KB L2 cache scales better for multiprocessor transaction processing. It follows, then, that a ProLiant 5000 with four 200-MHz processors, each with a 512-KB L2 cache, provides the best performance in all environments.

I/O

The ProLiant 5000 is optimized for PCI devices. There are several reasons for the shift from EISA to PCI technology:

First, I/O requirements for application servers are growing and will continue to grow into the future. PCI technology will support growth well beyond the capabilities of EISA. Because of their high levels of processing power, initial Pentium Pro quad-processor servers will require I/O bandwidth that is close to the maximum bandwidth available with EISA.

Increased use of higher bandwidth peripherals will continue to raise I/O performance requirements. At maximum bandwidth:

- Drive arrays could require up to 35 MB/s.
- FDDI requires up to 15 MB/s.
- Full duplex 100baseTX requires up to 20 MB/s.
- An ATM card requires up to 20 MB/s.

While the maximum bandwidth of an EISA bus is 33 MB/s, a single PCI bus provides a maximum bandwidth of 133 MB/s.

Second, future needs will require increased I/O expansion. Slot expansion requirements will grow with increasing processing power. While EISA is logically limited to 15 expansion slots, PCI can be logically expanded to hundreds of slots or virtual slots by means of hierarchical bus expansion designs using PCI-to-PCI bridges.

Third, PCI is a young and developing technology. Further performance enhancements are already being specified, including a 64-bit implementation.

The optimal time for customers to transition to the Pentium Pro/PCI-based ProLiant 5000 will vary, depending upon their computing needs and current technology. Customers using processor- and I/O-intensive applications such as databases will benefit most by a prompt transition to the ProLiant 5000.

Memory

Compaq believes that the memory subsystem significantly differentiates the ProLiant 5000 from competitors' Pentium Pro servers. In designing memory for the ProLiant 5000, Compaq is leading the industry in the move from Single Inline Memory Module (SIMM) to Dual Inline Memory Module (DIMM) technology.

The ProLiant 5000 contains 16 sockets for DIMMs. The system comes standard with 64 MB or 128 MB of RAM and is expandable to 4 GB of ECC-protected memory using 256 MB EDO DIMMs.

For midrange-class servers, DIMMs offer these advantages over SIMMs:

- **DIMMs offer a more flexible and cost effective upgrade path.** A 4-way interleaved system (which corresponds to a 256-bit-wide memory subsystem and is used in ProLiant 5000) provides maximum performance and maximum flexibility. Although using a 2-way interleaved system (which corresponds to a 128-bit-wide memory subsystem) will allow memory upgrade increments of 2 DIMMs or 4 SIMMs, it compromises performance.
- In a 4-way interleaved system, a memory upgrade requires either 8 SIMMs per bank or 4 DIMMs per bank. Upgrades using DIMMs, therefore, consume fewer memory slots and allow the highest total memory capacity.
- **DIMMs consume significantly less power.** SIMMs operate at 5 V. Because DIMMs operate at 3.3 V, they consume approximately 50% less power and dissipate significantly less heat than SIMMs, thereby improving overall system reliability and availability.
- **DIMMs provide larger memory capacities.** If the ProLiant 5000 had been designed using SIMM technology, the total system memory capacity would be limited to 1 GB. The switch to DIMMs increased total system memory capacity to 4 GB.

Some might argue that reuse of existing SIMMs for memory upgrades justifies using SIMMs in Pentium Pro servers. For optimal performance, however, 82450GX-based Pentium Pro servers must use DIMMs or SIMMs with an access time of 60 ns or faster. The access time of the majority of SIMMs now in the field, however, is 70 ns or slower.

System Architecture

Figure 1 illustrates the ProLiant 5000 Server architecture. The system features a Gunning Transceiver Logic processor bus (referred to in this paper as the host bus) that supports high-speed, low voltage transfers. It is a 64-bit split-transaction bus operating at 66 MHz, with a throughput capacity of 540 MB/s. Data integrity is ECC protected.

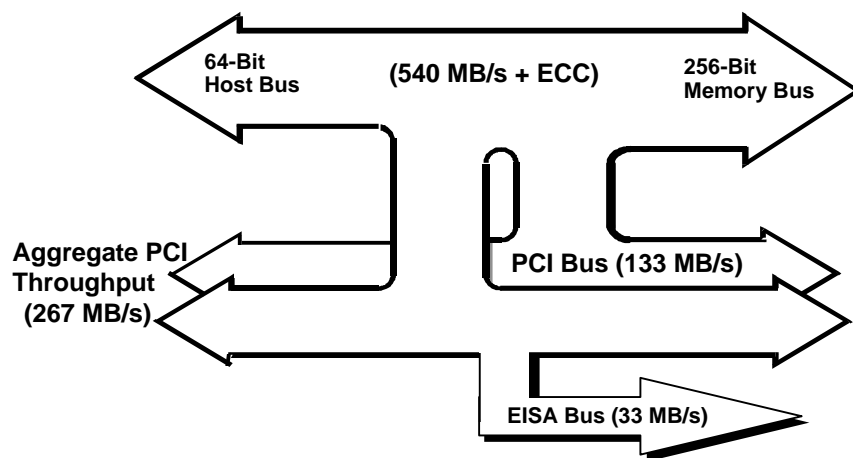


Figure 1. Compaq ProLiant 5000 Server Architecture

The ProLiant 5000 supports concurrent cache, memory, and I/O bus operation to optimize performance for a wide variety of processor- and I/O-intensive applications. The 4-way set associative L2 cache integrated into each Pentium Pro processor is accessed at full processor speed through a separate, dedicated processor-to-cache bus. These 4-way L2 caches minimize wait states by keeping data flowing to the processor. To fully exploit the power and scalability of Pentium Pro in four-processor-capable servers, the ProLiant 5000 also has a 256-bit memory bus.

The ProLiant 5000 design includes dual-peer PCI buses, each operating at 33 MHz and providing an aggregate bandwidth of 267 MB/s. The system also includes an EISA bus that provides a maximum bandwidth of 33 MB/s.

Storage

The ProLiant 5000 contains three removable media drive bays. A 3.5-inch, 1.44-MB diskette drive and a CD-ROM drive are standard. One half-height removable storage device, such as a tape drive, can be installed in the third bay. The server features hot-pluggable hard drive bays. In tower models using the 9.1-GB drives available today, four hot-pluggable drives provide maximum internal storage of 36.4 GB. In rack models, five hot-pluggable drives provide maximum internal storage of 45.5 GB. As larger capacity drives become available, maximum internal storage will increase.

Optional mass storage devices can be connected to the ProLiant 5000 via the External Fast-Wide SCSI-2 port on the back of the unit or the Fast-Wide SCSI-2 port on the optional Fast-Wide SCSI-2 Array Controllers. When combined with the dual bus Compaq ProLiant Storage

System/F, a complete ProLiant 5000 system provides a total storage capacity of more than 750 GB.

Power

The ProLiant 5000 contains a 490/540-W power supply that provides nominal output of 3.3 V, 5 V, and 12 V. The trend today is toward use of lower-voltage components. System memory in the ProLiant 5000, for example, operates at 3.3 V.

