

---

# Performance Report

---

## IBM Netfinity 7000

Version 1.3  
February 1998



---

## Executive Overview

---

*The performance of the IBM Netfinity\* 7000 server, announced worldwide in September 1997, was evaluated using the following benchmarks:*

- *TPC Benchmark C\*\**
- *SAP R/3 System Application Distributed and Centralized Sales and Distribution (SD) Benchmarks*
- *NotesBench for Lotus\*\* Domino\*\* Server Releases 4.51 and 4.6*
- *Ziff-Davis' ServerBench\*\* Version 4.0*
- *Ziff-Davis' WebBench\*\* Version 1.0*
- *Ziff-Davis' NetBench\*\* Version 5.01*

*For the Ziff-Davis benchmarks, the IBM Netfinity 7000 512KB and 1MB L2 cache models were measured. For comparison, the Compaq\*\* ProLiant\*\* 6000 and Hewlett-Packard\*\* NetServer\*\* LX Pro systems were measured in configurations similar to the Netfinity 7000 512KB L2 cache model.*

*This report version replaces Version 1.2, published in January 1998. This report provides updated price/performance results for the NotesBench measurements.*

---

### IBM Netfinity 7000

The IBM Netfinity 7000, available in rack and tower models, is a four-way symmetric multiprocessing (SMP) server featuring a 200MHz<sup>1</sup> Pentium Pro processor with either 1MB of L2 cache or 512KB of L2 cache, and up to 4GB of memory. The systems are ruggedly reliable enterprise servers that deliver exceptional power, functionality and value for handling complex, business-critical database or application server needs.

These servers, available in both tower and rack-ready models, are easy to install in either environment. Conversion kits are available to convert towers to racks, and vice versa, if requirements change. They offer configuration flexibility with such features as hot-swap, open-bay models; open systems to support conventional Ethernet

---

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

or single port- to multi-port Ethernet PCI adapters, and a full-high bay for internal tape units capable of backing up 70GB of data.

The Netfinity 7000 system is intended for enterprise customers whose networks depend on a reliable, high-performance server platform with excellent scalability for future growth. The system is ideal for handling business-critical database, groupware, messaging and intranet applications. The Netfinity 7000 can also function as a file server in LAN environments where there is a significant number of users requiring a large amount of file server resources. Standard features include:

- One to four 200MHz Pentium Pro processors, each with an integrated 1MB or 512KB L2 write-back cache
- Two processor cards with two ZIF sockets standard (a second, third, and fourth processor can be easily installed)
- 256MB ECC system memory, expandable to 4GB
- Advanced System Management Adapter for monitoring and controlling the Netfinity 7000 remotely or locally
- Advanced architecture: dual 132MB/second, I<sub>2</sub>O-ready PCI buses
- Room for adding 32-bit adapters: six PCI slots/four EISA slots
- 4.51GB<sup>2</sup> and 9.1GB Wide Ultra SCSI hot-swap drive options
- Optional IBM ServeRAID II Ultra SCSI Adapter and SSA RAID adapters for high-performance RAID 0, 1 and 5
- Twelve hot-swap hard disk drive bays, 8X-speed CD-ROM, 1.44MB diskette drive, and four open 5.25-inch, half-high bays
- Hot-swap, redundant cooling fans
- Two 400W worldwide power supplies standard (400W hot-swap, redundant power supply optional)
- Two serial ports (56Kbps) and one parallel port (peak transfer rate of 2MB/sec) that supports devices using ECP/EPP protocols adhering to the IEEE 1284 standard.
- ServerGuide\* for easy installation of network operating systems
- NetFinity Manager for comprehensive system and asset management
- Lotus Domino Server 4.6 (single-processor edition) for more efficient business communications

---

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

---

## 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 6.5 Enterprise Edition, the IBM Netfinity 7000 **delivered transaction processing results of 10,559.45 tpmC and price/performance of \$47.43/tpmC.**

The priced configuration consisted of one IBM Netfinity 7000 system with four 200MHz Pentium Pro processors (each with 1MB L2 cache) and 4GB of memory. Fifteen IBM Netfinity EXP10 Rack Storage Expansion Enclosures housing a total of 696.38GB of storage to accommodate database growth over a 180-day period. Five IBM PC Server 325 systems, each configured with one 266MHz Pentium II processor and 256MB of memory, were used as clients connected to the Netfinity 7000 via a 100BaseT Ethernet LAN. The five clients serviced the requests from a total of 8,750 emulated users.

The Executive Summary for the IBM Netfinity 7000 system is available on the World Wide Web at <http://www.tpc.org> (the TPC Web site) or at <http://www.pc.ibm.com/us/techlink/srvperf.html> (IBM's Web site).

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.

For a complete description of the benchmark, visit the TPC Web site at <http://www.tpc.org>.

---

**SAP R/3 SD Benchmarks**

In October 1997, the IBM Netfinity 7000 set a new performance record for running SAP R/3 **Distributed Sales and Distribution (SD) Benchmark** on 4-way, Intel-based servers. The IBM Netfinity 7000 **achieved 1,200 SD Benchmark users** with 1.75 seconds average dialog response time on Oracle 7.3 and Microsoft Windows NT 4.0. This represents a throughput of 123,000 fully processed order line items per hour.

In recent measurements, the IBM Netfinity 7000 has set another performance record for running SAP R/3 **Centralized SD Benchmark** on 4-way, Intel-based servers. The IBM Netfinity 7000 **achieved 130 SD Benchmark users** with 1.73 seconds response time on Oracle 7.3 and Microsoft Windows NT 4.0.

The IBM Netfinity 7000 now holds the number one position on both the Centralized and Distributed SAP R/3 Benchmarks on a 4-way, Intel-based server.

---

**NotesBench 4.51 and 4.6**

The IBM Netfinity 7000 system **demonstrated leadership performance and price/performance running Lotus Domino Server 4.51 (RAID-0)** using two NotesBench workloads:

Mail-only and Mail and Shared Database (MailDB). Configured with three 200MHz / 1MB L2 cache Pentium Pro processors and 1GB of memory for Mail, and four processors for MailDB, the IBM Netfinity 7000 system ran Lotus Domino Server Release 4.51 on Microsoft Windows NT Server 4.0 with Service Pack 3.

The IBM Netfinity 7000 system also **demonstrated leadership performance running Lotus Domino Server 4.6 (RAID-5)** configured as a partitioned server running the NotesBench Mail-only workload. For this test, the system was configured with four processors and 2GB of memory. The following table summarizes all results.

