WHITE PAPER

June 1996

Prepared By
Os Development

Compaq Computer
Corporation

CONTENTS

Performance of Microsoft Exchange Server 4.0 on Compaq ProLiant Servers

Electronic messaging is consistently becoming more pervasive in the business world. As people must process increasing amounts of information, there is a need for powerful software and hardware products to satisfy the requirements imposed.

To fill the need for a powerful messaging software product, Microsoft has created Microsoft Exchange Server 4.0. This product has been the focal point of extensive development and testing efforts involving the commitments of both Microsoft and its partners. During this time, Compaq and Microsoft have worked together to optimize Microsoft Exchange Server on Compaq System products.

Microsoft Exchange Server is enormously complex, with many features and capabilities. One important part of understanding the product is understanding its performance. For system administrators and potential purchasers of both Microsoft Exchange Server and Compaq products, this aspect of Microsoft Exchange Server is crucial to understand so the proper decisions can be made.

This white paper covers various performance aspects of Microsoft Exchange Server 4.0 on Compaq ProLiant Servers, specifically the ProLiant 4500 and the ProLiant 5000.

Fundamentally, this paper should provide the reader with an expectation of how Microsoft Exchange Server will perform on a Compaq ProLiant under various user loads. Furthermore, the paper should uncover other useful information pertaining to Microsoft Exchange Server configuration and performance tuning, as well as CPU, RAM, and disk utilization on Compaq ProLiant servers. Ultimately, this paper should help the potential purchaser make good decisions about which Compaq products to purchase in order to satisfy the requirements of his or her business needs.

COMPAQ

NOTICE

The information in this publication is subject to change without notice.

COMPAQ COMPUTER CORPORATION SHALL NOT BE LIABLE FOR TECHNICAL OR EDITORIAL ERRORS OR OMISSIONS CONTAINED HEREIN, NOR FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES RESULTING FROM THE FURNISHING, PERFORMANCE, OR USE OF THIS MATERIAL.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal, state or local requirements. Compaq does not warrant products other than its own strictly as stated in Compaq product warranties.

Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

Compaq, ProLiant, and NetFlex, registered United States Patent and Trademark Office.

Other product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

Microsoft, and Windows NT, are trademarks and/or registered trademarks of Microsoft Corporation.

Performance of Microsoft Exchange Server 4.0
on Compaq ProLiant Servers
First Edition (June 1996)

TEST CONFIGURATION

Introduction

Before diving right into the performance results, it is important to take a few minutes and review the tools and configuration used to arrive at these results.

There are two main tools used to prepare the data for this white paper:

- Microsoft Exchange Server LoadSim
- Microsoft Windows NT Performance Monitor

Naturally, your results will vary depending upon your exact conditions. In other words, the results shown probably won't be exactly what you will see. However, the trends should be the same.

Additionally, the configuration of test loads selected for this paper are a good overall mix of potential real world situations. If anything, the load is a bit on the heavy side, so it should illustrate a worst case scenario for most Exchange users. In other words, most likely your user load won't be heavier than the loads used in these tests, and as a result your performance should be at least as good as, if not better than, the performance shown here.

LoadSim

The main tool used in generating the performance data contained in this paper is the Microsoft Exchange Server User Load Simulation utility - called LoadSim.

As its name implies, LoadSim is a tool for simulating a client user load on an Exchange Server. Its purpose is to enable a single Windows NT machine - called a LoadSim client - to simulate multiple Microsoft Exchange client users.

For more information on the details of Loadsim operation and configuration, refer to the Microsoft white paper entitled, "LoadSim: Microsoft Exchange Server Load Simulation Tool for Microsoft Exchange 4.0". This paper is included with the LoadSim utility, and they can both be found on the Microsoft Exchange Server CD in the \SUPPORT\LOADSIM directory.

The operation of LoadSim users is governed by a LoadSim profile, which is stored on disk as a .SIM file. This profile controls things like how long a LoadSim 'day' is, how many email messages to send in a day's time, how many times to open and read existing email, whether to use distribution lists, whether to use public folders, etc.

Default User Profiles

There are three pre-configuredLoadSim profiles which are built intoLoadSim: *Light, Medium,* and *Heavy.*

DEFAULTLOADSIM USER PROFILES

LoadSim USER ATTRIBUTE	ATTRIBUTE DETAIL	LIGHT	MEDIUM	HEAVY
TEST DURATION	Length of a day (hours)	8	8	8
READING MAIL	New mail (times/day)	12	12	12
	Existing mail (times/day)	5	15	20
AFTER READING MAIL	% of Reply	5%	7%	15%
	% of Reply All	3%	5%	7%
	% of Forward	5%	7%	7%
	% of Move	20%	20%	20%
	% of Copy	0%	0%	0%
	% of Delete	40%	40%	40%
	% of Do nothing	27%	21%	11%
DISTRIBUTION LISTS	Minimum size	4	4	4
	Maximum size	50	50	50
	Average size	10	10	10
	Cover 100% of users (no overlap)	Yes	Yes	Yes
ATTACHMENTS	% to Run/Load Mail Attachment (if one exists)	25%	25%	25%
INBOX SIZE	Inbox Size Limit (# messages)	20	125	250
SENDING MAIL	New mail (times/day)	2	4	6
	Save a copy in Sent Mail Folder?	Yes	Yes	Yes
	Number of random recipients	3	3	3
	% of time to add a Distribution List	30%	30%	30%
	Message Priority	Normal	Normal	Normal
	Delivery Receipt?	No	No	No
	Read Receipt?	No	No	No
NEW MAIL MESSAGE CONTENT *Text-only, no attachment*	1K body (ups1K.msg)	90%	64%	50%
	2K body (ups2K.msg)	0%	17%	10%
	4K body (ups4K.msg)	0%	4%	5%
NEW MAIL MESSAGE CONTENT *1K mail body, with attachment*	10K attachment (ups10Kat.msg)	10%	5%	10%
	Embedded bitmap object (upsBMobj.msg)	0%	2%	5%
	Word attachment (upsWDatt.msg)	0%	2%	5%

LoadSim USER ATTRIBUTE	ATTRIBUTE DETAIL	LIGHT	MEDIUM	HEAVY
	Excel attachment (upsXLatt.msg)	0%	4%	5%
	Embedded Excel object (upsXLobj.msg)	0%	2%	10%
SCHEDULE+ CHANGES	Changes per day	1	5	10
	Update Free/Busy information?	No	No	No
	Average Schedule File Size	22KB	22KB	22KB
PUBLIC FOLDERS	Folder activity	None	None	None
CALCULATED DAILY LOAD (based on these defaults)	TOTAL MAIL RECEIVED PER DAY	22.94	66.30	118.89
CALCULATED DAILY LOAD (based on these defaults)	TOTAL MAIL SENT PER DAY	4.70	14.18	30.67
	Mail sent as New mail	2.00	4.00	6.00
	Mail sent as a Reply	1.05	3.76	13.03
	Mail sent as a Reply to All	0.60	2.67	5.82
	Mail sent as a Forward	1.05	3.76	5.82
CALCULATED DAILY LOAD (based on these defaults)	AVERAGE # RECIPIENTS FOR EACH MESSAGE	4.88	4.68	3.88

Table 1. DefaultLoadSim user profile definitions

User Profile Definition for This Paper

The actual LoadSim user profile used in these tests is based on the *Medium* user profile. In the following table, deviations from the standard *Medium* profile are in bold.

