

Performance Report IBM Netfinity 7000 with Microsoft Exchange Server

Version 1.0
January 1998

1. Executive Overview

IBM and Microsoft have partnered together to optimize the IBM Netfinity 7000 for the Microsoft Exchange Server version 5.5. This paper details the configuration used for performance testing of Exchange Server version 5.5 at IBM Kirkland Programming Center. An IBM Netfinity 7000, demonstrating excellent performance and scalability supported **10,000** simultaneous email users.

About IBM Netfinity 7000

The IBM Netfinity 7000, available in rack and tower models, is a four-way symmetric multiprocessing (SMP) server featuring a 200MHz Pentium Pro processor with either 1MB of L2 cache or 512KB of L2 cache, and up to 4GB of memory. The systems are ruggedly reliable enterprise servers that deliver exceptional power, functionality and value for handling complex, business-critical database or application server needs.

These servers, available in both tower and rack-ready models, are easy to install in either environment. Conversion kits are available to convert towers to racks, and vice versa, if requirements change. They offer configuration flexibility with such features as hot-swap, open-bay models; open systems to support conventional Ethernet or single port- to multi-port Ethernet PCI adapters, and a full-high bay for internal tape units capable of backing up 70GB of data.

The IBM Netfinity 7000 is intended for enterprise customers whose networks depend on a reliable, high-performance server platform with excellent scalability for future growth. The system is ideal for handling business-critical database, groupware, messaging and intranet applications. The IBM Netfinity 7000 can also function as a file server in LAN environments where there are a significant number of users requiring a large amount of file server resources.

More information on IBM Netfinity servers can be found at:

<http://www.pc.ibm.com/us/netfinity/netfinity.html>

The IBM Netfinity 7000 is a certified Microsoft Cluster Server configuration as listed on the Microsoft Hardware Compatibility List. For more information see

<http://www.microsoft.com/hwtest/hcl/>

2. Results Summary

Microsoft provided the Loadsim tools for Exchange email load Simulation. The medium MAPI user canonical profile was chosen which reflects the workload of a typical corporate email environment. Acceptable performance for this test is an average response time of less than one second, and prompt emptying of Exchange Server Information Store (IS) Send Queue after any momentary buildups. In this benchmark test, the IBM Netfinity 7000 was able to support an outstanding **10,000** Exchange email users under these conditions with an average response time of about half a second. The IS Send Queue was essentially zero. The measurements were taken from 4-hour steady state period of the benchmark test.

Number of Users	10,000 users (medium profile)
95th Percentile response time	649 milliseconds (ms)
Total message recipients delivered	807,744 messages for the 8-hour period
Messaging protocol	Exchange MAPI
Loadsim Canonical Profile	Medium (typical Corporate email user)
Microsoft Exchange Server	Version 5.5 (Build 1960.5)
Microsoft Loadsim tools	Version 5.5 (Build 2108.0)

3. Test Environment

Server Configuration

Server platform	IBM Netfinity 7000
Processors	Four 200 MHz Pentium Pro, 1MB L2 cache
System memory	3 GB
Disk Controllers	Two IBM ServeRAID II Ultra SCSI adapters Version 2.40 Firmware and device driver.
Disk Drives (total)	Internal: Twelve 4.51GB disk drives, Expansion: Ten 4.51GB disk drives. One EXP10 Rack Storage Expansion Enclosure is used. All disk drives are 7200rpm.
Logical drive C	NTFS, (drive 1), RAID0, write-back. For OS, pagefile, MTA, DS
Logical drive D	NTFS, (drive 2), RAID0, write-back. For Exchange Log files.
Logical drive E	NTFS, (drive 3 to 22), RAID0, write through. For Exchange Mail database.
Network adapters	Two 3COM 3C905 100TX Ethernet adapters
Operating Systems	Windows NT 4.0 Server Enterprise Edition + SP3
Microsoft Exchange Server	Version 5.5 (Build 1960.5)

Client Machine configuration

Client machines	IBM PC 350
Processors	Single 133 MHz Pentium.
System memory	80 MB
Network adapters	Single 3COM 3C905 100TX
Operating Systems	Windows NT 4.0 Workstation + SP3
Microsoft Loadsim tools	Version 5.5 (Build 2108.0)
Number of simulated users	Each machine simulates 300 users, except one control station at 100 users.
Total number of network segments	Two 100TX Ethernet segments
Total number of client machines	34 evenly divided into two Ethernet segments

