
Performance Report

IBM Netfinity 3500

Version 1.0
February 1998



Executive Overview

Powerful IBM Netfinity 3500 servers bring advanced levels of performance and functionality down to entry server price points. New models announced worldwide in February 1998 feature 233MHz¹, 266MHz and 333MHz Pentium[®] II processors with symmetric multiprocessing (SMP) capability.

The 233MHz system (Model 8644-10U) was evaluated using NotesBench for Lotus[®] Domino[®] Server Release 4.6.

The 266MHz system (Model 8644-20U) was evaluated using the following:

- *Ziff-Davis' ServerBench[®] Version 4.0*
- *Ziff-Davis' WebBench[®] Version 1.1*
- *Ziff-Davis' NetBench[®] Version 5.01*

For comparison, measurements were also conducted using the Ziff-Davis benchmarks with a Compaq[®] ProLiant[®] 1600R system.

IBM Netfinity 3500

The Netfinity 3500 models are powerful, entry-level, enterprise LAN servers. They combine leading-edge, industry-standard technologies, good internal data storage capacity, ease of installation and system management, and solid performance. These servers provide a cost-effective, high-performance platform to support networking applications where moderate amounts of adapter expansion and internal disk storage are required or anticipated for the future. They are ideal for small business environments for handling basic to moderately complex database and business application suites or as a file server. For large to medium-size businesses, they can run similar business applications in departmental or branch office environments.

The IBM Netfinity 3500 models are high-throughput Pentium II processor-based network servers that offer excellent scalability with the addition of a second Pentium II processor of equal speed and cache size. Featured are 233MHz, 266MHz and 333MHz Pentium II processor-based models that incorporate a 512KB ECC-protected L2 cache designed for server applications. Two

¹ MHz only measures microprocessor internal clock speed, not application performance. Many factors affect application performance.

Intel Slot 1 connectors, standard on the system board, accommodate easy SMP processor upgrades. On the 233MHz model, 32MB of ECC, synchronous dynamic random access memory is standard; 64MB of SDRAM is standard on the other models. On all models, memory can be expanded to a maximum of 512MB.

The IBM Netfinity 3500 is housed in a stylish, stealth-gray, mini-tower mechanical package designed to economically handle the expansion needs of small business and entry enterprise. PCI controllers for Wide Ultra SCSI and full-duplex 100/10Mbps Ethernet are located on the system board. An advanced SVGA adapter with 4MB of video memory, installed in the Accelerated Graphics Port (AGP) slot, provides excellent video resolution.

Some of the other standard features include:

- Dual-channel Wide Ultra SCSI controller that supports up to 132MB/sec data transfers over the PCI bus
- Five available expansion slots (three 32-bit PCI slots, a PCI/ISA combination slot, and an ISA slot)
- Six drive bays: two accessible 5.25/3.5-inch half-high bays; two accessible 3.5-inch slim-high bays; and two internal 3.5-inch slim-high bays, supporting up to 22.7GB² of internal hard disk storage
- 24X-speed IDE CD-ROM and a 1.44MB diskette drive
- A 330-watt worldwide, voltage-sensing power supply with auto restart to support maximum configurations while minimizing operator intervention after a temporary power outage
- ServerGuide* CD library, which includes online publications and utilities and drivers that provide assisted loading of popular network operating systems.

All Netfinity 3500 systems feature Netfinity Manager, a comprehensive systems management tool designed with a graphical interface, which is contained on a CD-ROM for installation convenience.

² When referring to hard disk capacity, GB equals one billion bytes. Total user-accessible capacity depends on operating environment..

Performance Highlights

Following are highlights of the benchmark results. Please review the more detailed information concerning competitive results later in this report.

NotesBench 4.6

The IBM Netfinity 3500, configured with one 233MHz Pentium II processor, **demonstrated outstanding price/performance running Lotus Domino Server 4.6 on Microsoft Windows NT Server 4.0** in recently conducted measurements using the NotesBench 4.6 Mail-only workload. The following table summarizes the results.

Test Script	Maximum Users	NotesMark (tpm)	Response Time (sec)	\$/User	\$/NotesMark
Mail-Only	820	1,041	2.389	\$7.44	\$5.86

Ziff-Davis Benchmarks

For these benchmarks, the IBM Netfinity 3500 and the Compaq ProLiant 1600R were configured with the 266MHz Pentium II processor.

