

---

# Performance Report

---

## **IBM Netfinity 7000 M10** with 450MHz Pentium II Xeon and 2MB L2 Cache

Version 1.0  
February 1999



---

## Executive Overview

---

*The performance of the IBM Netfinity\* 7000 M10 server (with 450MHz Pentium\*\* II Xeon\*\* processor with 2MB L2 cache), announced worldwide on January 19, 1999, was evaluated using the following benchmarks:*

- *TPC Benchmark C\*\**
- *TPC Benchmark D\*\**
- *SPECweb96\*\**
- *Ziff-Davis' ServerBench\*\* Version 4.02*
- *Ziff-Davis' NetBench\*\* Version 5.01*

*For the Ziff-Davis benchmarks, a similarly configured Netfinity 7000 M10 server using the 400MHz Pentium II Xeon processor with 1MB L2 cache was measured for comparison with the 450MHz / 2MB model.*

*The results of all the benchmarks mentioned are presented in this report.*

---

### Performance Highlights

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

#### **TPC Benchmark C**

Running Microsoft\*\* Windows NT\*\* 4.0 Enterprise Edition and Microsoft SQL Server 7.0 Enterprise Edition, the IBM Netfinity 7000 M10 **delivered a transaction processing performance result of 22,459.80 tpmC** at a price/performance ratio of \$25.70/tpmC. Total solution availability is January 19, 1999. (Results current as of February 9, 1999.)

#### **TPC Benchmark D**

Running Microsoft Windows NT Server 4.0 and IBM DB2 Universal Database Enterprise Edition 5.2.0, the IBM Netfinity 7000 M10 **delivered leadership performance results of 4,977.9 QppD@100GB and 747.8 QthD@100GB at a price/performance of \$123 per QphD@100GB.** Total solution

availability is June 30, 1999. (Results current as of February 9, 1999.)

Running IBM DB2 Universal Database Enterprise - Extended Edition 5.2.0 and Microsoft Windows NT Server 4.0, the Netfinity 7000 M10 / 450MHz **delivered leadership performance results of 36,872.0 QppD@1000GB and 8,166.9 QthD@1000GB at a price/performance of \$352 per QphD@1000GB.** Total solution availability is June 30, 1999. (Results current as of February 9, 1999.)

---

## SPECweb96 Benchmark

Continuing to set the pace for SPECweb96 performance on Windows NT Server 4.0 in four-way, two-way and one-way configurations, the IBM Netfinity 7000 M10 achieved:

- 7,818 operations per second (four-way)
- 5,996 operations per second (two-way)
- 4,579 operations per second (one-way).

All configurations used the IBM HTTP Server Version 1.3.3, with an integrated Web server accelerator for improved performance.

---

## Ziff-Davis Benchmarks

### ServerBench 4.02

ServerBench 4.02 was used to measure the performance of the IBM Netfinity 7000 M10 / 450MHz and Netfinity 7000 M10 / 400MHz systems as four-way application servers running Windows NT Server 4.0.

The peak level of throughput achieved by the IBM Netfinity 7000 M10 / 450MHz was **28 percent higher** than that of the Netfinity 7000 / 400MHz for RAID-0 and RAID-5.

### NetBench 5.01

NetBench 5.01 Disk Mix Windows for Windows NT Workstation 4.0 clients was used to measure the performance of the IBM Netfinity 7000 M10 / 450MHz and Netfinity 7000 M10 / 400MHz systems as single-processor file servers running Novell\*\* NetWare\*\* 4.11 with IWSP6 loaded.

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 7000 M10 provided network clients with **9 percent more throughput** than the Netfinity 7000 M10 / 400MHz for RAID-0 and **10 percent more throughput** than the Netfinity 7000 M10 / 400MHz for RAID-5.

---

## Test Environments and Results

---

### TPC Benchmark C

---

#### Results Summary

Running Microsoft\*\* Windows NT Server 4.0 Enterprise Edition and Microsoft SQL Server 7.0 Enterprise Edition, the IBM Netfinity 7000 M10 **delivered a transaction processing performance result of 22,459.80 tpmC at a price/performance ratio of \$25.70/tpmC**. Total solution availability is January 19, 1999. (Results current as of February 9, 1999.)