4. Measurement Analysis

Item/Counter (Averaged during steady state)	Measurement/Results
Processor Utilization (%)	89%
Context Switches / sec	5517
Processor Queue Length	3.9
Available Bytes	654 MB
Pages / sec	0.015
IS Send Queue length	26.35
Store Working Set Bytes	2.363 GB
Store Virtual Bytes	2.876 GB
IS Table Opens / sec	172
MTA Work Queue Length	22.17
IS Database Volume	
Disk Queue Length	7.35
Read I/Os sec	205.72
Write I/Os sec	152.76
IS Log Volume	
Disk Queue Length	0.059
Write I/Os sec	89.90
OS/MTA/DS Volume	
Disk Queue Length	0.023
Read I/Os sec	0.002
Write I/Os sec	21.15
Messages Delivered	381.72 per minute (183,808 for 8-hour period)
Message Recipients Delivered	1680.51 per minute (807,744 for 8-hour period)

5. Benchmark Vs production configuration

- In the benchmark test, we have selected the default medium user profile that Microsoft recommends. With this profile the Loadsim tool simulates MAPI email user access to the Private Store and Schedule+ activity. The profile also specifies a total of 30 distribution lists per site. In the production environment, the corporation may choose to implement additional Exchange features on the same server, which would increase the workload and hence limit the number of users that a server can support. Such features include but not limited to the following:
 - Public Folders access,
 - POP3 or other email interfaces,
 - Multiple Exchange Servers environment,
 - Links and replication to remote Exchange sites.
- In this benchmark configuration, we have chosen to configure the disk sub-system with RAID 0. In a production environment, however, it is essential to have some form of redundancy configuration in order to ensure fault tolerance and continuous availability. We recommend the configuration as follow:

Drives C and D, the boot drive and the Exchange log file drive, configure with RAID 1 to provide both fault tolerance and maximum response time. We do not recommend using RAID 5 for the drive storing the Exchange log file because the log file writes sequentially and has no read access during normal Exchange operation. RAID 5 is less efficient for handling this disk activity. We'd also recommend the log files be put on a dedicated drive.

Drive E, the Exchange database drive, configure in RAID 1 or RAID 5 across all the remaining disk drives provides redundancy. A RAID 5 configuration would be the best compromise between response time and cost effectiveness.

Technical Note:

RAID Technology

RAID (Redundant Array of Independent Disks) is a means to improve the reliability and performance of disk storage subsystems by configuring several disk spindles to appear as a single drive to the system. Two different levels of RAID were used for these tests: RAID Level 1 and RAID Level 5.

RAID Level 1 (or simply RAID 1) uses *disk mirroring* to provide high fault tolerance and enhanced read performance. In this configuration, data is written simultaneously to a number of disks in the array. Each mirrored disk in the set contains an exact duplicate of the data, so that data can be read simultaneously off any disk, and the data is preserved even if one of the disks should fail.

RAID Level 5 (RAID 5) stores data by *striping*, or organizing the virtual drive space in "stripes" spread across each of the physical disk spindles. This allows large files to be accessed more quickly, as they are broken up into segments across several drives, and each segment can be read simultaneously, thus decreasing access time. This configuration also provides fault tolerance by allocating one virtual drive for parity information, enabling the system to reconstruct the data on a failed drive.

For more information on RAID technology as implemented by IBM, see the Web site <http://www.pc.ibm.com/infobrf/raidfin.html>

- In a production environment, the maximum number of user-per-server could be constrained by the amount of storage space that the server can provide, as well as the per-user storage requirement. For information on IBM server storage overview, see <http://www.pc.ibm.com/us/options/storesvr.html>, <http://www.pc.ibm.com/us/infobrf/storexp.html>, and <http://ww.storage.ibm.com/storage/>
- Backup capacity and the duration to backup the database are also key factors to be considered.
- Microsoft Exchange Server includes a utility called the Performance Optimizer that allows the administrator to configure the server properly to take full advantage of the system memory and processors installed in the server machine. This function can be accessed via a Microsoft Wizard interface by running perfwiz.exe. It is necessary to run this program after every change to the configuration of the server machine (for example, when memory is added, or additional processors activated) so that Exchange Server will be aware of the changes to the system. For the tests performed in this paper, we ran perfwiz.exe after each configuration setup and accepted the default settings.