Test Script	Maximum Users	NotesMark (tpm)	Response Time (sec)	\$/User	\$/NotesMark
<b>Mail-Only - Partitioned Server</b>	4,650	6,071	1.453	\$20.73	\$15.88
<b>Mail-Only</b>	3,500	4,623	0.312	\$14.92	\$11.30
<b>MailDB</b>	2,900	6,294	1.042	\$19.82	\$9.13

---

## Ziff-Davis Benchmarks

For these benchmarks, two configurations of the IBM Netfinity 7000 server were measured:

- The 200MHz / 1MB L2 cache model configured with IBM Netfinity EXP10 4.51GB 10K Wide Ultra SCSI (10K rpm) hard disks attached via a Netfinity EXP10 Rack Storage Expansion Enclosure.
- The 512KB cache model configured similarly to the competitors' systems.

Neither the Compaq nor the HP system was available with the 1MB L2 cache or the 10K rpm drives at the time of testing. The Compaq and HP systems used in these benchmarks were shipped with a 200MHz / 512KB cache Pentium Pro processor and 7200 rpm drives.

### ServerBench 4.0

ServerBench 4.0 system test suite was used to measure the performance of the IBM Netfinity 7000 system and the competitors' systems as 4-way application servers running Windows NT Server 4.0.

The peak level of throughput achieved by the IBM Netfinity 7000 200MHz/1MB was:

- **63 percent higher** than the HP NetServer LX Pro
- **16 percent higher** than the Compaq ProLiant 6000
- **14 percent higher** than the IBM Netfinity 7000

### WebBench 1.0

WebBench 1.0 was used to measure the performance of the IBM Netfinity 7000 system and the competitors' systems as 4-way Web servers running Microsoft Internet Information Server 3.0 on Microsoft Windows NT 4.0.

At the peak of 12 WebBench clients, the IBM Netfinity 7000 200MHz/1MB system delivered:

- **26 percent more throughput** than the HP NetServer LX Pro
- **11 percent more throughput** than the Compaq ProLiant 6000
- **11 percent more throughput** than the IBM Netfinity 7000

At the peak of 12 WebBench clients, the IBM Netfinity 7000 200MHz/1MB serviced:

- **22 percent more requests per second** than the HP NetServer LX Pro
- **8 percent more requests per second** than the Compaq ProLiant 6000
- **8 percent more throughput** than the IBM Netfinity 7000

#### **NetBench 5.01**

NetBench 5.01 Disk Mix Windows for Workgroup Clients was used to measure the performance of the IBM Netfinity 7000 system and the competitors' systems as single-processor file servers running Novell\*\* NetWare\*\* 4.11 in a 100Mbps Ethernet environment.

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 7000 200MHz/1MB system provided network clients with:

- **59 percent more throughput** than the HP NetServer LX Pro
- **10 percent more throughput** than the Compaq ProLiant 6000
- **8 percent more throughput** than the IBM Netfinity 7000

---

## Test Environments and Results

---

### TPC Benchmark C

In 1988, the Transaction Processing Performance Council (TPC) was formed to fulfill the need for transaction processing benchmarks that emulate the workloads found on database servers. The council includes representatives from a cross-section of 45 hardware and software companies who meet to establish benchmark content. A primary goal of the council is to provide objective and verifiable performance data to the industry.

The TPC developed the TPC-C, an on-line transaction processing (OLTP) benchmark, which evaluates the performance of database systems. TPC-C generates a more complex workload than its predecessors TPC-A and TPC-B, whose workloads consisted of a single transaction type.

TPC-C incorporates the remote client in its definition of the benchmark environment. Thus, TPC-C results reflect the influence of client interaction and data transmission on the server under test. 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.

Tested Configuration		tpmC	\$ / tpmC
IBM Netfinity 7000 C/S		10,559.45	\$47.43 / tpmC
Database	Microsoft SQL Server 6.5 Enterprise Edition		
Operating System	Windows NT Server 4.0 Enterprise Edition		
TP Monitor	NCR TOP END 2.04		
Availability Date	Hardware: Now	RAID Controller Device Driver and Firmware: April 30, 1998	Software: Now
Configuration	Client / Server		
Five-Year Cost of Ownership	\$500,831		
Benchmark Revision	Rev. 3.3		



---

**Measurement  
Methodology**

The TPC-C specification defines a mix of transactions that make up a generic OLTP application to be used as a benchmark. The benchmark is based on the application environment of a wholesale supplier with many distributors, in which the workload consists of order-entry processing and the associated operations.

The TPC-C performance metrics are tpmC and dollars/tpmC. The throughput metric, tpmC, represents the average number of New-Order transactions completed per minute during the measurement interval. Four other transaction types are included because they contribute to the realism of the benchmark by producing a varied system workload; however, these four transaction types are not counted as part of the tpmC score. The New-Order transaction type was chosen as the performance indicator because the order-entry function is common in real-world OLTP applications.

The price/performance metric, dollars/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. These guidelines are described below.

**Specific Transaction Mix**

The transaction mix completed by the TPC-C application must conform to the following minimum percentages:

- 43 percent - Payment transaction
- 4 percent - Delivery transaction
- 4 percent - Order status
- 4 percent - Stock-Level transaction

The remaining transactions are of the New-Order type, which count toward the tpmC metric. Requiring a mix of other transaction types to be executed while the New-Order is being measured ensures that a realistic workload exists on the system while the measurement is being taken.

**Quick Response Times**

Ninety percent of all transactions must be completed within a specified response time. This ensures that the response times for all transactions remain reasonable. The application is not allowed to favor the New-Order transaction to achieve a higher tpmC measurement.

**Appropriate Benchmark Scale**

A benchmark's scale is defined by the number of users who access the system during the measurement interval and the size of the TPC-C database required to support them. The TPC-C specification requires that the scale of the benchmark increase with the performance of the server. (When the server and the application are appropriately matched, the benchmark is more realistic. Higher results may be achieved on a server by running an inappropriately small scale benchmark, but these results would not accurately represent a real-world application.)

**Emulated User Interface**

The emulated user interface includes think and keying times. The benchmark must be able to emulate the time that it takes for a customer to determine what to order (think time) and enter it into the application interface (keying time).

**Fault-Tolerant System**

The system under test must be able to withstand power outages and disk failures, while preserving the integrity of the database.

**Priced Configuration**

The benchmark sponsor must reveal the market price for all of the hardware and software used in the test as well as the hardware and software maintenance costs for a five-year period. The equipment must be in plan to become generally available. The priced configuration is used to calculate the dollars/tpmC metric.