ServerBench 4.0

ServerBench 4.0 was used to measure the performance of the IBM Netfinity 3500 system and the Compaq ProLiant 1600R system as dual-processor application servers running Windows NT Server 4.0 and providing services to Windows NT Workstation 4.0 clients.

The peak level of throughput achieved by the IBM Netfinity 3500 was **similar** to that of the Compaq ProLiant 1600R system.

WebBench 1.1

WebBench 1.1 was used to measure the performance of the IBM Netfinity 3500 system and the Compaq ProLiant 1600R system as dual-processor Web servers running Microsoft Internet Information Server 3.0 on Windows NT Server 4.0.

Under a high-end workload of 60 WebBench clients, the IBM Netfinity 3500 system delivered nearly **10 percent more throughput** than the Compaq ProLiant 1600R.

Under the high-end workload of 60 WebBench clients, the IBM Netfinity 3500 serviced **10 percent more requests per second** than the Compaq ProLiant 1600R.

NetBench 5.01

NetBench Version 5.01 Disk Mix for Windows for Workgroup Clients was used to measure the performance of the IBM Netfinity 3500 and the Compaq ProLiant 1600R as single-processor file servers running Novell** NetWare** 4.11.

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 3500 system provided **8 percent more throughput** than the Compaq ProLiant 1600R.

Test Environments and Results

NotesBench 4.6

Results Summary

The IBM Netfinity 3500, configured with one 233MHz Pentium II processor, **demonstrated leadership price/performance running Lotus Domino Server 4.6 on Microsoft Windows NT Server 4.0** in recently conducted measurements³ using the NotesBench 4.6 Mail-only workload. The following table summarizes the results.

Test Script	Maximum Users	NotesMark (tpm)	Response Time (sec)	\$/User	\$/NotesMark
Mail-Only	820	1,041	2.389	\$7.44	\$5.86

NotesBench generates a transactions-per-minute (tpm) throughput metric, called a NotesMark, for each test, along with a value for the maximum capacity (number of users) supported, and the average response time. The price/performance results for the IBM Netfinity 3500 are derived using a total system cost of \$6,098.⁴

Competitive Results

At the time of publication, there were no competitive results available for NotesBench 4.6. NotesBench Disclosure Reports for all published results are available on the World Wide Web at:

<http://www.notesbench.org>

Measurement Methodology

The IBM Netfinity 3500 system under test was configured with one 233MHz Pentium II processor with 512KB ECC L2 cache; 320MB of memory; two 4.51GB Wide Ultra SCSI hard disks (non-RAID). For these tests, a single 100Mbps Ethernet LAN

³ The NotesBench Disclosure Report for the IBM Netfinity 3500 was audited in January 1998 by KMDS Technical Associates, Inc., and was approved for publication. Highlights from the report are presented here.

⁴ Price/performance results are based on pricing provided by an IBM Business Partner. IBM resellers set their own prices, and actual prices may vary.

segment was used. The system under test, the destination servers, and the driver systems were connected to the LAN by one Asante 100BaseT Hubs. An IBM PC 350 computer was used as the source driver (parent) system; IBM PC 350 computers were used as the client driver (child) systems. Three IBM PC Server 720 systems were used as destination servers. Destination mail addresses were distributed across these three destination servers. The IBM Netfinity 3500 system under test ran Microsoft Windows NT Server Version 4.0 with Service Pack 3 and Domino Server Release 4.6.

To ensure that the test results were reproducible, all tests were repeated, and the results were compared and found to be consistent.

Measurement Analysis

A detailed description of the benchmarked configuration and analysis of the results are provided in the NotesBench Disclosure Report for the IBM Netfinity 3500 system, which is available on the World Wide Web at: <http://www.notesbench.org>

These results are based on running the IBM Netfinity 3500 as a dedicated Domino server; the addition of other application workloads will affect the number of users supported as well as the response time. Achieving optimum performance in a customer environment is highly dependent upon selecting adequate processor power, memory and disk storage as well as balancing the configuration of that hardware and appropriately tuning the operating system and Domino software.