The standard ProLiant 5000 power supply has a dual rating: 490 W at 100/115 V and 540 W at 208/230 V. When the system is fully configured with four processors, four internal drives, and more than six I/O boards, Compaq strongly recommends operating the ProLiant 5000 at 208/230 V.

Intel Pentium Pro processors have precise voltage requirements; therefore, Pentium Pro-based systems need to provide dedicated regulators for each CPU (and for the processor support chips). If a regulator fails, the processor fails, and the system fails.

A unique feature of the ProLiant 5000 server is Compaq Processor Power Modules. These modules determine the precise voltage requirement of processors and the support chip set. Then they ensure that the proper voltage is delivered to each component. The I/O system board comes standard with one of these modules but will accommodate two for redundancy. Each ProLiant 5000 server comes standard with one Processor Power Module per processor. An optional N+1 redundancy scheme supports up to three Processor Power Modules for every two processors and provides a highly available solution.

PROLIANT 5000 ARCHITECTURE

For all the reasons mentioned previously, Compaq chose to optimize the ProLiant 5000 for PCI. There are two primary bus expansion schemes for PCI systems: PCI-to-PCI bridge architecture and dual-peer PCI bus architecture.

Figure 2 is a block diagram showing PCI-to-PCI bridge architecture. This scheme provides extra slots for additional PCI devices and is the lower cost implementation. However, it has two significant limitations for high-end systems:

- It supports only one PCI bus directly and independently linked to the host bus.
- It provides no additional system level bus bandwidth.

Because of these limitations, optimizing the ProLiant 5000 I/O system for PCI meant using the dual-peer PCI bridge architecture instead. Dual-peer PCI bus architecture significantly increases the configuration flexibility of the system, allowing optimization levels that are not possible with PCI-to-PCI bridge architecture.

Figure 3 is a block diagram showing the dual-peer PCI bus architecture Compaq chose for the ProLiant 5000. This expansion scheme supports two PCI buses directly and independently linked to the high speed Pentium Pro processor bus via dual Host-to-PCI bridges. The two PCI buses provide aggregate I/O bandwidth totaling 267 MB/s and extra slots for PCI devices. The ProLiant 5000 also includes an EISA bus to maintain backward compatibility with installed EISA and ISA peripherals.

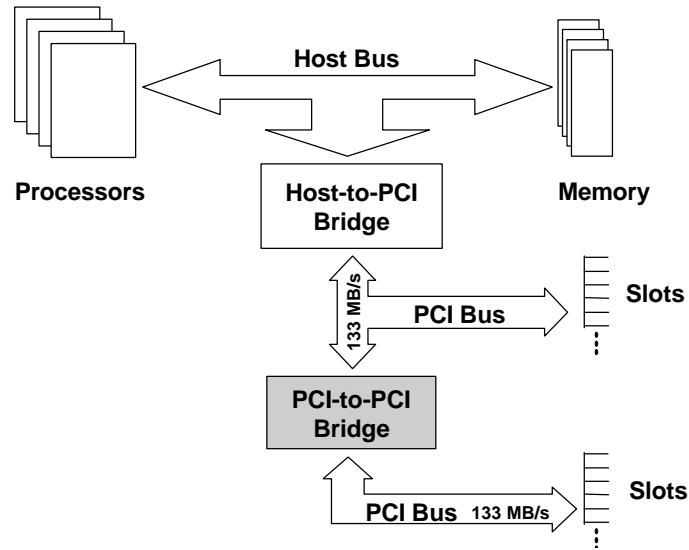


Figure 2. Block diagram showing PCI-to-PCI bridge architecture




Figure 3. Block diagram of the dual-peer PCI bridge architecture used in the Compaq ProLiant 5000 Server

Using high-bandwidth burst for all data transfers, PCI achieves a 133 MB/s peak throughput for reads and writes. Arbitration time in the ProLiant 5000 has been minimized so that more time can be spent effectively moving data. The high-performance PCI bridges support full concurrency of the host bus, the PCI bus, and the expansion bus in simultaneous operation. PCI also provides fast access. It is possible to access a PCI device in as little as 60 ns using a 33-MHz PCI clock.

Figure 4 depicts the I/O buses (primary and secondary PCI buses and the EISA bus) and slots in the ProLiant 5000. The primary PCI bus supports the following:

- PCI on dedicated slots 5, 6, 7, and 8.
- EISA bridge.

- EISA on dedicated slot 1 and on shared slots 2 and 3.
- Embedded EISA peripherals including video, timers, keyboard, and mouse.

The secondary bus supports the following:

- PCI on dedicated slot 4 and on shared slots 2 and 3.
- An embedded Fast-Wide SCSI controller.

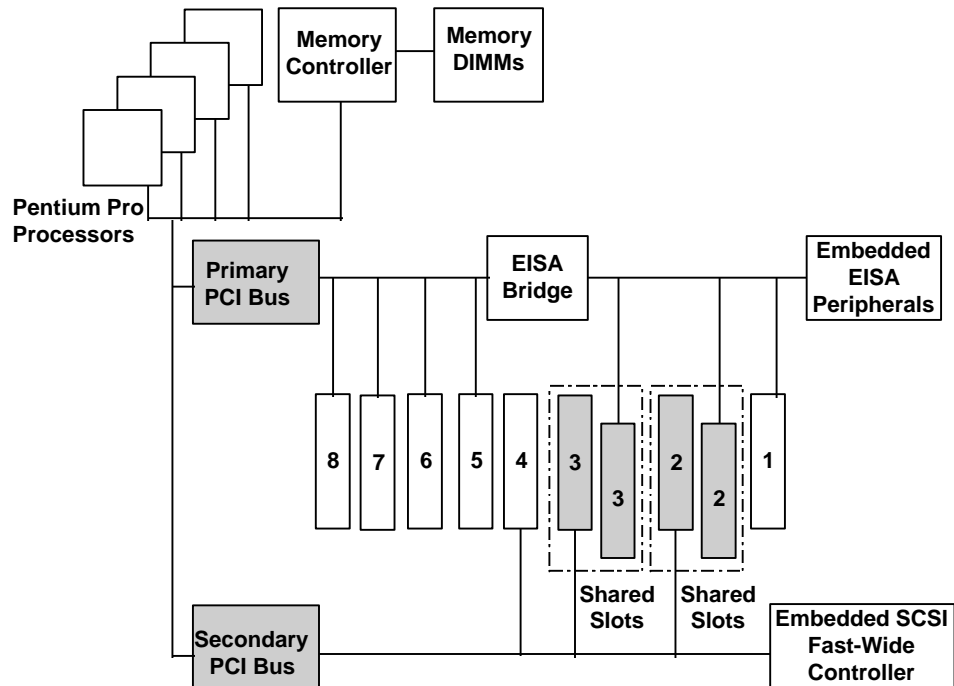


Figure 4. ProLiant 5000 I/O architecture

Performance Implications of the EISA Bus

Limitations of the Intel 82450GX chip set require that an EISA bus in a Pentium Pro server be bridged off a PCI bus, which precludes directly buffering EISA to the host bus or to memory. Essentially, EISA is one stage removed from the host bus. This can affect system performance in two ways:

- Because EISA is bridged off a PCI bus in Pentium Pro servers, EISA peripherals consume PCI bandwidth. In fact, *EISA peripherals can consume a disproportionate amount of PCI bandwidth* due to the indirect buffering to memory. A heavily loaded EISA peripheral can significantly limit the performance of the entire system. Low- to medium-bandwidth embedded EISA peripherals, however, do not significantly affect PCI performance.
- For all 82450GX-based Pentium Pro systems, the result of the architectural focus on PCI is lower EISA performance in Pentium Pro systems than in Pentium systems optimized for EISA (including the ProLiant 4500). Because Pentium Pro servers are optimized for PCI performance and EISA is bridged off PCI, low-level EISA performance in a Pentium Pro server using high-bandwidth EISA peripherals will be *one half or less* than in a system like the ProLiant 4500 that is optimized for EISA.