LOADSIM USER PROFILE FOR THIS PAPER

LoadSim USER ATTRIBUTE	ATTRIBUTE DETAIL	
TEST DURATION	Length of a day (hours)	8
READING MAIL	New mail (times/day)	12
	Existing mail (times/day)	15
AFTER READING MAIL	% of Reply	7%
	% of Reply All	5%
	% of Forward	7%
	% of Move	20%
	% of Copy	0%
	% of Delete	40%
	% of Do nothing	21%
DISTRIBUTION LISTS	Minimum size	4
	Maximum size	50
	Average size	**12**
	Cover 100% of users (no overlap)	Yes
ATTACHMENTS	% to Run/Load Mail Attachment (if one exists)	25%
INBOX SIZE	Inbox Size Limit (# messages)	125
SENDING MAIL	New mail (times/day)	4
	Save a copy in Sent Mail Folder?	Yes
	Number of random recipients	3
	% of time to add a Distribution List	30%
	Message Priority	Normal
	Delivery Receipt?	No
	Read Receipt?	No
NEW MAIL MESSAGE CONTENT	1K body (ups1K.msg)	64%
Text-only, no attachment	2K body (ups2K.msg)	17%
	4K body (ups4K.msg)	4%

LoadSim USER ATTRIBUTE	ATTRIBUTE DETAIL	
NEW MAIL MESSAGE CONTENT	10K attachment (ups10Kat.msg)	5%
1K mail body, with attachment	Embedded bitmap object (upsBMobj.msg)	2%
	Word attachment (upsWDatt.msg)	2%
	Excel attachment (upsXLatt.msg)	4%
	Embedded Excel object (upsXLobj.msg)	2%
SCHEDULE+ CHANGES	Changes per day	5
	Update Free/Busy information?	No
	Average Schedule File Size	22KB
PUBLIC FOLDERS	**Public folder tasks?**	**Yes**
	Number of times/day to open public folders	**5**
	% of time to open active folders	**60%**
	Read old item	**4%**
	Read new item	**15%**
	Change view	**7%**
	Delete item	**3%**
	Post items	**5%**
	Items to post	**1**
	Reply	**5%**
	Reply to author	**0%**
	Forward item	**0%**
	Copy item	**5%**
	Move item	**5%**
	Modify item	**5%**
	% of time to do resolve-as-all	**30%**
CALCULATED DAILY LOAD (based on these defaults)	**TOTAL MAIL RECEIVED PER DAY**	**76.76**
CALCULATED DAILY LOAD (based on these defaults)	**TOTAL MAIL SENT PER DAY**	**15.11**
	Mail sent as New mail	4.00
	Mail sent as a Reply	**4.10**
	Mail sent as a Reply to All	**2.91**
	Mail sent as a Forward	**4.10**
CALCULATED DAILY LOAD (based on these defaults)	**AVERAGE # RECIPIENTS FOR EACH MESSAGE**	**5.08**

Table 2. LoadSim user profile definition for this paper

LoadSim Score

The data point resulting from a LoadSim run is called the *Score*. The LoadSim Score represents a weighted average of the 95th percentile Exchange client response time (in milliseconds) of the various Exchange tasks. The read task is the highest weighted task — accounting for over half of the score.

Many of the figures in this paper use the Score as a means of implying performance, i.e. a lower Score indicates better Exchange Server performance.

Only you can determine the acceptable response time in your environment, but most administrators opt for one of two criteria: sub-second response time, or sub-3-second response time. Three seconds has been assumed as a safe upper limit for most graphs in this paper.

Calculating the LoadSim Score

The Score you see in the graphs of this paper are derived using the following process.

1. A LoadSim client simulates 100 *Medium* users
2. LoadSim is configured to log data to disk.
3. All Loadsim runs are for at least four hours.
4. The 95th percentile LoadSim Score is then calculated from the LoadSim logs using the LSLOG utility. (Refer to the Microsoft white paper, "LoadSim: Microsoft Exchange Server Load Simulation Tool for Microsoft Exchange 4.0" for more information on LSLOG)
5. When running LSLOG, the first hour is thrown away, and only the second, third, and fourth hours are used for calculating the Score.
6. The Scores returned by LSLOG for all the LoadSim clients in the test are averaged. The mean Score is used as a data point on the graph.

Performance Monitor

The main tool used in monitoring and collecting the performance data contained in this paper is the Windows NT Performance Monitor (PerfMon).

PerfMon monitors performance objects and counters within Windows NT, and it is these objects and counter which depict how the Exchange Server machine is performing under load.

For more information on how to use PerfMon, refer to the Windows NT Resource Kit for Windows NT 3.51, Volume 4, 'Optimizing Windows NT'. Contained in the volume is a thorough treatment of PerfMon and some good suggestions for detecting bottlenecks. The principles outlined are relevant to monitoring performance of Exchange Server.

Configuration of Test Facility

The test facility is configured on two isolated 10BaseT Ethernet networks - one network for the LoadSim clients and one network for the data collection. This is to isolate the network traffic imposed by data collection from the actual test network traffic. Figure 1 is a representation of the network topology.

There are 15 LoadSim client machines. As stated earlier, a LoadSim client is simply a Windows NT machine that is configured with the Microsoft Exchange client software and LoadSim. A single LoadSim client can simulate multiple users. In this case, each of the 15 LoadSim clients simulate 100 users. The load imposed by each user is based on the profile outlined previously in Table 2.

