

---

# Performance Report

---

## IBM Netfinity 7000 M10 (400MHz)

Version 1.0  
September 1998



---

## Executive Overview

---

*The performance of the IBM Netfinity\* 7000 M10 server, announced worldwide in September 1998, was evaluated using the following benchmarks:*

- *TPC Benchmark C\*\**
- *TPC Benchmark D\*\**
- *SAP R/3 System Centralized Sales and Distribution (SD) Benchmark*
- *SPECweb96\*\**
- *Ziff-Davis' ServerBench\*\* Version 4.0*
- *Ziff-Davis' WebBench\*\* Version 2.0*
- *Ziff-Davis' NetBench\*\* Version 5.01*

*For the Ziff-Davis benchmarks, the IBM Netfinity 7000 M10 400Mz<sup>1</sup>/1MB L2 cache model was measured. For comparison, a similarly configured Netfinity 7000 system was also measured.*

*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 Enterprise Edition 4.0 and Microsoft SQL Server Enterprise Edition 7.0, the IBM Netfinity 7000 M10 **delivered a leadership transaction processing performance result of 18,893.43 tpmC** at a price/performance ratio of \$29.09/tpmC.

#### **TPC Benchmark D**

When introduced June 29, 1998, the Netfinity 7000 M10 set a new performance record for TPC-D results on a 100GB database. Running Microsoft\*\* Windows NT Enterprise Edition 4.0 and DB2 Universal Database Enterprise Edition 5.2.0, the IBM Netfinity 7000 M10 **delivered leadership performance results of**

**1,871.1 QppD@100GB and 661.3 QthD @ 100GB, as well as price/performance of \$243 per QphD @ 100GB.**<sup>3</sup> (From June 29, through September 10, the Netfinity 7000 M10 server's TPC-D results were ranked number one in performance and price/performance on a four-way Intel-based system.)

---

#### **SAP R/3 System Centralized SD Benchmark**

In August 1998, the IBM Netfinity 7000 M10 set a new performance record running SAP R/3 Centralized System Sales and Distribution (SD) Benchmark on four-way, Intel-based servers. The IBM Netfinity 7000 M10 **achieved 240 Centralized SD Benchmark users** with 1.95 seconds average dialog response time on Oracle 8.0.4 and Microsoft Windows NT Enterprise Edition 4.0. This represents a throughput of 24,000 fully processed order line items per hour.

---

#### **SPECweb96 Benchmark**

With the best-ever SPECweb96 results on Windows NT Server 4.0, the IBM Netfinity 7000 M10 **achieved 5,277 operations per second with a response time of 8.3 milliseconds per operation**, running Apache HTTP Server Version 1.3.1 (and using IBM Research-developed technology that accelerates Web server performance).

---

#### **Ziff-Davis Benchmarks**

##### **ServerBench 4.01**

ServerBench 4.01 was used to measure the performance of the IBM Netfinity 7000 M10 and Netfinity 7000 systems as four-way application servers running Windows NT Server 4.0.

The peak level of throughput achieved by the IBM Netfinity 7000 M10 was **71 percent higher** than that of the Netfinity 7000 for RAID-0 and **69 percent higher** than that of the Netfinity 7000 for RAID-5.

##### **WebBench 2.0**

WebBench 2.0 was used to measure the performance of the IBM Netfinity 7000 M10 and Netfinity 7000 systems as four-way Web servers running Microsoft Internet Information Server 3.0 on Microsoft Windows NT 4.0.

The IBM Netfinity 7000 M10 achieved a peak level of throughput that was **1.4 times higher** than the Netfinity 7000 for RAID-0 and **1.6 times higher** than the Netfinity 7000 for RAID-5.

The IBM Netfinity 7000 M10 serviced a number of requests per second that was **1.4 times higher** than the Netfinity 7000 for RAID-0 and **1.6 times higher** than the Netfinity 7000 for 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 and Netfinity 7000 systems as single-processor file servers running Novell\*\* NetWare\*\* 4.11 with IWSPB4.

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 7000 M10 provided network clients with **1.4 times more throughput** than the Netfinity 7000 for RAID-0 and **two times more throughput** than the Netfinity 7000 for RAID-5.