In Pentium Pro servers, EISA is intended to be used only for low- to medium-bandwidth, low-duty-cycle adapters, such as one 10-Mbit network interface card (NIC) with low utilization, a modem, or a system management card. If the use of specific, high-bandwidth EISA peripherals is an overriding consideration for an application, a system optimized for EISA will tend to provide higher overall system performance. The ProLiant 4500, for example, is so highly optimized for EISA that it can outperform competitors' PCI-based Pentium servers—in file and print benchmarks and in database benchmarks.

Configuration Recommendations

Compaq has conducted extensive testing to determine optimal configurations for achieving high performance and ease of use at the initial installation of a ProLiant 5000 system and as additional peripherals are installed. This section contains Compaq recommendations for configuring the ProLiant 5000 to achieve optimal performance. Additional information about configuring the ProLiant 5000 is provided in the Appendix.

Although the PCI BIOS in the ProLiant 5000 allows for automatic device detection and configuration, Compaq strongly recommends that users run the ProLiant 5000 Configuration Utilities provided with the system.

PCI Peripherals

For top performance from a ProLiant 5000 server, follow these recommendations for configuring PCI peripherals:

- Use only high-performance PCI peripherals. Typical characteristics of high-performance PCI peripherals include bus mastering and burst mode support. For optimal performance, balance high-bandwidth PCI peripherals across the two PCI buses.
- Avoid using programmed I/O peripherals. Because programmed I/O peripherals use the bus inefficiently, they consume large amounts of the PCI bandwidth. This can significantly degrade performance of all peripherals on the same bus, and possibly of the entire system.
- Because of driver incompatibilities, some third-party PCI peripherals will not operate optimally on the secondary PCI bus. If that appears to be true of a specific card, contact the card manufacturer to request an updated device driver or move the card to the primary PCI bus.
- Ensure that the peripheral device being configured is using the most current driver. If a user experiences problems with third-party cards, the first diagnostic step should be to download the most current driver from the third-party card vendor's Web page.
- Parity checking is an important data integrity benefit that is enabled by default in the ProLiant 5000. However, certain PCI cards generate false parity errors on the PCI bus. Compaq strongly recommends that cards exhibiting this behavior not be used in the ProLiant 5000.

If a user does choose to install such cards, they can be supported in the ProLiant 5000 by disabling parity checking on the PCI bus. A PCI parity error will cause a non-maskable interrupt (NMI), which will result in a screen message notifying the system administrator that a PCI bus parity error has occurred. Upon receiving this message, the administrator has the option of turning off PCI parity NMI events. This can be done within the advanced mode of the ProLiant 5000 system configuration utility. This action should only be taken after checking with the card manufacturer to verify that the card does not support parity; otherwise, a real error may be masked.

EISA Peripherals

For best overall performance from a ProLiant 5000 system, minimize the use of EISA peripherals and avoid using high-bandwidth EISA peripherals such as NICs and disk controllers. High-bandwidth EISA peripherals will negatively impact overall system performance.

If the use of specific, high-bandwidth EISA peripherals is an overriding consideration for an application, a system optimized for EISA (for example, the ProLiant 4500) will tend to provide higher overall system performance.

Video Cards

The ProLiant 5000 has an embedded video controller capable of supporting resolutions of 1024 x 768 and 16 colors. The embedded controller is EISA-based, but the bus utilization by video activity in server environments is not high enough to impact performance.

Adding an optional video controller adapter on the ProLiant 5000 (or on any dual-peer PCI bus Pentium Pro server) is allowable only on the primary PCI bus (in slot 5, 6, 7, or 8 in the ProLiant 5000). The primary PCI bridge is the only bridge capable of decoding the video address space near the top of 1 MB. The second PCI bridge cannot address that space. This limitation is a function of the Intel 82450GX chip set.

While a PCI video adapter will perform equally well in slot 5, 6, 7, or 8, Compaq recommends installing it in slot 8. In each of the slot configurations recommended by Compaq, slot 8 has been left open for use with a video adapter card, should one be needed.

Use of high resolution EISA-based video adapter cards is not recommended in the ProLiant 5000, because they will consume significant bandwidth on the PCI bus.

Memory Upgrades

ProLiant 5000 memory can be upgraded in banks of four 60ns DIMM memory modules. Buffered Fast Page Mode (FPM) DIMMs are used for the 64, 128, 256, and 512MB memory options and buffered EDO DIMMs are used for the 1024MB memory option. All DIMMs in a bank must be of the same size, speed, and type. When an upgrade increment of four DIMMs is installed, two of the DIMMs must be installed on each of the two memory expansion boards. The procedure for installing DIMMs is graphically depicted in Chapter 2 of the *ProLiant 5000 Servers Reference Guide*. The 1024 MB memory option must be installed in memory boards that support either FPM or EDO memory. Systems with a serial number of x706 xxxx xxxx or lower must be upgraded with the FP/EDO Memory Board Conversion kit (PN 219481-B21). An attempt to use EDO on these systems without the proper memory board conversion kit could cause memory corruption problems. Systems with a serial number of x707 xxxx xxxx or greater already have the memory boards installed that will support either FPM or EDO memory.

The ProLiant 5000 supports 16-MB, 32-MB, 64-MB, and 128-MB 60-ns, buffered FPM DIMMs and 256-MB buffered EDO DIMMs for maximum memory expansion to 4 GB. All DIMMs in the ProLiant 5000 memory option kits have 60-ns access time for optimal performance. Slower third-party DIMMs can be used, but system performance will be reduced. ProLiant 5000 system read-only memory (ROM) detects speed differences between DIMMs and adjusts system timing to the lower of two speeds to ensure proper system operation.

For proper operation of a ProLiant 5000 server, users must add the right kind of DIMMs—buffered FPM or buffered EDO DIMMs—during memory upgrades. System ROM cannot differentiate between the multiple types of DIMMs (buffered FPM, unbuffered FPM, buffered EDO, and unbuffered EDO).