**Adequate Space for Growth**

The database must allow for the addition of records that would occur in a real-world environment over time. The test configuration must have enough disk space to allow the TPC-C benchmark to run for eight hours without error. The priced configuration must have enough disk space to support the addition of records that would occur in 180 days.

---

## Measurement Analysis

The TPC-C benchmark is as much a test of a server's ability to functionally support large memory and disk configurations as it is a measure of performance. To obtain competitive performance from a multiprocessor Pentium Pro server, the server must support at least 4GB of memory and approximately 700GB to 800GB of disk storage. In this benchmark in a Microsoft SQL Server environment, the Netfinity 7000 demonstrates best-of-breed performance and the flexibility to be configured and operate with enough memory and disk storage to achieve optimal performance from four 200MHz Pentium Pro processors.

The TPC-C benchmark employed a database scale of 875, which means that the server was configured with 875 warehouses. Each warehouse is associated with 10 districts; 30,000 customer entries; 30,000 history entries; 30,000 orders entries; 9,000 new-order entries; 300,000 order-line items; and 100,000 stock entries. In addition, the database was configured to have 100,000 different stock items for all warehouses. This large database was distributed across twelve 4.51GB 7200 rpm drives, seventy-nine 4.51GB 10K rpm drives, and thirty-one 9.1GB 7200 rpm drives. The primary objective of the benchmark was to demonstrate the ability of the Netfinity 7000 to be configured for and operate efficiently on a large database.

In addition to measuring transaction processing performance, which requires robust configuration capability, the TPC-C benchmark requires that the system under test support the ACID properties: Atomicity, Consistency, Isolation, and Durability. A series of ACID tests must be performed to ensure that these properties are supported by the system under test. The ability to withstand these rigorous tests demonstrates the integrity of the database and server hardware as well as the ability of a database server to operate under real-world conditions.

The Atomicity test ensures that all database transactions completed by the system under test are atomic; that is, the system either executes all parts of a transaction, or it executes none at all. This test is designed to ensure the integrity of the database.

The Consistency tests ensure that any operation by the database server takes the database from one consistent state to another consistent state. For example, the database begins with year-to-date sales of a warehouse equal to the total of year-to-date sales of all districts supported by that warehouse. Consistency demands that any transaction by a district should maintain the consistent relationship of district year-to-date sales and warehouse year-to-date sales. The TPC-C specification defines 12 consistency tests that are used to prove database consistency for the entire

database. Only four of these tests are required by the specification; that is the number used for this benchmark run.

The Isolation tests demonstrate the ability of a database to properly lock and serialize simultaneous transactions targeted at the same data elements. There are nine isolation tests that ensure the database can maintain accurate data for simultaneous updates, inquiries, and deletion of the same database fields.

Finally, the Durability tests are designed to illustrate the ability of the server database software and hardware to preserve committed transactions after power, disk, or memory failures. To prove durability requires a series of tests designed to simulate a database disk drive failure, a database log drive failure, and main memory failure. In each case the server must recover and not lose any committed transactions. This test is a most demanding demonstration of the ability of the Netfinity 7000 to safe-guard data under full-speed operating conditions.

The IBM Netfinity 7000 TPC-C Full Disclosure Report documents the performance, configuration capabilities, and robust database protection features of the Netfinity 7000 and Microsoft SQL Server database products. This document illustrates the abilities of the Netfinity 7000 to provide robust configuration options, database reliability and high performance in transaction processing environments.

The TPC-C Executive Summary for the IBM Netfinity 7000 is available in PostScript and Acrobat PDF formats on the TPC's World Wide Web site at <http://www.tpc.org>.

---

## SAP R/3 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 Council, established in April 1995, is made up of SAP and its business partners who conduct SAP benchmark tests. The council monitors SAP benchmarking activities and provides standards for reporting results, which must be verified by SAP prior to certification.

---

### Results Summary

#### SAP R/3 Distributed SD Benchmark

In October 1997, the IBM Netfinity 7000 system set a new performance record for running SAP R/3 Distributed Sales and Distribution (SD) Benchmark on 4-way, Intel-based servers. The IBM Netfinity 7000 **achieved 1,200 SD Benchmark users** with 1.75 seconds average dialog response time on Oracle 7.3 and Microsoft Windows NT Server 4.0. This represents a throughput of 123,000 fully processed order line items per hour. This record surpasses all previous competitive results using four Intel Pentium Pro 200MHz processors.

The environment for this benchmark was a 4-way Pentium Pro 200MHz IBM Netfinity 7000 database server with 2048MB of main memory and 234GB of total disk space running Windows NT Server 4.0, Oracle 7.3 and SAP's R/3 Rel. 3.1G. The measured throughput was 368,000 dialog steps per hour with an average CPU utilization of 96 percent for the database server. Thirteen 4-way Pentium Pro 200MHz IBM Netfinity 7000 systems running Windows NT 4.0 with Service Pack 3 were used as application servers with an average CPU utilization of 59 percent. Nine of

these IBM Netfinity 7000 systems acted as dialog servers, three as update servers and one as a message/enqueue server. The database server and all application servers ran Microsoft Windows NT Version 4.0 (with Service Pack 3) operating system. This benchmark fully complies with SAP's issued benchmark regulations and has been audited and certified by SAP. Details can be obtained from IBM and SAP AG.<sup>3</sup>

These benchmarks were completed at IBM's Server Performance Lab in Research Triangle Park, NC, by IBM engineers.

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

In recently conducted measurements, the IBM Netfinity 7000 system set a new performance record for running SAP R/3 Centralized SD Benchmark on 4-way, Intel-based servers. The IBM Netfinity 7000 **achieved 130 SD Benchmark users** with 1.73 seconds average response time on Oracle 7.3 and Microsoft Windows NT Server 4.0. This represents a throughput of 13,300 fully processed order line items per hour.

The environment for this benchmark was a 4-way Pentium Pro 200MHz IBM Netfinity 7000 database server with 1048MB of main memory and 162GB of total disk space running Windows NT Server 4.0, Oracle 7.3 and SAP's R/3 Rel. 3.1G. The measured throughput was 39,900 dialog steps per hour with an average CPU utilization of 99 percent for the central server. This benchmark fully complies with SAP's issued benchmark regulations and has been audited and certified by SAP. Details can be obtained from IBM and SAP AG.

The benchmark was completed at IBM's Server Performance Lab in Research Triangle Park, NC, by IBM engineers.