Tested Configuration	tpmC	\$ / tpmC
IBM Netfinity 7000 M10 C/S	22,459.80	\$25.70 / tpmC
Database	Microsoft SQL Server 7.0 Enterprise Edition	
Operating System	Microsoft Windows NT Server 4.0 Enterprise Edition	
TP Monitor	IBM TXSeries V4.2 for Windows NT	
Availability Date	January 29, 1999	
Configuration	Client / Server	
5-Year Cost of Ownership	\$577,117	
Benchmark Revision	Rev. 3.4	

The priced configuration consisted of one IBM Netfinity 7000 M10 server with four 450MHz Pentium II Xeon processors (each with 2MB L2 cache) and 4GB<sup>2</sup> of memory. The configuration included twenty-two IBM Netfinity EXP15 Rack Storage Expansion Enclosures, which housed a total of 1,992.9GB of storage to accommodate database growth over a 180-day period. Four IBM Netfinity 5000 servers, each configured with two 450MHz Pentium II processors and 384MB of memory, were used as clients connected to the Netfinity 7000 M10 via a 100BaseT Ethernet LAN. The four clients serviced the requests from a total of 18,360 emulated users.

## About TPC Benchmark C

Based on an order-entry workload, the TPC-C benchmark exercises the database components used to perform a wide range of tasks of varying complexity associated with an online transaction processing (OLTP) application. The benchmark application environment is that of a wholesale supplier with many distributors. The workload consists of order-entry processing and related transactions. The metrics used are throughput as measured by the number of New-Order transactions completed per minute (tpmC) and the associated system cost per tpmC, reported as \$/tpmC, and the total solution availability date.

The TPC-C performance metrics are tpmC and \$/tpmC. The throughput metric, tpmC, represents the average number of New-Order transactions completed per minute during the measurement interval. The price/performance metric, \$/tpmC, represents the cost of each tpmC and is determined by dividing the cost of the priced configuration by the average tpmC achieved on that system during the measurement interval. The price/performance metric encourages benchmark sponsors to test realistic system environments, and it enables readers of TPC-C results to make valid comparisons between published results. The two TPC-C metrics should be considered together.

The TPC-C specification sets forth specific guidelines to ensure that this benchmark successfully emulates a real-world application. Other factors besides the hardware platform influence the TPC-C results. These factors include the database software and operating system used, as well as the benchmark sponsor's tuning expertise. The reader should keep these factors in mind when comparing TPC-C results reported by different benchmark sponsors.

For a complete description of the benchmark and executive summaries for competitors' results, visit the TPC Web site at <http://www.tpc.org>. The TPC-C Executive Summary for the IBM Netfinity 7000 M10 is also available at IBM's Netfinity Web site at <http://www.pc.ibm.com/us/techlink/srvperf.html>.

---

## TPC Benchmark D

---

### Results Summary

#### Cluster Configuration with 32 Nodes

Running IBM DB2 Universal Database Enterprise - Extended Edition 5.2.0 and Microsoft\*\* Windows NT Server 4.0, the Netfinity 7000 M10 **delivered leadership performance results of 36,872.0 QppD@1000GB and 8,166.9 QthD@1000GB at a price/performance of \$352 per QphD@1000GB.** Total solution availability is June 30, 1999. (Results current as of February 9, 1999.)

Tested Configuration	QppD @ 1000GB	QthD @ 1000GB	\$ per QphD @ 1000GB
IBM Netfinity 7000 M10	36,872	8,166.9	\$352
Database	DB2 Universal Database Enterprise - Extended Edition 5.2.0		
Operating System	Microsoft Windows NT Server 4.0		
Availability Date	Netfinity 7000 M10: Currently Available Windows NT Server 4.0: Currently Available DB2 UDB 5.2.0: June 30, 1999 Total Solution Availability Date: June 30, 1999		
5-Year Cost of Ownership	\$6,104,897		
TPC-D Benchmark Rev.	Rev. 1.3.1		

The performance breakthroughs, which were verified by Transaction Processing Performance Council (TPC\*\*) auditors on January 11, 1999, ranked the 32-node Netfinity 7000 M10 cluster and DB2 Universal Database as the performance and price/performance leader among all hardware and database vendors for 1000GB TPC-D results.