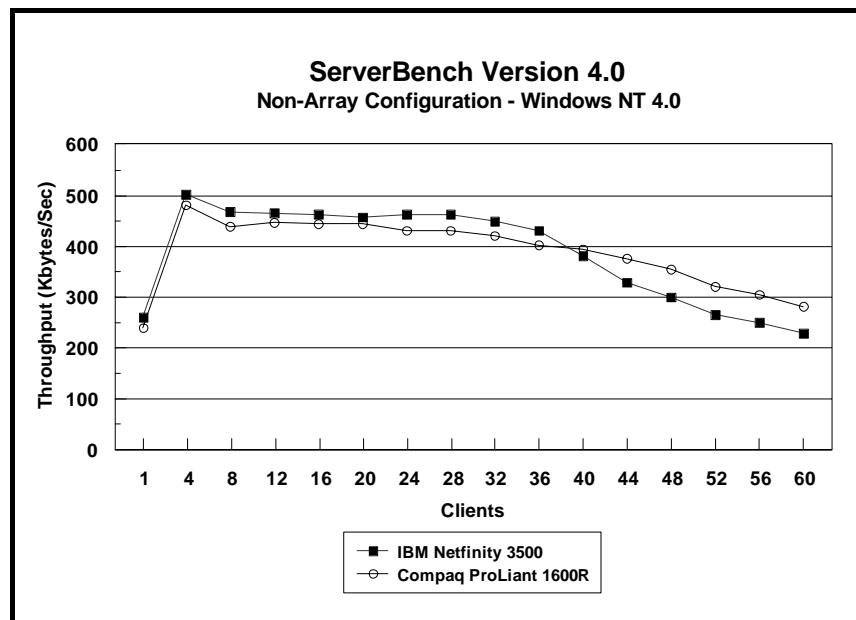
ServerBench 4.0

The ServerBench 4.0 system test suite SYS_60.TST was used to measure the performance of the IBM Netfinity 3500 and the Compaq ProLiant 1600R systems, configured as 2-way Pentium II-based application servers running Windows NT Server 4.0.

ServerBench 4.0 provides an overall transactions-per-second (TPS) score showing how well the server handles client requests for a variety of operations involving the server's processors, disk and network subsystems.

Results Summary

The peak level of throughput achieved by the IBM Netfinity 3500 was **similar** to that of the Compaq ProLiant 1600R system.



Measurement Methodology

The system test suite was performed using four 100Mbps Ethernet network segments with a total of 60 IBM PC 750 166MHz systems as client workstations attached to the server. Each workstation ran Windows NT 4.0 Workstation and executed the ServerBench 4.0 SYS_60.TST workload, which includes the client/server, processor, server/client, random read, and random write requests typically made in a client/server computing environment.

A transaction is a request issued by any one of the 60 clients; the TPS score is the number of transactions per second completed by the server under test. In the ServerBench environment, the server will not service the next request until it has finished the previous one. Thus, a higher number of transactions per second indicates better performance.

The clients randomly send requests for work to the server. These requests produce different types of loads on the server. The server performs the work by disk caching if system memory is available, or swapping mapped memory out to paged files if system memory is full.

The SYS_60.TST test suite contains a total of 16 test mixes. Measurements of transactions per second (TPS) were recorded as a weighted harmonic mean of the total TPS obtained by all clients in each test mix as clients were added. Clients were added in increments of four as follows: 1, 4, 8, 12, 16, 20, 24, 28, 32, 36, 40, 44, 48, 52, 56, 60.

Measurement Analysis

ServerBench's server application on Windows NT provides up to 47 service threads with 60 clients, plus one thread for each server processor. For this test, the servers were configured with two processors; therefore, a total of 49 service threads were used. A client workstation generates a request for the server to begin the next phase of a mix or to ask the server to perform some operation. The server creates a new service thread and passes that connection with the client to an I/O completion port.

As clients are added to the network, the I/O workload increases, requiring more service threads to be allocated to the clients. When all the service threads have been allocated, any new client requests cannot be serviced until an I/O completion port becomes available. Using four 100Mbps network adapters provided sufficient bandwidth to the application server.