In addition, it is important to understand this key point regarding Pentium Pro memory upgrades: Errata in the Pentium Pro memory controller itself prohibit mixing certain types of single-bank

and double-bank DIMMs (or SIMMs). If an 82450GX-based Pentium Pro server were allowed to operate with a mixture of these types of single-bank and double-bank DIMMs (or SIMMs), subtle data corruption errors might occur. This issue exists for all 82450GX-based Pentium Pro servers shipping in mid-1996, regardless of whether system memory contains SIMMs or DIMMs. ProLiant 5000 system ROM prevents operation with non-mixable DIMM bank types.

Using the Compaq memory option kits identified in Table 1 to upgrade memory in a ProLiant 5000 server will eliminate the mixing issues discussed above. Compaq memory upgrade kits for the ProLiant 5000 are fully compatible and can be mixed and matched at will. Use of third-party DIMMs is supported in the ProLiant 5000, but it is the user's responsibility to ensure that third-party DIMMs are of the right type and that they are mixable with memory already in the system. More details about mixable DIMMs is provided in Compaq Service Advisory No. 1017.

In a ProLiant 5000 server, Compaq memory, drives, and Pentium Pro processors are covered under the Compaq Pre-Failure Warranty. If Compaq Insight Manager detects operational degradation of memory, drives, or processors, Compaq will replace the component free of charge during its three-year warranty period, even if the component has not actually failed.

TABLE 1: COMPAQ MEMORY OPTION KITS FOR THE PROLIANT 5000 SERVER

Compaq Part No.	Total Memory	Kit Contains
219282	64 MB	Four 16-MB DIMMs (PN 169231)
219283	128 MB	Four 32-MB DIMMs (PN 221099)
219284	256 MB	Four 64-MB DIMMs (PN 185891)
219285	512 MB	Four 128-MB DIMMs (PN 169234)
241774*	1024 MB	Four 256-MB DIMMs (PN 228471)

*CPUs with a serial number of x706 xxxx xxxx or below require the addition of a FP/EDO Memory board conversion kit (PN 219481-B21).

Power

Compaq strongly recommends operating the ProLiant 5000 (like all Compaq high-end servers that are heavily loaded) at 208/230 V. Operating at 240 V increases the power supply rating from 490 W to 540 W.

If a ProLiant 5000 system is operated at 100/115 V, the limitations identified in Table 2 apply to the ProLiant 5000R system.

TABLE 2

ProLiant 5000R Configurations			
	Configuration A	Configuration B	Configuration C
Maximum number of EISA /PCI slots filled	6	7	7
Maximum number of hard drives	5	4	5
Maximum number of processors	4	4	3

Balancing Bus Use

Compaq Insight Manager supports bus use measurements on each PCI bus and on the EISA bus. The bus use measurement tools provide a good way to check the I/O configuration to ensure that peripherals are properly balanced across the two PCI buses. An indication of unbalanced bus use implies that the system may need to be reconfigured to optimize performance. Refer to Compaq Insight Manager documentation for more information about using these tools.

EXAMPLE CONFIGURATIONS

All Compaq ProLiant 5000 models ship from the factory with a Netelligent 10/100 TX Ethernet adapter in slot 2. Additionally, ProLiant 5000 array models ship with a SMART-2/P Array Controller installed in slot 5. This is the factory configuration *and* the recommended starting configuration when options are added later. This configuration has the advantage that the NIC and the SMART-2/P Array Controller reside on separate PCI buses.

Table 3 describes this factory configuration. Tables 4 - 8 show recommended configurations for installing as many as four SMART-2/P Array Controllers and two 10/100 TX PCI UTP Controllers to achieve maximum disk performance. Other I/O configurations may prove better for optimizing network throughput, for example.

Figure 5 provides a quick comparison of the six configurations described in Tables 3 - 8. Compaq recommends these configurations for two reasons:

- They provide the I/O load balancing necessary to achieve maximum overall system performance in a disk I/O-intensive environment.
- They have been verified in Compaq labs as optimal for ease of configuration.

TABLE 3

Factory Configuration, Array Models		
Device	Bus	Slot
10/100 TX PCI UTP Controller	Secondary	2
SMART-2/P Array Controller	Primary	5

TABLE 4

Recommended Three-Controller Configuration One 10/100 TX PCI UTP Controller, Two SMART-2 Array Controllers		
Device	Bus	Slot
10/100 TX PCI UTP Controller	Secondary	2
SMART-2/P Array Controller	Primary	5
SMART-2/P Array Controller	Secondary	3

TABLE 5

Recommended Three-Controller Configuration Two 10/100 TX PCI UTP Controllers, One SMART-2 Array Controller		
Device	Bus	Slot
10/100 TX PCI UTP Controller	Secondary	2
SMART-2/P Array Controller	Primary	5
10/100 TX PCI UTP Controller	Primary	6

TABLE 6

Recommended Four-Controller Configuration Two 10/100 TX PCI UTP Controllers, Two SMART-2 Array Controllers		
Device	Bus	Slot
10/100 TX PCI UTP Controller	Secondary	2
SMART-2/P Array Controller	Primary	5
10/100 TX PCI UTP Controller	Primary	6
SMART-2/P Array Controller	Secondary	3

TABLE 7

Recommended Five-Controller Configuration Two 10/100 TX PCI UTP Controllers, Three SMART-2 Array Controllers		
Device	Bus	Slot
10/100 TX PCI UTP Controller	Secondary	2
SMART-2/P Array Controller	Primary	5
SMART-2/P Array Controller	Secondary	3
10/100 TX PCI UTP Controller	Primary	6
SMART-2/P Array Controller	Secondary	4

TABLE 8

Recommended Six-Controller Configuration Two 10/100 TX PCI UTP Controllers, Four SMART-2 Array Controllers		
Device	Bus	Slot
10/100 TX PCI UTP Controller	Secondary	2
SMART-2/P Array Controller	Primary	5
SMART-2/P Array Controller	Secondary	3
10/100 TX PCI UTP Controller	Primary	6
SMART-2/P Array Controller	Secondary	4
SMART-2/P Array Controller	Primary	7

Config.	Slot 1	Slot 2		Slot 3		Slot 4	Slot 5	Slot 6	Slot 7	Slot 8
	EISA	EISA	Secondary PCI	EISA	Secondary PCI	Secondary PCI	Primary PCI	Primary PCI	Primary PCI	Primary PCI
1			NIC				A/C			
2			NIC		A/C		A/C			
3			NIC				A/C	NIC		
4			NIC		A/C		A/C	NIC		
5			NIC		A/C	A/C	A/C	NIC		
6			NIC		A/C	A/C	A/C	NIC	A/C	

Legend: NIC = Network Interface Card
A/C = Array Controller

Figure 5. Comparison of ProLiant 5000 configurations recommended in Tables 3 - 8

CONCLUSION

Compaq developed the new ProLiant 5000 Server to exploit the full performance advantages and unique characteristics of today's Pentium Pro processors, along with the latest in high-performance options. With the design of the ProLiant 5000, Compaq is leading the industry move toward the PCI architecture and memory technology required to meet customers' growing I/O requirements now and in the future.

The ProLiant 5000 embodies unique features specifically engineered to overcome Intel-documented limitations in the Intel 82450GX chip set. With proper system configurations (as recommended in this paper), the performance of the quad-processor ProLiant 5000 will be up to 50% greater in a database server environment than an equivalent Pentium-based system. Compaq continues to work with industry partners to enhance this breakthrough platform for the distributed enterprise.