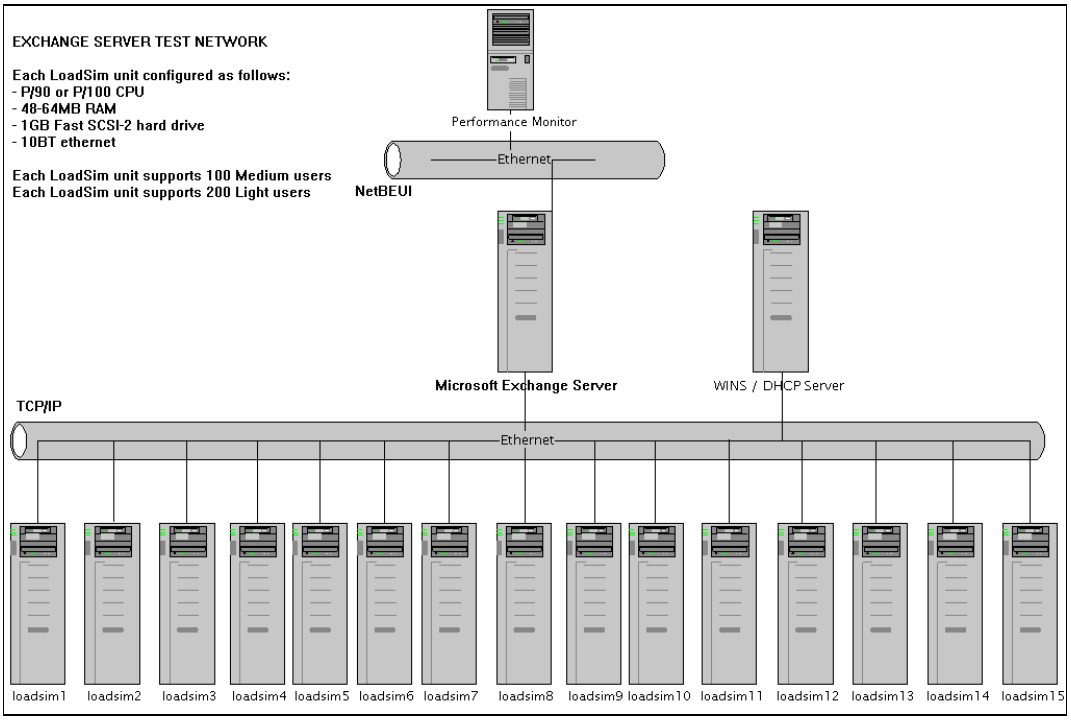

Figure 1. Exchange Server test network topology.

LOADSIM CLIENT HARDWARE CONFIGURATIONS

Machine Class	Compaq ProLiant 2000
System Processor	Single Pentium/100
System Memory	64 Mbytes
Disk subsystem	1GB Fast SCSI-2
Network Interface	Compaq NetFlex/2 (10BaseT)
Operating System	Windows NT Workstation v3.51 + SP2

Table 3. LoadSim client hardware configurations

PERFORMANCE RESULTS

Introduction

In this section of the paper the focus will be on the actual performance results of Microsoft Exchange Server.

There are three main areas of server resources to warrant attention:

- Processor subsystem - all the server's CPU resource, whether a uniprocessor system or a multiprocessor system
- Disk subsystem - all the server's disk storage resource, including controller type and number of drives in a RAID set

- System memory - all the server's memory resource, but not including cache memory on processor boards or drive arrays. This is the amount of RAM installed in the system.

The data shows that Exchange Server is primarily processor intensive. That is to say, the first resource to be consumed in an Exchange Server machine is the processor. With that in mind, most of the following data will be centered around an analysis of performance with varying CPU configurations.

After the processor subsystem, it is a tossup between the disk subsystem and system memory as to which is more crucial. It's a bit of a balance, i.e. if the server has a high performance disk subsystem, then the system memory is more likely to reach capacity first. Conversely, if the system has lots of memory, then the disk subsystem is likely to reach capacity first.

The cardinal rule to remember is this:*if a subsystem is not a bottleneck, then adding more of that resource will not help the problem*In all likelihood, it will merely amplify the existing weakness in the system. For example, if an Exchange Server exhibits poor response time - and the processor subsystem is not a bottleneck - then adding more processors to the server will not improve response time.

Following the processor subsystem section are some data pertaining to the disk subsystem and system memory. While there's enough variables in those to warrant a white paper of their own, these sections are included to provide a feel for how the Exchange Server system responds as different resources are adjusted.

Processor Subsystem

As mentioned above, the processor subsystem is usually the first to be exhausted in an Exchange Server system. Following are graphs showing how the different processor and server products from Compaq affect the performance of Microsoft Exchange Server.

The value on the vertical axis is theLoadSim Score. The value on the horizontal axis is the number ofLoadSim users.

1P indicates one processor installed in the server, 2P indicates two processors, etc. Each section will describe the individual configuration of the Exchange Server computer being tested.

In all these test, the Exchange Server computer is configured as shown, and Exchange Optimizer is run to set the Exchange system parameters. No manual tuning was done.

Pentium/100

PENTIUM/100 EXCHANGE SERVER HARDWARE CONFIGURATION

Machine Class	Compaq ProLiant 4500	
Processor Subsystem	1,2,4 Pentium/100, 512 KByte L2 cache	
System Memory	128 Mbytes	
Disk Subsystem	Integrated Fast Wide SCSI-2 controller	*Pagefile and executables volume*: 1x2GB Fast Wide SCSI-2 drive
	EISA SMART Controller	*Log volume*: 2x2GB Fast SCSI-2 drives on port 0, RAID 1 (Array Accelerator enabled) *IS volume*: 7x2GB Fast SCSI-2 drives on port 1, RAID 5 (Array Accelerator enabled)
Network Interface	Compaq NetFlex/3 (10BaseT)	
Operating System	Windows NT Server v3.51 + SP4	

