WHITE PAPER

October 1996

Prepared By
Intranet Engineering Group

Compaq Computer Corporation

CONTENTS

Performance Tuning Your Windows NT Web Server

Performance planning for Web servers is relatively new. Today there are very few tools available to judge how well a Web server may perform in a particular environment for a given load. This lack of available test tools and the information that they provide may result in installations that are under utilized or unable to meet load demand.

The focus of this paper is to introduce some of the methodologies employed to build a peak performing Web server. These methodologies can help you identify where potential bottlenecks may occur with a given load and thus avoid them all together. Also included are guidelines on what to expect for different platforms. Ultimately, the information provided will help the reader better plan and implement a Windows NT-based Web server.

COMPAQ

NOTICE

The information in this publication is subject to change without notice.

COMPAQ COMPUTER CORPORATION SHALL NOT BE LIABLE FOR TECHNICAL OR EDITORIAL ERRORS OR OMISSIONS CONTAINED HEREIN, NOR FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES RESULTING FROM THE FURNISHING, PERFORMANCE, OR USE OF THIS MATERIAL.

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 does not warrant products other than its own strictly as stated in Compaq product warranties.

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

Compaq, ProLiant, and NetFlex, registered United States Patent and Trademark Office.

Netelligent is a trademark and/or service mark of Compaq Computer Corporation.

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

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

Web Server Performance Tuning
First Edition (October 1996)

TESTING

Synthetic benchmarking, despite its limitations, can provide helpful information about your system including the identity and location of potential bottlenecks. Some benchmarking tools include the functionality to profile a Web site before introducing a real user load. With accurate profile data, you can then best configure your site for that type of load. While it may be impossible to predict your Web site's usage, if you know ahead of time what type of content will be served, this profiling information can be invaluable.

Tools

In order to generate some kind of load on a web server, one must consider the use of synthetic load generation tools. This section outlines some tools that will allow the tester to place a synthetic load on the server under test as well as evaluate it's performance.

WebCat

NOTE:
WebCat was designed as an engineering tool to help the Internet Information Server team at Microsoft determine the performance characteristics of their system during its design and implementation.

The Microsoft Web Capacity Analysis Toolkit (WebCat) is a tool for evaluating the performance characteristics for HTTP servers and the operating systems on which they run. Some useful applications of the WebCat analysis tool include:

- Capacity planning
- Performance evaluation of different product offerings
- Load simulation -- useful for system tuning and design

WebCat has three components which must be installed:

- ***Server***: The server contains content for any Web Server on any platform. You must install the content files of various sizes used to simulate the content of a Web site for WebCat. If you choose to run WebCat, the following applications are also included: Common Gateway Interface (CGI), Information Server Application Programming Interface (ISAPI), and Netscape Server Application Programming Interface (NSAPI).
- ***Controller***: The controller initiates and monitors an actual WebCat run. The controller must be running on Windows NT Workstation or Windows NT Server.
- ***Clients***: Each client simulates a Web browser by repeatedly sending requests to the Web server. Clients must be running either Windows 95 or Windows NT Workstation.

Webstone

This tool is a product of Silicon Graphics, Inc. (SGI). It is currently positioned as the de facto industry standard tool for performance analysis on Web servers. The source code for Webstone is freely available from SGI. Additional information about SGI and Webstone can be found on the Internet at: http://www.sgi.com

SPECweb

SPECweb96 for Windows NT is a Web server analysis tool that is still in beta form. SPECweb is a development effort by a coalition committee called the SPEC Committee. Additional information about this product and the SPEC Committee can be found on the Internet at: http://www.spec.com

Test Methodology and Test Configurations

Our test system was configured with eight to twelve clients running multiple threads to simulate up to a maximum of 264 clients. Tests were first conducted on a 10-megabit network where Compaq engineers noted a network bottleneck at light to moderate request loads. The tests were repeated over a 100-megabit network and again response proved to be inadequate. Finally, we configured a 100-megabit network where server and clients were split among two 100-BaseTX LANs using Compaq 100-BaseTX NICs as indicated in the chart below. All tests used Microsoft's WebCat analysis tool with the various file sets to measure the server throughput in megabits per second. Details of the testing results are provided in the sections that follow. For more information about analysis tools, refer to the "Tools" section provided later in this document.