Appendix

This appendix provides configuration information supplemental to that included in the first and second editions of this technology brief. The rules cited here provide a practical guideline for successfully configuring the ProLiant 5000 Server in a variety of operating environments.

GENERAL SYSTEM CONFIGURATION

The rules in this section pertain to general configuration of the ProLiant 5000 system.

RULE: Use System Configuration Utility Version 2.31 or later to configure the ProLiant 5000.

Version 2.31 is the earliest version of the Compaq System Configuration Utility that supports the ProLiant 5000. When adding any device or when changing the configuration of the ProLiant 5000, run the System Configuration Utility to identify the slot number and bus number for each configured PCI device and EISA device. The types of installed devices and where they reside in the PCI/EISA architecture will determine how well the ProLiant 5000 performs. To view device-specific information in the System Configuration Utility, from the System Configuration screen, select *Configure Hardware*, then *Review or Modify Hardware Settings*, and then *View or Edit Details*. Scan the list of configured devices.

Once installed in the ProLiant 5000, some devices prompt the user to run the Compaq System Configuration Utility if the devices are not yet configured. For example, some Compaq

Controllers generate a 175 POST error, indicating a PCI/Configuration Slot Mismatch and recommending that the System Configuration Utility be run to reconfigure the system. The POST error occurs if an unconfigured EISA device or unconfigured Compaq PCI device is added to the server.

RULE: Ensure that each I/O device installed in the ProLiant 5000 is properly seated in its slot.

When EISA devices are improperly seated, the System Configuration Utility requests a .CFG file with an invalid filename. When PCI devices are improperly seated, read operations of the configuration space header are corrupted.

EISA DEVICE CONFIGURATION

Documentation from Intel and Compaq recommends avoiding the use of EISA devices in the ProLiant 5000. If EISA or ISA devices are to be installed in the ProLiant 5000, however, follow the rules in this section to ensure that the server is properly configured and that the workload is appropriately balanced for the best possible performance.

RULE: If high performance and high bandwidth are required for your entire server implementation, configure only high-speed PCI devices in the ProLiant 5000 to achieve peak performance.

Use of EISA or ISA devices in the ProLiant 5000 can so reduce overall system performance that the value of using high-speed Pentium Pro processors may be negated. Because Pentium Pro servers are optimized for PCI performance and EISA is bridged off PCI, low-level EISA performance in a Pentium Pro server using high-bandwidth EISA peripherals will be *one half or less* than EISA performance in a system like the ProLiant 4500 that is optimized for EISA.

RULE: Use the high-speed PCI version of a server component when the performance of that component is critical to network and/or server use.

EISA devices consume a large portion of PCI bus bandwidth. EISA devices in high-use situations can starve other devices on the bus, causing degradation of performance. In some network implementations, however, there may be only one or two components in the server for which high performance is demanded. If performance of a specific component is critical, ensure that the component is a high-speed PCI device and that the load on the dual-peer buses is properly balanced.

For example, in the NetWare SFT III environment, the Mirrored Server Link (MSL) connects the primary server to the secondary server for resynchronization of the two servers. When either server is initialized, the resynchronization process occurs. Until resynchronization is complete, network services are not performed. Using EISA-based MSLs in two ProLiant 5000 Servers can substantially increase the time required for resynchronization; eventually, clients will begin to drop off the network. Furthermore, during operation the mirrored servers continually communicate via the MSLs. This communication consumes a large portion of PCI bus bandwidth and can eventually cause bus starvation. Compaq highly recommends use of Novell-approved PCI MSLs.

Note: Because EISA devices consume a large portion of PCI bus bandwidth, overall server performance degradation will occur when *any* EISA controllers are installed.

RULE: If EISA Network Interface Controllers are installed in the ProLiant 5000, ensure that PCI SCSI Controllers reside on the secondary PCI bus.

EISA bus starvation occurs when PCI SCSI controllers reside on the primary PCI bus along with high-use EISA devices, such as EISA network interface controllers. Since the EISA bus is bridged off of the primary PCI bus, installed EISA devices that consume a large portion of PCI bandwidth, such as programmed I/O devices, can starve other EISA or PCI devices on the primary PCI bus. Furthermore, once EISA devices relinquish control of the bus, other devices on the primary PCI bus can cause starvation of the EISA devices.

Because PCI SCSI hard drive and drive array controllers are high-use devices, they should be moved to the secondary PCI bus to allow EISA NICs enough bandwidth to operate and to avoid bus starvation. The Integrated Fast-Wide SCSI-2 Controller (a PCI device) resides on the secondary PCI bus for this reason.

RULE: Use the latest available E16 System ROM revision for the ProLiant 5000.

Some PCI configuration and compatibility issues have been corrected in the E16 system ROM. To minimize the risk of a configuration problem with third-party PCI devices, ensure that the E16 system ROM is dated June 4, 1996, or later.

Compaq System ROMs can be updated via ROMPaqs, which are available as SoftPaqs on the Compaq web site at www.compaq.com. Select *Service and Support* and then *Downloadable Software Files*. ROMPaqs are also available on the Compaq Support Software CDs that ship with Compaq server products.

PCI DEVICE CONFIGURATION

RULE: When adding a PCI device to the ProLiant 5000, ensure that the PCI device is assigned to the appropriate PCI bus.

Because of early driver incompatibilities, some PCI devices do not support the dual-peer bus architecture of the ProLiant 5000 properly and therefore may not operate properly on the secondary bus. When one of these devices is installed on the secondary PCI bus, the PCI device driver cannot find the device and therefore generates an error message similar to “device not found.” Although PCI device vendors have corrected this problem by upgrading their drivers or their controllers, some earlier versions of controllers may have to be configured on the primary PCI bus (slots 5, 6, 7, and 8) to operate properly in the ProLiant 5000.

RULE: Install PCI devices with small amounts of buffer space on the secondary PCI bus to minimize the risk of device starvation.

If starved, PCI devices with small amounts of buffer space generate errors. In NICs, these errors are manifested as Transmit and Receive underruns or overruns. For example, the Compaq Netelligent 4/16 TR PCI IBM UTP/STP Controller may exhibit this problem when configured on the primary PCI bus. This starvation problem occurs because the EISA bus is bridged off the primary PCI bus, and installed EISA devices consume a large portion of the primary PCI bus bandwidth, leaving little bandwidth for the PCI device. If the PCI device cannot gain access to the bus, the buffers fill quickly and cannot be flushed until the device gains access to the bus. This results in a steady increase in transmit and receive underruns/overruns while the device awaits bus access. Because of the limited buffer space, the device must request and gain access to the PCI bus more often, leaving fewer bus cycles for other devices on the same bus. Placing such PCI devices on the secondary PCI bus reduces the risk of bus starvation.

Intel recommends that when selecting and installing PCI devices, system administrators ensure that the devices (especially high-speed NICs) support bus mastering and burst mode and that the devices include enough on-board buffer space to support one or more full data packets.

RULE: Verify that the PCI device supports even parity generation.

