UNITED STATES

Performance Report

IBM PC Server 330

Version 3.0
September 1997

IBM®

Executive Overview

*The performance of the IBM PC Server Pentium** II-based models, announced worldwide in September 1997, was evaluated by the IBM PC Company's performance group using the following benchmarks:*

- *NotesBench for Lotus** Domino** Server Release 4.51*
- *Ziff-Davis' ServerBench** Version 4.0*
- *Ziff-Davis' WebBench** Version 1.0*
- *Ziff-Davis' NetBench** Version 5.01*

*For comparison, the PC Company's performance group also conducted the Ziff-Davis benchmarks with Compaq** ProLiant** 2500 and Hewlett-Packard** NetServer** LH Pro systems.*

IBM PC Server 330

The new IBM PC Server 330 models deliver greater processor power and enhanced functionality to enterprise or small to medium-size business. Two of the new models feature either a 266MHz[1] or a 233MHz Intel** Pentium II processor with integrated 512KB ECC L2 cache. The third model features a 200MHz Pentium Pro processor with integrated 512KB ECC L2 cache and up to 1GB of memory for optimum server performance in applications requiring more than 512MB of system memory. All models feature increased integration of server function on the motherboard, including a Wide Ultra SCSI RAID controller, a conventional Wide Ultra SCSI controller, and a high-performance, full-duplex 100/10Mbps Ethernet controller.

The IBM PC Server 330 models are intended for use as departmental file servers, database servers, and Internet/Intranet servers. They are ideal as high-performance branch-office file or database servers running applications such as e-mail and office productivity suites. The IBM PC Server 330 also features fault-tolerance in support of mission-critical business applications in either stand-alone or clustered configurations.

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

Standard features include:

- One- or two-way 200MHz Pentium Pro processors with dedicated 512KB L2 cache, or 266MHz or 233MHz Pentium II processor with 512KB ECC L2 cache
- One 64MB EDO dual in-line memory module (DIMM) standard, expandable to 512MB on Pentium II models and to 1GB[2] on the Pentium Pro model
- An advanced bus architecture: two 132MB-per-second PCI buses with secondary ISA bus for compatibility
- Ten status LEDs to allow monitoring of critical functions such as processor and Ethernet activity
- Room for adding 32-bit adapters, a total of six slots (four dedicated PCI slots and two PCI/ISA combination slots)
- Ten drive bays: six 3.5-inch slim-high, hot-swap bays; one 5.25-inch half-high drive bay; three 3.5-inch slim-high drive bays, all allowing access to removable media
- Up to 36.16GB of internal storage using six 4.51GB hot-swap disk drives and one 9.1GB conventional disk drive
- Integrated SVGA PCI controller with 1MB video memory, Wide Ultra SCSI PCI RAID controller, conventional Wide Ultra SCSI controller, and a full-duplex 100/10Mbps Ethernet PCI controller to provide server functionality without consuming valuable expansion slots
- 8X-speed CD-ROM, ServerGuide* and IBM Netfinity* Manager for easier setup and network management
- Lotus Domino Release 4.51 Server (single-processor): the tools for making decisions and communicating, locally or around the world

[2] When referring to hard disk capacity, GB stands for one thousand billion bytes. Total user-accessible capacity may vary dep ending on operating environments.

Performance Highlights

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

NotesBench

The IBM PC Server 330, configured with two 233MHz Pentium II processors, and the IBM PC Server 325, configured with one 266MHz Pentium II processor, **demonstrated leadership price/performance running Lotus Domino Server 4.51 on Microsoft Windows NT Server 4.0** in recently conducted measurements using two NotesBench workloads: Mail-only and Mail and Shared Database (MailDB). The following tables summarize the results.