Table 4. Pentium/100 Exchange Server hardware configuration

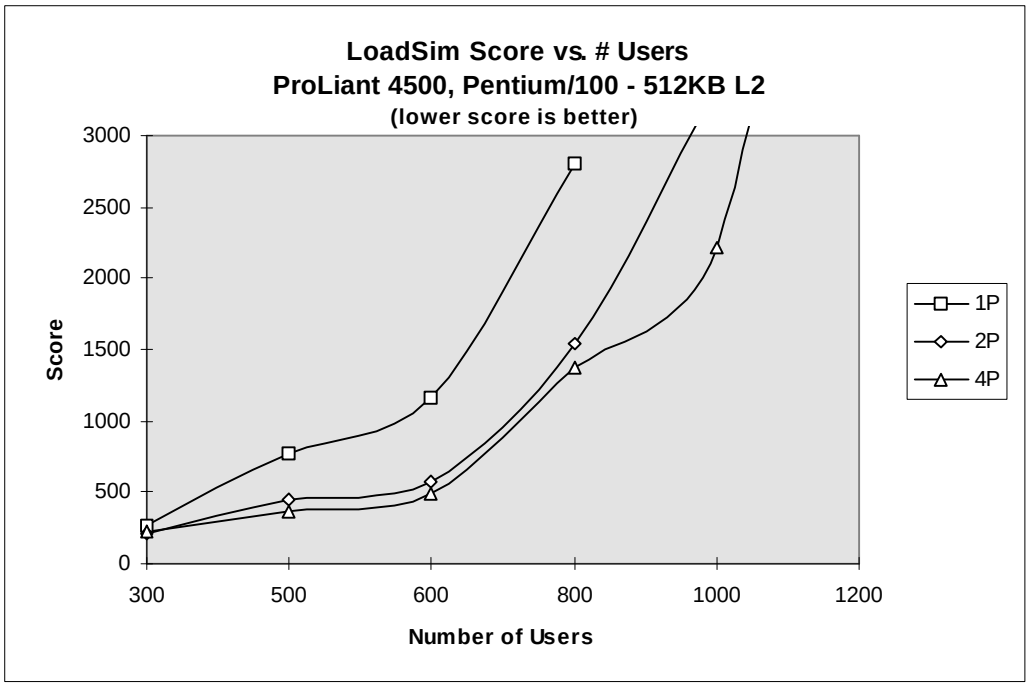

Figure 2. Relative Score performance of Pentium/100

- The server very quickly consumes processor resource with a single CPU. Only 500 users can be supported withsubsecond response time.
- In the single CPU data there is a huge jump in Score above 500 users. The trend carries over to the 2P and 4P data as well. This could be due to the L2 cache getting overrun at those numbers of users.
- This is a good example of showing the benefit of a second CPU because Exchange is becoming processor constrained after 600 users.

- This example also shows how there is little benefit in adding a CPU when the system *is not* processor constrained. In the case of going from 2P to 4P the response time improves very little until you get above 800 users, which is where 2P becomes processor constrained.
- There is a tradeoff in adding CPUs to the processor subsystem due to system overhead. System requests and context switches increase substantially as more CPUs are added.

Pentium/133

PENTIUM/133 EXCHANGE SERVER HARDWARE CONFIGURATION

Machine Class	Compaq ProLiant 4500	
Processor Subsystem	1,2 Pentium/133, 2 Mbyte L2 cache	
System Memory	128 Mbytes	
Disk Subsystem	Integrated Fast Wide SCSI-2 controller	*Pagefile and executables volume*: 1x2GB Fast Wide SCSI-2 drive
	EISA SMART Controller	*Log volume*: 2x2GB Fast SCSI-2 drives on port 0, RAID 1 (Array Accelerator enabled)
		IS volume: 7x2GB Fast SCSI-2 drives on port 1, RAID 5 (Array Accelerator enabled)
Network Interface	Compaq NetFlex/3 (10BaseT)	
Operating System	Windows NT Server v3.51 + SP4	

Table 5. Pentium/133 Exchange Server hardware configuration

NOTE:
These test results are based on a configuration that included Public Folders and Schedule Plus. We feel that this profile better represents an actual customer environment. In a similar report produced by Microsoft, testing did not include Public Folders or Schedule Plus and therefore produced significantly different results.

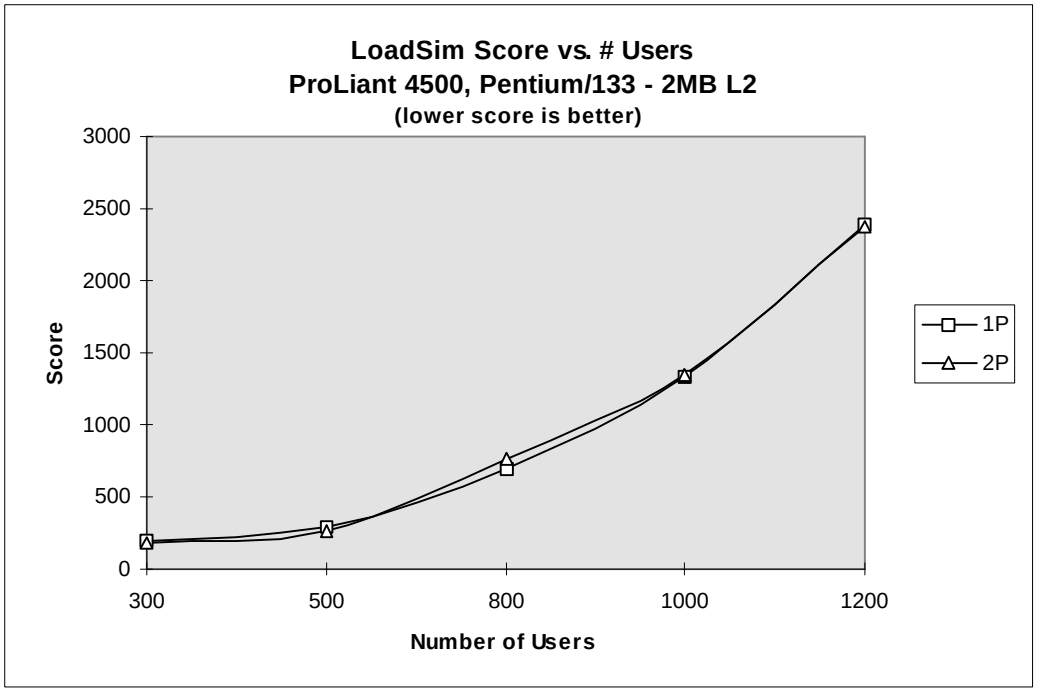

Figure 3. Relative Score performance of Pentium/133.

- This data perfectly illustrates the power of the L2 cache combined with the faster clock speed of the Pentium/133. After comparing the results shown in Figure2 with those of Figure3, it is clear that a single Pentium/133 easily outperforms two or four Pentium/100 CPUs.
- The server is less quick to consume processor resource. 800 users can easily be supported with subsecond response time.
- In this data there is not the huge jump in Score as exhibited by the Pentium/100. This is possibly because the L2 cache is not being overrun at these user loads.
- This example shows how there is little benefit in adding a CPU when the system is *not* processor constrained. Interestingly, there is no real apparent bottleneck in the server hardware itself. The gradual slowdown could be due to the Exchange software itself.
- This data implies that the system would be able to handle other tasks besides the email and public folder tasks in theLoadSim user profile. For example, the second CPU would likely be used more effectively if a gateway or some other Exchange process is added to the system.
- Four CPU data is omitted because it provides no additional benefit in this scenario.

