
Performance Report

IBM Netfinity 5500

Version 1.0
May 1998



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 April 1998 feature single or dual symmetrical multiprocessing (SMP) 400MHz¹ and 350MHz Pentium** II processors that support 100MHz operations to memory.*

*The 350MHz system (Model 8660-1RU) was evaluated using Ziff-Davis' benchmarks, ServerBench** Version 4.0 and WebBench** Version 1.1*

*For comparison, IBM also conducted these benchmarks with the Compaq** ProLiant** 3000/300MHz system, which was not available with a 350MHz Pentium II at the time of testing.*

IBM Netfinity 5500

Features include:

- Advanced 400MHz or 350MHz Pentium II processor with 512KB ECC L2 cache
- Planned upgrade path to Intel Pentium II processors for Slot 2
- Standard 400W, hot-swap power supply, optional hot-swap redundant power supply, and standard hot-swap, redundant fans
- 128MB of 100MHz system memory (1GB maximum) with error checking and correcting (ECC)
- Integrated system management processor providing a dedicated processor to perform remote management and monitoring functions
- Seven full-length slots: six PCI (four hot-swap) and one ISA
- Ten drive bays: standard diskette drive and 32X/14X² CD-ROM, six hot-swap bays, and two 5.25/3.5-inch half-high bays available for removable media devices
- Up to 54.6GB³ of internal hot-swap disk storage
- Industry standard 8 U, 19-inch rack-drawer models
- Tower models with innovative Netfinity NetBAY3 enclosure

- Integrated, dual-channel ServeRAID II, 100/10Mbps Ethernet, and SVGA video controllers
- Three serial ports (two 16550A-compatible), one parallel port, two USB ports, a mouse and a keyboard port
- ServerGuide* CD library, including Lotus** Domino** Server and online publications and utilities and drivers that provide assisted loading of popular network operating systems.

All Netfinity 5500 systems feature Netfinity Manager, a comprehensive systems management tool designed with a graphical interface, which is contained on a CD-ROM for installation convenience.

Performance Highlights

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

Ziff-Davis Benchmarks

For these benchmarks, the Compaq ProLiant 3000 with a 300MHz Pentium II processor was measured for comparison with the IBM Netfinity 5500. At the time these benchmarks were conducted, the Compaq ProLiant 3000 was not available with the 350MHz Pentium II processor.

ServerBench 4.0

ServerBench 4.0 was used to measure the performance of these 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 achieved a peak level of transactions per second that was:

- **25 percent higher** than the Compaq ProLiant 3000 in a RAID-0 configuration
- **25 percent higher** than the Compaq ProLiant 3000 in a RAID-5 configuration

WebBench 1.1

WebBench 1.1 was used to measure the performance of these systems as dual-processor Web servers running Microsoft Internet Information Server 3.0 on Windows NT Server 4.0.

Under a high-end workload of 60 WebBench clients, the IBM Netfinity 5500 system delivered:

- **25 percent more throughput** than the Compaq ProLiant 3000 in a RAID-0 configuration
- **20 percent more throughput** than the Compaq ProLiant 3000 in a RAID-5 configuration

Under the high-end workload of 60 WebBench clients, the IBM Netfinity 5500 system serviced:

- **25 percent more requests per second** than the Compaq ProLiant 3000 in a RAID-0 configuration
- **20 percent more requests per second** than the Compaq ProLiant 3000 in a RAID-5 configuration

Test Environments and Results

ServerBench 4.0


The ServerBench 4.0 system test suite SYS_60.TST was used to measure the performance of the IBM Netfinity 5500 and the Compaq ProLiant 3000 systems, configured as 2-way Pentium II-based 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


RAID-0 Configuration

The IBM Netfinity 5500 achieved a peak level of transactions per second that was **25 percent higher** than the Compaq ProLiant 3000.



RAID-5 Configuration

The IBM Netfinity 5500 achieved a peak level of transactions per second that was **25 percent higher** than the Compaq ProLiant 3000 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 4.0 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. (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 40 clients, the application became disk-bound.

WebBench 1.1


Ziff-Davis' WebBench 1.1 system test suite NT_SIMPLE_ISAPI20_V11.TST was used to measure the performance of the IBM Netfinity 5500 and the Compaq ProLiant 3000 systems, configured as 2-way Pentium II-based Web servers running Microsoft Internet Information Server 3.0 on Microsoft Windows NT Server 4.0 with Service Pack 3.