IBM PC Server 330 - Two-Way 233MHz Pentium II with 512MB Memory

Test Script	Maximum Users	NotesMark (tpm)	Response Time (sec)	$/User	$/NotesMark
Mail-Only	2,200	2,917	0.360	$13.45	$10.14
MailDB	1,700	3,914	1.427	$17.40	$7.56

IBM PC Server 325 - One-Way 266MHz Pentium II with 384MB Memory

Test Script	Maximum Users	NotesMark (tpm)	Response Time (sec)	$/User	$/NotesMark
Mail-Only	1,795	2,356	0.704	$10.84	$8.26
MailDB	1,400	3,203	1.670	$13.90	$6.07

Ziff-Davis Benchmarks

For the these benchmarks, the IBM PC Server 330 was configured with the 266MHz Pentium II processor. The Compaq and HP systems were not available with the Pentium II processor at the time of testing. These two systems are configured with the 200MHz Pentium Pro. Otherwise, the configuration of and the performance tuning for the three systems were the same for all Ziff-Davis benchmark tests conducted.

ServerBench 4.0

ServerBench 4.0 was used to measure the performance of the IBM PC Server 330 system and the competitors' systems as dual-processor application servers running Windows NT Server 4.0 and providing services to Windows NT Workstation 4.0 clients.

The peak level of throughput achieved by the IBM PC Server 330 was **15 percent higher** than that of the HP NetServer LH Pro system and slightly higher than that of the Compaq ProLiant 2500 system.

WebBench 1.0

WebBench 1.0 was used to measure the performance of the IBM PC Server 330 system and the competitors' systems as dual-processor Web servers running Microsoft Internet Information Server 3.0 on Windows NT Server 4.0.

At the peak of 12 WebBench clients, the IBM PC Server 330 system delivered **22 percent more throughput** than the HP NetServer LH Pro and **7 percent more throughput** than the Compaq ProLiant 2500.

At the peak of 12 WebBench clients, the IBM PC Server 330 serviced **25 percent more requests per second** than the HP NetServer LH Pro. The IBM PC Server 330 and the Compaq ProLiant 2500 serviced a similar number of requests per second at the peak of 12 WebBench clients.

NetBench 5.01

NetBench Version 5.01 Disk Mix for Windows for Workgroup Clients was used to measure the performance of the IBM PC Server 330 system and the competitors' systems as single-processor file servers running Novell** NetWare** 4.11.

Under a high-end workload of 60 NetBench clients, the IBM PC Server 330 system provided **30 percent more throughput** than the HP NetServer LH Pro and a level of throughput similar to that of the Compaq ProLiant 2500.

Test Environments and Results

NotesBench

Results Summary

The IBM PC Server 330 and the IBM PC Server 325 systems **demonstrated leadership price/performance running Lotus Domino Server 4.51 on Microsoft Windows NT Server 4.0** in recently conducted measurements using two NotesBench workloads: Mail-only and Mail and Shared Database (MailDB).[3] The following tables summarize the results.

IBM PC Server 330 - Two-Way 233MHz Pentium II / 512MB Memory

Test Script	Maximum Users	NotesMark (tpm)	Response Time (sec)	$/User	$/NotesMark
Mail-Only	2,200	2,917	0.360	$13.45	$10.14
MailDB	1,700	3,914	1.427	$17.40	$7.56

IBM PC Server 325 - One-Way 266MHz Pentium II / 384MB Memory

Test Script	Maximum Users	NotesMark (tpm)	Response Time (sec)	$/User	$/NotesMark
Mail-Only	1,795	2,356	0.704	$10.84	$8.26
MailDB	1,400	3,203	1.670	$13.90	$6.07

NotesBench generates a transactions-per-minute (tpm) throughput metric, called a NotesMark, for each test, along with a value for the maximum capacity (number of users) supported, and the average response time. The price/performance results for the IBM PC Server 330 are derived using a total system cost of $29,581. For the IBM PC Server 325, the price/performance results are based on a total system price of $19,454.[4]

[3] The NotesBench Disclosure Reports for the IBM PC Server 330 and the IBM PC Server 325 were audited in September 1997 by KMDS Technical Associates, Inc., and were approved for publication. Highlights from the PC Server 330 audited report are presented here.

[4] Price/performance results are based on pricing provided by an IBM Business Partner. IBM resellers set their own prices, and actual prices may vary.

