Purpose Built N440BX DP Server Board

Performance Brief

SPECweb96 Performance

350 MHz/512KB Pentium® II processor
400 MHz/512KB Pentium II processor

intel®

Revision 1.0
April, 1998

THIS TEST REPORT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, reference www.intel.com/procs/perf/limits.htm or call (U.S.) 1-800-628-8686 or 1-916-356-3104.

SPECweb96* benchmark test provides performance data helpful to those looking into purchasing a Web server. In its initial release, the benchmark focuses on server performance for static Web pages, measuring a server's ability to serve HTTP requests. SPEC officials are quick to emphasize that SPECweb96 is strictly a Web server benchmark. It does not measure Web client, Web client/server or WAN performance. For more information about SPECweb96, including a description of the systems and other information about microprocessor and system performance, visit SPEC's World Wide Web site at www.specbench.org.

Information in this document is provided in connection with Intel® products. No license, express or implied, by Intel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

The hardware vendor remains solely responsible for the design, sale and functionality of its product, including any liability arising from product infringement or product warranty.

Table of Contents

INTRODUCTION

The NC440BX DP server platform features dual Pentium® II Processors. The Intel® N440BX server board and the 82440BX chip set are the key ingredients you need to build robust servers. Designed for dual processing using the latest high-performance Intel® Pentium® II processors with 100MHz System Bus, the N440BX server board is the foundation for creating purpose-built servers designed for the most demanding business environments. It comes with everything you need to quickly build your servers and install at your customers fast: drivers, documentation, management software, diagnostics, and tremendous flexibility and genuine Intel quality.

With the N440BX Server board as your foundation, you'll get a wealth of competitive features. Intel's new Emergency Management Port, which provides remote server access for quick troubleshooting and repair; extra-large memory support — up to 1 GB; on-board dual-channel ultra SCSI, high performance 10/100Mb Ethernet LAN, 2 MB video; and five slots for huge flexibility in customizing solutions. All the latest technology is here, for fast, easy integration and fast time to market. One and two CPU measurements were taken at 350 and 400MHz

This NC440BX platform report is comprised of results from the SPECWeb96* benchmark. The configuration used for this benchmark was put to maximum stress to achieve reasonably performing results. We call these baseline numbers. A customer should be able to achieve at least this performance given that the configuration is duplicated.

DESCRIPTION

SPECweb96 is a standardized benchmark for comparing web server performance. The benchmark is designed to provide comparable measures of how well systems can handle HTTP GET requests. The SPECweb96 workload is based upon analysis of server logs from web sites ranging from a small personal server up through some of the Internet's most-popular servers. During a benchmark session, one or more clients are used to send HTTP requests to the server. The SPECweb96 software then measures the response time for each request. At the end of a benchmark run, which takes about three to four hours, SPECweb96 calculates a metric based on overall throughput, measured as maximum operations per second (OPS/SEC).

SPECweb96 workloads are based on analyses of server logs from a variety of popular Internet servers and some small Web sites. SPEC further validates workload information by comparing data from its analyses to logs from Netscape and CommerceNet, and to those of its sponsoring vendors. SPEC divides workload files into four categories according to size, ranging from less than 1KB to less then 1MB. Access patterns to the files are rated according to the analyses of server logs, and the access analyses approximate "real-world usage."

SPECweb96 works by having one or more client systems generate a load of HTTP GET requests against a selected server system. The SPEC committee provides the code, which executes simultaneously these clients. SPECweb96 reports its overall web server performance in operations per second (ops/sec). The operation per second metric used in SPECweb96 is defined as the maximum number of HTTP operations per second that a Web-server can sustain. The workload mix is built out of files in four classes: files less than 1KB account for 35% of all requests, files between 1KB and 10KB account for 50% of requests, 14% between 10KB and 100KB, and finally 1% between 100KB and 1MB. The configuration and implementations of the HTTP Server on the

N440BX server under test (SUT) are up to the user performing the benchmark.

A valid result consists of ten evenly spaced measurements. The spacing of the request size is determined by the maximum operations per second selected by the user. All results, to be valid, must be within 5% of the requested size.

