Performance Report

IBM Netfinity 5000 (550MHz)

Version 2.0
June 1999

IBM®

Executive Overview

The new model of the IBM Netfinity 5000, announced worldwide in June 1999, is a high-throughput, two-way SMP-capable server that features the Intel** 550MHz[1] Pentium** III processor, which supports 100MHz operations to memory. This new model expands IBM's Netfinity line of midrange servers, providing solid network performance for business-critical applications.*

*The new 550MHz system (Model 8659-51Y) was evaluated using Ziff-Davis' file server benchmark NetBench** Version 5.01.*

For comparison, the IBM Netfinity server performance laboratory also conducted the benchmarks with the IBM Netfinity 5000 configured with the 500MHz Pentium III processor.

All results obtained from this benchmark are presented in this report.

Performance Highlights

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

NetBench 5.01

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 5000 550MHz system, configured with RAID-5, delivered a level of network throughput that was **22 percent higher** than that of the 500MHz system.

Test Environments and Results

NetBench 5.01

The NetBench 5.01 Disk Mix test suite was used to measure the performance of the IBM Netfinity 5000 systems as single-processor file servers running Windows NT Server 4.0 and servicing requests from Windows NT Workstation 4.0 clients.

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

Results Summary

RAID-5 Configuration

Under a high-end workload of 60 NetBench clients, the IBM Netfinity 5000 550MHz system delivered a level of network throughput that was **22 percent higher** than that of the 500MHz system.

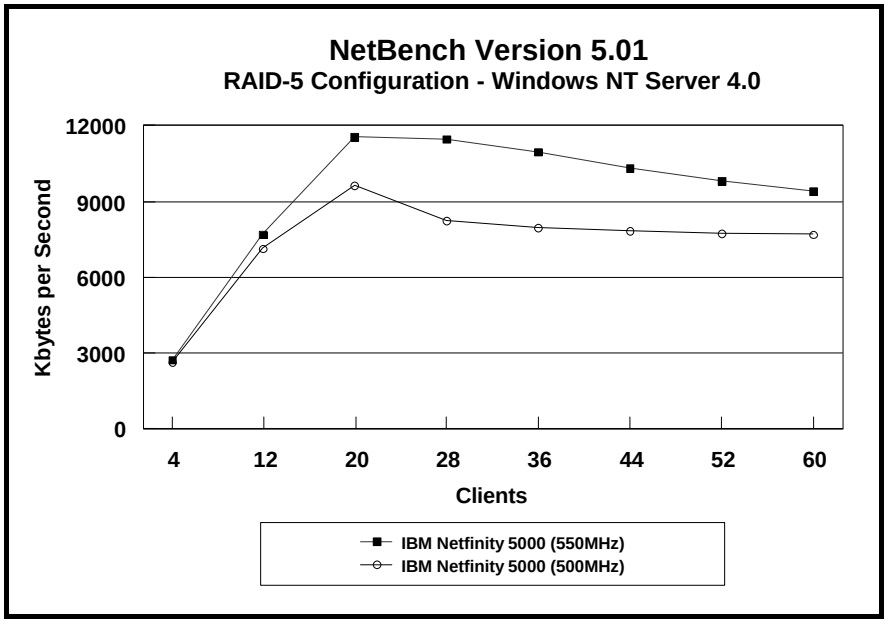

Measurement Methodology

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 20MB. 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 1GB of memory, in this configuration of 60 clients, the cache-hit ratio was more than 95 percent, and the cache buffer in the resource utilization was more than 90 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

NetBench 5.01

Features	IBM Netfinity 5000 550MHz/512KB	IBM Netfinity 5000 500MHz/512KB
Processor	One 550MHz Pentium III	One 500MHz Pentium III
Memory	1024MB ECC SDRAM	1024MB ECC SDRAM
L2 Cache	512KB (Write-Back)	512KB (Write-Back)
RAID Level	RAID-5	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	Integrated Wide Ultra SCSI PCI Controller	Integrated Wide Ultra SCSI PCI Controller
Disk Driver	AIC7870.DSK V 4.20	AIC7870.DSK V 4.20
Network Adapter	Three IBM 10/100 Ethernet PCI Adapters V3.1 One integrated AMD 10/100 Ethernet Controller	Three IBM 10/100 Ethernet PCI Adapters V3.1 One integrated AMD 10/100 Ethernet Controller
Bus	PCI	PCI
Network Driver	E100B.LAN V3.63	E100B.LAN V3.63
Network Operating System	Windows NT Server 4.0 with Service Pack 3	Windows NT Server 4.0 with Service Pack 3
NetBench Version / Test Suite	NB5.01 / Disk Mix	NB5.01 / Disk Mix

Test Disclosure Information

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: 20MB
- Total runtime: 11 minutes
- Ramp up and down: 30 seconds

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

NOS Parameters

- Network Speed = 100Mbps
- Duplex Mode = Full
- Receive Buffers = 16
- TCBs = 32
- Map Registers = 64
- Coalesce Buffers = 8
- Transmit Threshold = 16

Testbed Disclosure

The Netfinity 5000 550MHz system is planned to be available June 22, 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
Switches	IBM 8275 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
Cache	L2 = 512KB
Controller Software	Microsoft 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 and Netfinity are trademarks or registered trademarks of International Business Machines Corporation.

**Intel and Pentium are trademarks or 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 internal/external 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.