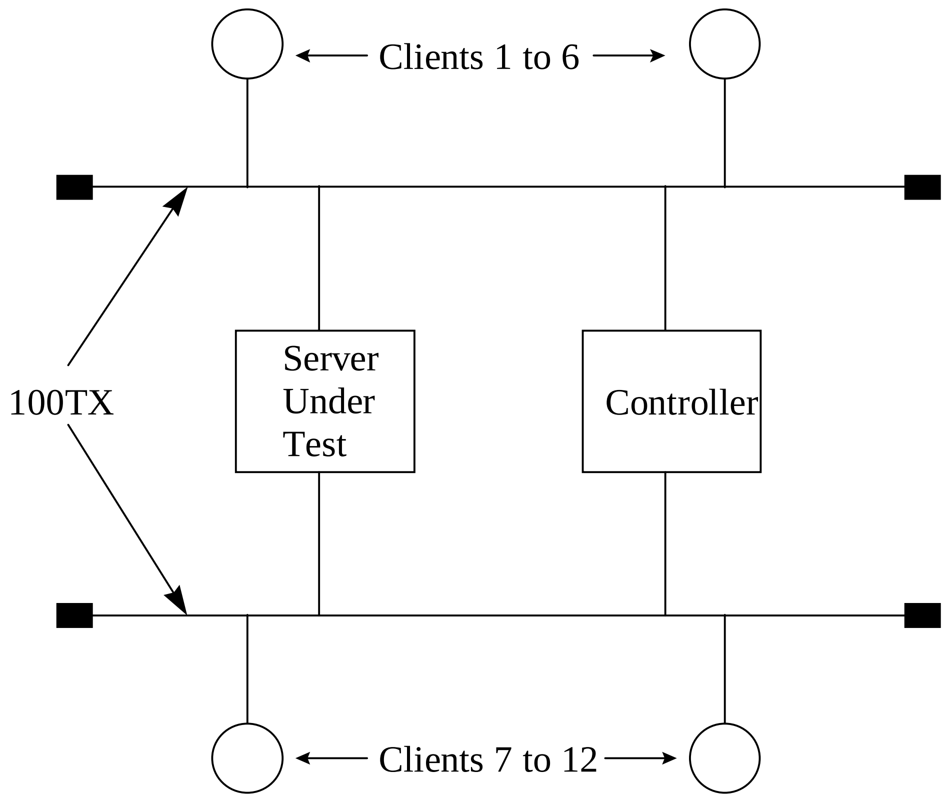

Webstone File Set Sizes

Web server performance will vary based on the type of workload. It is, therefore, important that the test workload closely represents the expected production environment for the server. Compaq engineers determined that the Webstone file set should be used in our tests, since it is the de facto standard for Web server performance comparison at this time. The WebCat analysis tool supports the Webstone file set. This compatibility enabled the use of WebCat to generate a load that approximates the results one would get using Webstone.

CGI/HTML File Distribution

Percentage of time requested	Random file size/selection
2%	256K
2%	512 Bytes
2%	1K
3%	2K
3%	3K
2%	4K
3%	6K
1%	8K
3%	16K
2%	32K
1%	64K
1%	256K
75%	10K
100.00%	total

While Webstone is useful in comparing performance across platforms, actual customer use will include requests for pages containing both Hypertext Markup Language (HTML) and Common Gateway Interface (CGI) content. To highlight the processing power of Compaq ProLiant servers, the CGI file set was selected as the mechanism for testing the server's ability to process requests. CGI is a term that is used for Web servers to denote an application that requires processing and is requested by a Web client but executes on the Web server. The CGI test we used distributed the load between CGI and text files according to the CGI/HTML File Distribution table shown here.

Webstone File Distribution

Percentage of time requested	Random file size/selection
40%	2K, 3K
25%	1K, 5K
15%	4K, 6K
5%	7K
4%	8K, 9K, 10K, 11K
4%	12K, 14K, 15K, 17K, 18K
6%	33K
1%	200K
100.00%	total

Large Filemix File Distribution

Percentage of time requested	Random file size/selection
8%	256 bytes
9%	512 Bytes
8%	1K
17%	2K
12%	3K
7%	4K
10%	6K
5%	8K
10%	16K
7%	32K
4%	64K
1%	256K
1%	512K
1%	1M
100.00%	total

Compaq Server Configuration

- ProLiant 1500
  - One to two 166-MHz Pentium processors
  - 32 MB, 64 MB, and 128 MB of RAM
  - Compaq SMART-2 Array Controller
  - Five 2.1-GB Fast-Wide SCSI-2 hard drives (RAID5)
  - Compaq Netelligent 10/100 TX PCI UTP Controller
  - Microsoft Windows NT 4.0 Server
  - Microsoft Internet Information Server (IIS) 2.0
  - Netscape Enterprise Server 2.0
- ProLiant 5000
  - One to four 166-MHz Pentium Pro processors
  - 32 MB, 64 MB, 128 MB, and 512 MB of RAM
  - Compaq SMART-2 Array Controller
  - Five 2.1-GB Fast-Wide SCSI-2 hard drives (RAID5)
  - Compaq Netelligent 10/100 TX PCI UTP Controller
  - Microsoft Windows NT 4.0 Server
  - Microsoft Internet Information Server (IIS) 2.0
  - Netscape Enterprise Server 2.0

Web Client Configuration

- Twelve ProLiant 2000 clients configured with:
  - Two 90-MHz Pentium processors
  - 32 MB RAM
  - 2.1-GB disk
  - Compaq SMART SCSI Array Controller
  - Netelligent 10/100 NIC configured for 100 Mb/s
  - Microsoft Windows NT 3.51 Service Pack 4, SSD 1.18
  - WebCat Client software

Performance Characteristics

Using the previously described systems and tools, Compaq engineers conducted performance testing on each of the following subsystems:

- Processor
- Memory
- Bus
- Disk
- Network

On the pages that follow, various charts accompanied by descriptions provide details of the subsystem testing and results.