### **Competitive Results**

The tables below summarize some of the published results for the SAP R/3 Distributed and Centralized SD Benchmarks run on 4-way, Intel-based systems. (For a complete list of SAP R/3 SD Benchmark results, go to SAP Web site at the URL provided below.)

The results in the table are provided for comparison. All competitive results shown are based on the benchmark measurements conducted by the respective companies. IBM did not test or in any way verify the test results obtained by these companies. The configuration of the server under test as well as the test environment may vary. Readers are encouraged to examine the companies' audited disclosure reports for details concerning

---

<sup>3</sup> Configuration details are provided in a report that is available at <http://www.ibm.com/pc/us/techlink/srvperf.html>.

the server configuration and methodology used to obtain the published results.<sup>4</sup>

	<b>Distributed SD Users</b>	<b>Response Time</b>	<b>Certification Number*</b>	<b>Software Platform</b>	<b># of Servers</b>
<b>IBM Netfinity 7000</b>	<b>1,200</b>	<b>1.75</b>	<b>#1997030</b>	<b>NT 4.0 Oracle 7.3 R/3 3.1g</b>	<b>14</b>
<b>IBM Netfinity 7000</b>	<b>1,125</b>	<b>1.77</b>	<b>#1997027</b>	<b>NT 4.0 Oracle 7.3 R/3 3.1g</b>	<b>14</b>
<b>Compaq ProLiant 5000</b>	1,116	1.92	*	NT 4.0 Informix 7.2 R/3 3.0e	14
<b>Compaq ProLiant 7000</b>	1,011	1.84	#1997021	NT 4.0 SQL Server 6.5 R/3 3.0f	14
<b>DG Aviion 4900</b>	936	1.91	*	NT 3.51 Oracle 7.2 R/3 3.0d	10
<b>Compaq ProLiant 7000</b>	918	1.76	#1997021	NT 4.0 SQL Server 6.5 R/3 3.0f	14
<b>HP NetServer LX Pro</b>	900	1.89	*	NT 3.51 Oracle 7.2 R/3 3.0c	10
<b>IBM PC Server 704</b>	<b>870</b>	<b>1.97</b>	<b>#1996024</b>	<b>NT 3.51 Oracle 7.2 R/3 3.0d</b>	<b>10</b>
<b>HP NetServer LX Pro</b>	774	1.92	*	NT 4.0 SQL Server 6.5 R/3 3.0c	10

---

<sup>4</sup> Data on competitive products obtained from publicly available information and is subject to change without notice. Contact the manufacturer for the most recent information.

	Centralized SD Users	Response Time	Certification Number*	Software Platform
IBM Netfinity 7000	130	1.73	#1997039	NT 4.0 Oracle 7.3 SAP R/3 3.1g
Bull Escala Series T	121	1.84	#1997041	AIX 4.2.1 Oracle 7.2.3 SAP R/3 3.1g
SNI Primergy 760	118	1.6	*	NT 4.0 Adabas SAP R/3 3.0e
Bull Escala T104	85	1.64	*	AIX Oracle SAP R/3 3.1g

\* For full configuration information, see R/3 Standard Application Benchmarks - Published Results at the SAP Benchmark Web site:  
<http://www.sap.com/products/techno/media/index.htm>

---

#### Methodology for the SD Benchmark

A SAP R/3 installation supports a client/server environment made up of three major software components:

- The Database Server, which handles all database transactions. The entire SAP R/3 software, ABAP/4 application programs, and user data are stored on this server. To date, only one physical database server is supported within an R/3 installation.
- The Application Server, which processes the input from R/3 work processes. It can read data from the local memory or initiate database requests. An R/3 installation can consist of one or more application servers.
- The Presentation Server, which provides the interface to the R/3 application. This software resides on a PC and is generally a Windows-based graphical user interface (GUI).

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 SAP R/3 benchmark consists of executing scripts that simulate a user on the system. Each script contains a common scenario for a given type of user working on a predefined set of data provided by SAP for benchmarking activities.



Typical activities performed by an SD user include creating an order, setting up its delivery, displaying the customer order, changing the delivery date, listing orders already present in the systems, and finally, creating an invoice. Each specific task represents a dialog step within the R/3 system. (A dialog step is a SAP screen change and all its entry criteria.) A 10-second “think time” is included between each dialog step. The SD benchmark consists of 15 distinct dialog steps (not including logging on and off, since this only occurs once during the benchmark run).

At the start of the benchmark run, the user can define how many times to loop through the 15 dialog steps. The goal is to keep a constant peak load for at least 15 minutes in order to produce statistically significant results.

---

## NotesBench

*Performance results obtained using NotesBench 4.6 on a system configured with RAID-5 should not be compared with results from previous versions of NotesBench, which used RAID-0. While RAID-5 implementations provide substantial data redundancy, the increased I/O overhead consumes disk subsystem bandwidth, which could otherwise be used to support additional users.*

For the measurements conducted with NotesBench 4.51 (RAID-0), the IBM Netfinity 7000 server was configured with three 200MHz Pentium Pro processors (with 1MB of L2 cache each) and 1GB of memory for the Mail-only test; for the Mail and Shared Database test, the IBM Netfinity 7000 server was configured with a fourth processor. For both configurations, the IBM Netfinity 7000 ran Lotus Domino Server Release 4.51 on Microsoft Windows NT Server Version 4.0 with Service Pack 3.

For the measurements conducted with NotesBench 4.6 (RAID-5), the IBM Netfinity 7000 server was configured with four 200MHz Pentium Pro processors (with 1MB of L2 cache each) and 2GB of memory for the Mail-only test. The IBM Netfinity 7000 ran Lotus Domino Server Release 4.6 on Microsoft Windows NT Server Version 4.0 with Service Pack 3.

---

## Results Summary

### NotesBench 4.51

The IBM Netfinity 7000 system **demonstrated leadership performance and price/performance as a Lotus Domino 4.51 Server**. The following table provides the results for the IBM Netfinity 7000 running two NotesBench workloads: Mail-only and Mail and Shared Database (MailDB).<sup>5</sup>

Test Script	Maximum Users	NotesMark (tpm)	Response Time (sec)	\$/User	\$/NotesMark
Mail-Only	3,500	4,623	0.312	\$14.92	\$11.30
MailDB	2,900	6,294	1.042	\$19.82	\$9.13

NotesBench provides an objective method for evaluating the performance of different platforms running Lotus Domino Server. NotesBench generates a transactions-per-minute (tpm) throughput metric, called a NotesMark, for each test, along with a value for

---

<sup>5</sup> The NotesBench Disclosure Report for the IBM Netfinity 7000 was audited in September 1997 by KMDS Technical Associates, Inc., and was approved for publication. Highlights from the report are presented here.

the maximum capacity (number of users) supported, and the average response time. The price/performance results are derived using a total system cost of \$52,227 for the Mail-only test and \$57,472 for the MailDB test.<sup>6</sup>

#### NotesBench 4.6

The IBM Netfinity 7000 system **demonstrated leadership performance as a Lotus Domino 4.6 Server**. The following table provides the results for the IBM Netfinity 7000 running the Mail-Only workload as a partitioned server.<sup>7</sup> The price/performance results are derived using a total system cost of \$96,423.