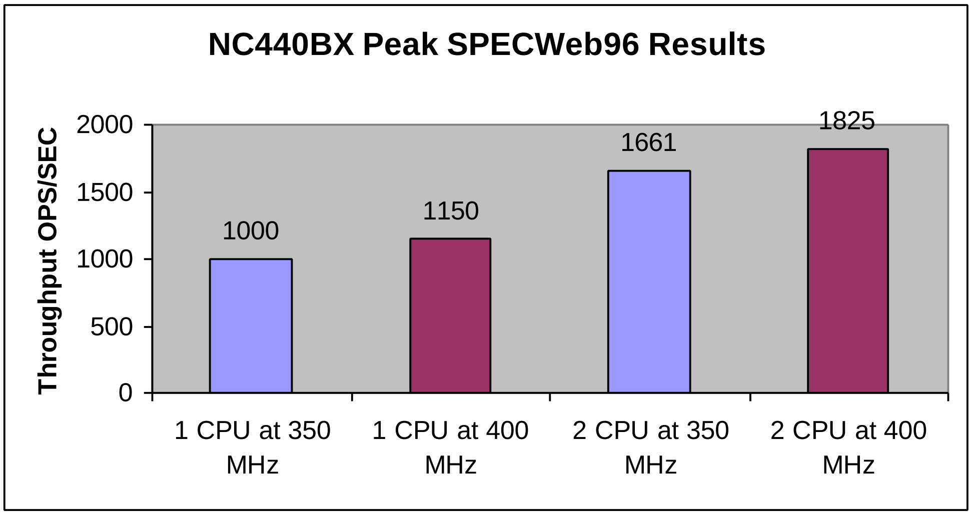

CLIENT/SERVER TEST ENVIRONMENT

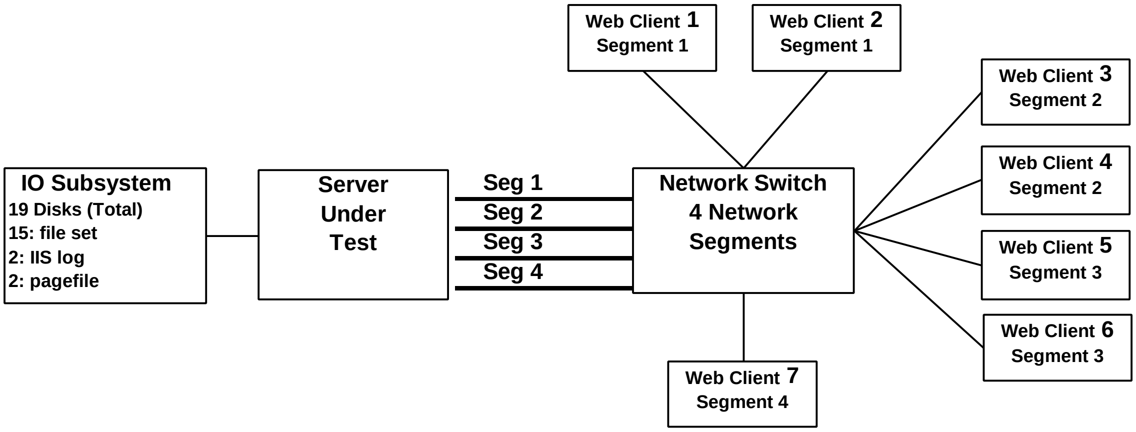

SYSTEM CONFIGURATION TUNING; I/O SUBSYSTEM

The IDE disk subsystem was used for the boot device. The AMI 434 RAID controller was configured using 15 disks across two I/O channels (7+8 disks) to represent a logical device. This logical device was the web server's home page that contained the file set and was configured as RAID 0, 8k-stripe size, write through mode, with direct IO. A Windows NT* 1.5 GB NT NTFS file system was created on this logical device. Two other logical devices were created on the third channel of the AMI RAID controller. Each consisted of two hard drives each, also configured as RAID 0, 8k-stripe size, write through mode, with direct IO. One of these logical devices was used for the IIS log device using a similar Windows NT (1.5 GB, NTFS) file system. The other logical device was used for the pagefile using a 2.5 GB NTFS file system.

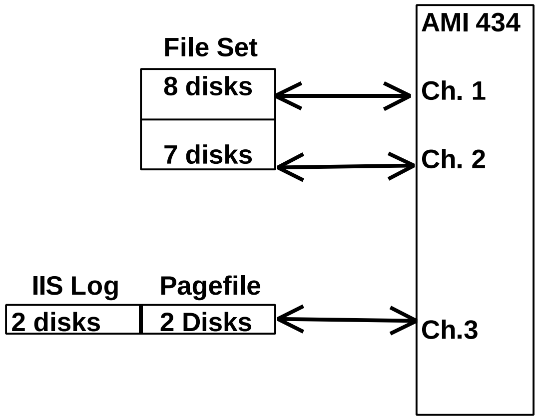

SYSTEM CONFIGURATION TUNING; NETWORKING

The server's networking environment consisted of four 100Mb/s, switched full duplex network segments (three Intel® EtherExpress™ PRO100+ and the onboard Intel 82558). TCP/IP network protocol was used for connectivity between all clients and the web server. The network was split into four segments. This is done to balance the network load. Three segments consisted of two clients and one connection to the web server. The fourth segment only had one client and a connection to the web server.