---

## Test Environments and Results

---

### TPC Benchmark C

---

#### Results Summary

Running Microsoft\*\* Windows NT Enterprise Edition 4.0 and Microsoft SQL Server Enterprise Edition 7.0, the IBM Netfinity 7000 M10 **delivered a leadership transaction processing performance result of 18,893.43 tpmC** at a price/performance ratio of \$29.09/tpmC. (Results current as of September 22, 1998.)

Tested Configuration		tpmC	\$ / tpmC
IBM Netfinity 7000 C/S		18,893.43	\$29.09 / tpmC
Database	Microsoft SQL Server Enterprise Edition 7.0		
Operating System	Microsoft Windows NT Server Enterprise Edition 4.0		
TP Monitor	TUXEDO 6.3 CFS		
Availability Date	Hardware: September 25, 1998	RAID Controller Device Driver and Firmware: September 29, 1998	Software: December 29, 1998
Configuration	Client / Server		
5-Year Cost of Ownership	\$549,551		
Benchmark Revision	Rev. 3.4		

The priced configuration consisted of one IBM Netfinity 7000 M10 server with four 400MHz Pentium II Xeon processors (each with 1MB L2 cache) and 4GB<sup>2</sup> of memory. Nineteen IBM Netfinity EXP15 Rack Storage Expansion Enclosures housed a total of 1,674.4GB of storage to accommodate database growth over a 180-day period. Three IBM Netfinity 3500 servers, each configured with two 333MHz Pentium II processors and 384MB of memory, were used as clients connected to the Netfinity 7000 M10 via a 100BaseT Ethernet LAN. The three clients serviced the requests from a total of 15,330 emulated users.

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.

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

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 1,871.1 QppD @ 100GB and 661.3 QthD @ 100GB at a price/performance of \$243 per QphD @ 100GB.** (From June 29, through September 10, the Netfinity 7000 M10 server's TPC-D results were ranked number one in performance and price/performance on a four-way Intel-based system.)

Tested Configuration	QppD @ 100GB	QthD @ 100GB	\$ per QphD @ 100GB
IBM Netfinity 7000 M10	1,871.1	661.3	\$243
Database	DB2 Universal Database Enterprise Edition 5.2.0		
Operating System	Microsoft Windows NT Server 4.0 Enterprise Edition 4.0		
Availability Date	Hardware: September 25, 1998		Software: October 31, 1998
5-Year Cost of Ownership	\$270,526		
Benchmark Revision	Rev. 1.3.1		

The priced configuration consisted of one IBM Netfinity 7000 M10 server with four 400MHz Pentium II Xeon processors (each with 1MB L2 cache) and 2GB of memory. Eight Netfinity EXP10 Rack Storage Expansion Enclosures housed 746.04GB of storage.

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 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>.

---

## SAP R/3 Centralized 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: centralized and distributed. For the centralized benchmark, the database server and application server run on the same physical machine. For the distributed benchmark, the database server and the application server run on different physical machines. The results presented in this report are based on the centralized benchmark only.

---

## Results Summary

### SAP R/3 Centralized SD Benchmark

In August, the IBM Netfinity 7000 M10 system set a new performance record for running the SAP R/3 System Centralized SD Benchmark on four-way, Intel-based servers. The IBM Netfinity 7000 M10 **achieved 240 SD Benchmark users** with 1.95 seconds average response time and a throughput of 24,000 fully processed order line items per hour.

The environment for this benchmark was an IBM Netfinity 7000 M10 database server using four 400MHz Pentium II Xeon processors with 1MB L2 cache each, 4GB of main memory, and 455GB of total disk space running Windows NT Server Enterprise Edition 4.0 (SP3), Oracle 8.0.4 and SAP's R/3 Rel. 3.1 H (3.1 I kernel). The measured throughput was 72,000 dialog steps per hour with an average CPU utilization of 99 percent for the central server. 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 on August 8, 1998, by IBM at its Netfinity Server Performance Laboratory in Research Triangle Park, NC, USA, by IBM engineers, and the results were certified with the information detailed above.

A complete list of SAP R/3 SD Benchmark results is available at the SAP Web site; see R/3 Standard Application Benchmarks - Published Results at:

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

---

## SPECweb96 Benchmark

---

### Results Summary

With the best-ever SPECweb96 results on Windows NT Server 4.0, the IBM Netfinity 7000 M10, configured with four 400MHz Pentium II Xeon processors and 1GB of memory, **achieved 5,277 operations per second with a response time of 8.3 milliseconds per operation**, running Apache HTTP Server Version 1.3.1 (and using IBM Research-developed technology that accelerates Web server 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.

