COMPAQ

Lotus Domino 4.5 Performance and Capacity Planning on Compaq Platforms

TechNote

First Edition (January 1998)
Document Number ECG048/0198
Compaq Computer Corporation

Notice

Lotus Domino Server 4.5 Performance and Capacity Planning on Compaq Platforms

First Edition (January 1998)
Document Number ECG048/0198

Contents

Subsystem Performance Comparison *continued*

Chapter 4
Capacity Planning

Appendix A

Microsoft Windows NT NotesBench Results

Appendix B

IBM OS/2 Warp Server Advanced NotesBench Results

Appendix C

Novell NetWare 4.11 NotesBench Results

Chapter 1
About This TechNote

This Compaq TechNote provides the results of a performance analysis conducted by Compaq engineers on the Lotus Domino Server for Microsoft Windows NT. The information presented here is based on technical knowledge acquired by Compaq engineers while testing these products in a closely controlled environment.

This TechNote is for system integrators and network administrators with a knowledge of Compaq Server products, Lotus Domino, and Windows NT. It is a supplement to the *Compaq Hardware Reference* document and the *Lotus Domino 4.5* documentation. The results and conclusions of this TechNote provide:

- An understanding of how individual Lotus Domino user and server tasks impact overall system performance.
- Suggestions for improving your Lotus Domino Server for Windows NT performance.
- Recommendations for selecting the appropriate server hardware for your Lotus Domino Server for Windows NT.

Objective

One objective of this TechNote is to provide customers running Lotus Domino 4.5 for Windows NT on Compaq servers information to assist them in optimally configuring their server(s) to achieve the highest possible performance from their hardware and software. Information is also provided to assist customers in making configuration upgrade decisions based on an anticipated return in performance gains.

Another objective of this TechNote is to provide information to future customers to assist them in selecting the appropriate server hardware configuration for their operating environment. This data illustrates performance and system utilization that can be expected for various processor types, server memory configurations, and disk subsystem options. Customers can use this data to determine which configuration would best suit their business needs considering price and performance information.

Notation Conventions

Table 1-1 lists the conventions this TechNote uses to distinguish elements of text found within this document.

Table 1-1
Notation Conventions

Convention	Use
Screen selections, *variables*, *and new terms*	These items always appear in italics.
FILENAMES	Names of files appear in uppercase italics in DOS and in other environments.
COMMANDS, DIRECTORY NAMES, DRIVE NAMES, and PROGRAMS	These items appear in uppercase in DOS and in other environments.
USER INPUT	Information you type exactly as it appears is shown in uppercase.
NOTE:	Presents commentary, sidelights, or interesting points of information.
type	When instructed to type information, do so without pressing the **Enter** key.
Select item → item → item	Items separated by arrows indicate items you select in a sequence.

Additional Resources

Consult the following resources for additional information on obtaining the best possible performance and throughput with Lotus Domino Server for Windows NT:

- Lotus Domino 4.5x documentation

  Provides a comprehensive set of documents covering installation and reference, including an administrator's guide with detailed information on Lotus Domino 4.5x.

- World Wide Web on the Internet

  http://www.compaq.com/support/techpubs/ for Compaq technical publications

  http://www.compaq.com/products/servers/platforms.html for Compaq system information for Compaq server offerings

  http://www.notesbench.org for published Compaq NotesBench audit reports

  http://www.lotus.com. See Domino Server under Products for Domino software family information.

- *Compaq Hardware Reference* documentation.

  Provides information similar to that available on the website for all Compaq server and option offerings.

- *Optimizing Windows NT volume of Microsoft Windows NT Resource Kit*

Helps to determine bottlenecks in networks and servers. Provides an understanding of how various activities affect the performance of computer hardware. Perform capacity planning to determine your future equipment needs for performance or capacity purpose.

NOTE: This list is not intended to be all inclusive of the materials available, however these materials will be of benefit to the reader.

Chapter 2
Performance Management

Successful performance management is achieved by fully understanding the performance impact that system resources — such as the system processor, memory, and disk subsystems components — have on the overall operation of your entire system. By changing the configuration of these components, performance is affected in some way. The goal of this chapter is to help you better understand the relationship between system resources and Lotus Domino Server performance so you can make informed server configuration decisions when upgrading or purchasing a new system.

This chapter includes:

- Defining two perceptions of performance
- Describing performance analysis
- Discussing standard and customized benchmarks as a performance measuring tool
- Describing the testing methodology used during the study while focusing on Lotus NotesBench as the benchmark tool used for measuring performance of the CPU, memory, and disk subsystem

Data gathered from Lotus NotesBench testing is presented and configuration recommendations are provided based upon data analysis and the experience of Compaq engineers.

Performance Characteristics

The term performance can be viewed in either of two ways. To a system administrator, performance means effective management of system resources. Therefore a system administrator is concerned with system throughput and utilization. To an end user, however, performance is measured by system response time. In practice, it is necessary to balance the two perspectives, understanding that a change made to improve response time may require more system resources.

This chapter provides information on how Lotus Domino Server performed under various test configuration scenarios or benchmarks. Based on these tests, information is provided that can be used as a guideline for gauging the response time, throughput, and capacity expected of Lotus Domino Server running on a Compaq system.