The priced configuration consisted of 32 IBM Netfinity 7000 M10 servers. Each of the 32 nodes contained the following components:

- Four Pentium II Xeon 450MHz processors with 2MB L2 cache each
- 2GB of memory
- One internal SCSI 4.51GB disk drive
- Two IBM Serial Storage Architecture (SSA) disk controllers attached to two SSA disk enclosures, each containing 14 SSA 9.1GB disk drives (plus two inactive spare drives)
- One GigaNet\*\* network interface adapter
- One Ethernet adapter
- Windows NT Server 4.0

- DB2 UDB EEE 5.2.0

Total storage was 8298GB. The network interconnect consisted of 32 GigaNet PCI network interface adapters and 12 GigaNet cLAN\*\* switches.

An executive summary of the TPC-D report can be obtained on the Internet from IBM Netfinity at

<http://www.pc.ibm.com/us/techlink/perform.html>.

### Single-Node, Single-User Configuration

Running IBM DB2 Universal Database Enterprise Edition 5.2.0 and Microsoft\*\* Windows NT Enterprise Edition 4.0, the Netfinity 7000 M10 **delivered leadership performance results of 4,977.9 QppD@100GB and 747.8 QthD@100GB at a price/performance of \$123 per QphD@100GB.** Total solution availability is June 30, 1999. (Results current as of February 9, 1999.)

Tested Configuration	QppD @ 100GB	QthD @ 100GB	\$ per QphD @ 100GB
IBM Netfinity 7000 M10	4,977.9	747.8	\$123
Database	DB2 Universal Database Enterprise Edition 5.2.0		
Operating System	Microsoft Windows NT Server Enterprise Edition 4.0		
Availability Date	Netfinity 7000 M10: Currently Available Windows NT Server 4.0: Currently Available DB2 UDB 5.2.0: June 30, 1999 Total Solution Availability Date: June 30, 1999		
5-Year Cost of Ownership	\$238,250		
TPC-D Benchmark Rev.	Rev. 1.3.1		

The priced configuration consisted of one IBM Netfinity 7000 M10 server with four 450MHz Pentium II Xeon processors (each with 2MB L2 cache) and 2.5GB of memory. Nine Netfinity EXP15 Rack Storage Expansion Enclosures housed 837.04GB of storage.

### About TPC Benchmark D

TPC Benchmark D is a Decision Support benchmark. It is a suite of business-oriented queries and concurrent updates. The performance metrics reported by the TPC-D measure multiple aspects of the capability of the system to process queries. These aspects include the selected database size against which the queries are executed, the TPC-D query processing power at the selected size (QppD@Size), and the TPC-D throughput at the selected size (QthD@Size) when queries are submitted by one or more concurrent users. The TPC-D Price/Performance metric is

expressed as  $\$/QphD@Size$  and is based on a composite query-per-hour rating derived from the QppD and QthD. To be compliant with the TPC-D standard, all references to TPC-D results for a given configuration must include all required reporting components.

For a complete description of the benchmark and executive summaries for competitors' results, visit the TPC Web site at <http://www.tpc.org>. The TPC-D Executive Summary for the IBM Netfinity 7000 M10 is also available at IBM's Netfinity Web site at <http://www.pc.ibm.com/us/techlink/srvperf.html>.



---

## SPECweb96 Benchmark

---

### Results Summary

The IBM Netfinity 7000 M10 server continues to deliver industry-leading SPECweb96 results on Windows NT Server 4.0 in four-way, two-way and one-way Intel-based configurations.

The IBM Netfinity 7000 M10, configured with four 450MHz Pentium II Xeon processors with 2MB of L2 cache and 2GB of memory, **achieved 7,818 operations per second with a response time of 12.5 milliseconds per operation.**

The IBM Netfinity 7000 M10, configured with two 450MHz Pentium II Xeon processors with 2MB of L2 cache and 2GB of memory, **achieved 5,996 operations per second with a response time of 15.8 milliseconds per operation.**

The IBM Netfinity 7000 M10, configured with one 450MHz Pentium II Xeon processor with 2MB of L2 cache and 2GB of memory, **achieved 4,579 operations per second with a response time of 20.7 milliseconds per operation.**

These latest Netfinity performance milestones were achieved using IBM HTTP Server Version 1.3.3 and Netfinity Web server acceleration on NT, and Alteon Networks' ACEnic Gigabit Ethernet Adapter with Jumbo Frame support, which improves bulk data transfer performance and minimizes packet-processing overhead on servers. Also used was the ACEswitch 180, a per-port-selectable 10/100/1000 Mbps switch.