Competitive Results

The tables 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. These results are provided for comparison. All competitive results shown are based on the tests 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' NotesBench Disclosure Reports (**http://www.notesbench.org** on the World Wide Web) for details concerning the server configuration and methodology used to obtain the published results.[5]

IBM PC Server 330 vs. Competitive Systems (SMP)	**$ / User**	**$ / NotesMark**	**Maximum Users**	**NotesMark (tpm)**
	Mail / MailDB	**Mail / MailDB**	**Mail / MailDB**	**Mail / MailDB**
IBM PC Server 330 (2 x 233MHz Pentium II, 512MB) Sept. 1997	**$13.45 / $17.40**	**$10.14 / $7.56**	**2,200 / 1,700**	**2,917 / 3,914**
IBM PC Server 704 (2 - 200MHz Pentium Pro, 768MB) June 1997	**$15.18 / $21.11**	**$11.51 / $9.21**	**2,850 / 2,050**	**3,761 / 4,697**
IBM PC Server 704 (3 - 200MHz Pentium Pro, 1GB) July 1997	**$16.28 / $23.36**	**$12.41 / $10.25**	**3,300 / 2,300**	**4,331 / 5,242**
IBM Netfinity 7000 (3 - 200MHz Pentium Pro, 1GB for Mail; 4 - 200MHz Pentium Pro, 1GB for MailDB) Sept. 1997	**$18.91 / $24.70**	**$14.32 / $11.38**	**3,500 / 2,900**	**4,623 / 6,294**
HP NetServer LX Pro (4 - 200MHz Pentium Pro, 1GB) May 1997	$19.81 / NA	$15.49 / NA	3,175 / NA	4,061 / NA
DEC AlphaServer 4100 5/400 (2 - 400MHz Alpha CPU, 1GB) Jan. 1997	$29.38 / NA	$22.49 / NA	3,100 / NA	4,050 / NA
DEC AlphaServer 4100 5/466 (2 - 466MHz Alpha CPU, 1GB) Feb. 1997	$31.07 / NA	$24.61 / NA	3,350 / NA	4,230 / NA

IBM PC Server 325 vs. Competitive Systems (Single-Processor)	**$ / User**	**$ / NotesMark**	**Maximum Users**	**NotesMark (tpm)**
	Mail / MailDB	**Mail / MailDB**	**Mail / MailDB**	**Mail / MailDB**
IBM PC Server 325 (1 x 266MHz Pentium II, 384MB) Sept. 1997	**$10.84 / $13.90**	**$8.26 / $6.07**	**1,795 / 1,400**	**2,917 / 3,914**
Compaq ProLiant 800 (1 x 200MHz Pentium Pro, 256MB) Jan. 1997	$13.48 / $16.05	$10.12 / $6.82	1,500 / 1,260	1,997 / 2,955
Compaq ProLiant 5000 (1 x 200MHz Pentium Pro, 512MB) June 1997	$19.97 / $26.63	$15.06 / $11.80	1,920 / 1,440	2,546 / 3,250
DEC AlphaServer 800 5/400 (1 x 400MHz Alpha, 576MB) April 1997	$23.11 / $27.37	$16.78 / $11.96	1,800 / 1,450	2,479 / 3,317
DEC AlphaServer 1000A 5/500 (1 x 500MHz Alpha, 768MB) June 1997	$25.78 / $31.50	$19.71 / $14.52	2,200 / 1,800	2,878 / 3,904

[5] Data on competitive products obtained from publicly available information and is subject to change without notice. Contact the manufacturer for the most recent information.

Measurement Methodology

The IBM PC Server 330 system under test was configured with two 233MHz Pentium II processors each with 512KB ECC L2 cache; 512MB of memory; nine 4.51GB Wide Ultra SCSI Hot-Swap hard disks configured as a RAID-0 array, using the IBM ServeRAID II Ultra SCSI Adapter; and one EtherJet 100/10 PCI Adapter. (The integrated IBM ServeRAID SCSI controller and the integrated 100/10Mbps Ethernet adapter were not used.)