ServerBench requires a large amount of system memory (e.g., 512MB) in order to produce a meaningful result. When workload increases gradually, the processor subsystem (processor and system memory) provides adequate service to all requests by caching them in the system memory, which is the primary factor affecting the TPS throughput.

As workload continued to increase (i.e., more clients joined the test mixes), system memory was exhausted, and the server had to rely on the disk subsystem for virtual memory. When this happened, the bottleneck shifted to the disk subsystem, and the

application became disk-bound. Running ServerBench with Windows NT may result in a low cache-hit ratio because some NT system threads (e.g., cache manager's lazy writer thread, memory manager's mapped page writer thread) will automatically move some mapped memory into paged files. If a client happens to request that paged-out data again, a cache-hit-miss will result.

The exact number of clients required to move the bottleneck from the processor to the disk subsystem depends on the amount of installed system memory. In our measurements, the application was processor-bound when running from 4 to 8 clients; with more than 40 clients, the application became disk-bound.

With any application, the balance of price and performance is always a key issue. Depending on the application environment, it may be well worth purchasing a second processor to alleviate a processor bottleneck. However, if the disk subsystem is the bottleneck, increasing the size of memory, rather than adding a second processor, can improve the performance.

The default values were used for all NT registry variables. The NT default is 'Max throughput for file sharing'.

WebBench 1.1


Ziff-Davis' WebBench 1.1 system test suite NT_SIMPLE_ISAPI20_V11.TST was used to measure the performance of the IBM Netfinity 3500 and the Compaq ProLiant 1600R systems, configured as 2-way 266MHz Pentium II-based systems as Web servers running Microsoft Internet Information Server 3.0 on Microsoft Windows NT Server 4.0 with Service Pack 3.

This system test suite performs both static HTML page requests and dynamic Common Gateway Interface (CGI) scripts requests, which are the two primary functions of an enterprise Web server.

Results Summary


Throughput

Under a high-end workload of 60 WebBench clients, the IBM Netfinity 3500 system delivered nearly **10 percent more throughput** than the Compaq ProLiant 1600R.



Requests per Second

Under a high-end workload of 60 clients, the IBM Netfinity 3500 serviced **10 percent more requests per second** than the Compaq ProLiant 1600R.



Measurement Methodology

The system test suite was performed using four 100Mbps Ethernet network segments with a total of 60 IBM PC 750 systems as client workstations attached to the server. Two kinds of performance data were collected:

- Static HTML pages requests, which demonstrates server throughput as each of the 60 clients, simulating an actual Web browser, fetched predesigned HTML pages using the HTTP protocol from the server. The HTML pages of different sizes (from 0.25KB to 128KB) were spread out in the home and 10 other directories (one of which has subdirectories two levels deep) in the server. The unit of this throughput measurement is bytes per second (shown here as Kbytes per second), indicating the number of bytes of HTML pages per second that were moved to the clients.
- Dynamic Common Gateway Interface requests, which demonstrates the number of requests per second completed by the server.

Each workstation ran Windows NT Workstation 4.0 and executed the WebBench 1.1 NT_SIMPLE_ISAPI20_V11.TST workload,

which includes HTML pages requests and Common Gateway Interface (CGI) requests, two of the primary functions of a web server. Each client randomly issued these requests to the web server according to a workload file that specifies each request a client makes and how frequently the client makes that request. The workload file associates a request percentage with each HTTP request and CGI request. The request percentage tells the client the number of requests it issues during a mix and what the percentage of requests should be for that particular mix. If all clients requested the same file at the same time, the results could be adversely affected; therefore, each client's request access patterns are randomized.

Clients were added incrementally to each mix as follows: 1, 4, 8, 12, 16, 20, 24, 28, 32, 36, 40, 44, 48, 52, 56, 60.

The NT_SIMPLE_ISAPI20_V11.TST test suite contains a total of 16 mixes. Each mix uses 30 seconds as ramp-up time, 30 seconds as ramp-down time, during which periods measurements were not done. Each mix ran for 300 seconds.

After the 16 mixes in the test suite were finished, WebBench created two graphs: one that plots the requests per second against each mix and another that plots the throughput against each mix. Also supplied is the amount of time how it took the clients to connect to the server and receive data from the server, and the number of connections per second made by each client.