Performance Analysis

Performance analysis is an ongoing, interactive process that is necessary for determining whether or not your server is performing as it should. Performance analysis required as a part of performance management includes:

- Understanding your user requirements
- Monitoring your server and network load patterns
- Making appropriate modifications to your configuration to achieve optimal use of resources

For the performance analysis investigation, Compaq engineers used a standard benchmark tool to examine the following Lotus Domino Server system resource areas:

- System Processor (CPU) Performance
- Memory
- Disk Subsystem
- Bus Architecture
- File Systems (NTFS vs. FAT)

Also examined during the performance analysis investigation were the following Domino Advanced Server Options:

- Partitioned Servers
- Clustering

Standard Benchmark Tool

A standard benchmark tool provides the ability to run the exact same test scenario under various operating environments to allow the comparison of one environment to another. For example, Test A executes a test script which initiates the execution of a fixed set of database or file operations for a consistent period of time on a hardware configuration, followed by the identical Test A running on another hardware configuration. The hardware configuration change implies that the processor, total system memory, network card, or disk subsystem configuration has been changed. To accurately measure the effect of configuration changes to one of these subsystems, all other variables are held constant except for the one under test.

Customized Benchmark Tool

A customized benchmark is simply an extension of the standard benchmark tool. The customized benchmark provides the capability for test engineers to pick the type of workload from a number of provided profiles which most closely matches their real world operating environment. Thus one engineer's test results with a customized set of profiles should only be compared to other tests that used the same workloads. The output of the benchmark tools is raw data which must be analyzed before any conclusions can be made.

Capacity Planning Tool

A capacity planning tool is similar to a benchmark tool in functionality, yet different in that it provides the capability to more accurately reflect "real world" system utilization by introducing the ability to customize the tests to reflect peak and low load times. This allows the test to be configured to reflect high utilization during the peak load time or times during the work day, and lower utilization during the period of the day when the system experiences less of a workload. A capacity planning tool allows these peaks and low times to be configured into the test as appropriate for any company. Rather than simply providing raw data as the output like a benchmarking tool, the capacity planning tool uses built-in intelligence that takes input provided and returns useable information as the output. For example, after running the capacity planning tool under a given scenario, the tool provides the recommended number of users as the output of the run. No analysis of raw data has to be performed by the engineer; this intelligence factor is built into the tool.

Test Methodology

NotesBench was used as the Workload Generator with the NT Performance Monitor Tool.

NotesBench

Lotus NotesBench is an implementation of a standard benchmark tool. Lotus NotesBench is a collection of benchmarks and documentation for evaluating the performance of Lotus Domino 4.x servers. The benchmarks (usually referred to as tests in this paper) model the behavior of Domino workstation-to-server or server-to-server operations. The benchmark returns measurements that enable evaluation of server performance in relation to the server system's cost.

This section discusses:

- Benchmarking Basics
- What is NotesBench?
- NotesBench Workload Basics
- Users and Threads

Benchmarking Basics

A *benchmark* is a software application that tests the performance of a computer system. Benchmarks can test the following:

- Two software applications running on the same hardware
- Different hardware platforms from the same vendor running the same software
- Different releases of software on the same machine
- Systems from different vendors running the same software

Below are some terms and definitions/descriptions that are associated with benchmarking.

- **Workloads**

  A benchmark is equivalent to a *workload* that is presented to a system under test by another system called a *driver.* Workloads consist of transactions that are executed by the software being used in the testing. Workloads typically provide a means for performing customized benchmarks intended to model various types of user activity.

- **Performance**

  A benchmark is run on several systems and the performance of each is measured and recorded. The benchmark's performance is a throughput metric — usually in units of work/time period. For example, the performance of full text search software on each benchmarked system is typically search transactions completed per second.

- **Price**

  Along with its performance metric, the price of the system under test is an integral part of a benchmark. The price is usually a metric that represents the cost-of-ownership of the system. The benchmark provides guidelines for calculating the system price.

- **Price:Performance**

  Together a benchmark's price and performance define a price:performance ratio — price divided by performance. Price:performance lets you decide which system carries out the work done by the software with the lowest cost.

- **Scalability**

  A benchmark should apply to both large and small computer systems. A benchmark should maintain a constant relationship between the workload presented to a system and the capacity of the system. As the capacity increases the benchmark's workload increases proportionally. This allows the benchmark user to scale the workload up or down to the size of the system under test.

- **Typical Results**

  Benchmark results are typically presented in the form shown in the following table. The systems under test are listed along with the performance (tps or transactions per second) and the price:performance ratio ($/tps). These benchmark results apply to different systems running the same application software.

Table 2-1
Typical Benchmark Results

System Under Test	tps	$/tps
System A	52	7.7
System B	56	19.0
System C	51	13.5

- **Benchmark Standards**

  There is a strong trend toward standard benchmarks within software areas of interest or domains. Groups of vendors define standard benchmarks for their domains. For example, the Transaction Processing Performance Council (TPPC) approves benchmarks in transaction processing and database applications; and the Standard Performance Evaluation Corporation (SPEC) defines benchmarks in the workstation and scientific areas. Lotus Development Corporation has developed a performance benchmark standard for Domino Server called NotesBench.

What is NotesBench?

NotesBench is an industry standard performance benchmark for Domino server. NotesBench falls into the last benchmark category mentioned above — different systems running the same application software. NotesBench is a benchmark developed by Lotus Development Corporation to provide a means for customers to make an “apples to apples” comparison of Domino Server running on different hardware and under various operating systems. NotesBench gives customers an objective way of evaluating the performance of different platforms running Notes. It is intended for use by hardware and software test engineers and integration consultants. NotesBench requires that vendors run the tests or workloads in the same manner. The published tests must be audited if vendors want to publish their performance results.