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

Results Summary


Throughput for RAID-0 Configuration

Under a high-end workload of 60 WebBench clients, the IBM Netfinity 5500 system delivered **25 percent more throughput** than the Compaq ProLiant 3000 system.




Throughput for RAID-5 Configuration

Under a high-end workload of 60 WebBench clients, the IBM Netfinity 5500 system delivered **20 percent more throughput** than the Compaq ProLiant 3000 system.




Requests per Second for RAID-0 Configuration

Under the high-end workload of 60 WebBench clients, the IBM Netfinity 5500 system serviced **25 percent more requests per second** than the Compaq ProLiant 3000 system.



Requests per Second for RAID-5 Configuration

Under a high-end workload of 60 clients, the IBM Netfinity 5500 serviced **20 percent more requests per second** than the Compaq ProLiant 3000 system.



Measurement Methodology

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. Two kinds of performance data were collected:

- 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 (shown here as Kbytes 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.

Each workstation ran Windows NT Workstation 4.0 and executed the WebBench 1.1 NT_SIMPLE_ISAPI20_V11.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's 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_ISAPI20_V11.TST test suite contains a total of 16 mixes. Each mix uses 30 seconds as ramp-up time, 30 seconds as ramp-down time, during which periods measurements were not done. Each mix ran for 300 seconds.

After the 16 mixes in the test suite were finished, WebBench created two graphs: one that plots the requests per second against each mix and another that plots the throughput against each mix. Also supplied is the amount of time 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 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.

WebBench 1.1 expands its set of dynamic test suites to cover more operating systems and more types of interfaces. The new workload file used in these measurements is the Internet Server Application Program Interface (ISAPI) dynamic link library (.DLL) combined with the SIMCGI.EXE in the CGI-BIN directory on the server. Because this executable file runs on the server, it uses processing resources on the server.

To get a valid measure of the server's performance, the request-per-second score and the throughput score should reach a point where they flatten out. This "flattening out" indicates that the server has been saturated, or fully loaded. In these measurements, adding clients increased the total requests-per-second and throughput scores. The curves increased from 1 to 28 clients, peaked at 32 clients, and then flattened out, indicating that the server had reached its saturation point. Ideally, the curves after the saturation point should remain at the same level where the server's resources (e.g., 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.

Server Configurations

ServerBench 4.0

RAID-0 Configuration

Features	IBM Netfinity 5500 350MHz/512KB	Compaq ProLiant 3000 300MHz/512KB
Processor	Two 350MHz Pentium II	Two 300MHz Pentium II
Memory	512MB ECC Synchronous DRAM	512MB ECC EDO
L2 Cache	512KB (Write-Back)	512KB (Write-Back)
RAID Level	RAID-0	RAID-0
Disk Drive	Five IBM 4.51GB Wide Ultra SCSI Drives (10K rpm)	Five 4.3GB Wide Ultra SCSI Drives (7200 rpm)
Disk Drive Adapter	ServeRAID II Ultra SCSI PCI Bus on Planar	SMART-2/P Array Controller
Disk Driver	IPSRAIDN.SYS	CPQARRAY.SYS
Network Adapter	Four IBM EtherJet 100/10 PCI Adapters	Four Netelligent 10/100 TX PCI UTP Controller
Bus	PCI	PCI
Network Driver	E100BNT.SYS	NETFLEX3.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 4.0 / SYS_60.TST	ServerBench 4.0 / SYS_60.TST

RAID-5 Configuration

Features	IBM Netfinity 5500 350MHz/512KB	Compaq ProLiant 3000 300MHz/512KB
Processor	Two 350MHz Pentium II	Two 300MHz Pentium II
Memory	512MB ECC Synchronous DRAM	512MB ECC EDO
L2 Cache	512KB (Write-Back)	512KB (Write-Back)
RAID Level	RAID-5	RAID-5
Disk Drive	Five IBM 4.51GB Wide Ultra SCSI Drives (10K rpm)	Five 4.3GB Wide Ultra SCSI Drives (7200 rpm)
Disk Drive Adapter	ServeRAID II Ultra SCSI PCI Bus on Planar	SMART-2/P Array Controller
Disk Driver	IPSRAIDN.SYS	CPQARRAY.SYS
Network Adapter	Four IBM EtherJet 100/10 PCI Adapters	Four Netelligent 10/100 TX PCI UTP Controller
Bus	PCI	PCI
Network Driver	E100BNT.SYS	NETFLEX3.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 4.0 / SYS_60.TST	ServerBench 4.0 / SYS_60.TST