Appendix A: Testing Methodology

LoadSim

LoadSim is an Exchange client emulator provided by Microsoft to simulate client messaging workload to the server. The server perceives its actions as that of an actual client, because it sends the same messages to the server as a client would when the user of the client software is taking some sort of action. Actions performed by LoadSim and the Microsoft Exchange Client are very similar in terms of server load; the user simulation features in LoadSim are designed specifically to mimic the Microsoft Exchange Client's use of Mail API.

Of equal significance, LoadSim is a multi-client emulator, in that it allows for simulating large-scale client-server scenarios without actually involving a multitude of client machines. LoadSim takes advantage of multiprocessing, multithreading, and shared memory features in Windows NT to create and manage up to several hundred simulated users on a single client computer. Our testing relied on LoadSim for the following purposes:

- To test performance and capacity of sever hardware.
- To determine network bandwidth requirements given varying server configurations.
- To conduct performance impact tests on various server configurations, without affecting numerous actual users.

User behavior models define how many times per day a user sends a message, how often the user reads a message, and what percentage of messages are deleted, saved in folders, or forwarded to others. LoadSim allows user behavior models to be customized to a quite detailed level. LoadSim also provides a simple means of characterizing user loads according to predefined usage levels. These predefined user profiles, called *canonical user profiles*, provide a workload standard that can be used to compare data in common scenarios. LoadSim provides canonical user profiles for light, medium, and heavy users. For our testing, we used the medium canonical user.

As LoadSim runs a scenario, it probabilistically generates user activities according to the model provided. For example, if the model specifies users to send 22 messages in an 8-hour day, LoadSim will have each user in the scenario generate messages at an average rate of 2.75 messages per hour, though the intervals between messages will not be uniform. Random variations are included to simulate the behavioral patterns of actual users. A user action is simulated by LoadSim's generating the same MAPI calls to the server that a client would.

Generating Scores Using LsLog

When LoadSim runs a scenario, it creates two files. One of these, named *loadsimsim.out*, contains a history of all generated actions. Any errors that occurred during the run are logged here, in the form of a screen dump from the LoadSim application window during the scenario run. This file is checked (in addition to NT Performance Monitor results) for errors that might possibly invalidate the entire run. Another file, named *lsperf.log*, contains a collection of all the individual action timings that result from the LoadSim run. This data is used to generate scores — that is, numerical values (measured in milliseconds) that represent the responsiveness of the server

under the specified load. Results measuring the impact of component or other system variations, as well as variations in the number of users, are derived by averaging these scores.

The LoadSim Performance Log Parser is a Windows NT command-line executable file designed to manipulate data from the LoadSim performance log files. After each LoadSim test run, LsLog is used to merge the logged data from all test clients, then truncate unwanted data in order to extract the relevant time window. The result is used to generate the 95th percentile result for all actions. The default 95th percentile response time for a given action means that 95 percent of all response times recorded for that action were equal to or below that value. The average response time, measured in milliseconds, provides an indication of how responsive the server is to the user. A small number means quicker (hence better) response. Microsoft recommends 1000 milliseconds (one second) or less as the suggested response time for Exchange Server.

Windows NT Performance Monitor

LoadSim scores represent the end result of a scenario run. In order to properly understand what actually occurred during the course of a run, the Windows NT Performance Monitor (PerfMon) is used constantly during the test to collect performance data, which must be analyzed when the test is complete. Besides the operating system-specific counters included as a default with Windows NT, the Exchange Server also includes counters specific to Exchange to provide further means to track Exchange performance. PerfMon's record indicates bottlenecks and problem areas. Studying the results logged by PerfMon ensures that the results reflected in the LoadSim scores are not corrupted by factors outside the scope of the testing methodology.

Performance Criteria

LoadSim was configured to simulate 10,000 medium MAPI email users for eight hours. The test environment reached a stable state after two hours, so the results from the first two hours of the test run were discarded. The readings from the remaining six hours were sampled to provide the average response time over four hours of steady-state period.