NotesBench provides:

- A command-line user interface for running the benchmarks
- Scripts used in the benchmarks
- A listing of parameter values for *NOTES.INI* files
- A checklist of disclosure information
- A checklist for auditors

NotesBench Basics

What Does NotesBench Test?

NotesBench evaluates different operating systems and different hardware running the same software application — Domino Server 4.5. You can also use it to compare different versions (or code builds) of Domino running on the same operating system and hardware.

NotesBench Workloads

The NotesBench software consists of a suite of benchmarks. Each benchmark maps to a workload or test, and each workload models Notes workstation-to-server or server-to-server operations. The following NotesBench workloads are available for benchmarking:

- **Mail** - A server for mail users — a workload modeling sites that rely only on mail for communication.
- **MailDB** - A server for mail and shared database users — a workload that models active users who are only performing mail and simple shared database operations.
- **GroupWare** - A server for power users who are sending large mail messages, adding documents with attachments to shared databases, performing full-text searches, and replicating changes from their local machine to the server.
- **DiscDB** - A server for active users who are performing heavy shared database operations, applying to sites that heavily utilize the collaborative features of Domino.
- **Repl_Hub** - A replication hub — a server that exists to propagate changes among a collection of other servers.
- **Mail Routing Hub** - A server that exists to route messages to other servers (a "pure" router) and possibly also to deliver messages to local users.
- **Cluster Mail and Shared Database** - A server in a cluster for power users who are sending large mail messages, adding, updating, deleting documents to shared databases with replica copies throughout the cluster.
- **Idle** - An idle usage workload that establishes an upper bound on the number of sessions (which do nothing) that a Domino Server can support. You can use this metric to aid in setting up the other NotesBench tests.

A workload is specified by running the NotesBench command NOTEBNCH followed by workload parameters. Each NOTEBNCH command executes a script consisting of a simple procedural language that presents transactions to the system under test. Each statement in the language corresponds to one or more Notes API functions. After the test has completed, the NotesBench command NOTESNUM is used to generate the performance metrics.

NotesBench Performance Metrics

NotesBench generates the same throughput metric for each of its workloads (the value of the metric changes from test to test). This metric is called a *NotesMark* and has the units transactions per minute (tpm).

Along with a NotesMark value, each workload produces a value for the maximum users supported by the test as well as the average response time.

The Results Metrics section discusses NotesBench performance results in more detail.

NotesBench Pricing

To calculate the weighted performance of a Domino Server (the price:performance ratio) you must calculate the cost of the system under test. The cost of the system under test includes all hardware components as well as the operating system and application software that was used to achieve the reported workload performance.

NotesBench Price:Performance

Price:performance measurements for NotesBench include the price:NotesMark and the price/user. NotesMark and supported users are the performance values for a NotesBench test, generated by the NOTESNUM command. The system under test price is then divided by NotesMark to provide $/transaction per minute or by users to provide $/user.

NotesBench Scaling

NotesBench tests maintain a proportional relationship between the capacity of the system under test and the workloads presented to the system. To accomplish scalability for the workloads, you set Notes environmental variables to values outlined in the NotesBench manual.

NotesBench Users and Threads

NotesBench executes its tests (workloads) by assigning Notes users on driver systems to threads in the NotesBench process. Each thread is the equivalent of one Notes user.

Each thread executes the entire NotesBench script for its workload process. Each thread executes many iterations of the same NotesBench script. If you assign 100 users to a NotesBench driver, there are 100 threads simultaneously executing the workload script.

Test Configuration and Procedure

The testbed used for the NotesBench performance testing consists of a system under test, a number of child drivers that generate the client threads, and destination servers for the workloads that require mail to be routed from the system under test to a destination server. The following diagram is an example of a NotesBench testbed.