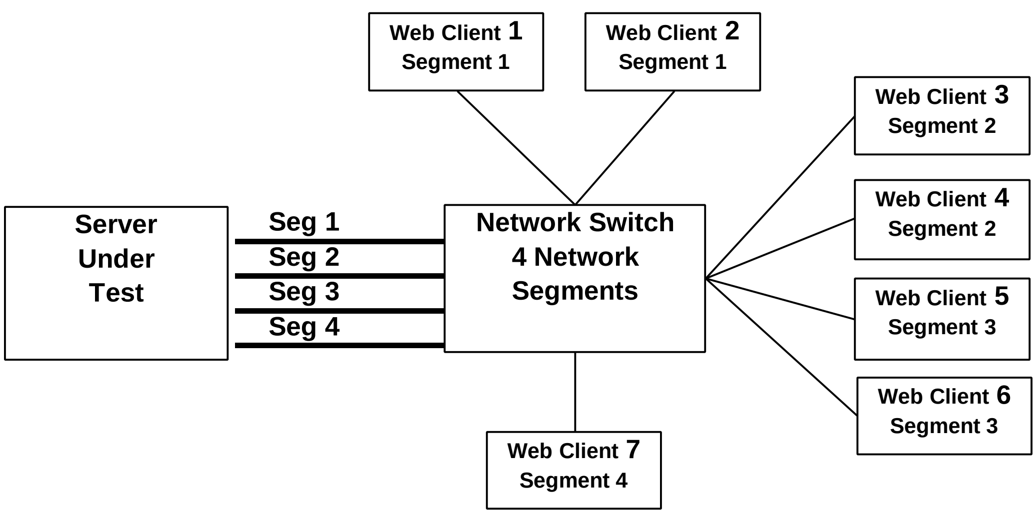

THROUGHPUT:

1 CPU at 350 MHz	2 CPU at 350 MHz	1 CPU at 400 MHz	2 CPU at 400 MHz
100	166	115	182
200	332	230	365
300	498	345	547
400	664	460	730
500	830	575	912
600	996	690	1095
700	1162	805	1277
800	1328	920	1460
900	1494	1035	1642
1000	1661	1150	1825

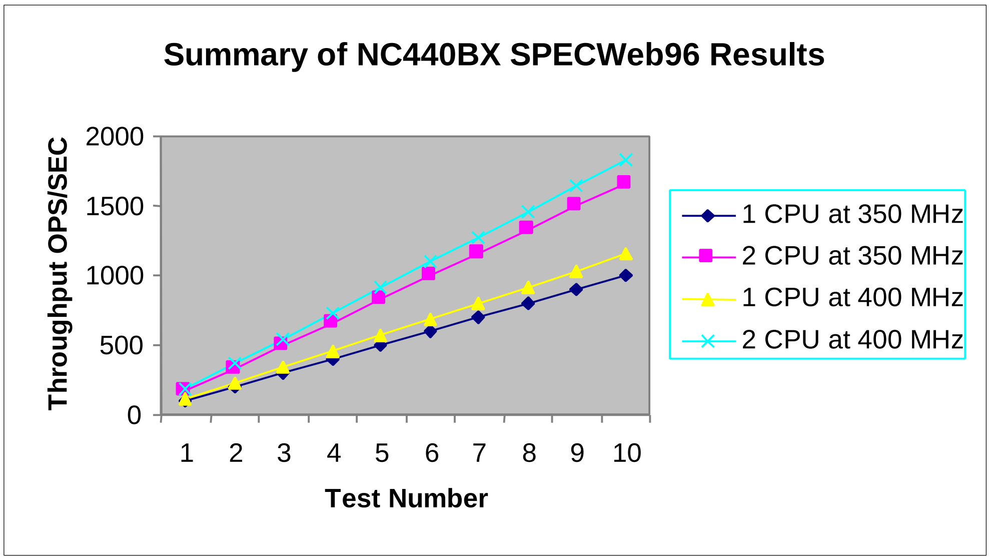

At 350MHz, two processor peak results increases by 60% compared to 1-processor results. Scaling at 350MHz is 1.6x. At 400MHz, two processor peak results increases by 59% compared to 1-processor results. Scaling at 400MHz is 1.59x. Scaling was calculated using the maximum ops/sec generated during the 10th and final test.

CONCLUSION

During the actual benchmark, in both one and two processor measurements, CPU utilization had reached 100% during the peak SPECweb96 results. The network utilization indicated was measured at the four network interface ports on the server system. Total Network utilization on the network switch was 23% and the switch CPU utilization was 30%. Networking did not appear to be a bottleneck as average utilization for each Intel® EtherExpress™ PRO100B in the clients was 32% for the one CPU measurement and 45% for the two CPU measurements. However, this is not network utilization, but instead equates to Mbit/sec using Network monitor agent. For one processor, the total Mbit/sec for all four NICs for transmit and receive was approximately 112 Mbit/sec. Two-processor network utilization for all four NICs was 161 Mbit/sec. Theoretical bandwidth is 400Mbits/sec, but actual bandwidth is 360 Mbit/sec maximum for four NICs.