As part of the PCI standards, the generation of even parity is mandatory for all PCI devices that drive address or data information onto the bus; however, the reporting of parity errors is not mandatory. Since parity checking has not been made mandatory, only some systems actually perform parity checking. If a PCI device incorrectly calculates parity, systems that do not perform parity checking cannot detect the incorrect calculation.

The Compaq ProLiant 5000 monitors parity checking on the PCI bus to ensure data integrity. If a parity error occurs, a system error (SERR#) NMI is generated during POST and an error message is displayed. In some operating environments, such as DOS, the generation of the NMI can cause the system to lock up or hang.

Parity errors that occur during POST indicate that the device has generated incorrect parity or has incorrectly calculated parity during initialization. Parity errors that occur during server operation rather than during POST may indicate a different problem. In either case, the controller manufacturer should be contacted. If the controller is known to generate parity errors, Compaq highly recommends that the controller not be used in the ProLiant 5000. However, if the controller must be used, the system administrator can disable the parity-checking feature in the ProLiant 5000 by using the “Advanced Mode” option of System Configuration Utility Version 2.31B.

PCI NETWORK INTERFACE CONTROLLERS

Because a variety of vendors develop NICs, NIC configuration provides the most challenges for integration with the ProLiant 5000. Some network application-specific integrated circuit (ASIC) vendors manufacture and OEM their chips to board vendors such as Compaq, who build a complete NIC product or integrate the chip in an existing PC or server product. Because PCI technology is new, these chip vendors and board vendors each address configuration and compatibility issues independently. Integrating these products into PCs or servers such as the ProLiant 5000 requires special consideration of the products used, the location of the product in the system (slot number and bus number), and the other peripherals installed in the system.

RULE: Use the latest version of the device driver from the NIC vendor, rather than the generic device driver from the chip vendor.

As part of the OEM agreement, OEM ASIC manufacturers provide an OEM developers' kit with the ASIC that usually includes core code from the vendor. This core code is required to provide the basic functionality of the driver for the NIC device, but it is usually incorporated in the board vendor's customized device drivers. This customization can sometimes result in compatibility issues as PCI technology and standards continue to develop and change. Some generic drivers (core code) provided by ASIC vendors will work adequately for NIC devices based on the ASIC vendors' chip sets. Others will not load or will not work properly, due to the unique implementation of the chip set in the NIC device.

Another issue concerns identification of these boards within the PC or server in which they reside. All PCI network chips include two fields, VendorID and DeviceID, in the chip's configuration space register. However, some of these chips do not support the additional fields SubsystemID and Subsystem_VendorID. These fields allow the board manufacturer to provide a unique identifier for each controller. Without support for these additional fields, compatibility and configuration problems can arise when multiple devices based on the same chip are used in the same PC or server. Consider the example of two controllers based on the same chip set and using the same core driver code. As the device drivers are loaded, the potential exists for the drivers to misidentify associated devices and attempt to load based on the chip, not on the device itself.

Most OEM board vendors provide a workaround for this problem. All chips provide support for a serial ROM that includes a board-vendor-unique MAC address for each controller. Part of this MAC address includes the board vendor's unique Ethernet identifier. Therefore, the vendor can customize the driver to load and scan for the MAC address in the board's serial ROM, rather than just for the VendorID and DeviceID fields in the configuration space register.

Note: Ensure that you load the latest version of the driver. Configuration issues, including support for the dual-peer bus architecture of the ProLiant 5000, may have been addressed in the latest version of the driver. For example, earlier versions of some PCI drivers may not recognize the NIC on the secondary PCI bus of the Compaq ProLiant 5000 and may cause the system to hang. Later versions of the same driver may correct the problem.

RULE: If the PCI board includes more than a single port or more than a single PCI device, always use the board vendor's device drivers, because the generic drivers typically will not support multiple ports.

The current generation of chips assumes a simple device design of a single chip and a single port with a single MAC address; however, some board vendors include multiple ports attached to a single chip. As with other issues, each board vendor addresses the need for multiple MAC addresses for multiple ports in a different manner. For example, a four-port controller that uses only one serial ROM may read the first address and calculate the next three sequential addresses; therefore, the board vendor's drivers are required in order to configure the four-port controller properly.

Note: If you are configuring NICs with multiple ports, be sure to use the board vendor's drivers and follow configuration instructions carefully.

RULE: PCI devices that use memory space below the 1-megabyte boundary must be installed on the primary PCI bus.

Some NICs have a unique hardware requirement: Memory below the 1-megabyte boundary must be allocated to the controller. The Intel 82450GX chip set architecture dictates that this memory area be available only to devices on the primary PCI bus. Installing such controllers on the secondary PCI bus will disable them, because no memory can be allocated to them. When the driver loads, it cannot locate the disabled PCI device and therefore displays an error message.

PCI VIDEO CONTROLLERS

RULE: Install PCI video controllers on the primary PCI bus (slots 5, 6, 7, and 8).

Some early versions of PCI video controllers require allocation of memory below the 1-megabyte boundary. However, the Intel 82450GX chip set architecture dictates that this memory area be available only to devices on the primary PCI bus. Therefore, installing these video controllers on the secondary PCI bus will leave them disabled. Installing PCI video controllers on the primary PCI bus (slot 5, 6, 7, or 8) is highly recommended to avoid this issue, even though some of these early PCI controllers may have been updated to support alternate memory allocations.

RULE: Install the QVision 1280/P Controller in slot 8 of the ProLiant 5000.

The Compaq QVision 1280/P Controller supports boundary scanning, which allows in-circuit testing of the controller. In-circuit testing is enabled when the Test Data In (TDI) signal is driven low during the controller's reset process. When the controller goes into test mode, it does not respond and appears to fail.

Test mode in the QVision 1280/P Controller is enabled only when both these conditions are met:

- The QVision 1280/P Controller is installed in primary PCI bus slot 5, 6, or 7.
- A controller that drives the TDI signal low comes before the QVision 1280/P Controller in the PCI scan chain.

In the ProLiant 5000, the TDI signal is always driven high in slot 8. To ensure that the QVision 1280/P Controller does not go into in-circuit test mode, install the QVision 1280/P Controller in slot 8 of the ProLiant 5000.

CONFIGURATION FOR SPECIFIC OPERATING SYSTEMS

This section provides configuration rules specific to the following operating environments:

- Microsoft Windows NT 3.5x or 4.0
- NetWare 3.x, NetWare 4.x, and IntranetWare
- SCO Open Server Release 3.0
- SCO OpenServer Release 5.0
- IBM OS/2 for SMP v2.11

Microsoft Windows NT 3.5x or 4.0

RULE: When adding any PCI bridge device to or removing any PCI bridge device from the primary PCI bus, reconfigure all NICs residing on the secondary PCI bus that do not load during the Windows NT boot process.

When any PCI bridge device, such as the Compaq SMART-2/P SCSI Array Controller, is added to or removed from the primary PCI bridge in the ProLiant 5000, NICs may not load for the following reason: In a Microsoft Windows NT environment, each PCI device and the bus number on which it resides are registered in the Windows NT registry. Adding or deleting PCI devices that include PCI-to-PCI bridges from the primary PCI bus causes the bus numbering to change; however, the previous bus numbering is not automatically changed in the registry.