For these tests, a single 100Mbps Ethernet LAN segment was used. The system under test, the destination servers, and the driver systems were connected to the LAN by two Asante 100BaseT Hubs. An IBM PC 350 computer was used as the source driver (parent) system; 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. The IBM PC Server 330 system under test ran Microsoft Windows NT Server Version 4.0 with Service Pack 3 and Domino Server Release 4.51.

To ensure that the test results were reproducible, all tests were repeated, and the results were compared for consistency.

Measurement Analysis

A detailed analysis of the measurements is provided in the NotesBench Disclosure Reports for the IBM PC Server 330 and the IBM PC Server 325 systems. In addition, the details of the benchmarked configuration are provided in the reports, which are available on the World Wide Web at **http://www.us.pc.ibm.com/techlink/srvperf.html**.

These results are based on running the IBM PC Server 330 and the IBM PC Server 325 as dedicated Domino servers; 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 configuration of that hardware and appropriately tuning the operating system and Domino software.

ServerBench 4.0

The ServerBench 4.0 system test suite SYS_60.TST was used to measure the performance of the IBM PC Server 330 266MHz Pentium II-based system and the competitors' 200MHz Pentium Pro-based systems as dual-processor 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 PC Server 330 was **15 percent higher** than that of the HP NetServer LH Pro system and slightly higher than that of the Compaq ProLiant 2500 system.

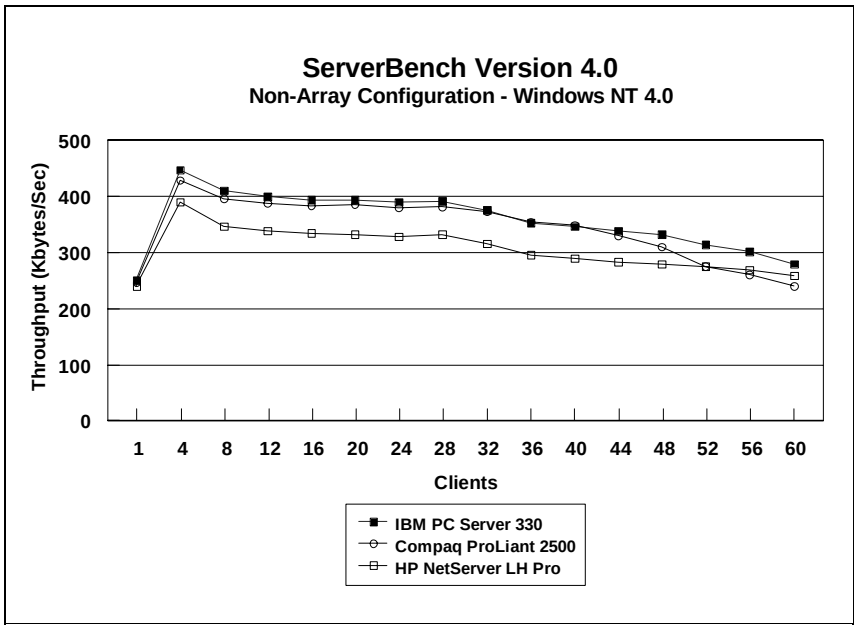

Measurement Methodology

The system test suite was performed using two 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.

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 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 (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 a second processor to alleviate a processor bottleneck. However, if the disk subsystem is the bottleneck, increasing the size of memory, rather than adding a second processor, can improve the performance.

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 PC Server 330 Pentium II-based system and the competitors' 200MHz Pentium Pro-based systems as 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

At the peak of 12 WebBench clients, the IBM PC Server 330 system delivered **22 percent more throughput** than the HP NetServer LH Pro and **7 percent more throughput** than the Compaq ProLiant 2500.