Processor

The figure below shows the processor comparison on the Compaq ProLiant 1500 running Windows NT 4.0 and IIS 2.0. The test results showed an average of 50 percent performance gain going from a single 166-MHz Pentium to a dual 166-MHz Pentium processor system. The test was run using the WebCat Webstone file set to simulate up to 264 client workloads that exercised static HTML requests. The throughput drop from the 96- to 192-client workloads were caused by client saturation on the network. Therefore, performance throughput stayed flat from the 192- to 264-client workloads.

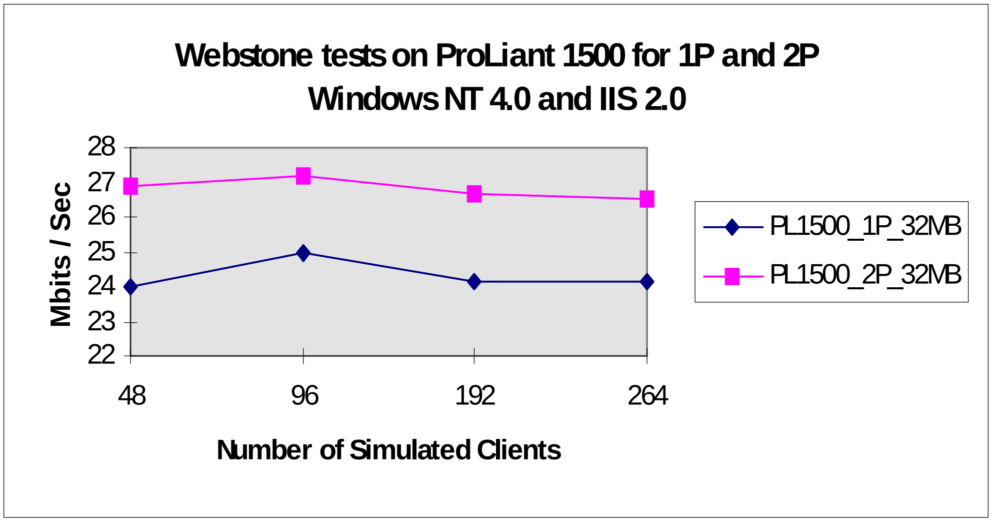

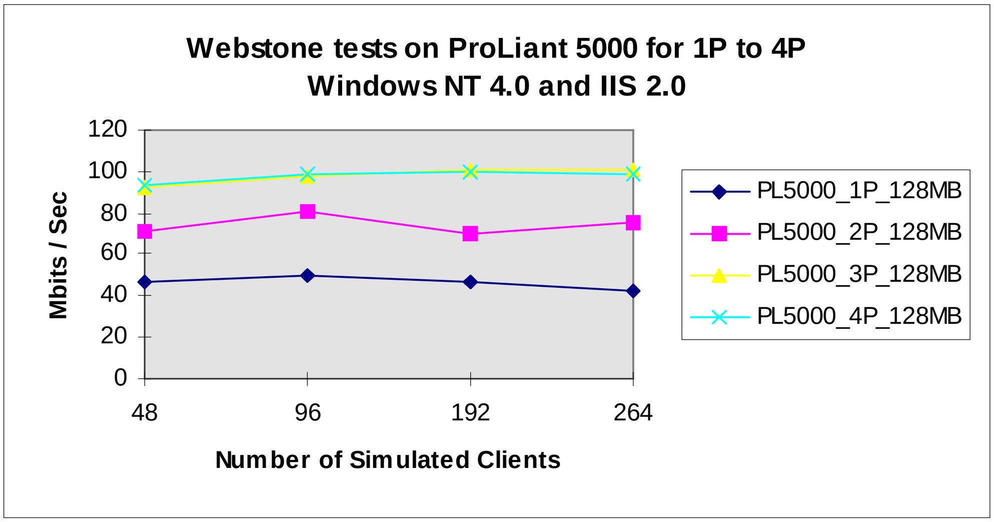

The ProLiant 5000 with dual 166-MHz Pentium Pro processors is 70 percent faster than the single 166-MHz Pentium Pro processor system. With three 166-MHz Pentium Pro processors, it is about 80 percent faster than the single 166-MHz Pentium Pro processor system and 10 percent faster the dual 166-MHz Pentium Pro processors system. The system does not scale beyond the third processor, therefore performance throughput on the third and fourth processors is flat. This is caused by the operating system's slow response to the interrupt requests (IRQs) across the processors. IRQs are distributed amongst all available processors in a Multi-Processor Specification (MPS) system. When a hardware device interrupts the processor, the Interrupt Handler may elect to execute the majority of its work in a deferred procedure call (DPC) which is executed in Privileged Mode. This is typical behavior of the current driver models.

