Performance Report

IBM Netfinity 5500 M10 (450MHz)

Version 1.0
February 1999

IBM ®

Executive Overview

The IBM Netfinity* 5500 M10, which features the Intel** 450MHz[1] Pentium** II Xeon** processor, was announced worldwide in January 1999. These new models, which are two-way SMP-capable servers, expand IBM's Netfinity line of midrange servers, providing solid network performance for business-critical applications.

The 450MHz system with 1MB of L2 cache (Model 8661-4RY) was evaluated using the following Ziff-Davis** benchmarks:

- *ServerBench** Version 4.02*

- *NetBench** Version 5.01*

For comparison, the IBM Netfinity server performance laboratory also conducted measurements with the IBM Netfinity 5500 system configured with a 450MHz Pentium II processor with 512KB of L2 cache.

All results from these benchmarks 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.

For these benchmarks, the IBM Netfinity 5500 M10 was configured with a 450MHz/100MHz Pentium II Xeon processor with 1MB of L2 cache; the Netfinity 5500 was configured with a 450MHz/100MHz Pentium II processor with 512KB of L2 cache. Both systems support two-way symmetrical multiprocessing (SMP) operations.

ServerBench 4.02

ServerBench 4.02 was used to measure the performance of the IBM Netfinity 5500 M10 and the IBM Netfinity 5500 as dual-processor application servers running Microsoft** Windows** NT Server 4.0 and providing services to Windows NT** Workstation 4.0 clients.

The IBM Netfinity 5500 M10 achieved a peak level of transactions per second in a RAID-0 configuration that was **33 percent higher** than the IBM Netfinity 5500.

The IBM Netfinity 5500 M10 achieved a peak level of transactions per second in a RAID-5 configuration that was **31 percent higher** than the IBM Netfinity 5500.

NetBench 5.01

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 5500 M10 delivered a level of network throughput in a RAID-0 configuration that was **11 percent higher** than that of the IBM Netfinity 5500**.**

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 5500 M10 delivered a level of network throughput in a RAID-5 configuration that was **16 percent higher** than that of the IBM Netfinity 5500**.**

Test Environments and Results

ServerBench 4.02

The ServerBench 4.02 test suite SYS_60.TST was used to measure the performance of the IBM Netfinity 5500 M10 and the IBM Netfinity 5500 systems, configured as two-way, application servers running Windows NT Server 4.0.

ServerBench 4.02 provides an overall transactions-per-second (TPS) score showing how well the server handles client requests for a variety of operations involving the server's processors, disk and network subsystems.
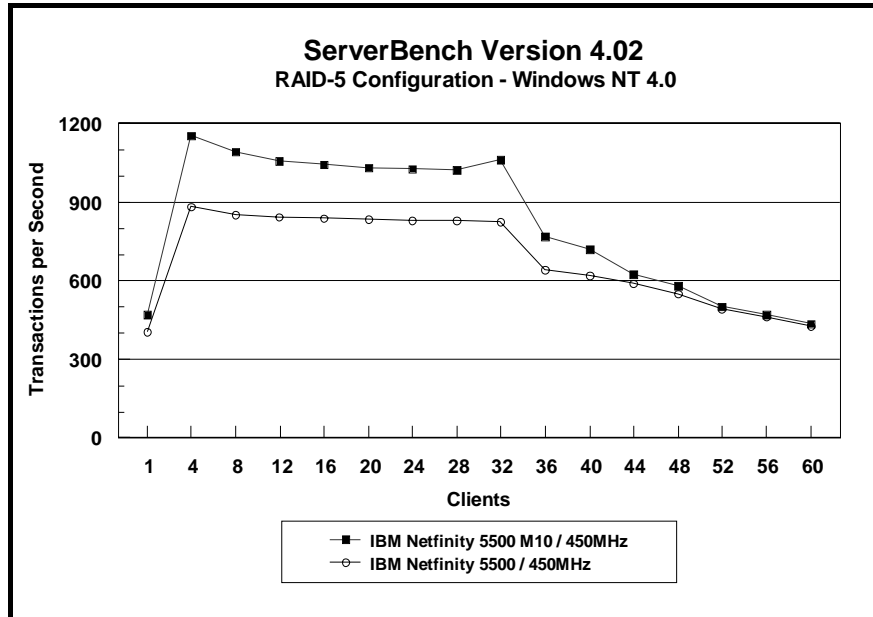
Results Summary