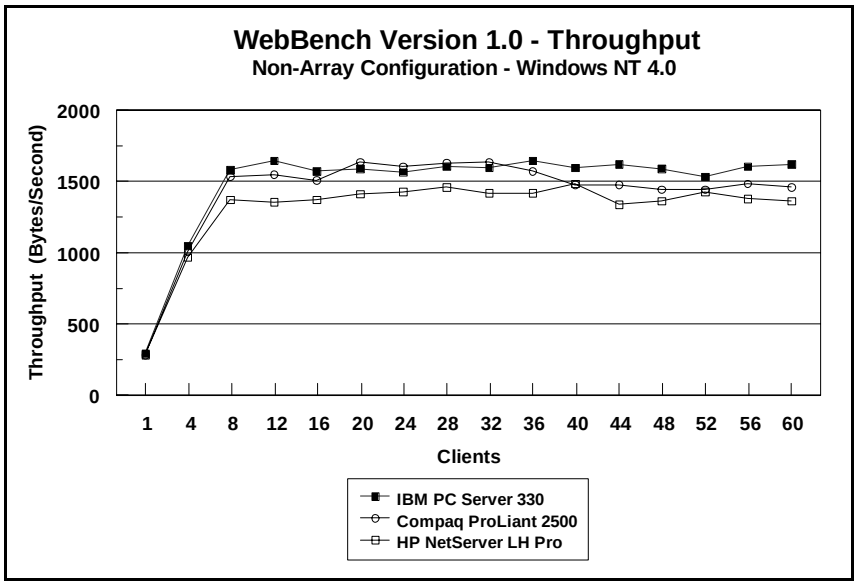

Requests per Second

At the peak of 12 WebBench clients, the IBM PC Server 330 serviced **25 percent more requests per second** than the HP NetServer LH Pro. The IBM PC Server 330 and the Compaq ProLiant 2500 serviced a similar number of requests per second at the peak of 12 WebBench clients.

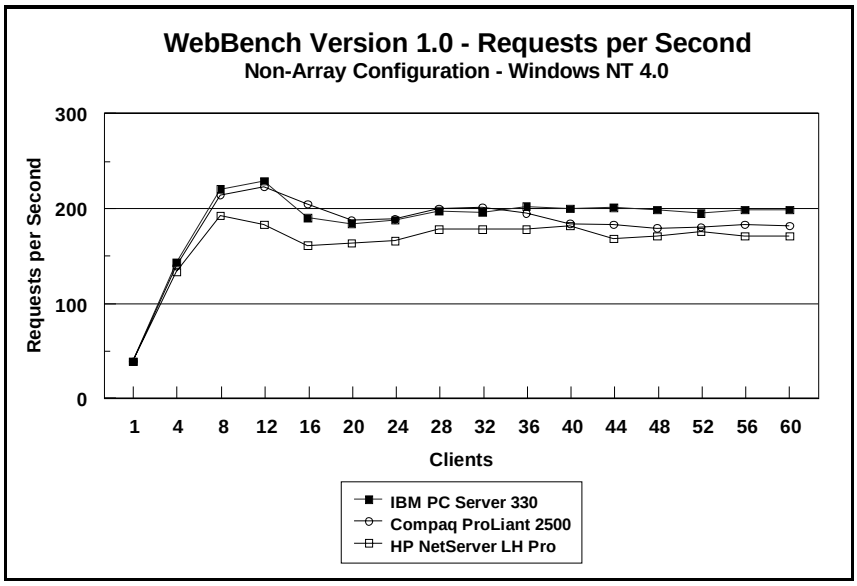

Measurement Methodology

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 1.0 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'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_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 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, the test suites can be run with a relatively small number of clients and still provide an accurate measure of a server's performance.

To get a valid measure of the server's performance, the requests-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 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 the server's memory manager policy to be tuned 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, the system was tuned 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

Ziff-Davis' NetBench 5.01 Disk Mix test suite was used to measure the performance of the IBM PC Server 330 266MHz Pentium II-based system and the competitors' 200MHz Pentium Pro-based systems as 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 PC Server 330 system provided **30 percent more throughput** than the HP NetServer LH Pro and a level of throughput similar to that of the Compaq ProLiant 2500.

