

---

# Performance Report

---

## IBM Netfinity 5500 M20 (500MHz)

Version 1.1  
April 1999



---

## Executive Overview

---

*The new IBM Netfinity\* 5500 M20 servers, announced worldwide in March 1999, advance the popular Netfinity 5500 line with affordable, four-way SMP processing power and scalability to enterprise network computing. The new 5500 M20 server features the 500MHz Intel\*\* Pentium\*\* III Xeon\*\* processor with up to 1MB L2 cache.*

*The new server (Model 8662-41Y) was evaluated using the following benchmarks:*

- *SAP\*\* R/3 Two-Tier Sales and Distribution (SD) Benchmark*
- *NotesBench for Domino\*\* R5*
- *SPECweb96\*\* Benchmark*
- *Ziff Davis' ServerBench\*\* Version 4.02 and NetBench\*\* Version 5.01*

*The results of all the benchmarks mentioned are presented in this report. This report replaces Version 1.0 from March 1999.*

---

### Performance Highlights

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

#### **SAP R/3 Two-Tier Sales and Distribution Benchmark**

The IBM Netfinity 5500 M20, configured with four 500MHz Pentium III Xeon processor and running the SAP R/3 Two-Tier SD Benchmark, Release 4.0 B, **achieved 200 SD Benchmark users** on Oracle 8.0.4 and Microsoft\*\* Windows NT\*\* Server 4.0 Enterprise Edition. This is the first published result using the 1MB L2 cache version of the new 500MHz Pentium III Xeon processor.

#### **NotesBench**

The IBM Netfinity 5500 M20, configured with two 500MHz Pentium III Xeon processors, **demonstrated leadership performance by supporting 8,250 concurrent active Mail users at a price/performance of \$5.51 per user**, running Lotus Domino R5 Server on Microsoft Windows NT Server 4.0 in recently conducted measurements using the NotesBench 164.1 R4 Mail-only workload.

### **SPECweb96 Benchmark**

The IBM Netfinity 5500 M20, configured with four 500MHz Pentium III Xeon processors and 1GB of memory, **achieved 7,796 operations per second**, running IBM HTTP Server Version 1.3.3, which integrates IBM's patented Cache Accelerator with Apache Web server technology. In two-way and one-way configurations, the Netfinity 5500 M20 achieved peak results of 6,684 and 5,958 operations per second, respectively.

### **ServerBench 4.02**

ServerBench 4.02 was used to measure the performance of the IBM Netfinity 5500 M20 500MHz and the Compaq ProLiant 5500 450MHz systems as four-way application servers running Windows NT Server 4.0.

In a RAID-0 configuration, the Netfinity 5500 M20 500MHz achieved a peak level of transactions per second at eight clients that was **13 percent higher** than that of the Compaq ProLiant 5500 450MHz system.

In a RAID-5 configuration, the Netfinity 5500 M20 500MHz achieved a peak level of transactions per second at eight clients that was **14 percent higher** than that of the Compaq ProLiant 5500 450MHz system.

### **NetBench 5.01**

NetBench 5.01 Disk Mix for Windows NT Workstation 4.0 clients was used to measure the performance of the IBM Netfinity 5500 M20 500MHz and the Compaq ProLiant 5500 450MHz systems as single-processor file servers running Novell\*\* NetWare\*\* 4.11.

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 5500 M20 500MHz system provided network clients with a level of throughput that was **49 percent higher** than that of the Compaq ProLiant 5500 450MHz system for RAID-0 and **10 percent higher** for RAID-5.

---

# Test Environments and Results

---

## SAP R/3 Two-Tier SD Benchmark

SAP AG developed its standard application benchmarks primarily to assure the quality of its R/3 System distributed application suite, which consists of seven system components: FI for Financial, MM for materials management, SD for Sales and Distribution, PP for Production and Planning, WM for Warehouse Management, PS for Project System, and HR for Human Resources. The benchmark suite is designed to test the performance of a system of hardware, software and the R/3 System applications.

The SD module benchmark has become the most popular of the seven SAP modules for reporting the number of users supported. Because of its complexity, the SD application represents the best mix of users accessing the different modules across the database. SAP has developed a conversion factor, based on the SD benchmark score, from which a common throughput score can be derived for each of the other benchmarks.

The SAP benchmark consists of two scenarios: central, or tier-two, and distributed, or three-tier. For the two-tier benchmark, the database server and application server run on the same physical machine. For the three-tier benchmark, the database server and the application server run on different physical machines. The results presented in this report are based on the two-tier benchmark only.

---

## Results Summary

The IBM Netfinity 5500 M20 server, running the SAP R/3 Two-Tier Centralized SD Benchmark, **achieved 200 SD Benchmark users** with 1.92 seconds average dialog response time and a measured throughput of 60,000 dialog steps per hour, equivalent to 20,000 processed order lines per hour, with an average CPU utilization of 99 percent for the database server.