RAID-0 Configuration

The IBM Netfinity 5500 M10 achieved a peak level of transactions per second that was **33 percent higher** than that of the IBM Netfinity 5500.
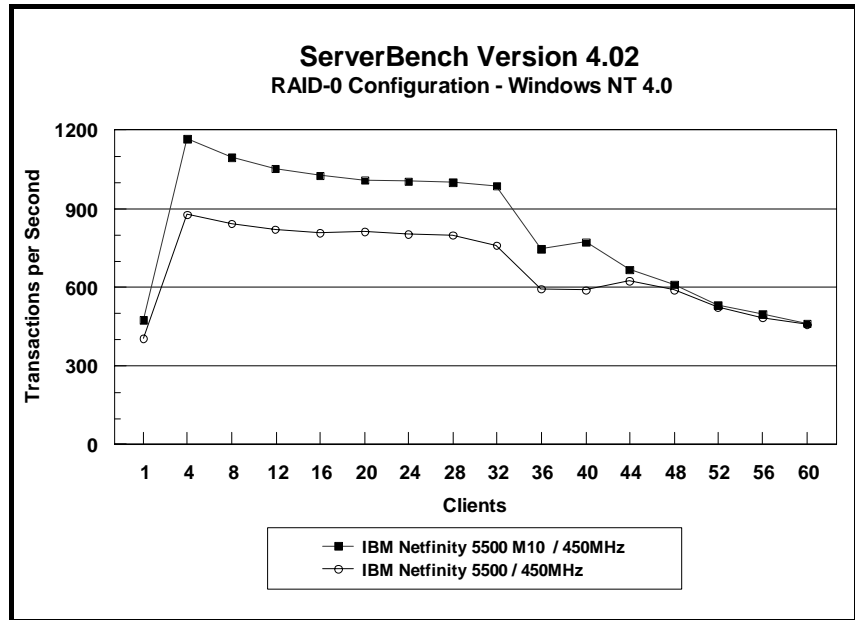
RAID-5 Configuration

The IBM Netfinity 5500 M10 achieved a peak level of transactions per second that was **31 percent higher** than that of the IBM Netfinity 5500.

ServerBench Version 4.02
RAID-5 Configuration - Windows NT 4.0

Transactions per Second

(Chart x-axis: Clients — 1, 4, 8, 12, 16, 20, 24, 28, 32, 36, 40, 44, 48, 52, 56, 60)

- ■ IBM Netfinity 5500 M10 / 450MHz
- ○ IBM Netfinity 5500 / 450MHz

Measurement Methodology

The system test suite was performed using four 100Mbps Ethernet network segments with a total of 60 IBM PC 750 166MHz systems as client workstations attached to the server. Each workstation ran Windows NT 4.0 Workstation and executed the ServerBench 4.02 SYS_60.TST workload, which includes the client/server, processor, server/client, random read, and random write requests typically made in a client/server computing environment. (The default values were used for all NT registry variables. The NT default is 'Max throughput for file sharing'.)

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

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

swapping mapped memory out to paged files if system memory is full.

The SYS_60.TST test suite contains 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 two processors; therefore, 49 service threads were used. A client workstation generates a request for the server to begin the next phase of a mix or to ask the server to perform some operation. The server creates a new service thread and passes that connection with the client to an I/O completion port.

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

ServerBench requires a large amount of system memory 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 these measurements, the application was processor-bound when running from 4 to 8 clients; with more than 40 clients, the application became disk-bound.
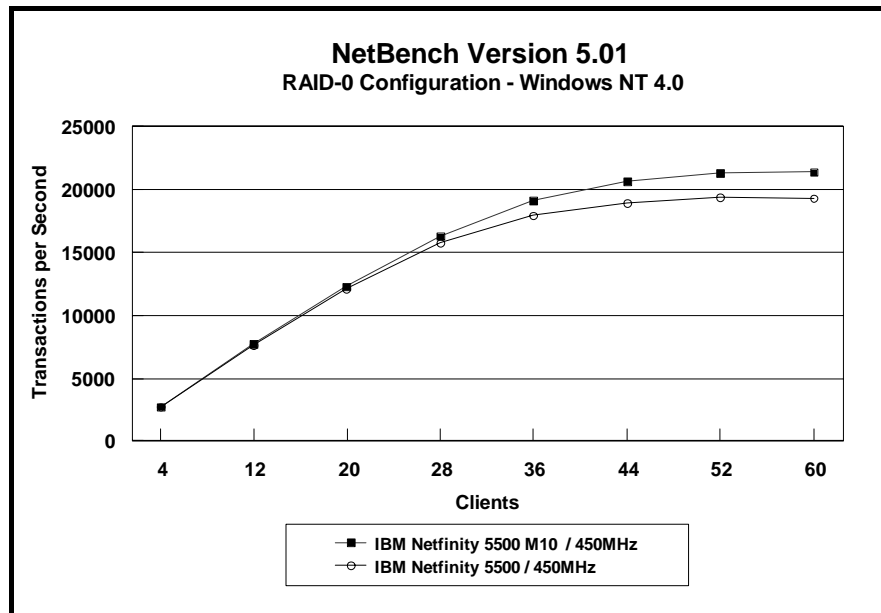
NetBench 5.01

The NetBench 5.01 Disk Mix test suite was used to measure the performance of the IBM Netfinity 5500 M10 and the IBM Netfinity 5500 systems as single-processor file servers running Novell NetWare 4.11 with Service Pack IWSP6. For these measurements, Windows NT Workstation 4.0 clients were used.