Response Time: Obviously, users do not want to experience a long wait for the program to respond to their actions. A response time of about one second is considered reasonable.

CPU Utilization: The greater the amount of processing power used by Exchange, the less efficient the server's performance will be. This ties in to response time; higher CPU utilization generally leads to slower response times.

Send Queue Size: The Exchange Server's send queue is used to store messages that have been sent by a client, but have not yet been processed by the server. This is needed to allow for the random nature of mail use: statistically, there will be times when many users are accessing Exchange at once, and other times when fewer users than normal access it. Ideally, the send queue should accumulate messages at times of heavy use, and then empty itself when usage falls below the average. Therefore, in a system with sufficient capacity, the send queue should be empty for a significant portion of the time (at these times, the server is processing messages as soon as they come in). This will result in an average send queue volume that is effectively zero.

Appendix B: LoadSim User Profile

Microsoft LoadSim provides three canonical user profiles to simulate the actions of a typical light, medium, or heavy corporate Exchange user. **Our test used the medium user profile.** For reference, the canonical profile characteristics are listed below.

LoadSim Test Attribute	Light	Medium	Heavy
Overall Test Configuration			
Logon per day	1	1	1
Logon immediately at the very beginning of test	yes	yes	yes
Always keep MAPI session connection	yes	yes	yes
Logoff at the end of each simulated day	no	no	no
Empty the Deleted Items folder on Logoff	yes	yes	yes
Public Folder Post	no	no	no
Browse Public Folders	no	no	no
Initial per user configuration (Test initialization)			
Number of messages in Inbox	1	4	9
Number of messages in Deleted Items	1	1	1
Number of new folders	20	40	60
Messages per new folder	5	5	5
Test Duration			
Length of day (hours)	8	8	8
Duration of simulation (hours)	8	8	8
Distribution Lists			
Use Distribution List	yes	yes	yes
DL per site	30	30	30
Minimum DL size	2	2	2
Average DL size	10	10	10
Maximum DL size	20	20	20
Send Mail			
Send new mail (times per day)	2	4	6
Number of recipients per msg.	3	3	3
Add a single distribution list to msg.	30%	30%	30%
Save a copy in Sent Items folder	yes	yes	yes
Message priority	normal	normal	normal
Delivery receipt	no	no	no
Read receipt	no	no	no
Type of messages sent			
1K message (ups1k.msg)	9	60	50
2K message (ups2k.msg)	0	16	10
4K message (ups4k.msg)	0	4	5
Message with 10K attachment (ups10kat.msg)	1	5	10
Message with Excel attachment (upsXLatt.msg)	0	4	5
Message with Word attachment (upsWDatt.msg)	0	2	5
Message with Bitmap object (upsBMobj.msg)	0	2	5
Message with Excel object (upsXLobj.msg)	0	2	10

Process Inbox			
Read new mail (times per day)	12	12	12
Reply	5%	7%	15%
Reply All	3%	5%	7%
Forward	5%	7%	7%
Delete	40%	40%	40%
Move	20%	20%	20%
Copy	0%	0%	0%
Load or execute message attachments	25%	25%	25%
Read Note delay			
Minimum	1s	1s	1s
Average	1s	1s	1s
Maximum	1s	1s	1s
Browse existing Mail (per day)	5	15	20
Schedule+			
Update schedule (times per day)	1	5	10
Update free busy information	no	no	No
Schedule file size (KB)			
Minimum	5	5	5
Average	22	22	22
Maximum	40	40	40
Approximate message traffic, per user, per day			
Total received	22.94	66.3	118.89
Total sent	4.7	14.18	30.67
Reply	1.05	3.76	13.03
Reply all	0.6	2.67	5.82
Forward	1.05	3.76	5.82
Average recipients per message (for all messages)	4.88	4.68	3.88
Approximate receipts requested per user per day			
Read Receipts	0	0	0
Delivery Receipts	0	0	0

The information contained in this document is distributed on an AS IS basis without any warranty either expressed 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, Netfinity and ServeRAID II are trademarks or registered trademarks of International Business Machines Corporation.

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

**Microsoft, Windows, Windows NT, and Exchange Server 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 Kirkland Programming Center, IBM Corp.

© 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.