Pentium/166

PENTIUM/166 EXCHANGE SERVER HARDWARE CONFIGURATION

Machine Class	Compaq ProLiant 4500	
Processor Subsystem	1,2 Pentium/166, 2 Mbyte L2 cache	
System Memory	128 Mbytes	
Disk Subsystem	Integrated Fast Wide SCSI-2 controller	*Pagefile and executables volume*: 1x2GB Fast Wide SCSI-2 drive
	EISA SMART Controller	*Log volume*: 2x2GB Fast SCSI-2 drives on port 0, RAID 1 (Array Accelerator enabled) *IS volume*: 7x2GB Fast SCSI-2 drives on port 1, RAID 5 (Array Accelerator enabled)
Network Interface	Compaq NetFlex/3 (10BaseT)	
Operating System	Windows NT Server v3.51 + SP4	

Table 6. Pentium/166 Exchange Server hardware configuration

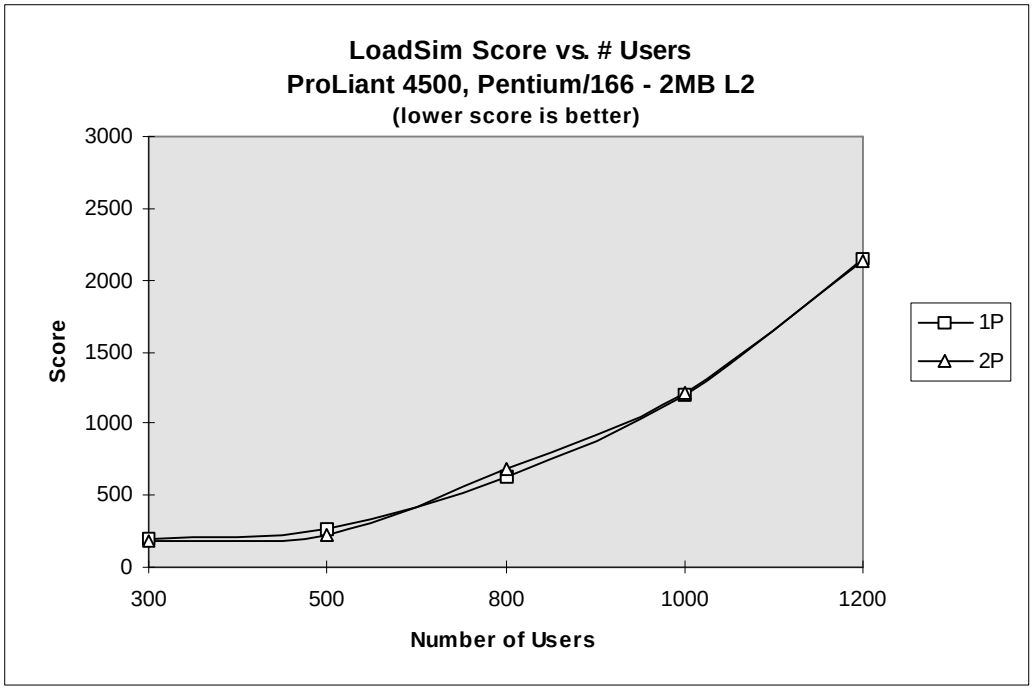

Figure 4. Relative Score performance of Pentium/166.

- The results and performance are very similar to that of the Pentium/133.
- This data again illustrates the power of the L2 cache combined with the faster clock speed of the Pentium/166. It is clear that a single Pentium/166 easily outperforms two or four Pentium/100 CPUs.
- The server is less quick to consume processor resource. 800 users can easily be supported with subsecond response time.

- In this data there is not the huge jump in Score as exhibited by the Pentium/100. This is possibly because the L2 cache is not being overrun at these user loads.
- This example shows how there is little benefit in adding a CPU when the system *is not* processor constrained. In the case of going from 1P to 2P the response time improves very little until the number of users exceeds 1000. Interestingly, there is no real apparent bottleneck in the server hardware itself. The gradual slowdown could be due to the Exchange software itself.
- For systems on which the workload is light (refer to Table 1 for workload comparisons), there is a tradeoff in adding CPUs to the processor subsystem due to system overhead. System requests and context switches increase substantially as more CPUs are added. This could also account for the marginal benefit in adding a second CPU, i.e. in this case the benefit could be offset by the additional overhead.
- This data implies that the system would be able to handle other tasks besides the email and public folder tasks in theLoadSim user profile. For example, the second CPU would likely be used more effectively if a gateway or some other Exchange process is added to the system.
- Four CPU data is omitted because it provides no additional benefit in this scenario.

Pentium Pro/166

PENTIUM PRO/166 EXCHANGE SERVER HARDWARE CONFIGURATION

Machine Class	Compaq ProLiant 5000	
Processor Subsystem	One Pentium Pro/166, 512 Kbyte L2 cache	
System Memory	128 Mbytes	
Disk Subsystem	Integrated Fast Wide SCSI-2 controller	*Pagefile and executables volume*: 1x2GB Fast Wide SCSI-2 drive
	PCI SMART-2 Controller	*Log volume*: 2x2GB Fast SCSI-2 drives on port 0, RAID 1 (Array Accelerator enabled) *IS volume*: 7x2GB Fast SCSI-2 drives on port 1, RAID 5 (Array Accelerator enabled)
Network Interface	Compaq NetFlex/3 PCI (10BaseT)	
Operating System	Windows NT Server v3.51 + SP4	