Test Script	Maximum Users	NotesMark (tpm)	Response Time (sec)	\$/User	\$/NotesMark
Mail-Only	4,650	6,071	1.453	\$20.73	\$15.88

#### Competitive Results

The table below summarizes some of the more recently published results for the Mail and MailDB tests that have been run using Domino Server 4.5x and Windows NT Server 4.0. At the time of publication of this report, there were no competitive results available for NotesBench 4.6 running on RAID-5 configurations.

The results are provided for comparison. All competitive results shown are based on the benchmark measurements conducted by the respective companies. IBM did not test or in any way verify the test results obtained by these companies. The configuration of the server under test as well as the test environment may vary. Readers are encouraged to examine the companies' audited disclosure reports for details concerning the server configuration and methodology used to obtain the published results.<sup>8</sup>

---

<sup>6</sup> The price/performance results are based on IBM's Estimated Reseller Price. IBM resellers set their own prices, and actual prices may vary.

<sup>7</sup> The NotesBench Disclosure Report for the IBM Netfinity 7000 was audited in January 1998 by KMDS Technical Associates, Inc., and was approved for publication. Highlights from the report are presented here.

<sup>8</sup> Data on competitive products obtained from publicly available information and is subject to change without notice. Contact the manufacturer for the most recent information.

	Maximum Users	NotesMark (tpm)	\$ / User	\$ / NotesMark
	Mail / MailDB	Mail / MailDB	Mail / MailDB	Mail / MailDB
<b>IBM Netfinity 7000</b> (3 - 200MHz Pentium Pro / Mail, 4 - 200MHz Pentium Pro - MailDB, 1GB) Sept. 1997	<b>3,500 / 2,900</b>	<b>4,623 / 6,294</b>	<b>\$14.92 / \$19.82</b>	<b>\$11.30 / \$9.13</b>
<b>DEC AlphaServer 4100 5/466</b> (2 - 466MHz Alpha, 1GB) Feb. 1997	3,350 / NA	4,230 / NA	\$31.07 / NA	\$24.61 / NA
<b>IBM PC Server 704</b> (3 - 200MHz Pentium Pro, 1GB) July 1997	<b>3,300 / 2,300</b>	<b>4,331 / 5,242</b>	<b>\$16.28 / \$23.36</b>	<b>\$12.41 / \$10.25</b>
<b>HP NetServer LX Pro</b> (4 - 200MHz Pentium Pro, 1GB) July 1997	3,175 / NA	4,061 / NA	\$19.81 / NA	\$15.49 / NA
<b>DEC AlphaServer 4100 5/400</b> (2 - 400MHz Alpha, 1GB) Jan. 1997	3,100 / NA	4,050 / NA	\$29.38 / NA	\$22.49 / NA
<b>IBM PC Server 704</b> (2 - 200MHz Pentium Pro, 1GB) June 1997	<b>2,850 / 2,050</b>	<b>3,761 / 4,697</b>	<b>\$15.18 / \$21.11</b>	<b>\$11.51 / \$9.21</b>
<b>DEC AlphaServer 1000A 5/500</b> (1 - 500MHz Alpha, 768MB) June 1997	2,200 / 1,800	2,878 / 3,904	\$25.78 / \$31.50	\$19.71 / \$14.52
<b>Compaq ProLiant 5000</b> (1 - 200MHz Pentium Pro, 512MB) June 1997	1,920 / 1,440	2,546 / 3,250	\$19.97 / \$26.63	\$15.06 / \$11.80
<b>DEC AlphaServer 800 5/400</b> (1 - 400MHz Alpha, 576MB) April 1997	1,800 / 1,450	2,479 / 3,317	\$23.11 / \$27.37	\$16.78 / \$11.96

---

## Measurement Methodology

### NotesBench 4.51

For the Mail-only test, the IBM Netfinity 7000 system under test was configured with three 200MHz Pentium Pro processors (1MB of L2 cache per processor); 1GB of memory; ten Netfinity EXP10 4.51GB 10K Wide Ultra SCSI hard disk drives configured as a RAID-0 array, using an IBM ServeRAID II Ultra SCSI Adapter; and two IBM EtherJet 100/10 PCI Adapters. For the MailDB test, a fourth processor was added.

In addition to the IBM Netfinity 7000 system under test, the benchmarked configurations used three IBM PC Server 720 systems destination servers, 14 client driver systems for Mail and 16 for MailDB, and one controller system. All systems were connected via two 100Mbps Ethernet LAN segments, using the TCP/IP network protocol.

An IBM PC 350 computer was used as the source driver (parent) system; 15 IBM PC 350 computers were used as the client driver (child) systems. Three IBM PC Server 720 systems were used as

destination servers. Destination mail addresses were distributed across these three destination servers.

To ensure that all results were reproducible, each test was run at least two times, and the results were compared for consistency.

#### **NotesBench 4.6**

The IBM Netfinity 7000 was configured as a partitioned server with four 200MHz Pentium Pro processors, 2GB of memory, three RAID-5 arrays with eight 4.51GB 10K rpm hard disk drives per array for full data redundancy. The server under test supported a Mail-only workload of 4,650 active mail users (see price/performance results<sup>9</sup> above). Three Domino 4.6 Partitioned Servers ran on the IBM Netfinity 7000 system, using identical copies of the Public Name & Address Book. Mail messages generated by users connected to the three partitioned servers were all routed to three destination servers.

In addition to the server under test, the benchmarked client/server network used three destination servers, 27 client driver systems, and one controller client system. All systems were connected via three 100Mbps Ethernet LAN segments, using the TCP/IP network protocol.

To ensure that all results were reproducible, the test was run at least two times, and the results were compared and found to be consistent.

---

### **Measurement Analysis**

The test results demonstrate that an IBM Netfinity 7000 configured as described within the respective NotesBench Disclosure Reports can support:

- 3,500 NotesBench 4.51 Mail-only users and 2,900 MailDB users with a response time well below what the test criteria permit
- 4,650 NotesBench 4.6 Mail-only users in a RAID-5 partitioned server environment with a response time well below what the test criteria permit.