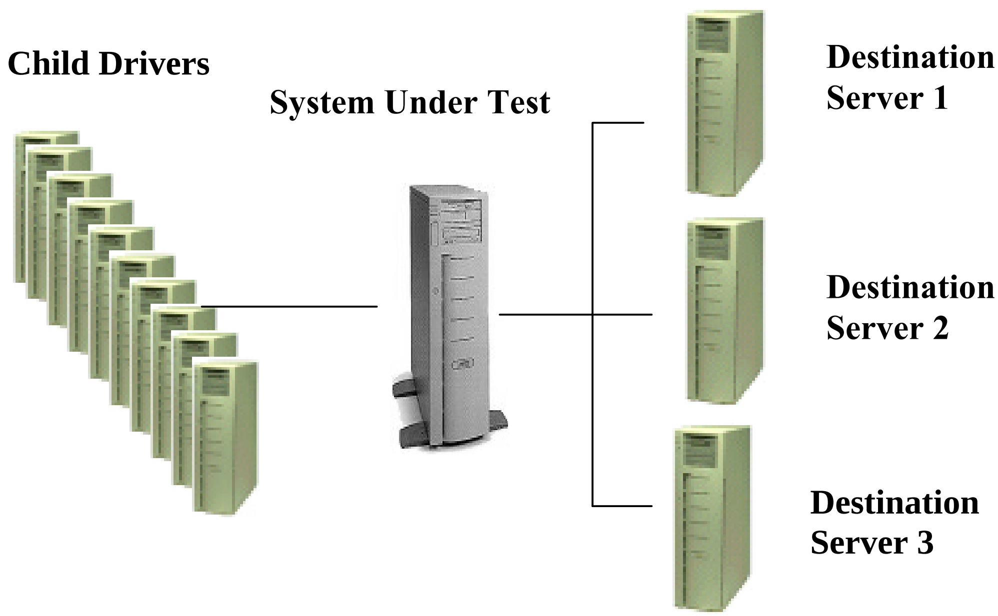

6 Driving Systems (Clients)	System Under Test	3 Destination Servers
ProLiant 2000	ProLiant 5000	ProLiant 4000
1xP5/66MHz	2xP6/166MHz	4xP5/133MHz
32-MB Memory	512-MB Memory	192-MB Memory
2x1-GB Disk	1x2-GB System Volume	1x1-GB Systems Volume
	7x2-GB Data Volume	5x1-GB Data Volume
	RAID 0	

Figure 2-1. NotesBench Test Configuration

Software Versions

- Client Systems: Windows NT Server with Domino Server 4.5 and NotesBench (Notes version 138 - Windows/32)
- System Under Test: Windows NT Server 3.51 with Domino Server 4.5 and NotesBench (Notes version 138 - Windows/32)
- Destination Servers: Windows NT Server with Domino Server 4.5 and NotesBench (Notes version 138 - Windows/32)

NotesBench Test Procedure

During NotesBench testing, Compaq Internet Solutions Engineers typically perform several trial runs to determine the best test duration and confirmation of steady state for a given test. Both test duration and steady state are determined using real time monitor utilities from Windows NT. During the trial runs, Windows NT Performance Monitor is used for monitoring system resources and for logging during the test process. The resulting data is presented as appropriate throughout this document. The results are compared with the NotesBench specification for conformity.

Actual testing began with the clients being added incrementally using a NotesBench Childstagger value setting of 5-10 minutes as specified in the parent *NOTES.INI*. The Childstagger setting caused a delay of 5-10 minutes between clients during ramp-up. Client systems sometimes used a Threadstagger value setting of 5 seconds specified in each client *NOTES.INI*. The Threadstagger setting caused a delay of 5 seconds between threads. The Childstagger and Threadstagger settings were used because the system under test typically showed tremendous stress during the workload initialization stage. Performance Monitor ran for the duration of the tests to log performance information.

Steady state was determined to be achieved during the test run by monitoring server output, the connected user threads as well as mail routing activity when appropriate. For example during a Mail workload test run that supported 1,000 users, the Domino server console displayed 1,000 users from the time the last client thread connected to the system during the ramp-up phase at the beginning of the test until the test duration was achieved.

The NOTESNUM utility results are used in the data presentation areas of this paper. Performance Monitor log files are also used when presenting resource utilization information.

Performance Monitor Tool

During the ramp-up time, the system's CPU usually reached the maximum utilization, and the system continued to be stressed until all the users were connected to the server. The processor utilization then dropped off and settled into a lower, steady rate of utilization depending on how many users were configured to run during the test. The available memory decreased steadily as the system ramped up, while the amount of system cache used increased steadily. The NT Performance Monitor parameter, DiskPerf, provides useful information about the disk usage, but it also has a significant impact on the system under test. The following Performance Monitor chart shows the typical resource usage of the system under test during a test workload of 2,850 Mail users. The ProLiant 5000 system was configured with 768-MB of memory and 2xP6/200MHz CPU with a 512-KB L2 cache.

Processor utilization averaged 83% while peaking at 97% during the test ramp-up. The chart also illustrates the symmetric multiprocessing capability of the NT operating system. The processor utilization is very balanced between the two Pentium Pro processors in the system.

Figure 2-2 is an example of the Performance Monitor results.