SPECweb96, with its standardized workload and implementation, measures a system's ability to perform as a World Wide Web server for static pages. The workload simulates the accesses to a Web service provider, where the server supports multiple pages for a number of different organizations. This benchmark is useful in evaluating systems that handle millions of hits per day and multiple hits per second. SPECweb96 provides the most objective, most representative benchmarks for measuring Web server performance.

SPECweb96 reports for results cited in this article, as well as all other published results, are available on the World Wide Web at **<http://www.specbench.org/osg/web96>**.

More information about these results is available at IBM's Netfinity Web site at **<http://www.pc.ibm.com/us/news/netfinity/index.html>**.

---

## ServerBench 4.02

Ziff-Davis' ServerBench 4.02 test suite SYS\_60.TST was used to measure the performance of the IBM Netfinity 7000 M10 / 450MHz and the Netfinity 7000 M10 / 400MHz systems as four-way application servers running Windows NT Server 4.0.


ServerBench 4.02 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

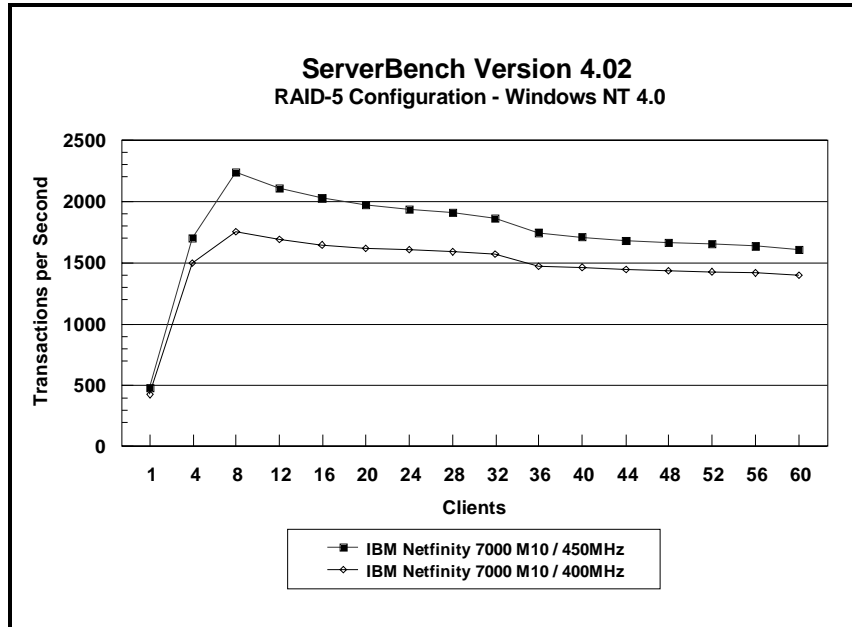
### RAID-0 Configuration

The Netfinity 7000 M10 / 450MHz achieved a peak level of throughput at eight clients that was **28 percent higher** than the Netfinity 7000 M10 / 400MHz.



## RAID-5 Configuration

The Netfinity 7000 M10 / 450MHz achieved a peak level of throughput at eight clients that was **28 percent higher** than the Netfinity 7000 M10 / 400MHz.



---

## 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 Workstation 4.0 and executed the ServerBench 4.02 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. (The default values were used for all NT registry variables. The NT default is 'Max throughput for file sharing'.)

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 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 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 incrementally 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 four processors; therefore, a total of 51 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 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.

---

## NetBench 5.01

The NetBench 5.01 Disk Mix was used to measure the performance of the IBM Netfinity 7000 M10 / 450MHz and the Netfinity 7000 M10 / 400MHz systems as single-processor file servers running Novell NetWare 4.11 with Service Pack IWSP6. For these measurements, NetBench 5.01 Windows NT Workstation 4.0 clients were used.