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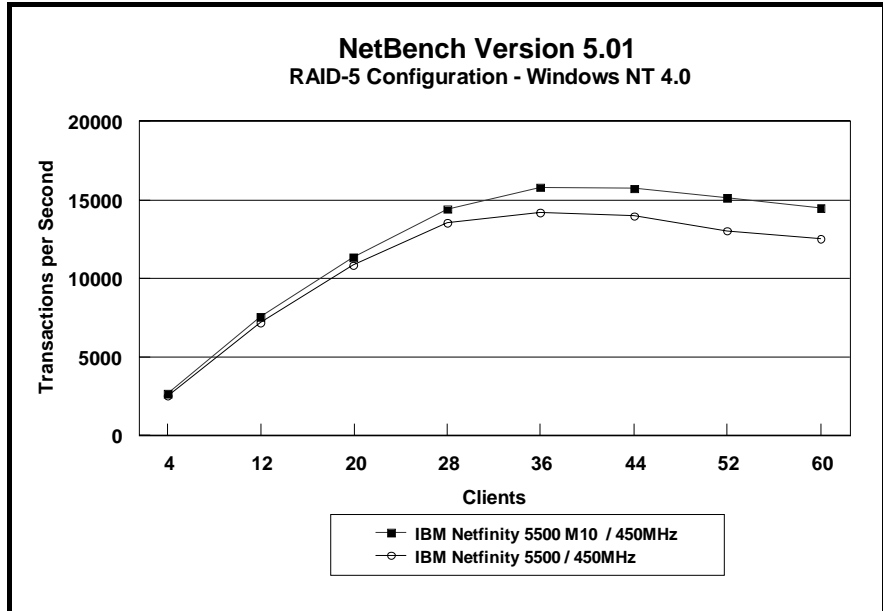
Results Summary

RAID-0 Configuration

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 5500 M10 provided **11 percent more throughput** than the IBM Netfinity 5500.

NetBench Version 5.01
RAID-0 Configuration - Windows NT 4.0

RAID-5 Configuration

Under a high-end workload of 60 NetBench clients, the IBM
Netfinity 5500 M10 provided **16 percent more throughput** than
the IBM Netfinity 5500.

NetBench Version 5.01
RAID-5 Configuration - Windows NT 4.0

A line chart titled "NetBench Version 5.01, RAID-5 Configuration - Windows NT 4.0" showing Transactions per Second (y-axis, 0 to 20000) versus Clients (x-axis: 4, 12, 20, 28, 36, 44, 52, 60). Two series: IBM Netfinity 5500 M10 / 450MHz and IBM Netfinity 5500 / 450MHz.

Legend:
- ■ IBM Netfinity 5500 M10 / 450MHz
- ○ IBM Netfinity 5500 / 450MHz

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 512MB memory in a configuration of 60 clients, the cache-hit ratio was more than 91 percent, and the cache buffer in resource utilization was more than 89 percent. These statistics indicate a

higher level of throughput. Also, the 100Mbps network adapter provided sufficient bandwidth to allow the server's subsystems (i.e., memory, disk and processor complex) to be saturated. This is important because in most production environments, the number of users is dynamic, and the server bottleneck may change several times daily. Showing both the network adapter and disk subsystem bottlenecks provides more useful information about how the server will perform in production environments.