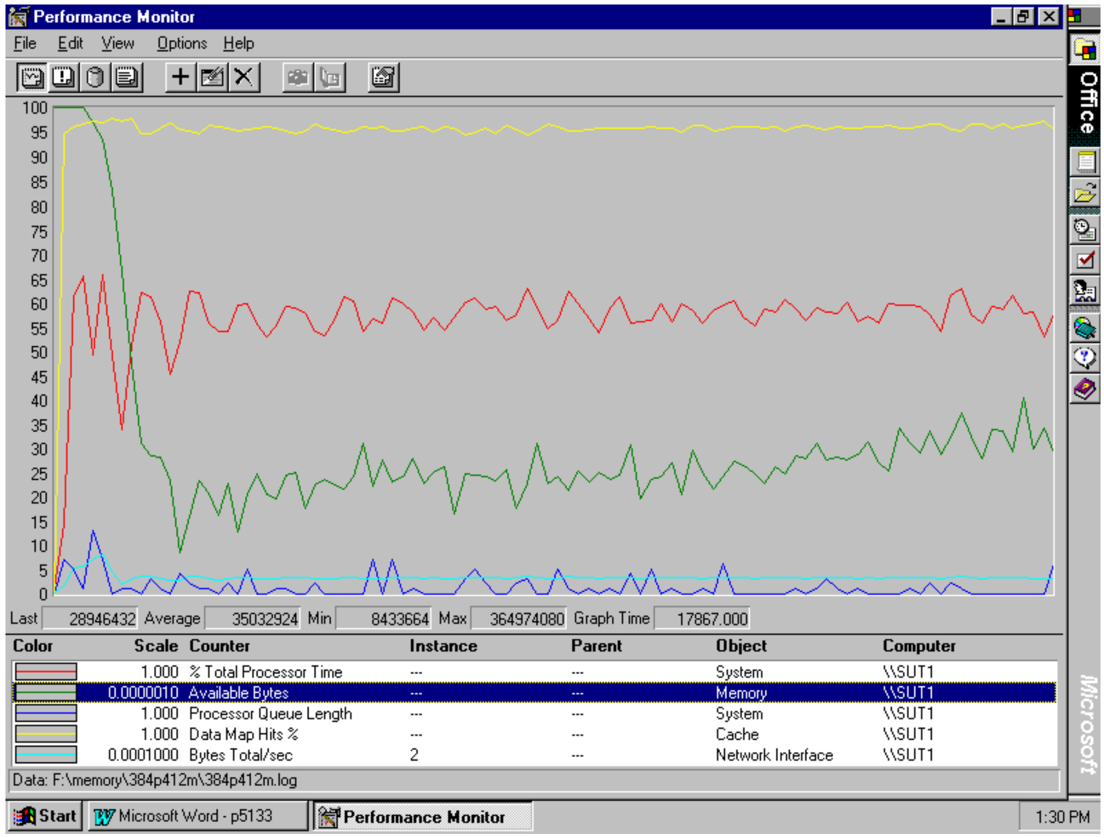

Figure 2-2. Performance Monitor Chart Showing System Resource Utilization

Chapter 3
Subsystem Performance Comparison

This section offers guidelines for obtaining the optimum value and performance from your Compaq server. These guidelines are based on tests designed by Compaq. The tests are based on the analysis of the data gathered from NotesBench testing. This section will contain a description of each of the subsystems, the data collected from testing, and recommendations for the configuration of your Compaq server. The subsystems to be discussed include:

- System Processor (CPU)
- Memory
- Disk

Note: See “Performance Summary and Recommendations” on page 3-42 for performance conclusions. Detailed information about the system processor, memory, and the disk subsystem’s impact on system performance and scalability can be found from this page through page 3-41.

System Processor (CPU)

In contrast to a resource sharing (file server) environment, a faster processor in a Lotus Domino Server for Windows NT Server yields faster client response times. In a resource sharing environment, the system processor plays a less important role in performance tuning than does the memory, disk, and network interface card. However, for Domino Server, the processor is the most important subsystem for high performance.

In the testing performed by the Compaq team, the performance of the Pentium 100MHz, 133MHz, and 166MHz processors were compared to the performance of the Pentium Pro 166MHz and 200MHz processors. As the test results will illustrate, the type of processor and its associated architecture features has as much of an impact on performance as processor rated clock speed. For example, the Pentium Pro processor offers outstanding performance that is partially attributed to the incorporation of dynamic execution features such as:

- A superscalar architecture gives the processor the ability to execute multiple instructions per clock cycle.

Writer: Krystal G. Williams Project: Lotus Domino Server 4.5 Performance and Capacity Planning on Compaq Platforms Comments: ECG048/0198

- Internal register renaming supports the execution of concurrent instructions.
- Speculative execution of branches is supported via the processor's branch target buffer which means that the processor is able to predict the correct branch in most instances, thus increasing the number of instructions that can be executed out of order.
- The processor fetches and decodes numerous instructions that are sent to an instruction pool which schedules instructions that have no dependencies on prior instructions to be executed even if the instruction is out of order.

Single CPU Performance Results

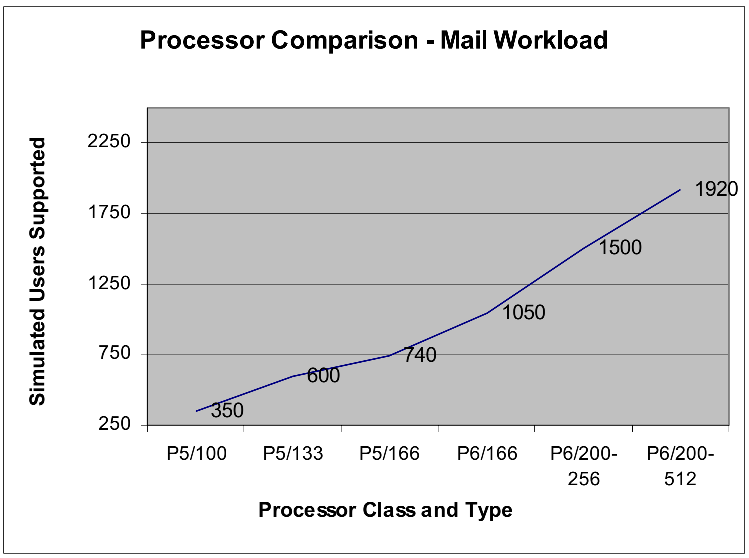

Figure 3-1. CPU performance comparison up to P6/200-512-KB L2 cache

The data charted above validates the conclusion regarding the importance of the processor. A relationship exists between number of users supported and the clock speed and cache size of particular processor, illustrating that the greater the processing power of the server, the higher number of users it can support. For example, the users supported can be increased by three times if upgrading from a P5/100 to a P6/166 processor. The results also show that 42 percent more users can be supported by the Pentium Pro processor than by a Pentium processor rated at the same clock speed of 166 MHz. The results also quantified a 20-30 percent performance gain for Domino mail activity when upgrading a server from a P6/200 MHz 256-KB L2 cache processor to a P6/200 MHz 512-KB L2 cache processor.