Measurement Analysis

In a typical two-tier Internet/Intranet environment, the Web browser is usually the user front end that makes requests to the Web server. The Web server functions either as a large HTML document store directly returning the HTML documents to the browser or as a back-end logic unit building a dynamic HTML document based on calculation of input fields from the Web browser. In a three-tier Internet/Intranet environment, the Web server usually functions as middleware directing Web browser requests to the appropriate business unit (e.g., database) to retrieve information for the user.

WebBench is designed to benchmark a Web server in a two-tier Internet/Intranet environment. In calculating the scores, WebBench counts only completed requests. A completed request consists of four steps:

- The client connects to the server.
- The client issues an HTTP request (either HTML or CGI) to the server.

- The server responds to the request. This response usually results in the server sending to the client an HTML file associated with the URL specified by the client.
- The client disconnects from the server.

In a single mix, the request begins with each client connecting to the server and ends with the client disconnecting from the server, followed immediately by another repeating the process. The cycle continues until the mix is completed.

WebBench 1.1 expands its set of dynamic test suites to cover more operating systems and more types of interfaces. The new workload file used in these measurements is the Internet Server Application Program Interface (ISAPI) dynamic link library (.DLL) combined with the SIMCGI.EXE in the CGI-BIN directory on the server. Because this executable file runs on the server, it uses processing resources on the server.

To get a valid measure of the server's performance, the request-per-second score and the throughput score should reach a point where they flatten out. This "flattening out" indicates that the server has been saturated, or fully loaded. In these measurements, adding clients increased the total requests-per-second and throughput scores. The curves increased from 1 to 28 clients, peaked at 32 clients, and then flattened out, indicating that the server had reached its saturation point. Ideally, the curves after the saturation point should remain at the same level where the server's resources (e.g., CPU, memory subsystem, disk subsystem) are used optimally. However, due to heavy network traffic and the need to balance each client request load, the curve may dip slightly, reducing the server load.

NetBench 5.01


Ziff-Davis' NetBench 5.01 Disk Mix test suite was used to measure the performance of the IBM Netfinity 3500 and the Compaq ProLiant 1600R 266MHz Pentium II-based systems as single-processor file servers running Novell NetWare 4.11.

For these measurements, NetBench 5.01 Windows for Workgroup Clients was used. The results for NetBench 5.01 in a Windows environment should not be compared with results from NetBench 3.01 or 3.0, which used DOS-only clients.

The Disk Mix test results are shown as the number of kilobytes (Kbytes) per second obtained by the server under test.

Results Summary

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 3500 system provided **8 percent more throughput** than the Compaq ProLiant 1600R.



Measurement Methodology

The Disk Mix test suite was performed using two 100Mbps Ethernet network segments with a total of 60 IBM PC 350 133MHz Pentium-based systems as client workstations attached to the server. Each workstation ran Windows for Workgroups

Version 3.11 and executed the NetBench 5.01 Disk Mix workload, which is based on leading Windows applications.

Each client randomly simulated the Windows for Workgroups application workloads, accessing shared and unshared data files located on the server. Each client used a workspace of 80MB. Clients were added incrementally as follows: 4, 12, 20, 28, 36, 44, 52 and 60. Measurements were recorded each time clients were added. See “Test Disclosure Information” for testbed details.

Measurement Analysis

The NetBench Version 5.01 workload exercises the server in a manner similar to actual Windows applications executing on a networked-attached PC; that is, the NetBench Version 5.01 Disk Mix emulates the actual I/O operations performed by leading Windows applications, placing a diverse load on the server by using multiple files, different request sizes and different network file operations.

As clients are added to the network, the I/O workload (i.e., the number of I/O requests to the server) increases, requiring more server resources, such as network adapter transfers, processing power, memory and disk operations. Initially, with a small number of clients, server resources are adequate to handle requests. During this time, the server’s network adapter becomes the bottleneck.

The Disk Mix test requires each client to have its own directory and also to be able to access the shared directory in the server. As the number of clients increases, any workload involving non-shared data files creates a burden on the disk subsystem. As a result, competition for caching user data in server memory causes the bottleneck to migrate from the network adapter to the disk subsystem.