We noticed the PRO100B #2 throughput was lower. The reason is this network segment had only one client while all others had two. Disk Latency was within reasonable limits for both one and two processor results. Disk Transfers on the IIS log device (E:) were 100% writes. Disk Transfers on the data device (D:) were 100% reads. One-processor results do not show transfer activity on the data device, as the file set was small enough to fit into file cache. This was not the case for two processors as disk read activity occurred. The request size was larger which meant more data needed to be cached. The data size for the request of 1825 ops/sec is approximately 954 Mbytes. The tested N440BX server system had 1GB of system memory installed. With the overhead of Windows NT plus 128MB assigned to the IIS cache (web cache), the data is too large to completely fit into file cache. This accounts for the disk activity of 6524 Bytes/transfer with 6ms latency. The size of the disk read transfer varied depending on the class of data selected from the web server. We also noticed the response time drops slightly during the second run. The reason is that the data is cached in memory.

Scaling is calculated by using peak results measured to compare one and two processor performances. Scaling from one CPU to two CPUs at 350MHz was 1.6x. Scaling from one CPU to two CPUs at 400MHz was 1.59x. The slightly lower scaling at two processor levels at 400Mhz is due to increased I/O activity as the benchmark request size was 165 ops/sec higher that the request size for two processors at 350 MHz

PERFORMANCE SUMMARY WITH CONCLUSIONS

- SMP scaling from 1 to 2 processors was 1.6 for 350MHz. The peak SPECweb96 result was 1000 ops/sec for one CPU and 1661 ops/sec for two CPUs. Adding a second CPU saw a 60% gain in performance.
- SMP scaling from 1 to 2 processors was 1.59 for 400MHz. The peak SPECweb96 result was 1150 OPS/sec for one CPU and 1825 ops/sec for two CPUs. Adding a second CPU saw a 59% gain in performance.
- The 400MHz results scaled slightly lower than 350MHz due to increased I/O activity caused by a larger benchmark request size.
- Performance gain from 350MHz to 400MHz with one-processor results is 13%.

- The performance gain from 350MHz to 400MHz with two processor results is 9%.

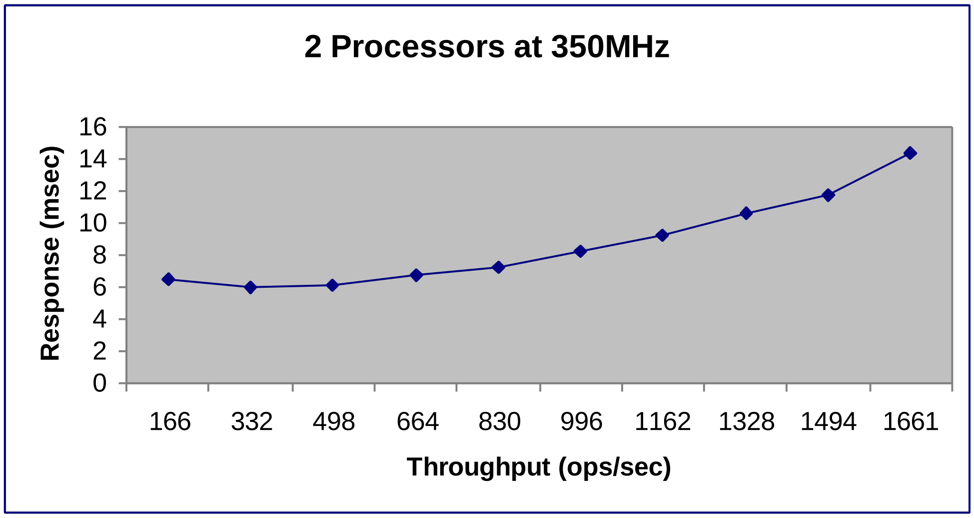

(2 CPU at 350MHz)
Maximum throughput was 1661 operations per second (ops/sec)

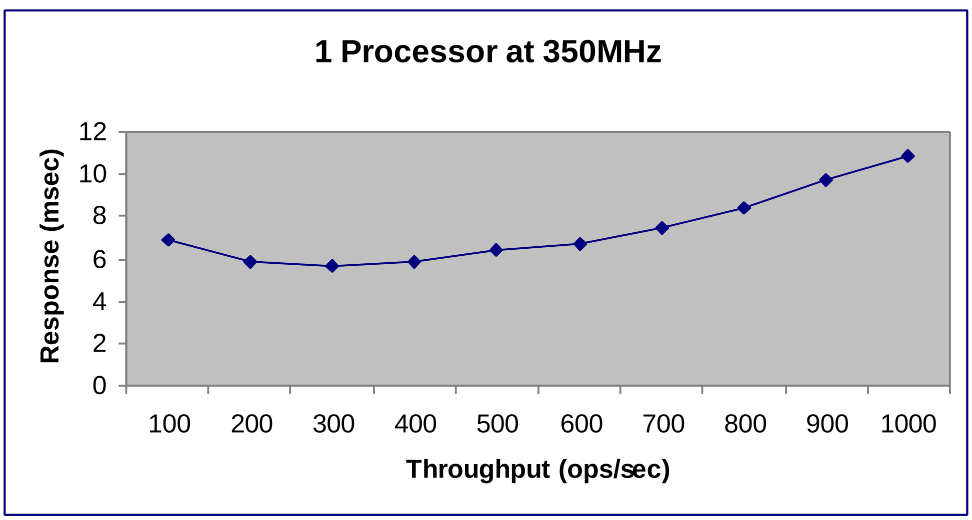

(1 CPU at 350MHz)
Maximum throughput was 1000 operations per second (ops/sec)