Table 3-1
Mail Workload Response Time Relating to Figure 3-1

Processor Type	Response Time	Number of Users
P5/100	2.392	350
P5/133	.204	600
P5/166	.145	740
P6/200-256-KB L2 cache	.24	1500
P6/200-512-KB L2 cache	.382	1920

Notice that the Pentium Pro 200 MHz processor is available with a 256-KB L2 cache and a 512-KB L2 cache. The differing cache size results in a 28% performance increase of the 512-KB cache over the 256-KB cache when looking at the number of users supported. The great performance gain from P5/100 to P5/133 is due not only to the CPU clock speed, but also to differences in the second level or L2 cache. Retired Compaq servers that featured the Pentium 133 and 166 processor board provided a 2 MB ServerCache-2 (second level cache) per processor board. Retired Compaq servers using the Pentium 100 processor provided a 512-KB ServerCache-2 (second level cache) with an optional Transaction Blaster or third level cache available. The 350 users in the testing were supported without using the third level cache option.

Processor Scalability Performance Results

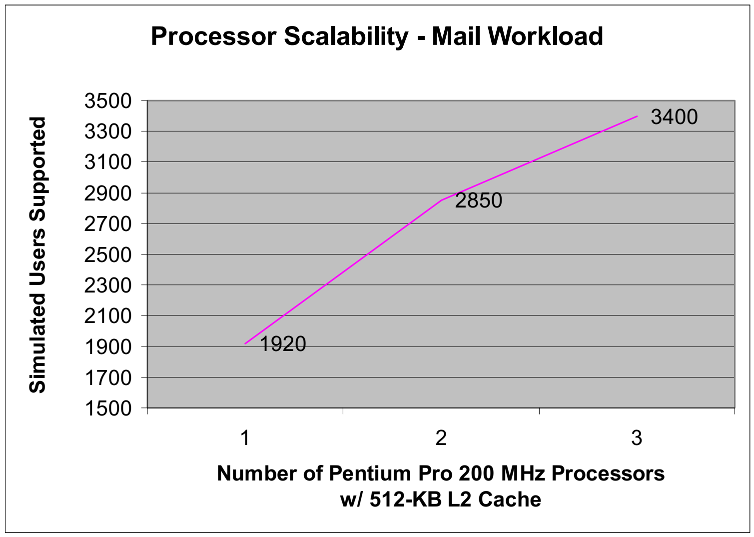

Figure 3-2. Mail Workload SMP Performance Comparison for the P6/200-512

The data charted above introduces the twist of multiple processors (Pentium Pro 200 MHz – 512-KB cache) into the performance picture. While the chart shows the number of users supported increasing as you move from 1P through 4P, the largest performance gain — a 48 percent increase in the number of supported simulated users — can be seen moving from 1P to 2P. By upgrading the system to three processors, 550 additional users can be supported. Thus the move from 2P to 3P resulted in a 19 percent increase in the number of users supported.

Note: The 4P maximum Mail user test could not be run due to problems with the benchmark testing software.

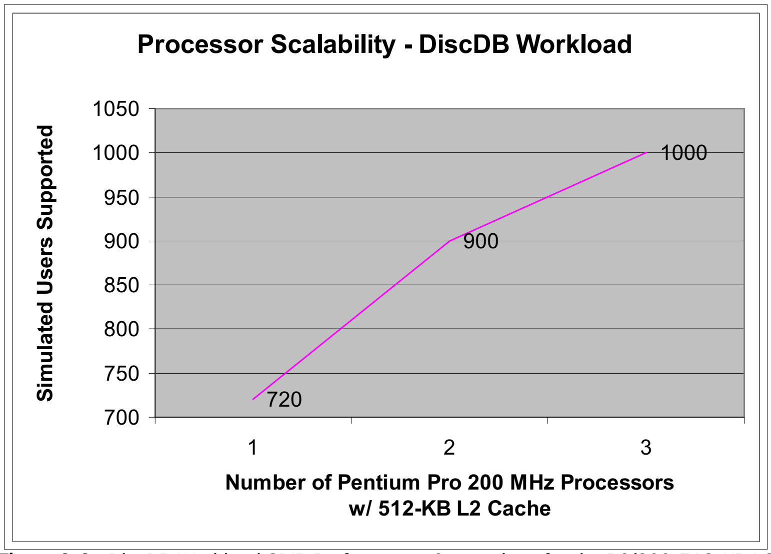

Figure 3-3. DiscDB Workload SMP Performance Comparison for the P6/200-512-KB L2 Cache

While the chart above shows the number of users supported increasing as you move from 1P through 3P, the most significant increase in the number of DiscDB users - 25% - can be seen moving from 1P to 2P. The performance gain from 2P to 3P is an increase of 11 percent for DiscDB users.

Note: The 4P maximum DiscDB user test could not be run due to problems with the benchmark testing software.