If PCI NIC devices do not load when the server is booted into Windows NT, perform either or both of the following procedures, depending on your configuration.

To reconfigure Compaq NICs, complete the following steps:

1. From the Control Panel in Windows NT, double-click on the Network icon, and then record pertinent information about each NIC in the ProLiant 5000.
2. Close the Network Settings dialog box.
3. From the Control Panel, double-click on the Compaq Net icon to launch the Compaq Advanced Network Control Utility.
4. Highlight the NIC that did not load and select the *Uninstall* option.
5. Highlight the NIC again and select the *Install* option.
6. Close the Compaq Advanced Network Control Utility.
7. In the Network Settings dialog box, enter the information recorded in Step 1 (IP address, etc.).
8. Shut down and reboot the server.

To reconfigure third-party NICs, complete the following steps:

1. From the Control Panel in Windows NT, double-click on the Network icon, and record pertinent information about each NIC in the ProLiant 5000.
2. Highlight the NIC that did not load, and select the *Remove* option.
3. Select the *Add Adapter* option.
4. From the Adapter List, select the NIC that did not load, and then select *Continue*.

5. In the Network Settings dialog box, enter the information recorded in Step 1 (IP address, etc.).
6. When the NIC configuration is complete, shut down and reboot the server.

RULE: To enhance ProLiant 5000 performance in a Windows NT 3.51 or 4.0 environment, use the *Advanced Mode* option in the System Configuration Utility Version 2.31B (or later) to increase the In-Order Queue (IOQ) depth to a value of 8.

Compaq testing of the ProLiant 5000 in a Windows NT 3.51 environment indicates that the overall server performance can be noticeably increased by:

- Increasing the In-Order-Queue (IOQ) depth parameter in the System Configuration Utility to a value of 8.
- Installing the Compaq Support Software (SSD) Version 1.18 (or later).

Be aware, however, that changing the IOQ depth to increase server performance results in loss of ECC protection for the Pentium Pro processor bus. Due to errata in the Intel 32450 PCI chip set, ECC protection of the Pentium Pro processor bus must be disabled when the IOQ depth is set to 8.

If ECC protection of the processor bus is an important feature in a server environment, do not change the default IOQ depth value of 1. If it is more desirable to improve performance by 10 to 15 percent, make the change.

To increase the IOQ depth in System Configuration Utility Version 2.31B (or later), complete the following steps:

1. Boot the server into System Configuration Utility Version 2.31 Revision B (or later).
2. Select the *System Configuration* option.
3. Highlight the *Configure Hardware* option, and press **CTRL + A** to enable Advanced Mode.
4. When a pop-up window indicates that Advanced Mode is enabled, press **ENTER** to continue.
5. After the System Configuration Utility completes the hardware configuration, select the *Review or modify hardware settings* option.
6. Select *Step 3: View or edit details*.
7. Scroll down to locate the Advanced Features section, and select *IOQ setting*.
8. Change the IOQ setting to IOQ=8.
9. Press **F10** to return to the Main Menu.
10. Select the option *Step 5: Save and exit*.
11. When the save process is complete, power-cycle the server.

IMPORTANT: The System Configuration Utility does *not* prompt you to power-cycle the server at the end of this procedure. If you do not power-cycle the server, the changes just made to the configuration will not be saved and the IOQ setting will return to the default value of 1.

RULE: When loading Windows NT support software for the ProLiant 5000, use the latest version of SSD for Windows NT.

SSD for Windows NT provides extensive support for the ProLiant 5000 in a Windows NT 3.51 or 4.0 environment, including the latest drivers and Compaq Hardware Abstraction Layers.

RULE: When installing TCP/IP, install multiple Compaq Netelligent 4/16 TR PCI IBM UTP/STP Controllers on the same PCI bus.

The TCP/IP protocol installation fails when multiple Compaq Netelligent 4/16 TR PCI IBM UTP/STP Controllers are installed over multiple PCI buses. For example, if two Compaq Netelligent 4/16 TR PCI IBM UTP/STP Controllers are installed, one on the primary bus and one on the secondary bus, the same IP address is used for both devices. This results in “invalid IP or subnetmask address” errors.

NetWare/IntranetWare

RULE: When loading NetWare/IntranetWare software support for the ProLiant 5000, use the latest version of Compaq support software for Novell products.

Novell Support Software Diskettes provide support for all Compaq servers, including the ProLiant 5000.

RULE: When configuring the ProLiant 5000 for NetWare/IntranetWare, do not allow devices that perform different functions (such as NICs and SCSI controllers) or devices of different architectures (PCI and EISA devices) to share interrupts.

NetWare/IntranetWare drivers support interrupt sharing among devices that perform the same function or devices of the same architecture (EISA or PCI). Sharing interrupts between dissimilar devices can cause the server to hang at the first real-mode transaction, such as when loading an NLM from drive C. For example, when a PCI NIC shares an interrupt with a PCI SCSI Controller, the server hangs at approximately 95% complete during the SmartStart installation of NetWare (during creation of the AUTOEXEC.NCF file). This occurs when the IPX protocol is bound to the NIC driver and a file copy operation from the NetWare CD (a real-mode transaction) occurs.

To reduce the risk of hanging an EISA/PCI server in a NetWare/IntranetWare environment:

- Do not configure EISA devices to share interrupts with PCI devices. If you must configure devices for shared interrupts, configure PCI devices to share interrupts with other PCI devices, and configure EISA devices to share interrupts with other EISA devices.
- Avoid configuring devices with different functions to share interrupts. If you must configure devices for shared interrupts, configure PCI SCSI devices to share interrupts with other PCI SCSI devices, and configure PCI NICs to share interrupts with other PCI NICs. When multiple identical devices are configured in the server and shared interrupts are required, configure the identical devices to share interrupts.

RULE: For AMD devices, ensure that the BUSTYPE=PCI1 parameter is defined in the LOAD command of the driver.

NetWare 4.1 drivers for AMD-based controllers require the BUSTYPE=PCI1 parameter to be included with the LOAD command as follows:

```
LOAD PCNTNW FRAME=ETHERNET_802.3 BUSTYPE=PCI1
```

RULE: To load the driver for a specific AMD-based device, include the PORT= parameter, and use the PORT address from the System Configuration Utility to specify the device in the server.

To load a driver for a specific AMD-based PCI device, you must use the PORT address to specify the device. The PORT address can be obtained by running the System Configuration Utility and reviewing the Used Resources List. This list provides the PORT address range for each device in the system. Record the PORT address values for each AMD device, especially if you are installing multiple AMD-based devices from multiple board vendors.

For example, to load an AMD-based controller in slot 8 of the ProLiant 5000, where the controller is configured at port address 6420h, use the following command:

```
LOAD PCNTNW PORT=6420 FRAME=ETHERNET_802.3 BUSTYPE=PCI1
```

Note: When new devices are added to the configuration, the PORT addresses can change. Always run the System Configuration Utility when devices are added or removed from the configuration. Obtain the new PORT addresses for all AMD-based controllers by running the System Configuration Utility and reviewing the Used Resources List.