Table 7. Pentium Pro /166 Exchange Server hardware configuration

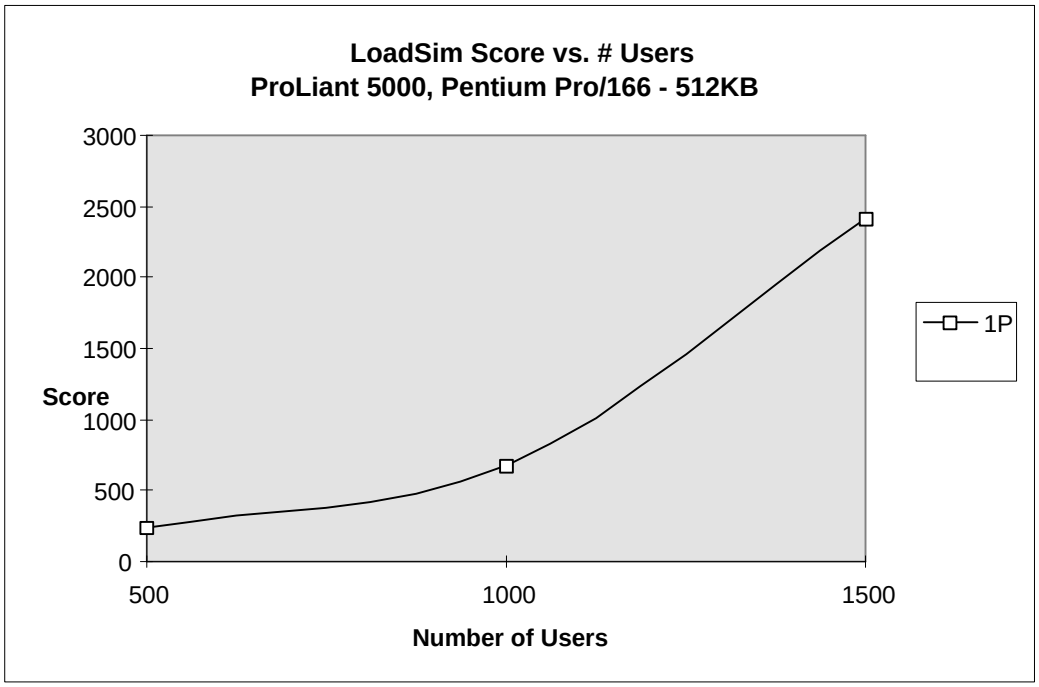

Figure 5. Relative Score performance of Pentium Pro/166.

- This data illustrates the increased power of the Pentium Pro processor with 32-bit Windows NT applications.
- Note that this system uses the PCI SMART-2 Array Controller rather than the EISA SMART Array Controller as in previous tests. Some of the performance benefit here could be attributed to better performance of the disk subsystem.
- The server is less quick to consume processor resource. 1000 users can easily be supported with subsecond response time.
- Interestingly, there is no real apparent bottleneck in the server hardware itself. However, the gradual slowdown trend continues as with Pentium/133 and Pentium/166. This could be due to the Exchange software itself, although it is faster overall than the Pentium-class processor subsystems.
- This data implies that the system would be able to handle other tasks besides the email and public folder tasks in theLoadSim user profile. For example, the second CPU would likely be used more effectively if a gateway or some other Exchange process is added to the system.

Single CPU Processor Subsystem Comparison

Figures 6 and 7 show a comparison of the performance of the different processor subsystems examined. Figure 6 compares single CPU configurations. Figure 7 compares dual CPU configurations.

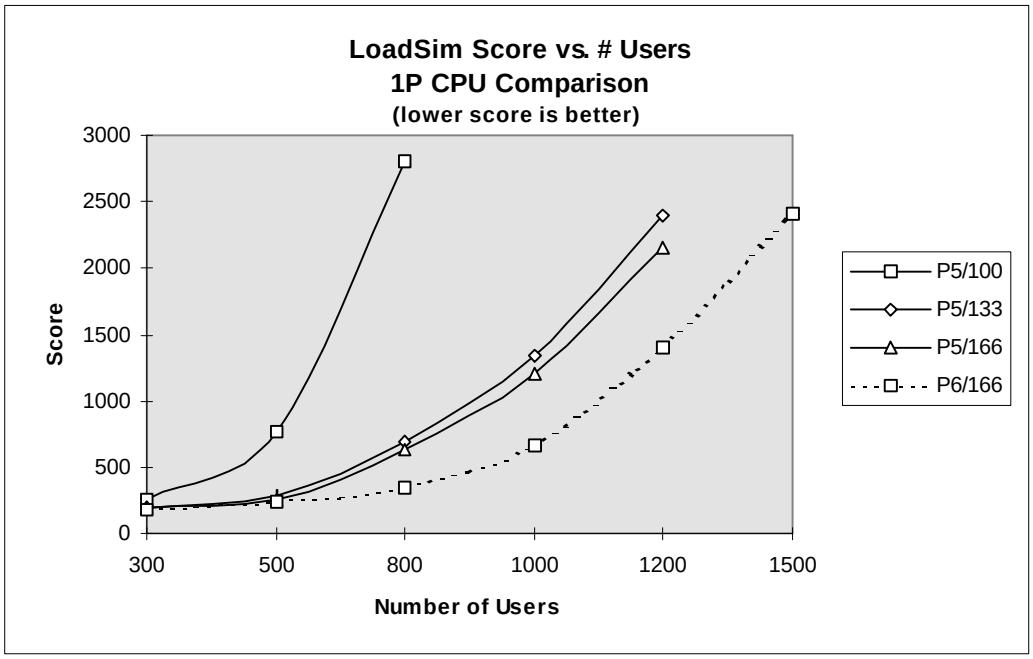

Figure 6. Single CPU performance comparison.

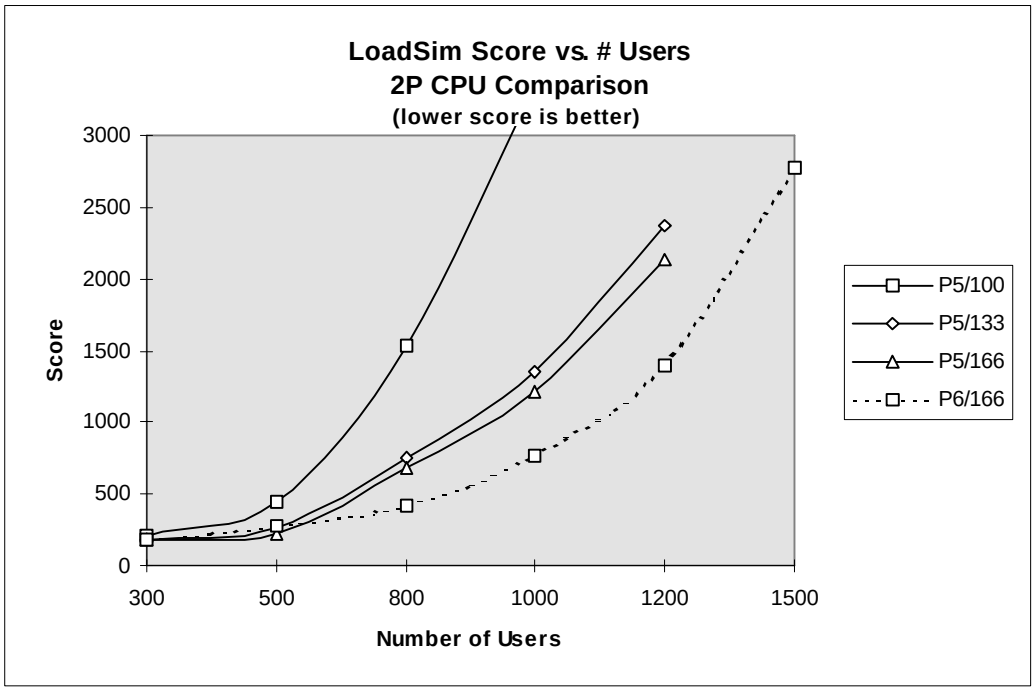

Figure 7. Dual CPU performance comparison.