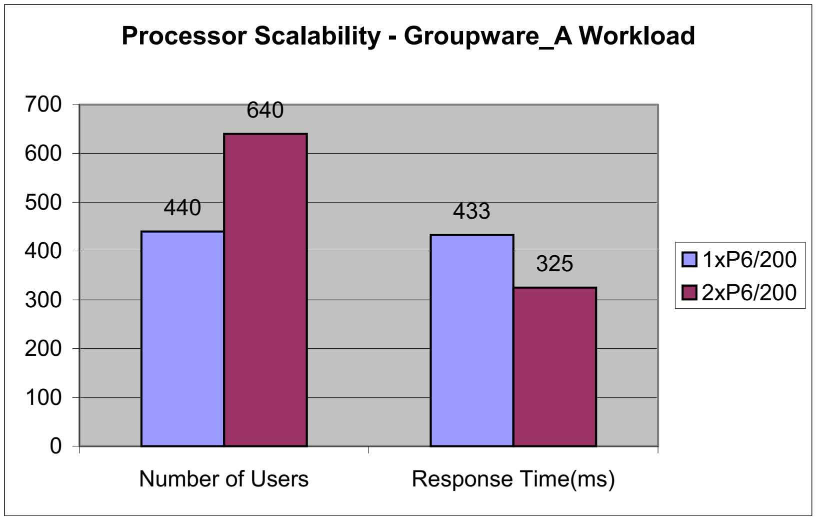

Figure 3-4. Groupware_A Workload SMP Performance Comparison for the P6/200-512KB L2 Cache.

The P6/200 chart above shows the number of users supported increasing by 45 percent as you move from 1P to 2P for Groupware_A applications.

Note: Due to limitations within the test software at the time this data was collected, only 1P and 2P scenarios could be thoroughly tested. Testing for 3P and 4P scenarios was performed, but limitations were encountered with the test software. A corrected software version that removes these limitations has recently become available and will be used in future GroupWare testing.

Memory Performance Results

Memory is a parameter that is dependent upon the total functional responsibilities of your Lotus Domino Server. At all times you must consider what background tasks your Domino Server may be performing when determining the optimal memory configuration. Also keep in mind that more memory than necessary simply means a larger pool of memory resource that your server and the processor have to manage. Therefore providing more memory than is required by Domino, the system and other applications you are running may hinder and not help performance. At the same time, you will see the performance of a server that has insufficient memory suffer due to disk thrashing as the system moves pages into virtual memory at a high rate.

Memory subsystem testing was performed for five NotesBench workloads where the number of users supported was held constant while the memory configuration varied. The purpose of the testing was to help customers better understand the performance impact of adding memory to their configurations. Tests show that the optimal memory configuration varied from one workload to the next. Another test conclusion is that more memory is not necessarily always better. However, testing showed that a system without proper memory configuration not only experiences performance degradation, but also begins to drop user connections. The proper memory configuration is therefore very important to the optimal operation of the system. Performance Monitor's report of available memory was also used in this memory analysis. The workloads included in this testing are as follows:

- Mail
- MailDB
- Groupware_A
- WebWalker
- Repl_Hub

Mail Workload Memory Testing

During the Mail Workload tests, the number of mail users was held to 1200 on a PL800 P6/200 256-KB cache. The memory configuration varied from 128-MB through 256-MB. The PL800 was configured using a Smart 2/P disk array controller. A RAID 0 array of 7 drives striped was used for Domino data. A 1024-MB system paging file was used for all PL800 Mail workload memory tests.

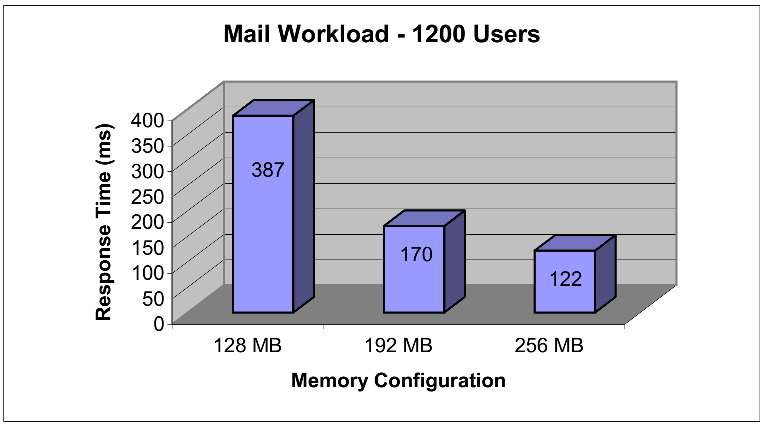

Figure 3-5. Memory Performance - Mail Workload on a PL800 1P6/200 256-KB Cache.

Notice that the 1200 user workload test shows adding memory improves response time. The response time improved, decreasing 56 percent as 64-MB was added to the 128-MB system. The response time improved, decreasing 68.5 percent as 128-MB was added to the 128-MB system, bringing total memory to 256-MB.

MailDB Workload Memory Testing

During the MailDB Workload tests, the number of MailDB users was held to 1000 on a PL5000 P6/200 512-KB cache. The memory configuration varied from 128-MB through 512-MB. The PL5000 was configured using a Smart 2/P disk array controller. A RAID 0 array of drives striped was used for Domino data. A 1024-MB system paging file was used for all PL5000 MailDB memory tests.

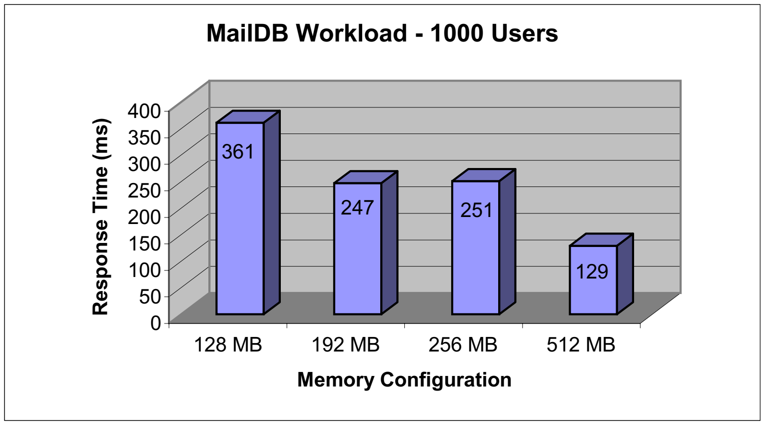

Figure 3-6. Memory Performance - MailDB Workload on PL5000 1P6/200 512KB Cache.