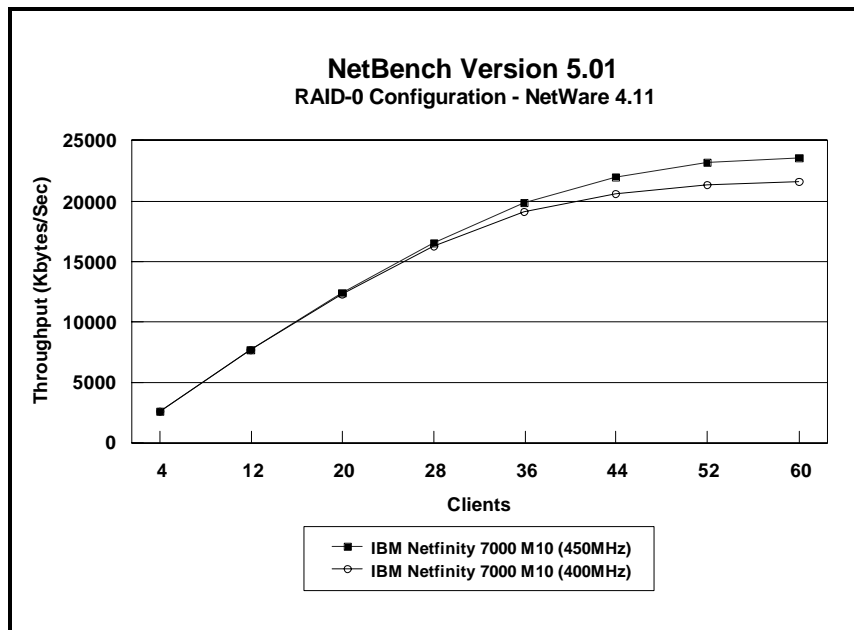
The Disk Mix test results are shown as the number of kilobytes (Kbytes) per second obtained by the server under test.

---

## Results Summary


### RAID-0 Configuration

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 7000 M10 / 450MHz server provided network clients with **9 percent more throughput** than the Netfinity 7000 M10 / 400MHz server.



## RAID-5 Configuration

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 7000 M10 / 450MHz provided network clients with **10 percent more throughput** than the Netfinity 7000 M10 / 400MHz.



---

### Measurement Methodology

The Disk Mix test suite was performed using four 100Mbps Ethernet network segments with a total of 60 IBM PC 350 200MHz Pentium-based systems as client workstations attached to the server. Each workstation ran Windows NT Workstation 4.0 and executed the NetBench 5.01 Disk Mix workload, which is based on leading Windows applications.

Each client randomly simulated the Windows 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, 60. Measurements were recorded each time clients were added.

---

### Measurement Analysis

The NetBench 5.01 workload exercises the server in a manner similar to actual Windows applications executing on a networked-attached PC; that is, the NetBench 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, using

1GB of memory, in this configuration of 60 clients, the cache-hit ratio was more than 95 percent, and the cache buffer in the resource utilization was more than 90 percent. These statistics indicate a higher level of throughput. Also, 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 4.02

Features	IBM Netfinity 7000 M10 / 450MHz	IBM Netfinity 7000 M10 / 400MHz
Processor	Four 450MHz Pentium II Xeon	Four 400MHz Pentium II Xeon
Memory	1GB ECC	1GB ECC
L2 Cache	2MB ECC (Write-Back)	1MB ECC (Write-Back)
RAID Level	RAID-0 and RAID-5	RAID-0 and RAID-5
Disk Drives	Nine 9.1GB 10K Wide Ultra SCSI Drives	Nine 9.1GB 10K Wide Ultra SCSI Drives
Disk Drive Adapter	One ServeRAID-3H Ultra2 SCSI Adapter	One ServeRAID-3H Ultra2 SCSI Adapter
Disk Driver	ISPRAIDN.SYS	ISPRAIDN.SYS
Network Adapters	Four IBM EtherJet 100/10 PCI Adapters	Four IBM EtherJet 100/10 PCI Adapters
Bus	PCI	PCI
Network Driver	IBMFENT.SYS	IBMFENT.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
Test Suite	ServerBench 4.02 SYS_60.TST	ServerBench 4.02 SYS_60.TST