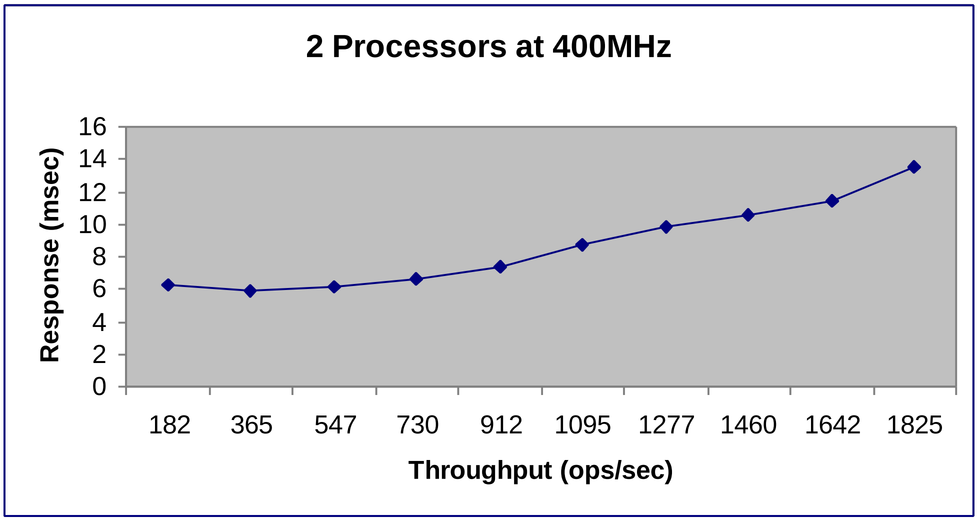

(2 CPU at 400MHz)
Maximum throughput was 1825 operations per second (ops/sec)

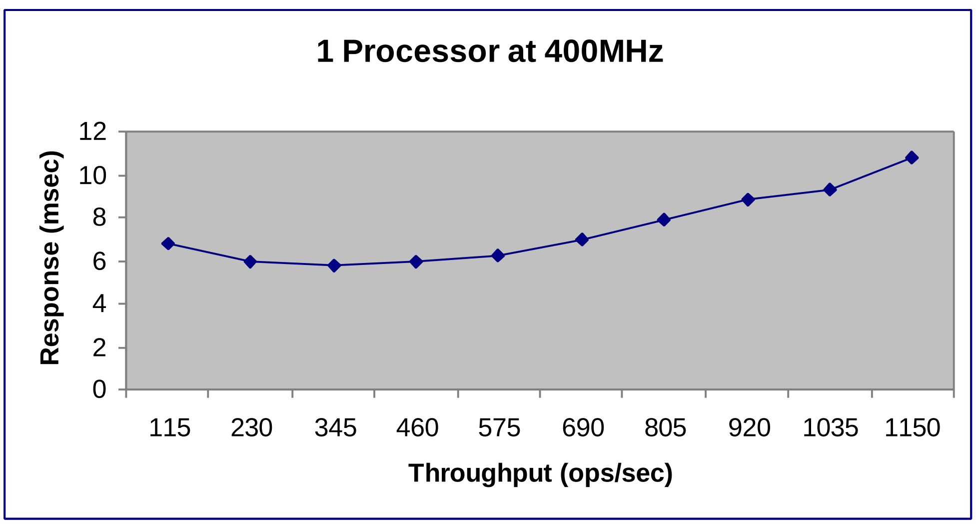

(1 CPU at 400MHz)
Maximum throughput was 1150 operations per second (ops/sec)

APPENDIX A: WORKLOAD PERFORMANCE METRICS

There are nine discrete sizes within each class (e.g. 1 KB, 2 KB, on up to 9KB, then 10 KB, 20 KB, through 90KB, etc.). However, accesses within a class are not evenly distributed; they are allocated using a Poisson distribution centered on the midpoint within the class. The resulting access pattern mimics the behavior where some files (such as "index.html") are more popular than the rest, and some files (such as "myfile.gif") are rarely requested.