---

## ServerBench 4.01

Ziff-Davis' ServerBench 4.01 test suite SYS\_60.TST was used to measure the performance of the IBM Netfinity 7000 M10 (400MHz Pentium II Xeon) and the Netfinity 7000 (200MHz Pentium Pro) systems as four-way application servers running Windows NT Server 4.0.

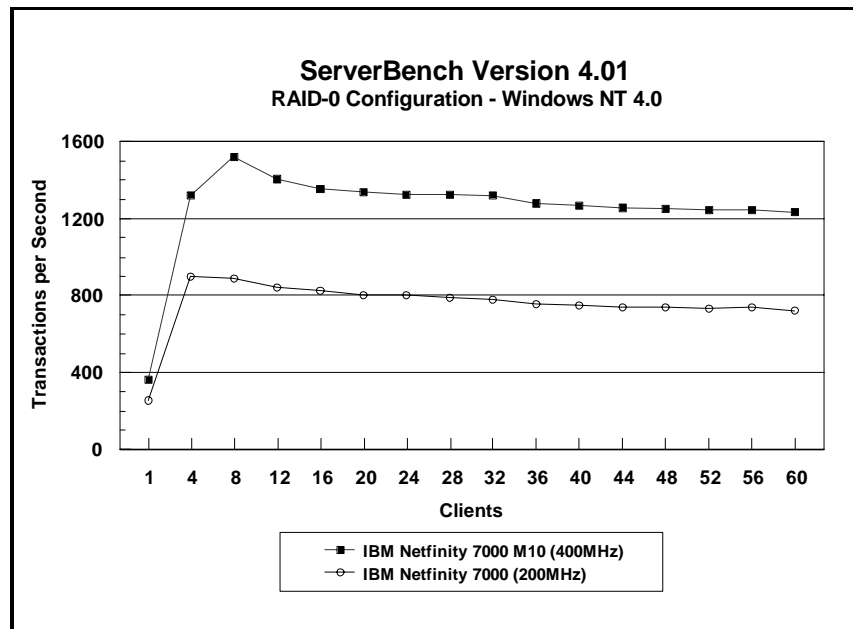
ServerBench 4.01 provides an overall transactions-per-second (TPS) score showing how well the server handles client requests for a variety of operations involving the server's processors, disk and network subsystems.

---

## Results Summary


### RAID-0 Configuration

The Netfinity 7000 M10 achieved a peak level of throughput at eight clients that was **71 percent higher** than the Netfinity 7000.



## RAID-5 Configuration

The Netfinity 7000 M10 achieved a peak level of throughput at eight clients that was **69 percent higher** than the Netfinity 7000.



---

### 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 4.01 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 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 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.

---

## WebBench 2.0

The WebBench 2.0 test suite NT\_SIMPLE\_CGI20\_V1.TST was used to measure the performance of the IBM Netfinity 7000 M10 and Netfinity 7000 systems as four-way Web servers running Microsoft Internet Information Server 3.0 on Windows NT Server 4.0 with Service Pack 3.


This test suite performs both static HTML page requests and dynamic Internet Server API (ISAPI) requests, which are the two primary functions of an enterprise Web server running Microsoft Internet Information Server.

---

## Results Summary


### Throughput for RAID-0 Configuration

Under a high-end workload of 60 WebBench clients, the IBM Netfinity 7000 M10 delivered **1.4 times more throughput** than the Netfinity 7000.