WebBench 1.1

RAID-0 Configuration

Features	IBM Netfinity 5500 350MHz/512KB	Compaq ProLiant 3000 300MHz/512KB
Processor	Two 350MHz Pentium II	Two 300MHz Pentium II
Memory	512MB ECC Synchronous DRAM	512MB ECC EDO
L2 Cache	512KB (Write-Back)	512KB (Write-Back)
RAID Level	RAID-0	RAID-0
Disk Drive	Five IBM 4.51GB Wide Ultra SCSI Drives (10K rpm)	Five 4.3GB Wide Ultra SCSI Drives (7200 rpm)
Disk Drive Adapter	ServeRAID II Ultra SCSI PCI Bus on Planar	SMART-2/P Array Controller
Disk Driver	IPSRAIDN.SYS	CPQARRAY.SYS
Network Adapter	Four IBM EtherJet 100/10 PCI Adapters	Four Netelligent 10/100 TX PCI UTP Controller
Bus	PCI	PCI
Network Driver	E100BNT.SYS	NETFLEX3.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
WebBench Version / Test Suite	WebBench 1.1 / NT_SIMPLE_ISAPI20_V11.TST	WebBench 1.1 / NT_SIMPLE_ISAPI20_V11.TST
Web Server	Microsoft Internet Information Server 3.0	Microsoft Internet Information Server 3.0

RAID-5 Configuration

Features	IBM Netfinity 5500 350MHz/512KB	Compaq ProLiant 3000 300MHz/512KB
Processor	Two 350MHz Pentium II	Two 300MHz Pentium II
Memory	512MB ECC Synchronous DRAM	512MB ECC EDO
L2 Cache	512KB (Write-Back)	512KB (Write-Back)
RAID Level	RAID-5	RAID-5
Disk Drive	Five IBM 4.51GB Wide Ultra SCSI Drives (10K rpm)	Five 4.3GB Wide Ultra SCSI Drives (7200 rpm)
Disk Drive Adapter	ServeRAID II Ultra SCSI PCI Bus on Planar	SMART-2/P Array Controller
Disk Driver	IPSRAIDN.SYS	CPQARRAY.SYS
Network Adapter	Four IBM EtherJet 100/10 PCI Adapters	Four Netelligent 10/100 TX PCI UTP Controller
Bus	PCI	PCI
Network Driver	E100BNT.SYS	NETFLEX3.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
WebBench Version / Test Suite	WebBench 1.1 / NT_SIMPLE_ISAPI20_V11.TST	WebBench 1.1 / NT_SIMPLE_ISAPI20_V11.TST
Web Server	Microsoft Internet Information Server 3.0	Microsoft Internet Information Server 3.0

Test Disclosure Information

ServerBench 4.0

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

Version: ServerBench 4.0

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
- Delay Time: 0
- Think Time: 0
- Ramp up: Default setup
- Ramp down: Default setup

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

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	Asante 100Mbps Ethernet
Clients per Segment	15
CPU / Memory	166MHz Pentium / 64MB
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

WebBench 1.1

The measurements were conducted using Ziff-Davis' WebBench 1.1 running the NT_SIMPLE_ISAPI20_V11.TST test suite with Windows NT Workstation 4.0 clients as described below:

Version: WebBench 1.1

Mixes

- NT_SIMPLE_ISAPI20_V11.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
- Web Server: Logon Disabled

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

Web Server: Microsoft Internet Information Server 3.0

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

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, ServerGuide and Netfinity are trademarks or registered trademarks of International Business Machines Corporation.

**Lotus, Lotus Notes and Domino are trademarks or registered trademarks of Lotus Development Corporation.

**Intel and Pentium are registered trademarks of Microsoft 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 Netfinity Server Performance Laboratory.

© 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.

Notes

¹ MHz denotes the clock speed of the microprocessor only, not application performance. Many factors affect application performance.

² 32X/14X CD-ROM variable read rate. Actual playback speed will vary and is often less than the maximum possible.

³ When referring to hard disk capacity, GB, or gigabyte, means one thousand million bytes. Total user-accessible capacity may vary depending on operating environment.