Server Configurations

ServerBench 4.02

RAID Configurations

Features	IBM Netfinity 5500 M10 450MHz/1MB	IBM Netfinity 5500 450MHz/512KB
Processor	Two 450MHz Pentium II Xeon	Two 450MHz Pentium II
Memory	512MB ECC SDRAM	512MB ECC SDRAM
L2 Cache	1MB (Write-Back)	512KB (Write-Back)
RAID Level	RAID-0 and RAID-5	RAID-0 and RAID-5
Disk Drive	Five IBM 9.1GB[2] Wide Ultra SCSI Drives (10K rpm)	Five IBM 9.1GB Wide Ultra SCSI Drives (10K rpm)
Disk Drive Adapter	ServeRAID-3DH Ultra2 SCSI Adapter	ServeRAID-3DH Ultra2 SCSI Adapter
Disk Driver	IPSRAID.SYS	IPSRAID.SYS
Network Adapter	Four IBM EtherJet 100/10 PCI Adapters	Four IBM EtherJet 100/10 PCI Adapters
Bus	PCI	PCI
Network Driver	IBMFENT.SYS	IBMFENT.SYS
Network Operating System	Windows NT Server 4.0 with Service Pack 3	Windows NT Server 4.0 with Service Pack 3
System Partition Size	1GB	1GB
File System	NTFS	NTFS
Allocation Unit Size	Predefined Default	Predefined Default
ServerBench Version / Test Suite	ServerBench 4.02 / SYS_60.TST	ServerBench 4.02 / SYS_60.TST

NetBench 5.01

RAID Configurations

Features	IBM Netfinity 5500 M10 450MHz/1MB	IBM Netfinity 5500 450MHz/512KB
Processor	One 450MHz Pentium II Xeon	One 450MHz Pentium II
Memory	512MB ECC SDRAM	512MB ECC SDRAM
L2 Cache	1MB (Write-Back)	512KB (Write-Back)
RAID Level	RAID-0 and RAID-5	RAID-0 and RAID-5
Disk Drive	Five IBM 9.1GB Wide Ultra SCSI Drives (10K rpm)	Five IBM 9.1GB Wide Ultra SCSI Drives (10K rpm)
Disk Drive Adapter	ServeRAID-3DH Ultra2 SCSI Adapter	ServeRAID-3DH Ultra2 SCSI Adapter
Disk Driver	IPSRAID.HAM V2.81	IPSRAID.HAM V2.81
Network Adapter	Four IBM EtherJet 100/10 PCI Adapters	Four IBM EtherJet 100/10 PCI Adapters
Bus	PCI	PCI
Network Driver	E100B.LAN V3.63	E100B.LAN V3.63
Network Operating System	NetWare 4.11 with IWSP6 Loaded	NetWare 4.11 with IWSP6 Loaded
NetWare Volume Block Size	16KB	16KB
File Compression	Off	Off
Block Allocation	On	On
Data Migration	Off	Off
NetBench Version / Test Suite	NB5.01 / Windows NT Workstation 4.0 Clients / Disk Mix	NB5.01 / Windows NT Workstation 4.0 Clients / Disk Mix

Test Disclosure Information

ServerBench 4.02

The ServerBench measurements were conducted using Ziff-Davis' ServerBench 4.02 running the SYS_60.TST test suite with Windows NT Workstation 4.0 as described below:

Version: ServerBench 4.02

Mixes

- System Test Mixes
- Clients: 1, 4, 8, 12, 16, 20, 24, 28, 32, 36, 40, 44, 48, 52, 56, 60
- Data Segment Size: 16MB
- Segment Access Ratio: 1
- Ramp up: Default setup
- Ramp down: Default setup

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

Testbed Disclosure

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

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

NetBench 5.01

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

Version: NetBench 5.01

Mixes

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

Network Operating System: NetWare 4.11 with IWSP6 loaded

NOS Parameters

- Immediate Purge of Deleted Files = ON
- Enable Disk Read after Write Verify = OFF
- Minimum Packet Receive Buffers = 700
- Maximum Packet Receive Buffers = 1400
- Set NCP Packet Signature Option = 0
- Maximum Physical Receive Packet Size = 1514
- Reserved Buffer Below 16MB = 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 = 150
- Maximum Number of Directory Handles = 30
- Dirty Disk Cache Delay Time = 5
- Directory Cache Buffer Non-Referenced Delay = 15
- Directory Cache Allocation Wait Time = 2.2 seconds
- Maximum Interrupt Events = 1000
- Set Enable File Compression = Off

If clients drop out, set the following:

- Number of Watchdog Packets = 50

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

To monitor the dropping out of clients, set:

- Console Display Watchdog Logouts = On

Testbed Disclosure

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

Network	100Mbps Ethernet
Clients	60
Hubs	3COM Super Stack II Hub 100TX
Clients per Segment	15
CPU / Memory	200MHz Pentium / 16MB
Network Adapter	IBM 100/10 PCI Ethernet Adapter (Bus 0)
Software	Windows NT Workstation 4.0
	Network Service - Client Service for NetWare
	Network Protocol - NWLink IPX/SPX Compatible Teansport -NWLink NetBIOS -TCP/IP Protocol
Cache	L2 = 256KB
Controller Software	Microsoft Windows 95 Novell Loginw 95

Notes

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

[2] When referring to hard disk capacity, GB, or gigabyte, means one thousand million bytes. Total user-accessible capacity may vary depending on operating environment.