In addition, when a server’s memory buffer space is exhausted, requests are forced to go directly to the disk; therefore, the performance bottleneck quickly migrates from the network adapter to the disk subsystem, resulting in a low, disk cache-hit-ratio. Moreover, if the disk subsystem cannot quickly write “dirty” (updated) data in memory to disk, thereby freeing memory for other I/O requests, memory fills up, creating a disk backlog.

The exact number of clients required to move the bottleneck from the network adapter to the disk subsystem is dependent upon many factors. However, the most significant contributors are the I/O workload, server memory, and server disk subsystem performance. Because the Disk Mix’s I/O workload is predefined, server

memory and server disk subsystem performance contribute most to the server's disk cache-hit-ratio.

Server hardware can be configured so that the results of the NetBench Disk Mix test highlight the performance of either the server network adapter or the server disk subsystem. For example, if a large amount of memory and a fixed number of 60 simultaneous clients are used, the bottleneck will always be on the server network adapter. If too little memory is used, the bottleneck will most likely occur at the disk subsystem. The ideal measurement configuration should utilize enough memory and simultaneous clients to demonstrate the performance of the server network adapter and the server disk subsystem. This was our goal for the Disk Mix test.

In evaluating the performance results of any measurement, it is important to understand the relationship between the server configuration and the workload generated by the benchmark. We experimented with several configurations. For these servers in this configuration of 60 clients, we found that 128MB of memory optimized the throughput and also stressed the server as the workload increased. The reason is that the 100Mbps network adapter provided sufficient bandwidth to allow the server's subsystems (i.e., memory, disk and processor complex) to be saturated. This is important because in most production environments, the number of users is dynamic, and the server bottleneck may change several times daily. Showing both the network adapter and disk subsystem bottlenecks provides more useful information about how the server will perform in production environments.

Server Configurations

ServerBench 4.0

Features	IBM Netfinity 3500 266MHz/512KB	Compaq ProLiant 1600R 266MHz/512KB
Processor	Two 266MHz Pentium II	Two 266MHz Pentium II
Memory	512MB ECC Synchronous DRAM	512MB EDO ECC
L2 Cache	512KB (Write-Back)	512KB (Write-Back)
RAID Level	Non-Array	Non-Array
Disk Drive	Three IBM 4.51GB Wide Ultra SCSI Drives (7200 rpm)	Three Compaq 4.3GB Wide Ultra SCSI Drives (7200 rpm)
Disk Drive Adapter	One Wide Ultra SCSI PCI Bus on Planar	One Wide Ultra SCSI PCI Bus on Planar
Disk Driver	AIC78XX.SYS	SYMC810.SYS
Network Adapter	Four IBM 100/10 Ethernet PCI Adapters	Four Netelligent 10/100 TX PCI Ethernet Adapters
Bus	PCI	PCI
Network Driver	E100B.SYS	NETFLX3.SYS
Network Operating System	Windows NT Server 4.0 with Service Pack 3	Windows NT Server 4.0 with Service Pack 3
System Partition Size	1GB	1GB
File System	NTFS	NTFS
Allocation Unit Size	Predefined Default	Predefined Default
ServerBench Version / Test Suite	ServerBench 4.0 / SYS_60.TST	ServerBench 4.0 / SYS_60.TST

WebBench 1.1

Features	IBM Netfinity 3500 266MHz/512KB	Compaq ProLiant 1600R 266MHz/512KB
Processor	Two 266MHz Pentium II	Two 266MHz Pentium II
Memory	512MB ECC Synchronous DRAM	512MB EDO ECC
L2 Cache	512KB (Write-Back)	512KB (Write-Back)
RAID Level	Non-Array	Non-Array
Disk Drive	Three IBM 4.51GB Wide Ultra SCSI Drives (7200 rpm)	Three Compaq 4.3GB Wide Ultra SCSI Drives (7200 rpm)
Disk Drive Adapter	One Wide Ultra SCSI PCI Bus on Planar	One Wide Ultra SCSI PCI Bus on Planar
Disk Driver	AIC78XX.SYS	SYMC810.SYS
Network Adapter	Four IBM 100/10 Ethernet PCI Adapters	Four Netelligent 10/100 TX PCI Ethernet Adapters
Bus	PCI	PCI
Network Driver	E100B.SYS	NETFLX3.SYS
Network Operating System	Windows NT Server 4.0 with Service Pack 3	Windows NT Server 4.0 with Service Pack 3
System Partition Size	1GB	1GB
File Compression	Off	Off
File System	NTFS	NTFS
Allocation Unit Size	Predefined Default	Predefined Default
WebBench Version / Test Suite	WebBench 1.1 / NT_SIMPLE_ISAPI20_V11.TST	WebBench 1.1 / NT_SIMPLE_ISAPI20_V11.TST
Web Server	Microsoft Internet Information Server 3.0	Microsoft Internet Information Server 3.0