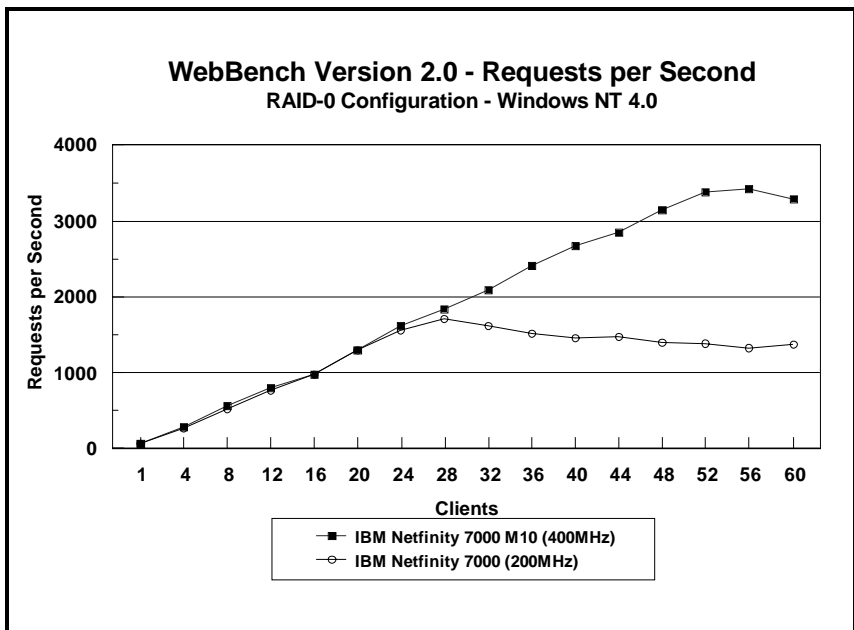
### Throughput for RAID-5 Configuration

Under a high-end workload of 60 WebBench clients, the IBM Netfinity 7000 M10 delivered **1.6 times more throughput** than the Netfinity 7000.




### Requests per Second for RAID-0 Configuration

Under a high-end workload of 60 WebBench clients, the IBM Netfinity 7000 M10 delivered **1.4 times more throughput** than the Netfinity 7000.



### Requests per Second for RAID-5 Configuration

Under a high-end workload of 60 WebBench clients, the IBM Netfinity 7000 M10 delivered **1.6 times more throughput** than the Netfinity 7000.



---

### Measurement Methodology

The servers were configured with four 100Mbps Ethernet adapters. Each adapter/segment contained 15 IBM PC 750 systems running Windows NT Workstation 4.0 for a total of 60 clients. Each client, using the WebBench 2.0 NT\_SIMPLE\_ISAPI20\_V20.TST workload, randomly issued two types of transactions:

- Static HTML page requests, which simulates an actual Web browser, fetching predesigned HTML pages from the server using the HTTP protocol.
- Dynamic Internet Server API requests, which execute on the server and create the HTML response data, thereby using the server's processing resources.

At the beginning of the test, the controller sets up a random pattern in each client, so that all 60 clients do not request the same file at the same time.

The NT\_SIMPLE\_ISAPI20\_V20.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 measurements were not made. Each mix runs for 300 seconds and consists of 80 percent static



and 20 percent dynamic requests. 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.

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 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.

To get a valid measure of the server's performance, the requests-per-second and the throughput scores 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 52 clients,

peaked at 56 clients, and then dropped off, indicating that the Netfinity 7000 M10 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., processor, 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

The NetBench 5.01 Disk Mix was used to measure the performance of the IBM Netfinity 7000 M10 and the Netfinity 7000 system as single-processor file servers running Novell NetWare 4.11 with Service Pack IWSP4B. 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 provided network clients with **1.4 times more throughput** than the Netfinity 7000.



## RAID-5 Configuration

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 7000 M10 provided network clients with **two times more throughput** than the Netfinity 7000.



---

### 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 4.01

Features	IBM Netfinity 7000 M10 400MHz	IBM Netfinity 7000 200MHz
Processor	Four 400MHz Pentium II Xeon	Four 200MHz Pentium Pro
Memory	1024MB ECC	1024MB ECC
L2 Cache	1MB (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 ServerRAID-3H Ultra2 SCSI Adapter	One ServerRAID II Ultra 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	E100BNT.SYS	E100BNT.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	1024MB	1024MB
File System	NTFS	NTFS
Allocation Unit Size	Predefined Default	Predefined Default
Test Suite	ServerBench 4.01 SYS_60.TST	ServerBench 4.01 SYS_60.TST