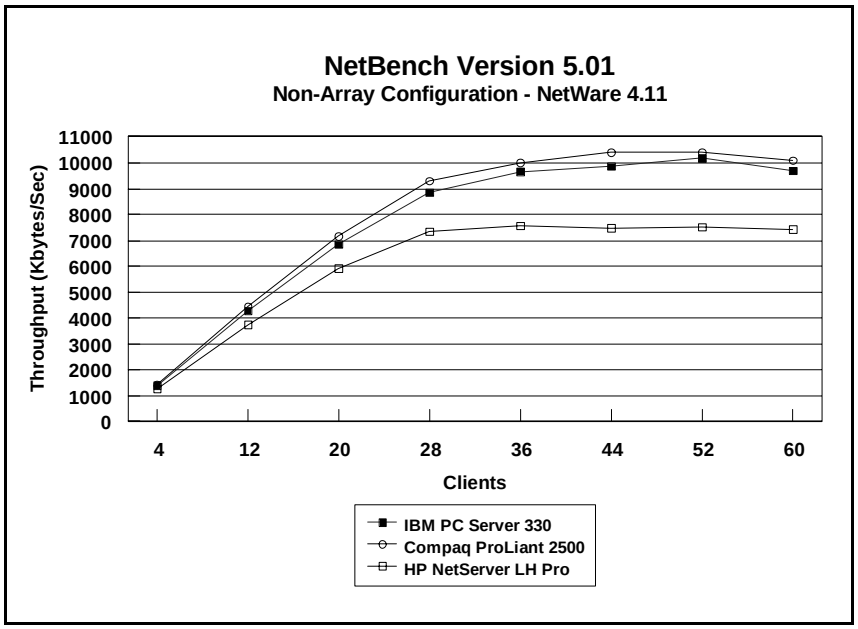

Measurement Methodology

The Disk Mix test suite was performed using two 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 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 for Workgroups 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 Version 5.01 workload exercises the server in a manner similar to actual Windows applications executing on a networked-attached PC; that is, the NetBench Version 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 128MB of memory optimized the throughput and also stressed the server 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.

Server Configurations

ServerBench 4.0

Features	IBM PC Server 330 266MHz/512KB	Compaq ProLiant 2500 200MHz/512KB	HP NetServer LH Pro 200MHz/512KB
Processor	Two 266MHz Pentium II	Two 200MHz Pentium Pro	Two 200MHz Pentium Pro
Memory	512MB ECC	512MB ECC	512MB ECC
L2 Cache	512KB ECC (Write-Back)	512KB (Write-Back)	512KB (Write-Back)
RAID Level	Non-Array	Non-Array	Non-Array
Disk Drive	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)
Disk Drive Adapter	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
Disk Driver	AIC78XX.SYS	SYMC810.SYS	AIC78XX.SYS
Network Adapter	Four IBM 100/10 PCI Ethernet Adapters	Four Netelligent 10/100 TX PCI Ethernet Adapters	Four HP DeskDirect LAN Adapters
Bus	PCI	PCI	PCI
Network Driver	E100B.SYS	NETFLX3.SYS	E100B.SYS
Network Operating System	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
System Partition Size	1GB	1GB	1GB
File System	NTFS	NTFS	NTFS
Allocation Unit Size	Predefined Default	Predefined Default	Predefined Default
ServerBench Version / Test Suite	ServerBench 4.0 / SYS_60.TST	ServerBench 4.0 / SYS_60.TST	ServerBench 4.0 / SYS_60.TST