RULE: For NetWare SFT III configurations, use PCI Mirrored Server Links (MSLs).

Configuring EISA-based MSLs in two Compaq ProLiant 5000 Servers in a NetWare SFT III environment can substantially increase the time required for resynchronization. Until the resynchronization is complete, network services are not performed; eventually clients will begin to drop off the network. Compaq highly recommends use of Novell-approved PCI MSLs.

SCO UNIX

RULE: Ensure that the number of system processors installed in the ProLiant 5000 matches the number of processors licensed under SCO Open Server 3.0.

If the ProLiant 5000 is configured with more system processors than are currently licensed under SCO Open Server 3.0, an “Unexpected trap in kernel mode” occurs and the operating system reports a kernel mode trap type 0x00000006 or kernel mode trap type 0x0000000E. Ensure that the number of system processors in the ProLiant 5000 does not exceed the number of processors licensed under SCO Open Server 3.0.

RULE: For multiprocessor ProLiant 5000 Servers in an SCO Open Server 3.0 environment, first install SCO Advanced Hardware Support (AHS) 3.4, then install SCO Support Level Supplement (SLS) uod421A.

The Compaq ProLiant 5000 uses the Intel Advanced Programmable Interrupt Controller (APIC) and conforms to Intel Multiprocessor Specification (MPS) 1.4. SCO includes Intel MPS support in the SCO Advanced Hardware Supplement (AHS) and in SCO Support Level Supplement (SLS) uod421A. Install SCO Open Server 3.0 first, install SCO AHS uod421 for MPS support, and then install SCO SLS.

RULE: When installing SCO Open Server 3.0 on a multiprocessor ProLiant 5000, disable the multiprocessor Systempro/XL device driver.

The multiprocessor device driver for the Compaq Systempro/XL improperly detects the EISA ID of the ProLiant 5000 system board. This improper identification allows SCO Open Server 3.0 to use the Systempro/XL driver instead of the Intel MPS driver. If the Systempro/XL multiprocessor driver is not disabled, SCO Open Server 3.0 cannot detect more than one system processor. When the Systempro/XL driver is disabled, Open Server 3.0 correctly recognizes the Intel MPS driver.

Note: Before disabling the Systempro/XL driver, ensure that SCO AHS 3.4 is installed first and that SCO SLS uod421 is installed second.

To disable the Systempro/XL driver, modify the sysproxl file as follows:

1. Log in as root.
2. Change to the directory /etc/conf/sdevice.d as follows:

```
cd /etc/conf/sdevice.d
```


3. Modify the `sysproxl` file and change the Y in the second column to N as indicated below:

```
From:
sysproxl  Y      1      0      0      0      0      0      0      0
```

```
To:
sysproxl  N      1      0      0      0      0      0      0      0
```

4. Relink the kernel, and reboot the system using the following commands:

```
cd /etc/conf cf.d
./link_unix -y
shutdown -i6
```

RULE: For Compaq ProLiant 5000 Servers configured with three or four system processors and operating under extremely heavy I/O workload conditions, modify the APIC driver to eliminate a deadlock condition that may arise from multiple interprocessor interrupts being queued into a single processor.

Under extreme workload conditions, the ProLiant 5000 may, on rare occasions, experience system lockups in a SCO Open Server 3.0 or OpenServer 5.0 environment. The problem is caused by a deadlock condition that occurs when multiple interprocessor interrupts are queued onto a single system processor in the ProLiant 5000. When the APIC takes too long to service these interprocessor interrupts, the server lockup occurs.

Correcting the problem requires modifying the APIC driver to reduce the amount of time spent waiting for the APIC to service the interprocessor interrupts. Two procedures are provided here; follow the procedure for the SCO Open Server version being used.

Procedure For SCO Open Server Release 3.0

For Intel MPS support on Open Server Release 3.0 systems, SCO AHS 3.4 must be installed first and the SCO SLS uod421 must be installed second. Once these installations are done in the correct order, complete the following steps:

1. Log in as root.
2. Change the directory to `/etc/conf/pack.d/apic` as follows:
`cd /etc/conf/pack.d/apic`
3. Make a backup copy of the existing driver:
`cp Driver.o Driver.o.orig`
4. Invoke the SCO patch utility, `_fst`, with the write option:
`/etc/_fst -w Driver.o`
5. Verify that the original code has not already been modified:
`apic_await_intcmd+0x7?ia`

Note that the correct response should be:

```
apic_await_intcmd+0x7:  mov    [epb-0x4], 0x5f5e100
```

6. Change the delay value:
`apic_await_intcmd+0xa?W0x5f5e10`

Note that the correct response should be:

```
apic_await_intcmd+0xa:  0x5f5e100=0x5f5e10
```

7. Verify that the code has been properly modified:

```
apic_await_intcmd+0x7?ia
```

Note that the correct response should be:

```
apic_await_intcmd+0x7:  mov    [epb-0x4], 0x5f5e10
```

8. Exit the patch utility by pressing **CNTRL + D**.
9. To implement the changes, relink the kernel and reboot the system as follows:

```
cd /etc/conf/cf.d
./link_unix -y
shutdown -i6
```

Procedure For SCO OpenServer Release 5.0

If OpenServer Release 5.0 systems were installed via Compaq SmartStart Release 3.0 (or later), this procedure has already been applied. If not, complete these steps:

1. Log in as root.
2. Change the directory to /etc/conf/pack.d/apic as follows:

```
cd /etc/conf/pack.d/apic
```
3. Make a backup copy of the existing driver:

```
cp Driver.o Driver.o.orig
```
4. Invoke the SCO patch utility, `_fst`, with the write option:

```
/etc/_fst -w Driver.o
```
5. Verify that the original code has not already been modified:

```
apic_await_intcmd+0x10?ia
```

Note that the correct response should be:

```
apic_await_intcmd+0x10:  mov    edi, 0x5f5e100
```

6. Change the delay value:

```
apic_await_intcmd+0x11?W0x5f5e10
```
7. Verify that the code has been properly modified:

```
apic_await_intcmd+0x10?ia
```

Note that the correct response should be:

```
apic_await_intcmd+0x10:  mov    edi, 0x5f5e10
```

8. Exit the patch utility by pressing **CNTRL + D**.
9. Relink the kernel and reboot the system as follows to implement the changes:

```
cd /etc/conf/cf.d
./link_unix -y
shutdown -i6
```

⋮ **IBM OS/2 for SMP v2.11**

⋮ **RULE:** When installing OS/2 for SMP v2.11 on the Compaq ProLiant 5000, use Compaq OS2APIC.PSD Version 2.00 (or later), which ships on OS/2 Support Software from Compaq.

⋮ IBM provides an OS2APIC.PSD; however, when the IBM version is loaded instead of the Compaq OS2APIC.PSD Version 2.00, the installation hangs. Use the OS2APIC.PSD on OS/2 Support Software from Compaq Version 2.08 (or later).

⋮ **Note:** When SmartStart Release 3.00 (or later) is used to install OS/2 for SMP v2.11, the Compaq OS2APIC.PSD is automatically loaded during the SmartStart installation.