The environment for this benchmark was an IBM Netfinity 5500 M20 database server using four 500MHz Pentium III Xeon processors with 1MB L2 cache each, 4GB of main memory, and 655GB of total disk space running Windows NT Server 4.0

Enterprise Edition with Service Pack 4, Oracle 8.0.4 and SAP's R/3 Two-Tier Centralized SD Benchmark, Release 4.0 B. This benchmark fully complied with SAP's issued benchmark regulations and was audited and certified by SAP. Details can be obtained from IBM and SAP AG.

The SAP Standard Application Benchmark was performed at IBM's Netfinity Server Performance Laboratory in Research Triangle Park, NC, USA, by IBM engineers, and the results were certified with the information detailed above.

For a complete list of SAP R/3 Benchmark results, visit:

**<http://www.sap.com/products/techno>**

Click on Quickfinder and select Standard Application Benchmarks. Scroll to the bottom of the page, and click on R/3 Benchmarks - Results, Version 4.1.

---

## NotesBench for Domino R5

---

### Results Summary

The IBM Netfinity 5500 M20, configured with two 500MHz Pentium III Xeon processors, **demonstrated leadership price/performance running Lotus Domino R5 Server on Microsoft Windows NT Server 4.0** in recently conducted measurements using the NotesBench 164.1 R4 Mail-only workload. Results summarized below:

Test Script	Maximum Users	NotesMark (tpm)	Response Time (sec)	\$/User	\$/NotesMark
Mail-Only (RAID-1)	8,250	10,957	0.264	\$5.51	\$4.15

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 5500 is based on a total system cost of \$45,468.

At the time of publication, no other server manufacturer had published results for NotesBench on Domino R5.

---

### Measurement Methodology

The IBM Netfinity 5500 M20 system under test was configured with two 500MHz Pentium III Xeon processors with 1MB ECC L2 cache; 2GB of memory; twenty-three 4.51GB 10K Wide Ultra SCSI hard disks using RAID-1. For these tests, two 100Mbps Ethernet LAN segments were used. The system under test, the destination servers, and the client driver systems were connected to the LAN by one IBM PC Server 325 configured as a router. Destination mail addresses were distributed across six destination servers. The IBM Netfinity 5500 M20 system under test ran Microsoft Windows NT Server Version 4.0 with Service Pack 4 and Domino R5 Server.

---

### Measurement Analysis

These results are based on running the IBM Netfinity 5500 M20 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.

A detailed description of the benchmarked configuration and analysis of the certified results are provided in the NotesBench Disclosure Report for the IBM Netfinity 5500 M20 system, which is available on the World Wide Web at:

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

NotesBench Disclosure Reports for all published results are also available on the World Wide Web at:

**<http://www.notesbench.org>**

---

## SPECweb96 Benchmark

---

### Results Summary

The IBM Netfinity 5500 M20, announced worldwide in March 1999, delivers industry-leading SPECweb96 results on Windows NT Server 4.0 in four-way, two-way and one-way configurations.

The IBM Netfinity 5500 M20, configured with four 500MHz Pentium III Xeon processors with 1MB of L2 cache and 1GB of memory, **achieved peak results of 7,796 operations per second with an average response time of 12.9 milliseconds per operation.**

In two-way and one-way configurations, the Netfinity 5500 M20 **achieved peak results of 6,684 operations per second with an average response time of 13.8 milliseconds, and 5,958 operations per second with an average response time of 15.2 milliseconds,** respectively.

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.

**Note:** SPECweb96 results based on Windows NT Server 4.0 should not be compared with SPECweb96 results based on beta versions of Windows 2000\*\*, which will not be generally available as a product until the fourth quarter of 1999. Windows 2000 contains network performance enhancements that significantly boost SPECweb96 performance.

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 defines two metrics: operations per second and average response time in milliseconds per workload. An "operation" is an HTTP request for an HTML file or an object referenced in an HTML file.



More information about these results is available at IBM's Netfinity  
Web site at:

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

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>**

---

## ServerBench 4.02

Ziff-Davis' ServerBench 4.02 test suite SYS\_60.TST was used to measure the performance of the IBM Netfinity 5500 M20 500MHz and the Compaq ProLiant 5500 450MHz 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 5500 M20 500MHz achieved a peak level of transactions per second at eight clients that was **13 percent higher** than that of the Compaq ProLiant 5500 450MHz system.