- This data again illustrates the increased power of the Pentium Pro processor with 32-bit Windows NT applications. It also demonstrates the benefit of large L2 cache in the Pentium/133 and /166 CPUs (as opposed to the smaller cache in Pentium/100 CPUs.)
- The overall indication is that 1500 users is about the practical upper limit for tests using this LoadSim user profile. Over 1500 users the response time exceeds three seconds.

- There is consistency in the relative performance of both the 1P and 2P data.
- Four CPU data is omitted because it provides no additional benefit in this scenario.

Disk Subsystem

Effect of Number of Spindles

The basic idea with a disk subsystem is to make disk I/O as fast as possible. One good way to do this is to have several spindles (disks) spinning at the same time connected to a Compaq SMART Array Controller.

For related information see the Compaq white paper, 'SMART-2 Array Controller Technology' or the Compaq technote, 'Configuring Compaq RAID Controllers for Database Servers'.

This concept applies to Microsoft Exchange Server. Table 8 describes the hardware used in this test. Figure 8 depicts the benefit of adding disks to the array set to increase response time.

PENTIUM PRO/166 EXCHANGE SERVER HARDWARE CONFIGURATION

Machine Class	Compaq ProLiant 5000	
Processor Subsystem	1 Pentium Pro/166, 512 Kbyte L2 cache	
System Memory	128 Mbytes	
Disk Subsystem	Integrated Fast Wide SCSI-2 controller	*Pagefile and executables volume*: 1x2GB Fast Wide SCSI-2 drive
	PCI SMART-2 Controller	*Log volume*: 2x2GB Fast SCSI-2 drives on port 0, RAID 1 (Array Accelerator enabled) *IS volume*: 7x2GB Fast SCSI-2 drives on port 1, RAID 5 (Array Accelerator enabled)
Network Interface	Compaq NetFlex/3 PCI (10BaseT)	
Operating System	Windows NT Server v3.51 + SP4	

Table 8. Pentium Pro /166 Exchange Server hardware configuration

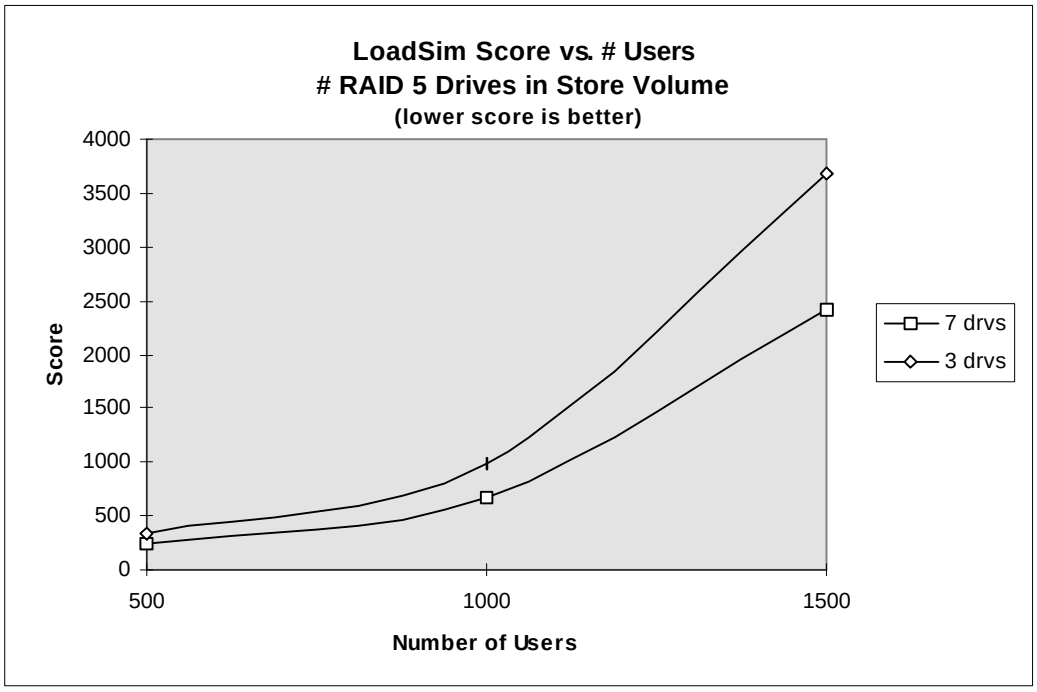

Figure 8. Performance benefit of adding disks to an array.

- Adding disk spindles to the RAID set produces a benefit in response time.
- The benefit increases as the number of users increases. This is because the number of disk I/Os required increases also, so the benefit of the array controller becomes more apparent.

System Memory

System Memory plays a crucial role in the performance of Windows NT and NT applications. It is very important to have enough system memory in an Exchange Server computer.

Effect of Increasing IS Buffers

Aside from the amount of memory required for running Windows NT and Exchange Server, system memory and disk performance are interrelated. This is because the IS buffers are allocated from system memory. A small IS buffer will constrain disk I/O whereas a large IS buffer will tend to relieve the disk I/O requirements on the disk subsystem.

PENTIUM PRO/166 EXCHANGE SERVER HARDWARE CONFIGURATION

Machine Class	Compaq ProLiant 5000	
Processor Subsystem	1 Pentium Pro/166, 512 Kbyte L2 cache	
System Memory	128/ 256 Mbytes	
Disk Subsystem	Integrated Fast Wide SCSI-2 controller	*Pagefile and executables volume*: 1x2GB Fast Wide SCSI-2 drive
	PCI SMART-2 Controller	*Log volume*: 2x2GB Fast SCSI-2 drives on port 0, RAID 1 (Array Accelerator enabled) *IS volume*: 3x2GB Fast SCSI-2 drives on port 1, RAID 5 (Array Accelerator enabled)
Network Interface	Compaq NetFlex/3 PCI (10BaseT)	
Operating System	Windows NT Server v3.51 + SP4	

Table 9. Pentium Pro /166 Exchange Server hardware configuration

In this case, the IS buffers were increased by adding 128 Mbytes of RAM to the configuration. Note that the entire additional 128 Mbytes of RAM are allocated to the IS buffer, i.e. the IS buffers were originally set to 9833 (39 MB), and they were increased to 42601 (167 MB). Each IS buffer is 4 Kbytes in size, so a net of 32768 buffers are added to allocate 128 Mbytes of RAM.