---

## NetBench 5.01

Features	IBM Netfinity 7000 M10 / 450MHz	IBM Netfinity 7000 M10 / 400MHz
Processor	One 450MHz Pentium II Xeon	One 400MHz Pentium II Xeon
Memory	1GB ECC	1GB ECC
L2 Cache	2MB (Write-Back)	1MB (Write-Back)
RAID Level	RAID-0 and RAID-5	RAID-0 and RAID-5
Disk Drives	Nine 9.1GB 10K Wide Ultra SCSI Drives	Nine 9.1GB 10K Wide Ultra SCSI Drives
Disk Drive Adapter	One ServeRAID-3H Ultra2 SCSI Adapter	One ServeRAID-3H Ultra2 SCSI Adapter
Disk Driver	ISPRAID.HAM V2.81	ISPRAID.HAM V2.81
Network Adapters	Four IBM EtherJet 100/10 PCI Adapters	Four IBM EtherJet 100/10 PCI Adapters
Bus	PCI	PCI
Network Driver	E100B.LAN V3.63	E100B.LAN V3.63
Network Operating System	NetWare 4.11 with IWSP6 Loaded	NetWare 4.11 with IWSP6 Loaded
NetWare Volume Block Size	16KB	16KB
File Compression	Off	Off
Block Allocation	On	On
Data Migration	Off	Off
Disk Mix	NB5.01/Windows NT Workstation 4.0	NB5.01/Windows NT Workstation 4.0

---

# Test Disclosure Information

---

## ServerBench 4.02

The measurements were conducted using Ziff-Davis' ServerBench 4.02 running the Windows NT Workstation 4.0 clients as described below:

**Version:** 4.02

### Mixes

- System Test Mixes
- Clients: 1, 4, 8, 12, 16, 20, 24, 28, 32, 36, 40, 44, 48, 52, 60
- Data Segment Size: 16MB
- Segment Access Ratio: 1
- Ramp up: Default setup
- Ramp down: Default setup

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

---

## 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.

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

---

## NetBench 5.01

The NetBench measurements were conducted using Ziff-Davis' NetBench 5.01 running the Disk Mix Windows NT Workstation 4.0 clients as described below:

**Version:** 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 IWSP6 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 Package Size = 1514
- Reserved Buffer Below 16MEG = 200
- Maximum Service Processes = 70
- Maximum Concurrent Directory Cache Write = 200
- Dirty Directory Cache Delay Time = 3
- Maximum Concurrent Disk Cache Write = 1000
- Maximum Directory Cache Buffers = 500
- Minimum Directory Cache Buffers = 150
- Minimum File Cache Buffers = 300
- Maximum Number of Directory Handles = 30
- Dirty Disk Cache Delay Time = 5
- Directory Cache Allocation Wait Time = 2.2 seconds
- Directory Cache Buffer Non-Referenced Delay = 15
- Maximum Interrupt Events = 1000
- Set Enable File Compression = Off

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.

<b>Network</b>	100Mbps Ethernet
<b>Clients</b>	60
<b>Hubs</b>	3COM Super Stack II Hub 100TX
<b>Clients per Segment</b>	15
<b>CPU</b>	200MHz Pentium
<b>Network Adapter</b>	IBM 100/10 PCI Ethernet Adapter (Bus 0)
<b>Software</b>	Windows NT Workstation 4.0
	Network Service - Client Service for NetWare
	Network Protocol - NWLink IPX/SPX compatible Teansport - NWLink NetBIOS - TCP/IP Protocol
<b>Cache</b>	L2 = 256KB
<b>Controller Software</b>	Microsoft Windows 95 Novell Loginw 95

---

THE INFORMATION CONTAINED IN THIS DOCUMENT IS DISTRIBUTED ON AN AS IS BASIS WITHOUT ANY WARRANTY EITHER EXPRESSED 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 and Netfinity are trademarks or registered trademarks of International Business Machines Corporation.

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

\*\*Intel, Pentium and Xeon are trademarks or registered trademarks of Intel Corporation.

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

\*\*SPECweb96 is a trademark of the Standard Performance Evaluation Corporation.

\*\*TPC, tpmC, \$/tpmC, TPC-D, QppD, QthD and QphD are trademarks of the Transaction Processing Performance Council.

\*\*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 Netfinity Server Performance Laboratory, IBM Corp.

© Copyright International Business Machines Corporation 1999 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.

#### **Notes**

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

<sup>2</sup> When referring to hard disk capacity, GB, or gigabyte, means one thousand million bytes; total user-accessible capacity may vary.