NetBench 5.01

Features	IBM Netfinity 3500 266MHz/512KB	Compaq ProLiant 1600R 266MHz/512KB
Processor	One 266MHz Pentium II	One 266MHz Pentium II
Memory	128MB ECC Synchronous DRAM	128MB EDO ECC
L2 Cache	512KB (Write-Back)	512KB (Write-Back)
RAID Level	Non-Array	Non-Array
Disk Drives	Three IBM 4.51GB Wide Ultra SCSI Drives (7200 rpm)	Three 4.3GB Wide Ultra SCSI Drives (7200 rpm)
Disk Drive Adapter	One Wide Ultra SCSI PCI Bus on Planar	One Wide Ultra SCSI PCI Bus on Planar
Disk Driver	AIC7870.DSK V2.13	CPQS710.DSK V2.11
Network Adapters	Two IBM 100/10 Ethernet PCI Adapters	Two Netelligent PCI 100/10 Ethernet Adapters
Bus	PCI	PCI
Network Driver	E100S.LAN V2.93	CPQNF3.LAN V2.23
Network Operating System	NetWare 4.11 with IWSP4 loaded	NetWare 4.11 with IWSP4 loaded
NetWare Volume Block Size	32KB	32KB
File Compression	Off	Off
Block Allocation	On	On
Data Migration	Off	Off
Disk Mix	NB5.01/Windows for Workgroup Clients	NB5.01/Windows for Workgroup Clients

Test Disclosure Information

ServerBench 4.0

The ServerBench measurements were conducted using Ziff-Davis' ServerBench 4.0 running the SYT_60.TST test suite with Windows NT Workstation 4.0 as described below:

Version: ServerBench 4.0

Mixes

- System Test Mixes
- Clients: 1, 4, 8, 12, 16, 20, 24, 28, 32, 36, 40, 44, 48, 52, 56, 60
- Disk Test Data File: 20MB
- Disk Test File I/O Range: 10MB
- Ramp up: Default setup
- Ramp down: Default setup

Network Operating System: Windows NT Server 4.0 with Service Pack 3

Testbed Disclosure

All products used for these measurements are shipping versions available to the general public. All measurements were performed without independent verification by Ziff-Davis.

Network	100Mbps Ethernet
Clients	60
Hubs	Asante 100Mbps Ethernet
Clients per Segment	15
CPU / Memory	166MHz Pentium / 64MB
Network Adapter	IBM 100/10 PCI Ethernet Adapter (Bus 0)
Software	Windows NT 4.0 Workstation
Cache	L2 = 512KB
Controller Software	Microsoft Windows NT Workstation 4.0

WebBench 1.1

The measurements were conducted using Ziff-Davis' WebBench 1.1 running the NT_SIMPLE_ISAPI20_V11.TST test suite with Windows NT Workstation 4.0 clients as described below:

Version: WebBench 1.1

Mixes

- NT_SIMPLE_ISAPI20_V11.TST
- Clients 1, 4, 8, 12, 16, 20, 24, 28, 32, 36, 40, 44, 48, 52, 56, 60
- Ramp up: 30 seconds
- Ramp down: 30 seconds
- Length: 300 seconds
- Think: 0 seconds
- Delay: 0 seconds
- Threads per client: 1
- Receive buffer size: 4KB
- Keep-alive: Off
- Web Server: Logon Disabled

Network Operating System: Windows NT Server 4.0 with Service Pack 3

Web Server: Microsoft Internet Information Server 3.0

Testbed Disclosure

All products used for these measurements are shipping versions available to the general public. All measurements were performed without independent verification by Ziff-Davis.