DPCs run at lower priority than Interrupts, thus permitting new Interrupts to occur while DPCs are being executed. The processor assigned to handle interrupts for a particular network interface is determined during initialization. Therefore, the addition of multiple NICs will distribute a heavy load of DPCs on an MPS system, like the ProLiant 1500 and ProLiant 5000 running Windows NT Server 4.0, to take advantage of multiple processors and help alleviate the possibility of a processor subsystem bottleneck.

All tests were run using the WebCat Webstone file set that generated static HTML requests. The server is configured with one to four 166-MHz Pentium Pro processors, Windows NT 4.0, and IIS 2.0 to test the processor scalability on the ProLiant 5000. In summary, the test results demonstrate that, in various workloads requests (light, medium, and heavy), the ProLiant servers scale from one processor to three processors very well.

Memory

The figures below show the memory comparison on the ProLiant 1500 and ProLiant 5000 running Windows NT 4.0 and IIS 2.0 with 32 MB, 64 MB, and 128 MB of RAM. The tests were run using the WebCat file mix distribution methodology with random file size selection ranging from 256 bytes to 1MB. The test results illustrated that adding additional memory to the server will improve server throughput. The 64-MB RAM configuration outperformed the 32-MB RAM configuration by an average of 50 percent while the 128-MB RAM configuration showed an average of 100 percent performance gain at light to moderate workloads. The spikes and performance drop on the charts at the 96 simulated clients to 264 simulated clients level were caused by a network bottleneck and client saturation.

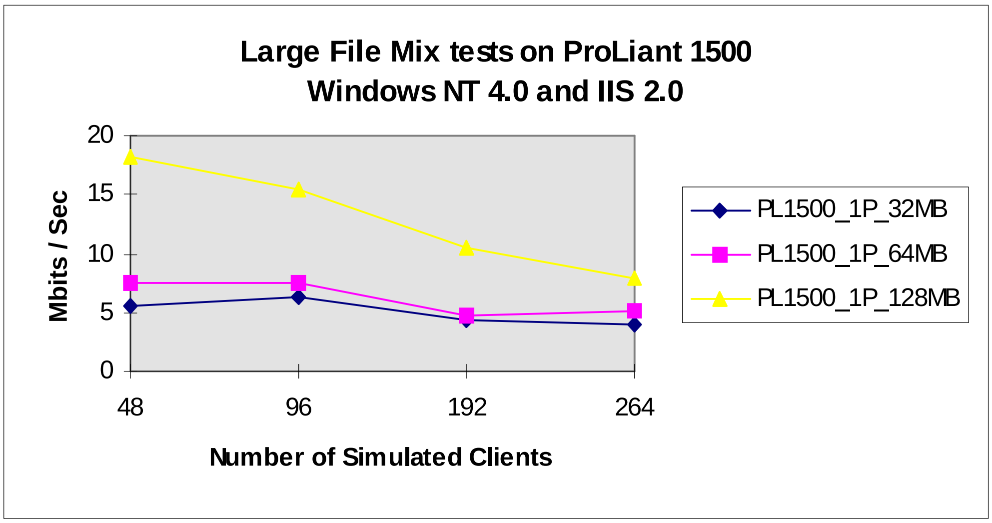

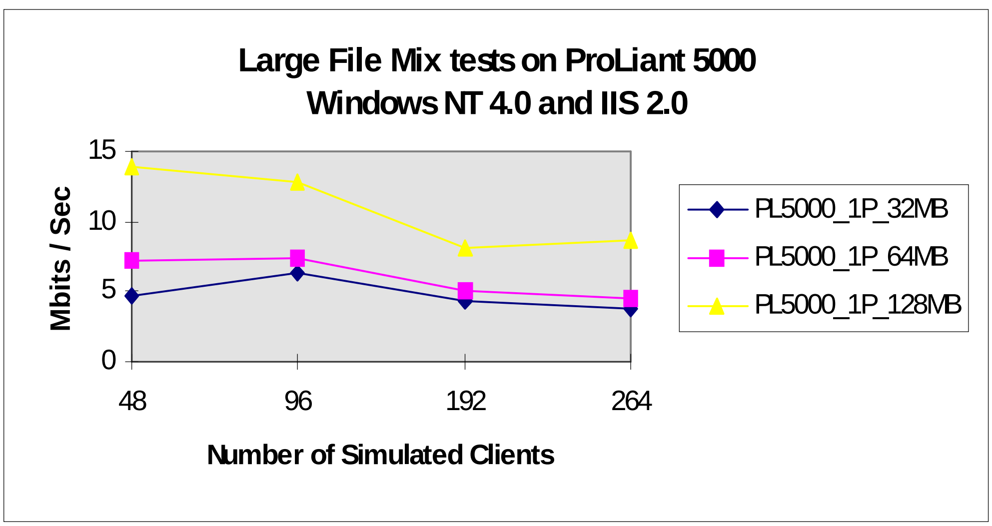

Memory is one of the most valuable resources in a Windows NT system for overall performance. The amount of memory installed and available in a system impacts on total system performance. Since memory is used for disk caching as well as program execution, additional memory in the system, in certain environments, will improve performance. The charts above show a steady performance gain from 32 MB to 64 MB and 128 MB. Based on observations and analysis of the test data, it is believed that performance will continue to improve with additional memory, up to

256 MB, for systems that serve data similar to the Webstone and Large file set tests where cache is important.

Bus Subsystem

The Compaq ProLiant 5000 and ProLiant 1500 were again used for our system under test. Each system was configured with a PCI disk and network controller and then an EISA disk and network controller. The Microsoft WebCat Analysis Tool was used to collect data which was analyzed to compare the PCI and EISA bus performance. The WebCat Webstone file set was the workload utilized in the test runs which generated the bus subsystem data presented in the following diagram. The Webstone file set testing was run for workloads varying from 48 to 264 simulated clients, exercising static HTML requests.