These results are based on running the IBM Netfinity 7000 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

---

<sup>9</sup> The price/performance results are based on IBM's Estimated Reseller Price. IBM resellers set their own prices, and actual prices may vary.

configuration of that hardware and appropriately tuning the operating system and Domino software.

A detailed analysis of the measurements for the all tests conducted with the IBM Netfinity 7000, along with details of the benchmarked configuration, is provided in the NotesBench Disclosure Reports for the system. The reports are available at:

- IBM's World Wide Web site at  
**<http://www.pc.ibm.com/us/techlink/srvperf.html>**
- Lotus' NotesBench World Wide Web site at  
**<http://www.notesbench.org>**

## ServerBench 4.0


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

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

### Results Summary

The peak level of throughput achieved by the IBM Netfinity 7000 200MHz/1MB was:

- **63 percent higher** than the HP NetServer LX Pro
- **16 percent higher** than the Compaq ProLiant 6000
- **14 percent higher** than the IBM 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 4.0 SYS\_60.TST workload, which includes the client/server, processor, server/client, random read, and random write requests typically made in a client/server computing environment.

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

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

The SYS\_60.TST test suite contains a total of 16 test mixes. Measurements of transactions per second (TPS) were recorded as a weighted harmonic mean of the total TPS obtained by all clients in each test mix as clients were added. Clients were added in increments of four as follows: 1, 4, 8, 12, 16, 20, 24, 28, 32, 36, 40, 44, 48, 52, 56, 60.

---

**Measurement Analysis**

ServerBench's server application on Windows NT provides up to 47 service threads with 60 clients, plus one thread for each server processor. For this test, the servers were configured with 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 (e.g., 512MB) in order to produce a meaningful result. When workload



increases gradually, the processor subsystem (processor and system memory) provides adequate service to all requests by caching them in the system memory, which is the primary factor affecting the TPS throughput.

As workload continued to increase (i.e., more clients joined the test mixes), system memory was exhausted, and the server had to rely on the disk subsystem for virtual memory. When this happened, the bottleneck shifted to the disk subsystem, and the application became disk-bound. Running ServerBench with Windows NT may result in a low cache-hit ratio because some NT system threads (e.g., cache manager's lazy writer thread, memory manager's mapped page writer thread) will automatically move some mapped memory into paged files. If a client happens to request that paged-out data again, a cache-hit-miss will result.

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

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

In a normal client/server environment, it is reasonable to install Windows NT Server as an application server. Ziff-Davis' setup instructions recommend that the server be set in file server mode to achieve better throughput.

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

## WebBench 1.0

Ziff-Davis' WebBench 1.0 system test suite NT\_SIMPLE\_CGI20\_V1.TST was used to measure the performance of the IBM Netfinity 7000 system and the competitors' systems as four-way Web servers running Microsoft Internet Information Server 3.0 on Windows NT Server 4.0 operating system.


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

## Results Summary

### Throughput

At the peak of 12 WebBench clients, the IBM Netfinity 7000 200MHz/1MB system delivered:


- **26 percent more throughput** than the HP NetServer LX Pro
- **11 percent more throughput** than the Compaq ProLiant 6000
- **11 percent more throughput** than the IBM Netfinity 7000



**Requests per Second**

At the peak of 12 WebBench clients, the IBM Netfinity 7000 200MHz/1MB serviced:

- **22 percent more requests per second** than the HP NetServer LX Pro
- **8 percent more requests per second** than the Compaq ProLiant 6000
- **8 percent more throughput** than the IBM Netfinity 7000



---

**Measurement  
Methodology**

Two kinds of performance measurements were conducted in this standard test suite:

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

The system test suite was performed using four 100Mbps Ethernet network segments with a total of 60 IBM PC 750 systems as client workstations attached to the server. Each workstation ran Windows NT Workstation 4.0 and executed the WebBench 1.0 NT\_SIMPLE\_CGI20\_V1.TST workload, which includes HTML pages requests and Common Gateway Interface (CGI) requests, two of the primary functions of a web server. Each client randomly issued these requests to the web server according to a workload file that specifies each request a client makes and how frequently the client makes that request. The workload file associates a request percentage with each HTTP request and CGI request. The request percentage tells the client the number of requests it issues during a mix and what the percentage of requests should be for that particular mix. If all clients requested the same file at the same time, the results could be adversely affected; therefore, each client request access patterns are randomized.

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

The NT\_SIMPLE\_CGI20\_V1.TST test suite contains a total of 15 mixes. Each mix uses 30 seconds as Ramp-Up time, 30 seconds as Ramp-Down time, during which periods measurements were not done. Each mix ran for 300 seconds.

After the 15 mixes in the test suite were finished, WebBench created two graphs: one that plots the requests per second against each mix and another that plots the throughput against each mix. Also supplied is the amount of time how it took the clients to connect to the server and receive data from the server, and the number of connections per second made by each client.

---

## Measurement Analysis

In a typical two-tier Internet/Intranet environment, the web browser is usually the user front end that makes requests to the web server. The web server functions either as a large HTML document store directly returning the HTML documents to the browser or as a back-end logic unit building a dynamic HTML document based on calculation of input fields from the web browser. In a three-tier Internet/Intranet environment, the web server usually functions as a 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.

Because each WebBench client generally stresses the server as much as several actual users do, you can run the test suites with a relatively small number of clients and still get an accurate measure of a server's performance.

To get a valid measure of the server's performance, ensure that the requests-per-second score and the throughput score reach a point where they flatten out. This "flattening out" indicates that the server has been saturated, or fully loaded. In our test, adding clients increased the total requests-per-second and throughput scores. The curves increased very sharply from 1 to 8 clients, peaked at 12 clients, and then flattened out, indicating that the server had reached its saturation point. Ideally, the curves after saturation point should remain at the same level where the server's resources (CPU, memory subsystem, disk subsystem) are used optimally. However, due to heavy network traffic and the need to balance each client request load, the curve may dip slightly, reducing the server load.

Our data shows that the saturation point occurred at 12 clients for all three servers. The scores at the saturation point, which are repeatable, indicate how well the server performs. The curves after 12 clients dipped and varied sometimes widely for each run for the reason mentioned above.