WebBench 1.0

Features	IBM PC Server 330 266MHz/512KB	Compaq ProLiant 2500 200MHz/512KB	HP NetServer LH Pro 200MHz/512KB
Processor	Two 266MHz Pentium II	Two 200MHz Pentium Pro	Two 200MHz Pentium Pro
Memory	512MB ECC	512MB ECC	512MB ECC
L2 Cache	512KB ECC (Write-Back)	512KB (Write-Back)	512KB (Write-Back)
RAID Level	Non-Array	Non-Array	Non-Array
Disk Drive	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)
Disk Drive Adapter	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
Disk Driver	AIC78XX.SYS	SYMC810.SYS	AIC78XX.SYS
Network Adapter	Four IBM 100/10 PCI Ethernet Adapters	Four Netelligent 10/100 TX PCI Ethernet Adapters	Four HP DeskDirect LAN Adapters
Bus	PCI	PCI	PCI
Network Driver	E100B.SYS	NETFLX3.SYS	E100B.SYS
Network Operating System	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
System Partition Size	1GB	1GB	1GB
File Compression	Off	Off	Off
File System	NTFS	NTFS	NTFS
Allocation Unit Size	Predefined Default	Predefined Default	Predefined Default
WebBench Version / Test Suite	WebBench 1.0 / NT_SIMPLE_CGI20_V1.TST	WebBench 1.0 / NT_SIMPLE_CGI20_V1.TST	WebBench 1.0 / NT_SIMPLE_CGI20_V1.TST
Web Server	Microsoft Internet Information Server 3.0	Microsoft Internet Information Server 3.0	Microsoft Internet Information Server 3.0

NetBench 5.01

Features	IBM PC Server 330 266MHz/512KB	Compaq ProLiant 2500 200MHz/512KB	HP NetServer LH Pro 200MHz/512KB
Processor	One 266MHz Pentium II	One 200MHz Pentium Pro	One 200MHz Pentium Pro
Memory	128MB ECC	128MB ECC	128MB ECC
L2 Cache	512KB ECC (Write-Back)	512KB (Write-Back)	512KB (Write-Back)
RAID Level	Non-Array	Non-Array	Non-Array
Disk Drives	Four IBM 4.51GB Wide Ultra SCSI Drives (7200 rpm)	Four 4.3GB Wide Ultra SCSI Drives (7200 rpm)	Four HP 4.2GB Fast/Wide SCSI-2 Drives (7200 rpm)
Disk Drive Adapter	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
Disk Driver	AIC7870.DSK V2.13	CPQS710.DSK V2.11	AIC7870.DSK V2.13
Network Adapters	Two Intel** EtherExpress** PRO 100/10 Smart Ethernet Adapters	Two Netelligent PCI 100/10 Ethernet Adapters	Four HP DeskDirect LAN Adapters
Bus	PCI	PCI	PCI
Network Driver	E100S.LAN V2.93	CPQNF3.LAN V2.23	HPTX.LAN V1.47
Network Operating System	NetWare 4.11 with IWSP3 loaded	NetWare 4.11 with IWSP3 loaded	NetWare 4.11 with IWSP3 loaded
NetWare Volume Block Size	32KB	32KB	32KB
File Compression	Off	Off	Off
Block Allocation	On	On	On
Data Migration	Off	Off	Off
Disk Mix	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 ServerBench measurements were conducted using Ziff-Davis' ServerBench 4.0 running the SYT_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
- Disk Test Data File: 20MB
- Disk Test File I/O Range: 10MB
- 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 100Mbp 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 1.0

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: WebBench 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 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

NetBench Version 5.01

The NetBench measurements were conducted using Ziff-Davis' NetBench 5.01 running the Disk Mix with Windows for Workgroup 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 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 Packet Size = 1514
- Reserved Buffer Below 16MB = 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 Buffer Non-Referenced Delay = 30
- Directory Cache Allocation Wait Time = 2.2 seconds

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	Asante 100Mbp Ethernet
Clients per Segment	30
CPU / Memory	133MHz Pentium / 16MB
Network Adapter	IBM 100/10 PCI Ethernet Adapter (Bus 0)
Software	IBM DOS 6.3
	NetWare DOS Requester
	LSL.COM (8-3-95)
	E100BODI (5-21-96)
	IPXODI (8-8-95)
	VLM.EXE (11-8-94)
Cache	L2 = 256KB
Controller Software	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

Published by the IBM Server Performance Laboratory, IBM Corp. The IBM Server Performance Laboratory publishes white papers and performance reports, which are available on the World Wide Web at the following URL:
http://www.us.pc.ibm.com/techlink/srvperf.html