Testing revealed that the bus subsystem has a significant impact on performance in some cases. The systems under test were ProLiant servers with single 166-MHz processors, 64 MB of RAM, and Windows NT 4.0 as the operating system . Both the ProLiant 1500 and the ProLiant 5000 were configured using the Compaq SMART-2/P Array Controller for the disk controller and two NetFlex-3/P network interface controllers (NICs) for the PCI configurations. The EISA test configurations consisted of the same ProLiant 5000 and ProLiant 1500 using a Compaq SMART-2/E Array Controller and two NetFlex-3/E NICs. An array of five 2.1 GB pluggable drives configured at a RAID 5 level of hardware fault tolerance was used for both PCI and EISA tests. The SMART-2 Controller was configured with a Read/Write ratio of 50% Read and 50% Write. The MaxReceives buffers for the NetFlex-3 NICs were set to 500. This tuning parameter and others will be discussed further in the "Performance Tuning" section.

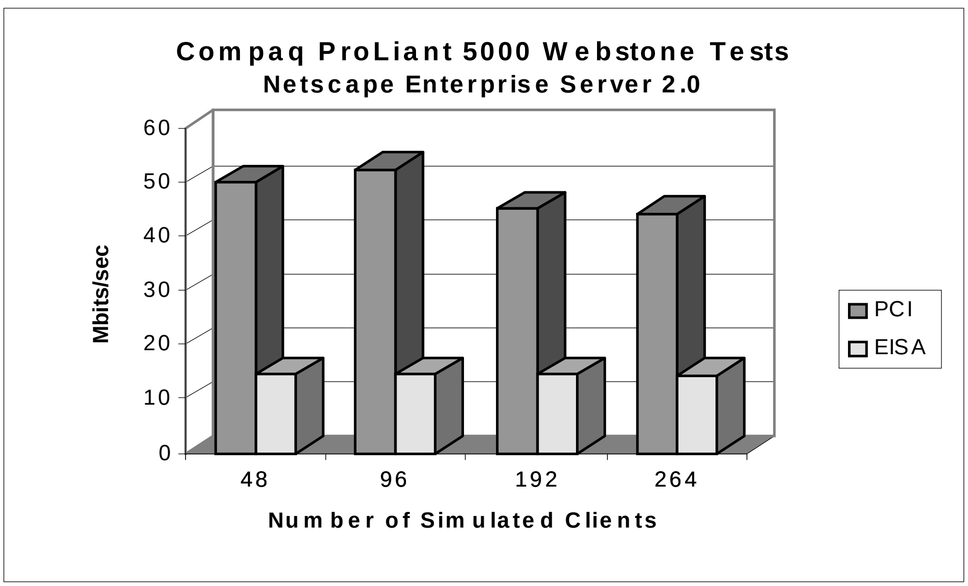

NOTE:
The Technology Brief titled "Configuring the Compaq ProLiant 5000 Server for Peak Performance" contains configuration reference charts that can help you ensure you are correctly configuring your PCI controllers. This document can be found on the Compaq Web site at "http://www.compaq.com".

The ProLiant 5000 is optimized for PCI controllers resulting in the superior performance of the PCI test configuration. The figure above shows the bus performance comparison on the ProLiant 5000 running Windows NT 4.0 and Netscape Enterprise Server 2.0. The PCI configuration outperformed the EISA configuration ranging from 310 to 363 percent. This large difference can be primarily attributed to the dual, peer PCI bus architecture of the ProLiant 5000. Two PCI buses are independently linked to the Pentium Pro processor bus by Host-to-PCI bridges. This dual, peer bus architecture supports aggregate I/O throughput as high as 267MB/s. To obtain the optimal load across both PCI buses, the system must be carefully configured.

Compaq engineers recommend PCI network and disk controllers over EISA controllers in the Compaq ProLiant 5000 to experience performance gains similar to those seen during WebCat testing. The performance improvement attainable by using PCI network and disk controllers justify the cost of the new controller in all environments where high throughput and rapid response time are requirements.