Sizes (in bytes) of Files in Each Class

Class 0	Class 1	Class 2	Class 3
102	1024	10,240	102,400
204	2048	20,480	204,800
		.	
922	9216	92,160	921,600

The total size of the file set should scale with the expected throughput of the server. The resulting workload can be thought of as what might be the behavior of a system supporting the "home pages" for a number of "members". There is a set of directories, one per member, with 36 files each, a complete set of nine files per each of the four classes. Requests are spread evenly across all applicable directories; the number of directories is set by the above-mentioned scaling function, which can be stated as: sqrt (throughput / 5) * 10.

Number of Directories (And Resulting Disk Space)

Based Upon the Target Throughput (ops/sec)

Ops/sec: 1	Dirs: 4	Size: 22 MB
Ops/sec: 2	Dirs: 6	Size: 31 MB
Ops/sec: 5	Dirs: 10	Size: 49 MB
Ops/sec: 10	Dirs: 14	Size: 69 MB

Ops/sec: 20	Dirs: 20	Size: 98 MB
Ops/sec: 50	Dirs: 31	Size: 154 MB
Ops/sec: 100	Dirs: 44	Size: 218 MB
Ops/sec: 200	Dirs: 63	Size: 309 MB
Ops/sec: 500	Dirs: 100	Size: 488 MB
Ops/sec: 1000	Dirs: 141	Size: 690 MB

The SPECweb96 benchmark's process structure involves several processes. The "prime" client process coordinates all of the test activity. The "client" process generates the workload on that driver. This "client" starts a configurable number of threads in the Windows NT* version which actually generates the workload, leaving the original "client" to do coordination and processing of results on that driver. The workload generation occurs within the client threads. Each of these threads generates an independent stream of HTTP requests, pausing in between requests so that on average it generates the specified number of requests per second. Each thread has a separate random deck of operation classes, initialized to match the percentages in the benchmark's defined operation mix. Each thread then works its way through the deck, using a Poisson distribution to select a file from the appropriate class, and then selecting at random a directory from which to fetch.

APPENDIX B: TESTBED CONFIGURATION

Server Hardware:
N440BX DP Server System with two Pentium® II Processors using BIOS version 1.0
Processor speeds tested: 350 & 400MHz.
Memory: 1024 MB (4x256mb Hitachi* Registered DIMMs)
IO subsystem: On board IDE and 1 AMI 434 controller BIOS version 2.44, firmware version TM76 with 4
Mb of cache
Disks: 1 - Western Digital* Caviar* 21000 IDE, 19 - 2 GB Quantum* XP 34300W fast wide SCSI disks.
Network Controller: Three Intel® EtherExpress™ PRO100+ PCI Ethernet Adapter and the on-board Intel 82558 NIC.

Client Hardware:
Quantity: Seven PCs with the Intel® 82440LX PCIset
Two Intel® Pentium® II Processor (300MHz)
Memory: 256MB
Network Controller: one Intel® EtherExpress™ PRO100B PCI Ethernet Adapter
IO Subsystem: On board IDE, 1 – Western Digital 22100

Server Software:
Operating system: NT 4.0 SP3.
IO subsystem driver: NT 4.0 ATAPI (Boot device), and AMI miniport Ver.3.51 (Target device)
Networking: Intel® EtherExpress™ PRO100+ and 82558 PRO set driver Version 3.0 (TCP/IP only)
Benchmarking Kit: SPECweb96 Version 5.0
Internet Information Server Version 4.0
NDIS, AFD and TCPIP Post SP3 drivers
(ftp://ftp.microsoft.com/bussys/winnt/winnt-public/fixes/usa/NT40/hotfixes-postSP3/)
Intbind Rev 0.2 (ftp://ftp.microsoft.com/bussys/winnt/winnt-public/tools/affinity/)
IIS Sequential write LOG fix. Available April 1998 at
ftp://ftp.microsoft.com/bussys/winnt/winnt-public/fixes/usa/NT40/hotfixes-postSP3/

Client Software:
Operating system: NT 4.0 SP3 (Server)
IO subsystem driver: ATAPI
Networking: Pro100B/82557 Version 2.51 (TCP/IP only)
Benchmarking Kit: SPECweb96* Version 5.0
Microsoft* Internet Explorer 2.0

Other Hardware:
Cisco* Catalyst* 5000 Ethernet switch

Under Control Panel, select system. Then select performance. Set Performance Boost for Foreground Application=NONE --Do not raise the priority of the foreground application, default is halfway.

Internet Service Manager: Select the web site. Then select properties, Directory Security, Edit.
Disable IIS Windows NT Challenge/Response --Disables IIS user information encryption, default is enabled.

Using regedt32 select HKEY_LOCAL_MACHINE/SYSTEM/CurrentControlSet/Services

Add Registry Entry InetInfo/Parameters/ObjectCacheTTL=0xea60 --Keeps IIS open descriptors active at least 16.6 hours, default is 10 minutes.