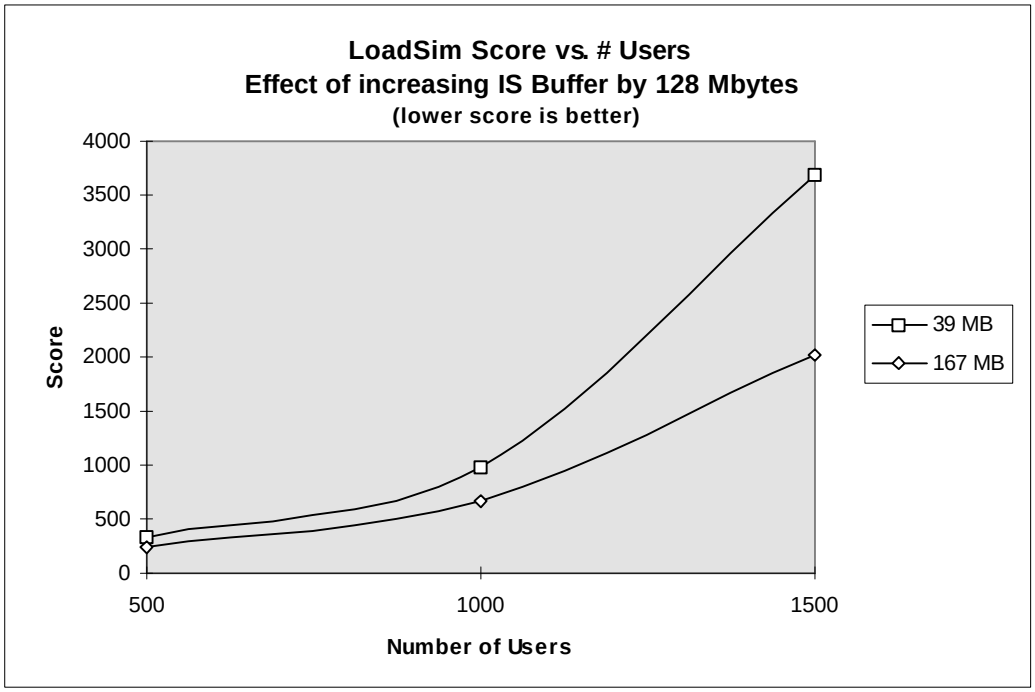

Figure 9. Performance benefit of increasing IS buffers.

- The benefit of adding extra IS buffers produces significant benefit, especially at higher user levels. This is because the extra IS buffers are relieving the disk subsystem of a certain amount of the I/O load.

- Most of the benefit from extra IS buffers will be from reads. All writes have to make it to the disk sooner or later, so the net amount of writes does not change that much. However, significant reductions in the amount of read I/O are observed, sometimes by as much as 60%.
- Depending on the cost of RAM vs. the cost of extra disks, adding RAM may be a cost effective alternative to adding disks to a RAID set in order to increase performance.

Effect of Increasing IS Buffers Too Much

If the IS buffers are increased by too much without adding extra RAM, you will cause the system to become memory constrained and start paging. This will defeat the entire purpose because response time will degrade as a result.

So if the IS buffers are increased based on free memory available in the system - rather than as a result of adding more RAM - be careful to not take away memory from system processes. There should be at least 10-15 Mbytes free memory in the system at all times.

It is always safe to run Exchange Optimizer to get a good recommendation for setting IS buffers based on system memory.

Conclusions and Recommendations

The purpose of this section is to provide general guidelines for configuring an Exchange Server computer. You should read these items and apply the data in this paper to your own configuration decisions. They are organized into four general categories: processor subsystem, disk subsystem, system memory and general.

Processor Subsystem

- Primarily, the fastest single CPU should be used for the processor subsystem. If the server will be supporting many users, and especially if the server will be handling many different non-IS processes such as gateways, multiple CPUs will be a benefit.
- The processor L2 cache provides a very significant benefit. The larger L2 caches on Pentium/133, Pentium/166, andPentiumPro CPUs help them perform exceptionally well.
- Although there is no Exchange-specific data to support the claim, the Pentium Pro/166 with 512 Kbyte L2 cache will likely outperform the Pentium Pro/200 with 256 Kbyte L2 cache.
- The Pentium Pro demonstrates its benefit with the full 32-bit code of Windows NT and Exchange Server.

Disk Subsystem

- Place the database logs on a separate physical volume from the Information Service databases. Using the two ports of a SMART controller for this works well as SMART can perform simultaneous I/O on both ports.
- The database store (IS) volume should at least be composed of an array of three disks, and it should be fault tolerant. Due to the random I/O nature of this volume, RAID 5 is a good fault tolerance configuration. Although RAID 5 imposes a performance penalty for disk writes, the ratio of disk reads to disk writes is about 2:1. This ratio can change if you have a very large IS buffer, but it generally holds true.

- The log volume must*always* be fault tolerant. Due to the sequential write-only nature of the log volume, many disks in a RAID set to not provide a significant benefit. Two high capacity disks mirrored (RAID 1) is a configuration that provides good sequential write performance and excellent fault tolerance. Note that the IS buffer has nothing to do with log volume performance.
- In most cases there is no performance degradation imposed by turning off the SMART Controller Array Accelerator cache on the Log volume. Leave it enabled on the database Store (IS) volume.

System Memory

- 128 Mbytes of RAM is a good amount of system memory to start with for the 500 user range. 256 Mbytes can provide additional benefit.
- Monitor the available memory on the system using NT Performance Monitor. If there is unused memory available on a consistent basis, allocate a portion of it to the IS buffers. However, do*not* over-allocate RAM. There should be at least 10-15 Mbytes free memory in the system at all times.
- Increasing RAM in the system and allocating it to the IS buffer can significantly improve client response time. In some cases, doing this may be more cost effective than adding disks to the drive subsystem.

General

- Based on the tests profiles used for purposes of this paper, the maximum practical number of users for a single Exchange Server is 1500. Depending upon your user load, that is most likely a worst case scenario. Results will vary with different cases.
- Always run Exchange Optimizer after the initial setup of Exchange Server. Also run it after changing configuration of the server.