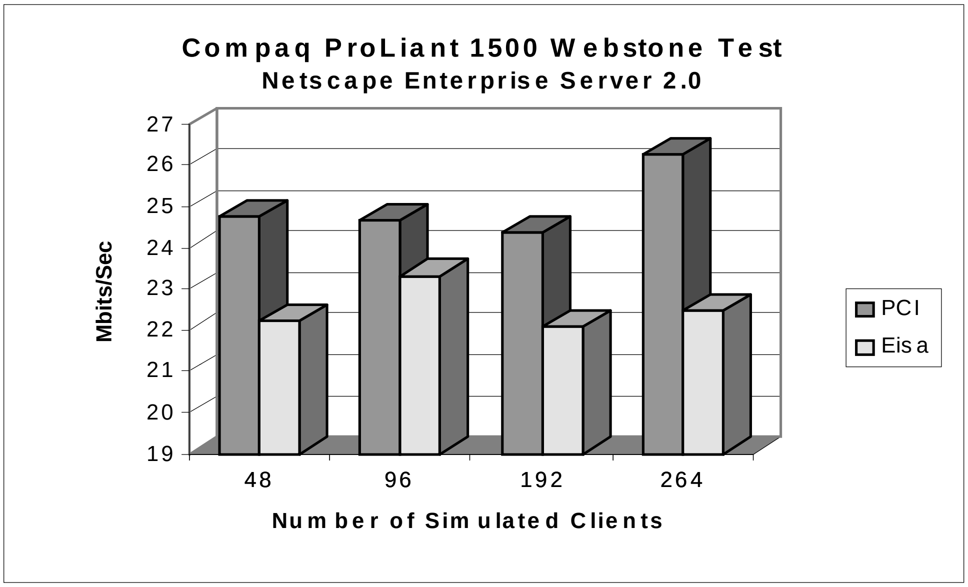

The figure above shows the bus performance comparison on the ProLiant 1500 running Windows NT 4.0 and Netscape Enterprise Server 2.0. The performance difference between the PCI and EISA configurations of the Compaq ProLiant 1500 is much less significant than previously shown for the Compaq ProLiant 5000. Tests for the ProLiant 1500 PCI configuration showed a 5 to 17 percent performance increase over those for the ProLiant 1500 EISA configuration. This much smaller difference between PCI and EISA can largely be attributed to the differences in architecture. The ProLiant 1500 does not have the ProLiant 5000's dual, peer PCI bus. The maximum transfer rate on the ProLiant 1500's EISA bus is 33 MB/s while the maximum transfer rate on the PCI bus (peak) is 133 MB/s.

Disk Subsystem

The disk subsystem has a performance impact on all applications. The amount of I/O required by your application determines the degree of impact on the disk subsystem performance. The higher the level of disk I/O, the greater the performance impact on the disk subsystem.

Customers have several options for configuring their Web server's hard disks and making a decision regarding fault tolerance. Frequently, fault tolerance is referred to by the level of Redundant Arrays of Inexpensive Disks (RAID) that is supported. RAID is a term used to refer to an array technology that provides data redundancy to increase the overall system reliability and performance. The Compaq SMART-2 Array Controller is required to configure Compaq ProLiant 5000 and ProLiant 1500 platforms with the following hardware fault tolerance options:

- RAID 0 - No fault tolerance support
- RAID 1 - Disk mirroring
- RAID 4 - Data guarding
- RAID 5 - Distributed data guarding

The disk subsystem performance testing of the ProLiant 5000 and ProLiant 1500 consisted of WebCat test runs which compared the performance of a single drive with no fault tolerance support, an array of drives configured with RAID 5 fault tolerance, and an array of drives configured with RAID 1 fault tolerance. RAID 4 was not included with the tests because it offers lower performance than RAID 5 while using a similar amount of available disk storage for fault tolerance. The WebCat Analysis Tool was used to collect the disk subsystem data presented in the charts which follow. The WebCat Large Filemix file set was used as the workload in test runs which generated the disk subsystem data presented below. These Large Filemix file set tests were run for workloads varying from 48 to 264 simulated clients, exercising static HTML requests.

RAID 1 or Disk Mirroring works as its name implies, storing duplicate data on a pair of disk drives. RAID 1, therefore, always requires an even number of disk drives in the array. For RAID 1 tests, four 2.1GB drives were set up using a single SMART-2/P Controller. The Array Configuration utility was used to set up two 2.1GB drives forming a logical drive which is then “mirrored” or duplicated to the other two 2.1GB drives which form a second logical drive. This form of fault tolerance has the highest associated cost due to 50 percent of the total drive storage capacity allocated for fault tolerance.

RAID 5 or Distributed Data Guarding is referred to as such because parity data is used to guard against the loss of data. This parity data is striped or distributed across all the drives in the array. RAID 5 provides very good data protection because this parity data can be used to reconstruct data on a failed drive. The usable storage depends on the total number of drives in the array. Up to fourteen drives can be configured in a single array. The more drives you set up in the array, the better the performance you can expect.

For RAID 5 tests, five 2.1GB pluggable drives were set up in an array and the SMART-2/P Controller was used for all tests. The SMART-2 Controller was configured with a Read/Write ratio of 50% Read and 50% Write.