---

## WebBench 2.0

<b>Features</b>	<b>IBM Netfinity 7000 M10 400MHz</b>	<b>IBM Netfinity 7000 200MHz</b>
<b>Processor</b>	Four 400MHz Pentium II Xeon	Four 200MHz Pentium Pro
<b>Memory</b>	1024MB ECC	1024MB ECC
<b>L2 Cache</b>	1MB (Write-Back)	1MB (Write-Back)
<b>RAID Level</b>	RAID-0 and RAID-5	RAID-0 and RAID-5
<b>Disk Drives</b>	Nine 9.1GB 10K Wide Ultra SCSI Drives	Nine 9.1GB 10K Wide Ultra SCSI Drives
<b>Disk Drive Adapter</b>	One ServerRAID-3H Ultra2 SCSI Adapter	One ServerRAID II Ultra SCSI Adapter
<b>Disk Driver</b>	ISPRAIDN.SYS	ISPRAIDN.SYS
<b>Network Adapters</b>	Four IBM EtherJet 100/10 PCI Adapters	Four IBM EtherJet 100/10 PCI Adapters
<b>Bus</b>	PCI	PCI
<b>Network Driver</b>	E100BNT.SYS	E100BNT.SYS
<b>Network Operating System</b>	Windows NT Server 4.0 with Service Pack 3	Windows NT Server 4.0 with Service Pack 3
<b>System Partition Size</b>	1GB	1GB
<b>File Compression</b>	Off	Off
<b>File System</b>	NTFS	NTFS
<b>Allocation Unit Size</b>	Predefined Default	Predefined Default
<b>Test Suite</b>	NT_SIMPLE_CGI20_V2.TST	NT_SIMPLE_CGI20_V2.TST
<b>Web Server</b>	Microsoft Internet Information Server 3.0	Microsoft Internet Information Server 3.0



---

## NetBench 5.01

<b>Features</b>	<b>IBM Netfinity 7000 M10 400MHz</b>	<b>IBM Netfinity 7000 200MHz</b>
<b>Processor</b>	One 400MHz Pentium II Xeon	One 200MHz Pentium Pro
<b>Memory</b>	1024MB ECC	1024MB ECC
<b>L2 Cache</b>	1MB (Write-Back)	1MB (Write-Back)
<b>RAID Level</b>	RAID-0 and RAID-5	RAID-0 and RAID-5
<b>Disk Drives</b>	Nine 9.1GB 10K Wide Ultra SCSI Drives	Nine 9.1GB 10K Wide Ultra SCSI Drives
<b>Disk Drive Adapter</b>	One ServerRAID-3H Ultra2 SCSI Adapter	One ServerRAID II Ultra SCSI Adapter
<b>Disk Driver</b>	ISPRAID.HAM V2.81	ISPRAID.HAM V2.81
<b>Network Adapters</b>	Four IBM EtherJet 100/10 PCI Adapters	Four IBM EtherJet 100/10 PCI Adapters
<b>Bus</b>	PCI	PCI
<b>Network Driver</b>	E100B.LAN V3.23	E100B.LAN V3.23
<b>Network Operating System</b>	NetWare 4.11 with IWSPB4 Loaded	NetWare 4.11 with IWSPB4 Loaded
<b>NetWare Volume Block Size</b>	16KB	16KB
<b>File Compression</b>	Off	Off
<b>Block Allocation</b>	On	On
<b>Data Migration</b>	Off	Off
<b>Disk Mix</b>	NB5.01/Windows NT Workstation 4.0	NB5.01/Windows NT Workstation 4.0

---

## Test Disclosure Information

---

### ServerBench 4.01

The measurements were conducted using Ziff-Davis' ServerBench 4.01 running the Windows NT Workstation 4.0 clients as described below:

**Version:** 4.01

#### 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

The Netfinity 7000 M10 is planned to be available September 25, 1998. All other products used for these measurements were 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

---

## WebBench 2.0

The measurements were conducted using Ziff-Davis' WebBench 2.0 running the NT\_SIMPLE\_CGI20\_V2.TST test suite with Windows NT Workstation 4.0 clients as described below:

**Version:** 1.0

**Mixes**

- NT\_SIMPLE\_CGI20\_V2.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

**Network Operating System:** Windows NT Server 4.0 (Build 1381) with Service Pack 3

**Web Server:** Microsoft Internet Information Server 3.0

---

## Testbed Disclosure

The Netfinity 7000 M10 is planned to be available September 25, 1998. All other products used for these measurements were 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 / Memory</b>	200MHz 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 IWSP4B 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

The Netfinity 7000 M10 is planned to be available September 25, 1998. All other products used for these measurements were 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.

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

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 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.

<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.