
Performance Report

IBM Netfinity 5500 (500MHz)

Version 1.0
March 1999



Executive Overview

*The IBM Netfinity 5500 servers offer new levels of power, scalability and serviceability for creating a reliable foundation for networked, business-critical applications. New models announced worldwide in March 1999 feature single or dual symmetrical multiprocessing (SMP) with the new 500MHz¹ Pentium** III processor, which supports 100MHz operations to memory.*

*The new 500MHz system (Model 8660-61U) was evaluated using Ziff-Davis' benchmarks, ServerBench** Version 4.02 and NetBench** Version 5.01*

For comparison, the IBM Netfinity server performance laboratory also conducted the benchmarks with the IBM Netfinity 5500 configured with a 450MHz Pentium II processor.

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.

ServerBench 4.02

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

The IBM Netfinity 5500 500MHz system achieved a peak level of transactions per second that was **12 and 13 percent higher** than that of the 450MHz system in both RAID-0 and RAID-5 configurations, respectively.

NetBench 5.01

Under the high-end workload of 60 NetBench clients, the IBM Netfinity 5500 500MHz system provided 5 percent higher throughput than the 450MHz system in the RAID-0 configuration and throughput similar to the 450MHz system in the RAID-5 configuration.

Test Environments and Results

ServerBench 4.02


The ServerBench 4.02 system test suite SYS_60.TST was used to measure the performance of the IBM Netfinity 5500 500MHz and 450MHz 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

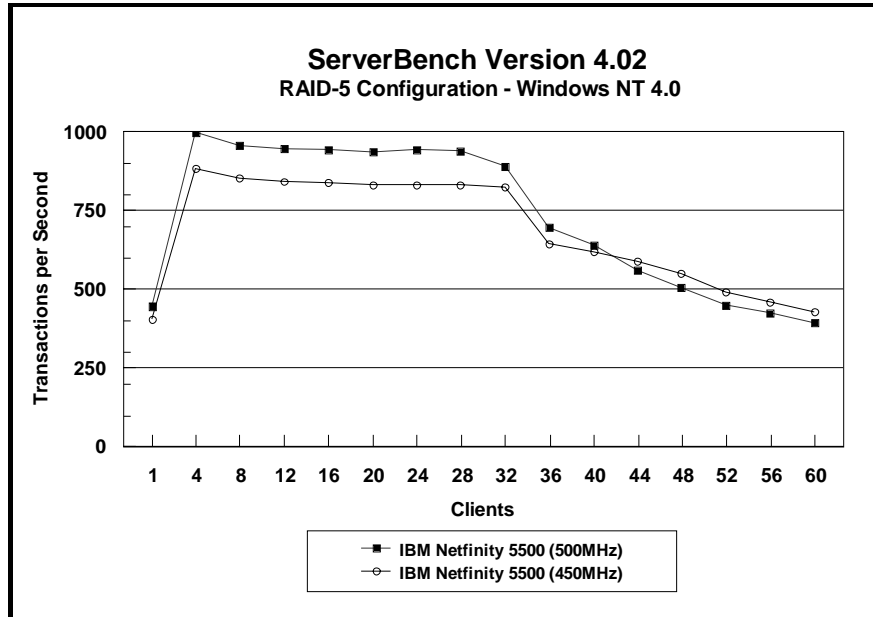
RAID-0 Configuration

The IBM Netfinity 5500 500MHz system achieved a peak level of transactions per second that was **12 percent higher** than that of the 450MHz system.



RAID-5 Configuration

The IBM Netfinity 5500 500MHz system achieved a peak level of transactions per second that was **13 percent higher** than that of the 450MHz system.



Measurement Methodology

The system test suite was performed using four 100Mbps Ethernet network segments with a total of 60 IBM PC 750 166MHz systems as client workstations attached to the server. Each workstation ran Windows NT 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. A higher TPS indicates better performance.

The clients randomly send requests to the server, producing 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 two processors; therefore, a total of 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 our measurements, the application was processor-bound when running from 4 to 8 clients; with more than 32 clients, the application became disk-bound.

NetBench 5.01

The NetBench 5.01 Disk Mix test suite was used to measure the performance of 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.

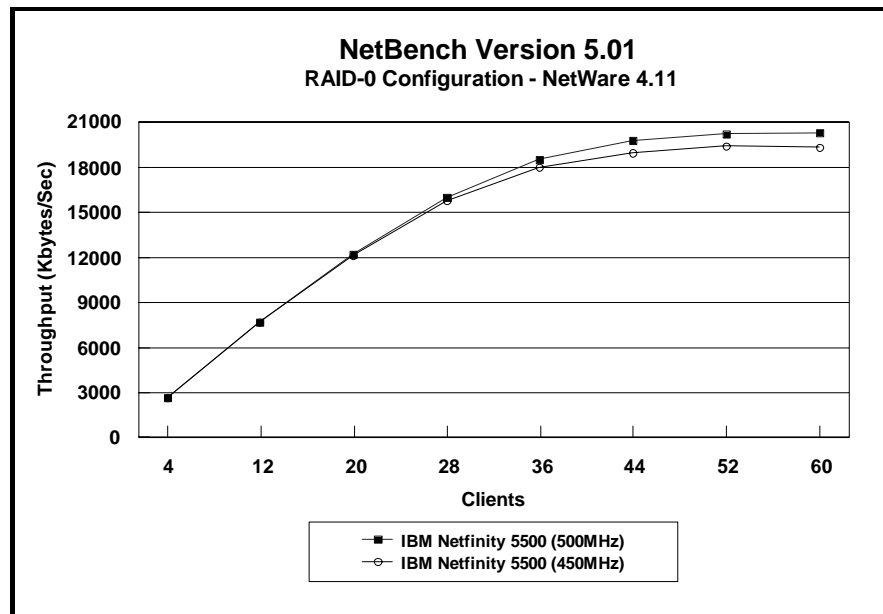
The Disk Mix test results are shown as the number of kilobytes (Kbytes) per second obtained by the server under test.