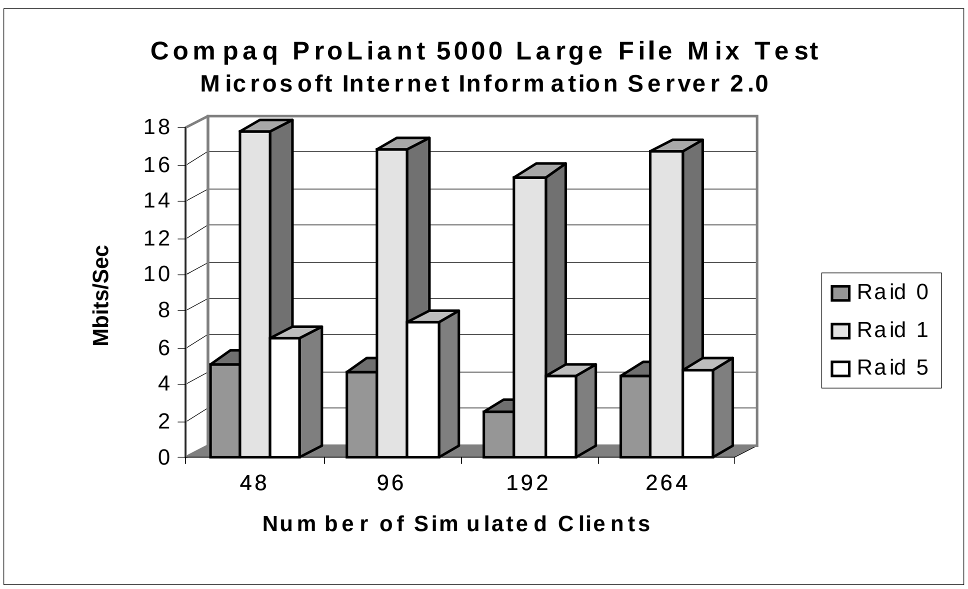

The chart above illustrates the disk subsystem impact on performance in the Compaq ProLiant 5000. Clearly RAID 1 outperforms the single drive spindle and the drive array configured with RAID 5 fault tolerance. RAID 1 outperforms RAID 5 providing from 2 - 3.5 times the throughput. RAID 5 outperforms the single drive to a less significant degree.

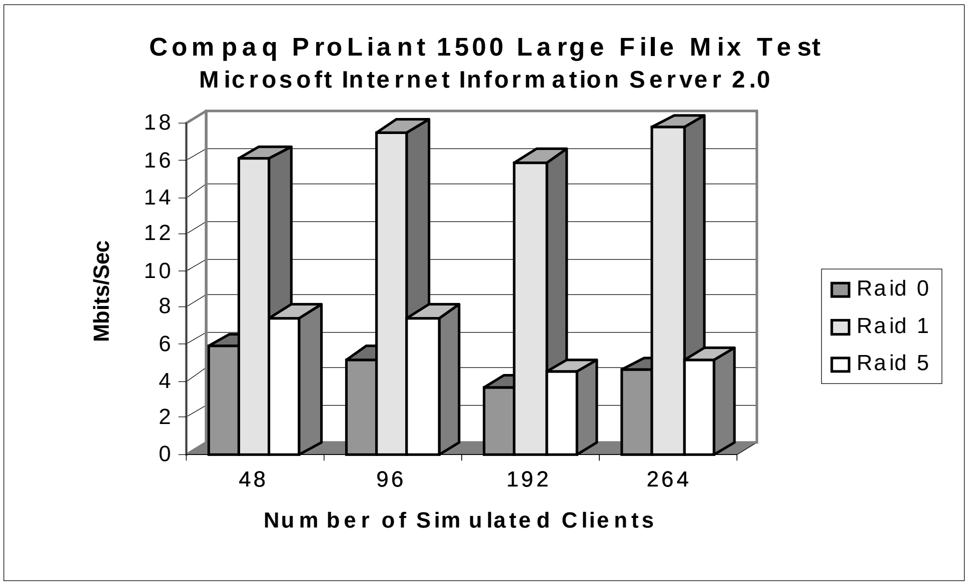

The chart above illustrates the disk subsystem impact on performance in the Compaq ProLiant 1500. Clearly RAID 1 outperforms the single drive spindle and the drive array configured with RAID 5 fault tolerance. RAID 1 outperforms RAID 5, providing a 30 to 50 percent increase in throughput. RAID 5 outperforms the single drive by approximately 20 percent.

RAID 1 is recommended when evaluating the optimal configuration from a purely performance perspective. Note, however, that RAID 1 configurations are more costly than RAID 5 because 50 percent of your RAID 1 configuration is used for fault tolerance.

Network

NOTE:
Factors such as whether the wire is atomic or divisible, client and server side TCP/IP stack over head, I/O subsystem overhead, network congestion, number of files/images composing the page, etc. will effect actual throughput.

In the networking tests run on 10-megabit Ethernet, we found that the network was easily the first bottleneck that we encountered. Splitting the network into several segments alleviated the bottleneck. Moving to a 100-megabit split network proved to be the best performing solution of all. By splitting the network into multiple segments, the load can be distributed across multiple NICs while distributing Deferred Procedure Calls (DPCs) across multiple processors.