## RAID-5 Configuration

The Netfinity 5500 M20 500MHz achieved a peak level of transactions per second at eight clients that was **14 percent higher** than that of the Compaq ProLiant 5500 450MHz 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 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 32 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 5500 M20 500MHz and the Compaq ProLiant 5500 450MHz 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

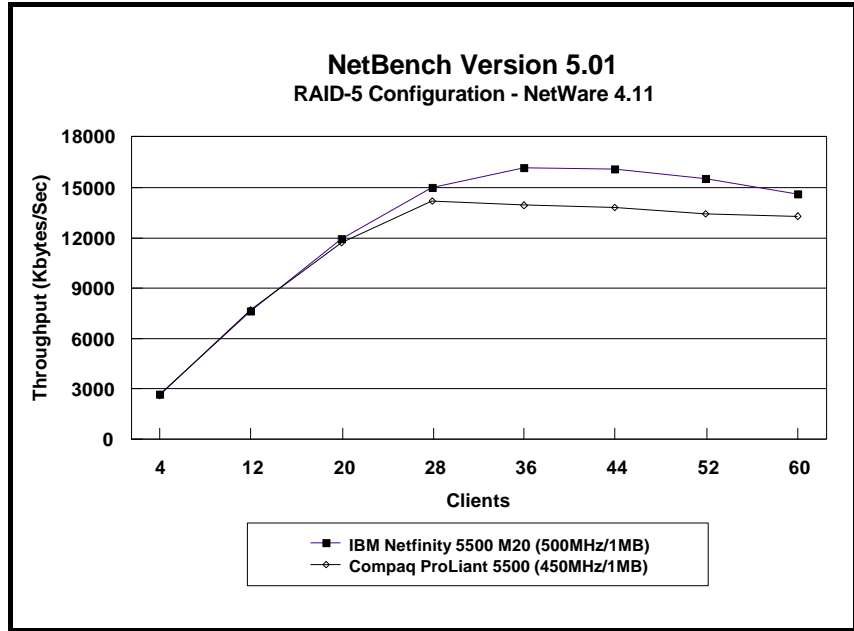
### RAID-0 Configuration

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 5500 M20 500MHz system provided network clients with a level of throughput that was **49 percent higher** than that of the Compaq ProLiant 5500 450MHz system.



## RAID-5 Configuration

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 5500 M20 500MHz server provided network clients with a level of throughput that was **10 percent higher** than that of the Compaq ProLiant 5500 450MHz server.



---

### Measurement Methodology

The Disk Mix test suite was performed using four 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 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 5500 M20 / 500MHz	Compaq ProLiant 5500 / 450MHz
Processor	Four 500MHz Pentium III Xeon	Four 450MHz Pentium II Xeon
Memory	1GB ECC	1GB ECC
L2 Cache	1MB ECC (Write-Back)	1MB ECC (Write-Back)
RAID Level	0 and 5	0 and 5
Disk Drives	Five 9.1GB 10K Wide Ultra SCSI Drives	Five 9.1GB 10K Wide Ultra SCSI Drives
Disk Drive Adapter	Integrated ServeRAID II Ultra SCSI Controller	Compaq Smart Array Controller 3200
Disk Driver	ISPRAIDN.SYS	CPQARRAY.SYS
Network Adapters	Four IBM EtherJet 100/10 PCI Adapters	Four Netelligent 10/100 Ethernet Adapters
Bus	PCI	PCI
Network Driver	IBMFENT.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
Test Suite	ServerBench 4.02 SYS_60.TST	ServerBench 4.02 SYS_60.TST

## NetBench 5.01

Features	IBM Netfinity 5500 M20 / 500MHz	Compaq ProLiant 5500 / 450MHz
Processor	One 500MHz Pentium III Xeon	One 450MHz Pentium II Xeon
Memory	1GB ECC	1GB ECC
L2 Cache	1MB (Write-Back)	1MB (Write-Back)
RAID Level	0 and 5	0 and 5
Disk Drives	Five 9.1GB 10K Wide Ultra SCSI Drives	Five 9.1GB 10K Wide Ultra SCSI Drives
Disk Drive Adapter	Integrated ServeRAID Ultra SCSI Controller	* Compaq Smart Array Controller 3200
Disk Driver	ISPRAID.HAM V2.81	CPQDA386.DSK V3.10
Network Adapters	Four IBM EtherJet 100/10 PCI Adapters	Four Netelligent 10/100 Ethernet Adapters
Bus	PCI	PCI
Network Driver	E100B.LAN V3.63	CPQNF3.LAN V2.42
Network Operating System	NetWare 4.11 with IWSP6 Loaded	NetWare 4.11 with IWSP6 Loaded
NetWare Volume Block Size	16KB	RAID-5 = 16KB, RAID-0 = 64KB
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

\* Adapter set at default values, for example, RAID-0 = 128KB, RAID-5 = 16KB

---

# 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
- Delay: 0
- Think: 0

**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>	Asante 100Mbps Ethernet
<b>Clients per Segment</b>	15
<b>CPU</b>	133MHz Pentium / 32MB
<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.

\*\*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.

\*\*Oracle is a registered trademark of Oracle Corporation.

\*\*SAP is a registered trademark of SAP AG.

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