Add Registry Entry InetInfo/Parameters/MemoryCacheSize=0x800000 --Sets the cache open descriptor table to 128MB, default is 4MB.

Add Registry Entry InetInfo/Parameters/ListenBacklog=0x19 --Specifies maximum active connections to hold in server queue, default is 15.

Add Registry Entry InetInfo/Parameters/LogFileFlushInterval=0x3c --Sets maximum active connections to hold in the server's queue to 60. Default is 30.

Add Registry Entry InetInfo/Parameters/MaxPoolThreads=2 --Specifies maximum network request threads per processor, default is 4.

Add Registry Entry InetInfo/Parameters/PoolThreadLimit=4 --Specifies maximum pool threads in system, default is 2 * #MB.

Add Registry Entry Tcpip/Parameters/MaxFreeTcbs=0x1bd50 --Increases the TCB timewait table to 114000 entries, default is 2000.

Add Registry Entry Tcpip/Parameters/MaxHashTableSize=0x10000 --Sets TCB hash table size to 65,536 entries, default is 512.

Add Registry Entry Tcpip/Parameters/TcpTimedWaitDelay=0x3c --Sets TIME_WAIT parameter to 60 seconds (non-RFC 1122), default is 30.

Set NDIS/Parameters/ProcessorAffinityMask to 0 --Forces processor that handled interrupt to also handle that DPC, default is any processor.

Using intbind Network interrupts affinity set to two network cards to each processor. (Note: This setting only applies to the to processor results).

CLIENT OPERATING SYSTEM TUNING

Set NDIS/Parameters/ProcessorAffinityMask to 0 --Forces processor that handled interrupt to also handle that DPC, default is any processor.

Number of processes for SPECweb workload generation. 4 (Set in SPECweb96* RC script)

NETWORK ADAPTER TUNING

	Server (PRO100+)	Clients (PRO100B)
Receive buffers:	128	1024
Transmit Control Blocks:	80	16
Coalesce Buffers:	32	32

APPENDIX D: RESOURCE UTILIZATION DATA

The following performance monitor data was collected during the third SPECweb test in a series of three runs. This was to ensure that the file set was cached in memory. Perfmon data was collected during the measurement that requested the peak SPECweb result that was achieved during the actual benchmark. This ensures the data is representative of the actual benchmark, which consists of ten evenly spaced runs. This perfmon data however is slightly lower in CPU utilization in the two-processor data. The reason is the SNMP and Network Monitor Agent services were turned on to collect network data. These two services were intrusive on the network, prohibiting the clients from driving the server as hard as they could during the actual benchmark. With SNMP and Network Monitor Agent turned off , CPU reached 100% utilization during the actual benchmark.