When choosing a network connection for a Web server, there are several factors that should be included in your comparison.

- Amount of simultaneous requests
- Available bandwidth
- Cost
- Existing infrastructure

The following table illustrates the theoretical, as well as practical application of common networking technologies available for comparison purposes.

Physical Channel	**100BaseT or FDDI**	**10BaseT**	**T1**	**ISDN (BRI)**
Bandwidth, Primary Business Use	100Mb, Backbone or Intranet	10Mb, Intranet	1.544Mb, Large business, Web Site	128Kbps, Med. - Small business, ISP Connection
Conn/Sec. @ 50K in theory	2048	204.8	32	2.5
Conn/Sec. @ 50K in practice	1536	153.6	24	2
Conn/Sec. @ 250K in theory	409.6	41	6.5	.5
Conn/Sec. @ 250K in practice	385.1	38.5	5.8	.4

TUNING PARAMETERS

Prior sections of this paper have discussed various hardware configurations to help you achieve the most out of a Compaq Web server. The recommendations below affect key software components which ultimately control hardware performance.

The tuning recommendations covered in General Tuning and Microsoft Internet Information Server Tuning are for the Windows NT Registry parameters. The tuning recommendations covered in Netscape FastTrack Server and Enterprise Server Tuning modify various non-Registry files.

General Tuning

The parameters specified for General Tuning are applicable to both Microsoft Internet Information Server and Netscape.

Compaq Network Controller Tuning:

- For the Windows NT Registry section:

  HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\cpqnf3(#)\Parameters

  Add the following parameter:

  MaxReceives = REG_DWORD 0x1F4 = 500

  – Increases the number of MaxReceives counters for Compaq Netelligent 10/100TX Network Controller to 500. (The default is 100.)

  – Specifies the maximum number of receive lists the driver allocates for receive frames

Microsoft Internet Information Server Tuning

- For the Windows NT Registry section:

  HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\InetInfo\Parameters

  Add the following parameter:

  ListenBackLog = REG_DWORD 300

  – This insures that the Web Server has enough listening endpoints for new connections coming in.

  – This avoids the connection failure errors on busy servers (caused due to lack of end-points).

- For the Windows NT Registry section:

  HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\W3SVC\Parameters

  Add the following parameter:

  AcceptExOutstanding = REG_DWORD 100

  – Specify the number of AcceptEx sockets that will be outstanding and active. The server maintains these incoming endpoints to pick up new connections. The larger you make this pool, the less amount of time the clients have to wait in the queue.

- For tests with large number of clients, keeping this number high will ensure that sufficient number of client connection contexts will be maintained.

Netscape FastTrack Server and Enterprise Server Tuning

- In the Netscape configuration file:

  MAGNUS.CONF

  Set MaxThreads equal to the number of megabytes of installed RAM.

- If you have a large number of files, or a large total content size, you can set some special parameters. In the Netscape configuration file:

  OBJ.CONF

  - for 1024 files: Init fn='cache-init' cache-size=1024
  - for 1024000kB(1000MB) of documents: Init fn='cache-init' mmap-max=1024000

NOTE:
Refer to the "Processor" section to understand why this is important in systems with multiple processors.

- If you have more than one CPU, perform the following steps:

  1. Determine which CPU is the IO Processor (Performance Monitor will show more usage on this CPU). For this example, assume that CPU 0 is the IO Processor.
  2. Create a file C:\WINNT\HTTP.INI containing:
     [HttpAccel]
     IOProcessor=0

SUMMARY

Web server tools that are available today can help determine potential bottlenecks in your Web server prior to investing resources on a permanent solution. Given the data presented in the charts as well as some of the tuning parameters, it's easy to see that the selection of a server and its components are critical to overall performance.

Performance testing thus far has shown the Web server to be very similar to a classic application server. The processor scales as expected with greater increases in performance experienced when moving from one processor to two than from any incremental processor additions beyond two processors. Data revealed that increases in memory help boost performance more in tests run using the Large Filemix file set. The amount of memory required for a Web server will be determined by the type of content the server is expected to support. For instance, memory will not be as big of a factor in smaller workloads as it is with larger workloads.

The Web server should be configured using PCI controllers such as the SMART-2 Controller if performance is a critical issue. The server should be configured with as many drives as possible to obtain the desired storage requirements and provide optimal performance. Lastly, monitor your network utilization and consider switching to a 100Mb/s network if a bottleneck is discovered.

Compaq offers high performance Web server solutions for the Windows NT environment with both Microsoft Internet Information Server and Netscape Enterprise Server software running on Compaq ProLiant servers. In addition, Compaq ProLiant servers are optimized to take advantage of the enhanced performance of Windows NT 4.0 as well as Microsoft Internet Information Server and Netscape Enterprise Server products.

Invest some time to planning any Web site, be it internal or external, for peak performance in preparation for future popularity and growth. By referencing the data available in this paper, you can optimize your configuration to match the type of content to be distributed. This should help you predict possible bottlenecks at peak access times and work to improve overall response time. If your site is too slow, your intended audience will only visit it once.