Results Summary

RAID-0 Configuration

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 5500 500MHz system provided **5 percent higher** throughput than the 450MHz system.


NetBench is designed to exercise the server's disk subsystem; therefore, the contribution of increased processor speed to overall throughput in a file-serving environment is insignificant.



RAID-5 Configuration

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 5500 500MHz system provided a level of throughput **similar** to that of the 450MHz system.

NetBench is designed to exercise the server's disk subsystem; therefore, the contribution of increased processor speed to overall throughput in a file-serving environment is insignificant.



Measurement Methodology

The Disk Mix test suite was performed using four 100Mbps Ethernet network segments with a total of 60 IBM PC 350 133MHz Pentium-based systems as client workstations attached to the server. Each workstation ran Windows NT Workstation 4.0 and executed the NetBench 5.01 Disk Mix workload, which is based on leading Windows applications.

Each client randomly simulated the Windows application workloads, accessing shared and unshared data files located on the server. Each client used a workspace of 80MB. Clients were added incrementally as follows: 4, 12, 20, 28, 36, 44, 52 and 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 of memory, in this configuration of 60 clients, the cache-hit ratio was more than 90 percent, and the cache buffer in the resource utilization was more than 90 percent. 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.

Server Configurations

ServerBench 4.02

Features	IBM Netfinity 5500 500MHz/512KB	IBM Netfinity 5500 450MHz/512KB
Processor	Two 500MHz Pentium III	Two 450MHz Pentium II
Memory	512MB ECC Synchronous DRAM	512MB ECC Synchronous DRAM
L2 Cache	512KB (Write-Back)	512KB (Write-Back)
RAID Level	0 and 5	0 and 5
Disk Drive	Five 9.1GB 10K Wide Ultra SCSI Drives	Five 9.1GB 10K Wide Ultra SCSI Drives
Disk Drive Adapter	ServeRAID II Ultra SCSI PCI Bus on Planar	ServeRAID II Ultra SCSI PCI Bus on Planar
Disk Driver	IPSRAIDN.SYS	IPSRAIDN.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

Features	IBM Netfinity 5500 500MHz/512KB	IBM Netfinity 5500 450MHz/512KB
Processor	One 500MHz Pentium III	One 450MHz Pentium II
Memory	512MB ECC SDRAM	512MB ECC SDRAM
L2 Cache	512KB (Write-Back)	512KB (Write-Back)
RAID Level	0 and 5	0 and 5
Disk Drives	Five 9.1GB 10K Wide Ultra SCSI Drives	Five 9.1GB 10K Wide Ultra SCSI Drives
Disk Drive Adapter	ServeRAID II Wide Ultra SCSI Controller on Planar	ServeRAID II Wide Ultra SCSI Controller on Planar
Disk Driver	ISPRAID.HAM V2.81	ISPRAID.HAM V2.81
Network Adapters	Four IBM EtherJet 100/10 PCI Adapters	Four IBM EtherJet 100/10 PCI Adapters
Bus	PCI	PCI
Network Driver	E100B.LAN V3.63	E100B.LAN V3.63
Network Operating System	NetWare 4.11 with IWSP6 loaded	NetWare 4.11 with IWSP6 loaded
NetWare Volume Block Size	16KB	16KB
File Compression	Off	Off
Block Allocation	On	On
Data Migration	Off	Off
Disk Mix	NB5.01 / Windows NT Workstation 4.0	NB5.01 / Windows NT Workstation 4.0

Test Disclosure Information

ServerBench 4.02

The 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
- Delay: 0
- Think: 0

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

Testbed Disclosure

The IBM Netfinity 5500 is planned to be available March 17, 1999. All other 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 4.0 Workstation
Cache	L2 = 512KB
Controller Software	Microsoft 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 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 = 300
- 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

The Netfinity 5500 500MHz model is planned to be available March 17, 1999. All other 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	Asante 100Mbps Ethernet
Clients per Segment	15
CPU / Memory	133MHz Pentium / 32MB
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 Transport - NWLink NetBIOS - TCP/IP Protocol
Cache	L2 = 256KB
Controller Software	Microsoft Windows 95 Novel Loginw 95

THE INFORMATION CONTAINED IN THIS DOCUMENT IS DISTRIBUTED ON AN AS IS BASIS WITHOUT ANY WARRANTY EITHER EXPRESS OR IMPLIED. The use of this information or the implementation of any of these techniques is the customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item has been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environment do so at their own risk.

This publication was produced in the United States. IBM may not offer the products, services, or features discussed in this document in other countries, and the information is subject to change without notice. Consult your local IBM representative for information on products and services available in your area.

*IBM and Netfinity are trademarks or registered trademarks of International Business Machines Corporation.

**Intel and Pentium are registered trademarks of Intel Corporation.

**Microsoft, Windows and Windows NT 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 Netfinity Server Performance Laboratory.

© Copyright International Business Machines Corporation 1999. All rights reserved.

Permission is granted to reproduce this document in whole or in part, provided the copyright notice as printed above is set forth in full text at the beginning or end of each reproduced document or portion thereof.

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Notes

¹ MHz denotes the clock speed of the microprocessor only, 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.