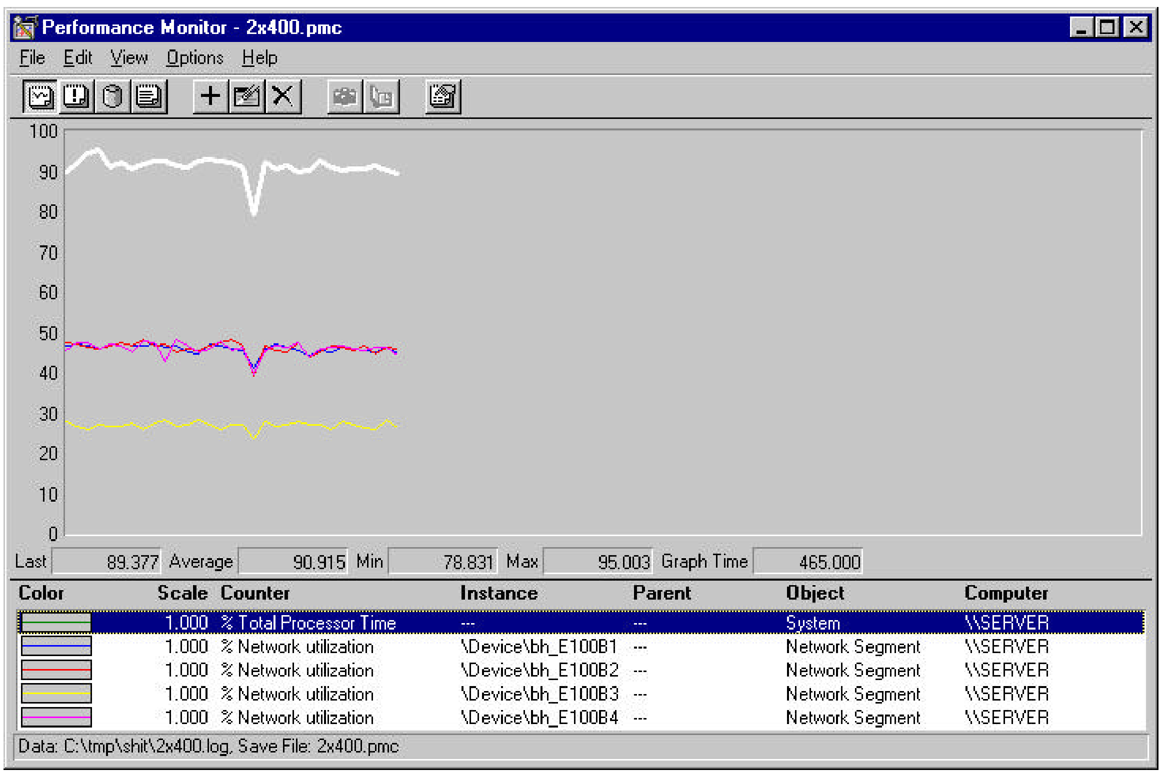

(2 CPU at 400MHz)

All Network segments are balanced with two clients each. The last (fourth) segment has one client. CPU utilization is 91% (See note above)

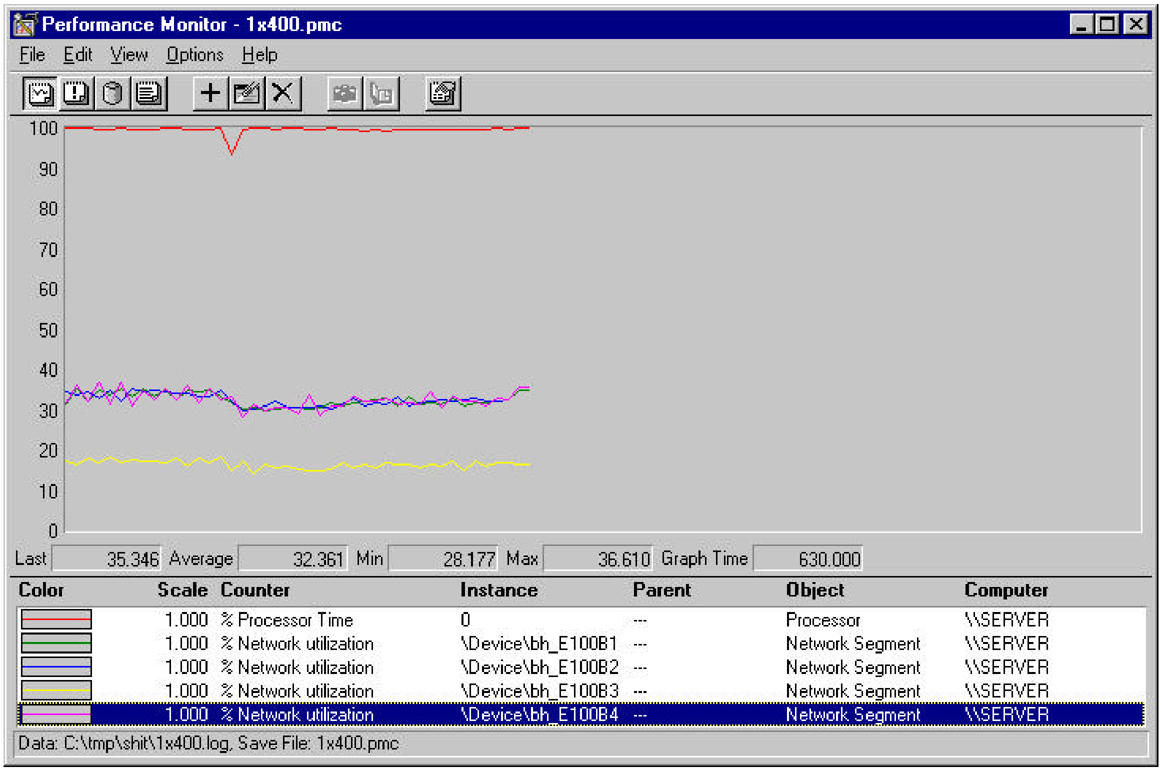

(1 CPU at 400MHz)

All Network segments are balanced with two clients each. The last (fourth) segment has one client. CPU utilization is 100%

PERFORMANCE MONITOR SUMMARY (1 AND 2 CPU AT 400MHZ)

400MHz		1 Processor	2 Processor
System statistics			
% Total Processor Time		99.4	90.9
% Total Privileged Time		85.24	80.5
% Total User Time		14.16	10.36
I/O statistics			
Avg. Disk Bytes/Transfer	Logical device D: (Data)	0	6524
	Logical device E: (IIS Log)	35518	38277
Avg. Disk sec/Transfer	Logical device D: (Data)	0ms	6ms
	Logical device E: (IIS Log)	5ms	7ms
Disk Transfers/sec	Logical device D: (Data)	0	13.01
	Logical device E: (IIS Log)	1.95	2.59
Network statistics			
% Network utilization	Device PRO100B1	32.29	45.6
	Device PRO100B2	16.3	26.69
	Device PRO100B3	32.4	45.8
	Device PRO100B4	32.36	45.6