File sharing mode/application server mode is a parameter in NT that allows us to tune the server's memory manager policy to favor either the file cache or the process' working set. In File sharing mode, the server does not trim file cache as often as it does in application server mode, and it allocates more memory resources to file cache. Because the system is configured with 512MB memory, we tuned the system in file sharing mode, which did not provide a higher score at the saturation point, but rather a smoother curve than was achieved in application server mode after the saturation point.

**NetBench 5.01**

NetBench 5.01 system test suite DM\_4GB\_TST was used to measure the performance of the IBM Netfinity 7000 system and the competitors' systems as single-processor file servers running Novell NetWare 4.11.


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

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

**Results Summary**

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 7000 200MHz/1MB system provided network clients with:

- **59 percent more throughput** than the HP NetServer LX Pro
- **10 percent more throughput** than the Compaq ProLiant 6000
- **8 percent more throughput** than the IBM Netfinity 7000



---

**Measurement  
Methodology**

The Disk Mix test suite was performed using two 100Mbps Ethernet network segments with a total of 60 IBM PC 350 systems as client workstations attached to the server. Each workstation ran Windows for Workgroups Version 3.11 and executed the NetBench 5.01 Disk Mix workload, which is based on leading Windows applications.

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

---

**Measurement Analysis**

The NetBench 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 in this configuration of 60 clients, we found that using 128MB (eight 16MB memory modules to enable four-way memory interleaving) of memory accounted for a high percentage of cache-hits. Also, the entire server was stressed as the workload increased. The reason is that the 100Mbps network adapter provided sufficient bandwidth to allow the server's subsystems (i.e., memory, disk and processor complex) to be saturated. This is important because in most production environments, the number of users is dynamic, and the server bottleneck may change several times daily. Showing both the network adapter and disk subsystem bottlenecks provides more useful information about how the server will perform in production environments.

NetBench is designed to measure the ability of a file server to respond to client file open, close, read and write I/O operations. Our measurements show that increasing the processor speed alone provides little, if any, increase in throughput for most file server environments. However, additional CPU power can provide improvements for servers that are supporting a very large number of users.

## Server Configurations

### ServerBench 4.0

Features	IBM Netfinity 7000 200MHz/1MB	IBM Netfinity 7000 200MHz/512KB	Compaq ProLiant 6000 200MHz/512KB	HP NetServer LX Pro 200MHz/512KB
<b>Processor</b>	Four 200MHz Pentium Pro	Four 200MHz Pentium Pro	Four 200Hz Pentium Pro	Four 200MHz Pentium Pro
<b>Memory</b>	512MB ECC	512MB ECC	512MB ECC	512MB ECC
<b>L2 Cache</b>	1MB (Write-Back)	512KB (Write-Back)	512KB (Write-Back)	512KB (Write-Back)
<b>RAID Level</b>	Non-Array	Non-Array	Non-Array	Non-Array
<b>Disk Drives</b>	Four IBM Netfinity EXP10 4.51GB 10K Wide Ultra SCSI Drives (10K rpm) <sup>10</sup>	Four IBM 4.51GB Wide Ultra SCSI Drives (7200 rpm)	Four Compaq 4.3GB Wide Ultra SCSI Drives (7200 rpm)	Four HP 4.2GB Fast/Wide SCSI-2 Drives (7200 rpm)
<b>Disk Drive Adapter</b>	One Wide Ultra SCSI PCI Bus on Planar	One Wide Ultra SCSI PCI Bus on Planar	One Wide Ultra SCSI PCI Bus on Planar	One Ultra SCSI-2 Fast/Wide PCI Bus on Planar
<b>Disk Driver</b>	AIC78XX.SYS	AIC78XX.SYS	SYMC810.SYS	AIC78XX.SYS
<b>Network Adapters</b>	Four IBM 100/10 PCI Ethernet Adapters	Four IBM 100/10 PCI Ethernet Adapters	Four Netelligent 10/100 TX PCI Ethernet Adapters	Four HP DeskDirect LAN Adapters
<b>Bus</b>	PCI	PCI	PCI	PCI
<b>Network Driver</b>	E100B.SYS	E100B.SYS	NETFLX3.SYS	E100B.SYS
<b>Network Operating System</b>	Windows NT Server 4.0 with Service Pack 3	Windows NT Server 4.0 with Service Pack 3	Windows NT Server 4.0 with Service Pack 3	Windows NT Server 4.0 with Service Pack 3
<b>System Partition Size</b>	16GB	16GB	16GB	16GB
<b>File System</b>	NTFS	NTFS	NTFS	NTFS
<b>Allocation Unit Size</b>	Predefined Default	Predefined Default	Predefined Default	Predefined Default
<b>Test Suite</b>	ServerBench 4.0 SYS_60.TST	ServerBench 4.0 SYS_60.TST	ServerBench 4.0 SYS_60.TST	ServerBench 4.0 SYS_60.TST

<sup>10</sup> For all Ziff-Davis benchmarks run with the IBM Netfinity 7000 1MB cache model, the 10K rpm drives were attached via a Netfinity EXP10 Rack Storage Expansion Enclosure.

---

**WebBench 1.0**

<b>Features</b>	<b>IBM Netfinity 7000 200MHz/1MB</b>	<b>IBM Netfinity 7000 200MHz/512KB</b>	<b>Compaq ProLiant 6000 200MHz/512KB</b>	<b>HP NetServer LX Pro 200MHz/512KB</b>
<b>Processor</b>	Four 200MHz Pentium Pro	Four 200MHz Pentium Pro	Four 200Hz Pentium Pro	Four 200MHz Pentium Pro
<b>Memory</b>	512MB ECC	512MB ECC	512MB ECC	512MB ECC
<b>L2 Cache</b>	1MB (Write-Back)	512KB (Write-Back)	512KB (Write-Back)	512KB (Write-Back)
<b>RAID Level</b>	Non-Array	Non-Array	Non-Array	Non-Array
<b>Disk Drives</b>	Six IBM Netfinity EXP10 4.51GB 10K Wide Ultra SCSI Drives (10K rpm)	Six IBM 4.51GB Wide Ultra SCSI Drives (7200 rpm)	Six Compaq 4.3GB Wide Ultra SCSI Drives (7200 rpm)	Six HP 4.2GB Fast/Wide SCSI-2 Drives (7200 rpm)
<b>Disk Drive Adapter</b>	One Wide Ultra SCSI PCI Bus on Planar	One Wide Ultra SCSI PCI Bus on Planar	One Wide Ultra SCSI PCI Bus on Planar	One Ultra SCSI-2 Fast/Wide PCI Bus on Planar
<b>Disk Driver</b>	AIC78XX.SYS	AIC78XX.SYS	SYMC810.SYS	AIC78XX.SYS
<b>Network Adapters</b>	Four IBM 100/10 PCI Ethernet Adapters	Four IBM 100/10 PCI Ethernet Adapters	Four Netelligent 10/100 TX PCI Ethernet Adapters	Four HP DeskDirect LAN Adapters
<b>Bus</b>	PCI	PCI	PCI	PCI
<b>Network Driver</b>	E100B.SYS	E100B.SYS	NETFLX3.SYS	E100B.SYS
<b>Network Operating System</b>	Windows NT Server 4.0 with Service Pack 3	Windows NT Server 4.0 with Service Pack 3	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	1GB	1GB
<b>File Compression</b>	Off	Off	Off	Off
<b>File System</b>	NTFS	NTFS	NTFS	NTFS
<b>Allocation Unit Size</b>	Predefined Default	Predefined Default	Predefined Default	Predefined Default
<b>Test Suite</b>	NT_SIMPLE_CGI20_V 1.TST	NT_SIMPLE_CGI20_V 1.TST	NT_SIMPLE_CGI20_V 1.TST	NT_SIMPLE_CGI20_V 1.TST
<b>Web Server</b>	Microsoft Internet Information Server 3.0	Microsoft Internet Information Server 3.0	Microsoft Internet Information Server 3.0	Microsoft Internet Information Server 3.0