Network	100Mbps Ethernet
Clients	60
Hubs	Asante 100Mbps Ethernet
Clients per Segment	15
CPU / Memory	166MHz Pentium / 32MB
Network Adapter	IBM 100/10 PCI Ethernet Adapter (Bus 0)
Software	Microsoft Windows NT Workstation 4.0
Cache	L2 = 512KB
Controller Software	Windows NT Workstation 4.0

**NetBench Version
5.01**

The NetBench measurements were conducted using Ziff-Davis' NetBench 5.01 running the Disk Mix with Windows for Workgroup Clients as described below:

Version: NetBench 5.01

Mixes

- Disk Mix
- Clients 4, 12, 20, 28, 36, 44, 52, 60
- Client workspace: 80MB
- Total runtime: 11 minutes
- Ramp up and down: 30 seconds

Network Operating System: NetWare 4.11 with IWSP4 loaded

NOS Parameters

- Immediate Purge of Deleted Files = ON
- Enable Disk Read after Write Verify = OFF
- Minimum Packet Receive Buffers = 700
- Maximum Packet Receive Buffers = 1400
- Set NCP Packet Signature Option = 0
- Maximum Physical Receive Packet Size = 1514
- Reserved Buffer Below 16MB = 200
- Maximum Service Processes = 70
- Maximum Concurrent Directory Cache Write = 100
- Dirty Directory Cache Delay Time = 10
- Maximum Concurrent Disk Cache Write = 100
- Maximum Directory Cache Buffers = 700
- Minimum Directory Cache Buffers = 150
- Minimum File Cache Buffers = 150
- Maximum Number of Directory Handles = 30
- Dirty Disk Cache Delay Time = 5
- Directory Cache Buffer Non-Referenced Delay = 30
- Directory Cache Allocation Wait Time = 2.2 seconds

If clients drop out, set the following:

- Number of Watchdog Packets = 50
- Delay Between Watchdog Packets = 10 minutes
- Delay Before First Watchdog Packet = 20 minutes

To monitor the dropping out of clients, set:

- Console Display Watchdog Logouts = On

Testbed Disclosure

All products used for these measurements are shipping versions available to the general public. All measurements were performed without independent verification by Ziff-Davis.

Network	100Mbps Ethernet
Clients	60
Hubs	Asante 100Mbps Ethernet
Clients per Segment	30
CPU / Memory	133MHz Pentium / 16MB
Network Adapter	IBM 100/10 PCI Ethernet Adapter (Bus 0)
Software	IBM DOS 6.3
	NetWare DOS Requester
	LSL.COM (8-3-95)
	E100BODI (5-21-96)
	IPXODI (8-8-95)
	VLM.EXE (11-8-94)
Cache	L2 = 256KB
Controller Software	PC-DOS Version 6.3 Microsoft Windows for Workgroups 3.11

Clients NET.CFG

- Checksum = 0
- Large Internet Packet = On
- PB Buffers = 10
- PBurst Read Windows Size = 64
- PBurst Write Windows Size = 64
- Cache Buffers = 64

THE INFORMATION CONTAINED IN THIS DOCUMENT IS DISTRIBUTED ON AN AS IS BASIS WITHOUT ANY WARRANTY EITHER EXPRESS OR IMPLIED. The use of this information or the implementation of any of these techniques is the customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item has been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environment do so at their own risk.

This publication was produced in the United States. IBM may not offer the products, services, or features discussed in this document in other countries, and the information is subject to change without notice. Consult your local IBM representative for information on products and services available in your area.

*IBM, ServerGuide and Netfinity are trademarks or registered trademarks of International Business Machines Corporation.

**Lotus, Lotus Notes and Domino are trademarks of Lotus Development Corporation.

**Intel and Pentium are registered trademarks of Microsoft Corporation.

**Microsoft and Windows are trademarks or registered trademarks of Microsoft Corporation.

Other company, product, or service names, which may be denoted by two asterisks (**), may be trademarks or service marks of others.

Published by the IBM Server Performance Laboratory, IBM Corp.

© Copyright International Business Machines Corporation 1998. All rights reserved.

Permission is granted to reproduce this document in whole or in part, provided the copyright notice as printed above is set forth in full text at the beginning or end of each reproduced document or portion thereof.

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