---

**NetBench 5.01**

	<b>IBM Netfinity 7000 200MHz/1MB</b>	<b>IBM Netfinity 7000 200MHz/512KB</b>	<b>Compaq ProLiant 6000 200MHz/512KB</b>	<b>HP NetServer LX Pro 200MHz/512KB</b>
<b>Processor</b>	One 200MHz Pentium Pro	One 200MHz Pentium Pro	One 200Hz Pentium Pro	One 200MHz Pentium Pro
<b>Memory</b>	256MB ECC	256MB ECC	256MB ECC	256MB ECC
<b>L2 Cache</b>	1MB (Write-Back)	512KB (Write-Back)	512KB (Write-Back)	512KB (Write-Back)
<b>RAID Level</b>	Non-Array	Non-Array	Non-Array	Non-Array
<b>Disk Drives</b>	Six IBM Netfinity EXP10 4.51GB 10K Wide Ultra SCSI Drives (10K rpm)	Six IBM 4.51GB Wide Ultra SCSI Drives (7200 rpm)	Six Compaq 4.3GB Wide Ultra SCSI Drives (7200 rpm)	Six HP 4.2GB Fast/Wide SCSI-2 Drives (7200 rpm)
<b>Disk Drive Adapter</b>	One Wide Ultra SCSI PCI Bus on Planar	One Wide Ultra SCSI PCI Bus on Planar	One Wide Ultra SCSI PCI Bus on Planar	One Ultra SCSI-2 Fast/Wide PCI Bus on Planar
<b>Disk Driver</b>	AIC7870.DSK V2.13	AIC7870.DSK V2.13	CPQS710.DSK V2.11	AIC7870.DSK V2.13
<b>Network Adapters</b>	Two Intel EtherExpress PRO/100 Smart Adapters	Two Intel EtherExpress PRO/100 Smart Adapters	Two Netelligent 10/100 TX PCI Ethernet Adapters	Two HP DeskDirect LAN Adapters
<b>Bus</b>	PCI	PCI	PCI	PCI
<b>Network Driver</b>	E100S.LAN V2.93	E100S.LAN V2.93	CPQNF3.LAN V2.23	HPTX.LAN V1.47
<b>Network Operating System</b>	NetWare 4.11	NetWare 4.11	NetWare 4.11	NetWare 4.11
<b>NetWare Volume Block Size</b>	32KB	32KB	32KB	32KB
<b>File Compression</b>	Off	Off	Off	Off
<b>Block Allocation</b>	On	On	On	On
<b>Data Migration</b>	Off	Off	Off	Off
<b>Disk Mix</b>	NB5.01/Windows for Workgroup Clients	NB5.01/Windows for Workgroup Clients	NB5.01/Windows for Workgroup Clients	NB5.01/Windows for Workgroup Clients

---

## Test Disclosure Information

---

### ServerBench 4.0

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

**Version:** 4.0

#### 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 non-IBM 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>	30
<b>CPU / Memory</b>	166MHz Pentium / 64MB
<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 1.0

The measurements were conducted using Ziff-Davis' WebBench 1.0 running the NT\_SIMPLE\_CGI20\_V1.TST test suite with Windows NT Workstation 4.0 clients as described below:

**Version:** 1.0

**Mixes**

- NT\_SIMPLE\_CGI20\_V1.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

All non-IBM 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 / 64MB
<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 for Workgroup 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 IWSP3 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 = 100
- Dirty Directory Cache Delay Time = 10
- Maximum Concurrent Disk Cache Write = 100
- Maximum Directory Cache Buffers = 700
- Minimum Directory Cache Buffers = 150
- Minimum File Cache Buffers = 150
- Maximum Number of Directory Handles = 30
- Dirty Disk Cache Delay Time = 5
- Directory Cache Allocation Wait Time = 2.2 seconds
- Directory Cache Buffer Non-Referenced Delay = 30
- Maximum Interrupt Events = 100,000

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 non-IBM 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>	30
<b>CPU</b>	133MHz Pentium
<b>Network Adapter</b>	IBM 100/10 PCI Ethernet Adapter (Bus 0)
<b>Software</b>	IBM DOS 6.3
	NetWare DOS Requester
	LSL.COM (8-3-95)
	E100BODI (5-21-96)
	IPXODI (8-8-95)
	VLM.EXE (11-08-94)
<b>Cache</b>	L2 = 256KB
<b>Controller Software</b>	PC-DOS Version 6.3 Microsoft Windows for Workgroups 3.11

### Clients NET.CFG

- Checksum = 0
- Large Internet Packet = On
- PB Buffers = 10
- PBurst Read Windows Size = 64
- PBurst Write Windows Size = 64
- Cache Buffers = 64
- Cache Write = On
- Cache Buffers Size = 4096
- True Commit = Off
- Signature Level = 0



---

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, Netfinity and ServerGuide 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 and Pentium are registered trademarks of Intel Corporation.

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

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

Published by the IBM Server Performance Laboratory, IBM Corp. The IBM Server Performance Laboratory publishes white papers and performance reports, including audited disclosures for benchmarks such as TPC-C and NotesBench. These documents are available on the World Wide Web at the following URL: <http://www.pc.ibm.com/us/techlink/srvperf.html>

© 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.