Notice that the 1000 user workload test shows adding memory improves response time. The response time improved, decreasing by 30.5 percent as 128-MB was added to the 128-MB system, bringing total memory to 256-MB. The response time improved, decreasing by 64 percent as 384-MB was added to the 128-MB system, bringing total memory to 512-MB. Thus the chart illustrates the improvement in response time relative to 128 31.7-MB increases in memory.

Groupware_A Workload Memory Testing

During the Groupware_A Workload tests, the number of Groupware_A users was held to 300 on a PL5000 P6/200 512KB cache. The memory configuration varied from 128 MB through 512 MB. The PL5000 was configured using a Smart 2/P disk array controller. A RAID 0 array of 7 drives striped was used for Domino data. A 1024 MB system paging file was used for all PL5000 Groupware_A memory tests.

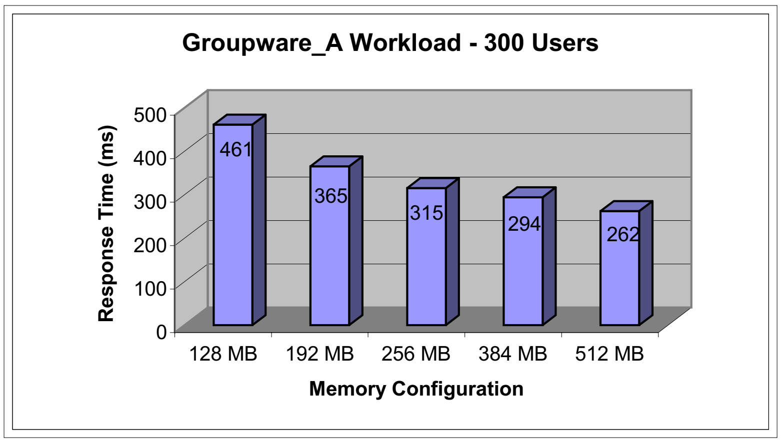

Figure 3-7. Memory Performance – Groupware_A Workload on PL5000 1P6/200 512KB Cache.

Notice that the 300 user workload test shows adding memory improves response time. The response time improved, decreasing by 31.7 percent as 128-MB was added to the 128-MB system, bringing total memory to 256-MB. The response time improved, decreasing by 43.2 percent as 384-MB was added to the 128-MB system, bringing total memory to 512-MB.

Repl_Hub Workload Memory Testing

During the Repl_Hub Workload tests, the number of spoke servers was held to 240 on a PL5000 P6/200 512-KB cache. The memory configuration varied from 64-MB through 192-MB. The PL5000 was configured using a Smart 2/P disk array controller. A RAID 0 array of 7 drives striped was used for Domino data. A 1024-MB system paging file was used for all PL5000 Repl_Hub memory tests.

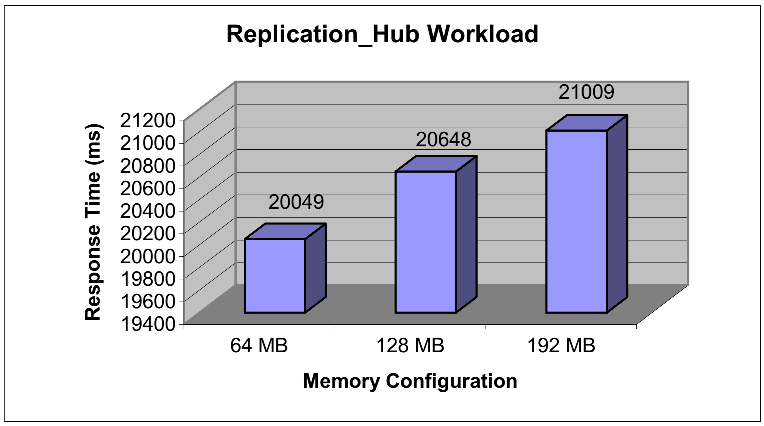

Figure 3-8. Memory Performance – Repl_Hub Workload on PL5000 1P6/200 512KB Cache.

Notice that the 240 spoke server workload test shows adding memory does not improve response time. The response time was degraded by three percent as 64-MB was added to the 64-MB system, bringing total memory to 128-MB. The response time degraded by five percent as 128-MB was added to the 64-MB system, bringing total memory to 192-MB.

While the Repl_Hub performance degraded slightly as memory was increased, the change was very small. The important information to gather from this testing is that Repl_Hub has lower memory requirements than some of the other more intensive workloads.

Memory Recommendations

Compaq engineers have provided the following memory recommendation guidelines. The tables below represent memory recommendations for mail and groupware application profiles. Domino Server configurations will obviously often need to provide support for a mixture of user profiles. The mail memory requirement represents the least memory resource intensive profile. The groupware memory requirement represents the most memory resource intensive profile. These tables can be used by the system administrator as a rule of thumb guideline for initially determining the system memory requirements. Once in production, the administrator can use a tool such as NT Performance Monitor to follow the memory resource utilization during operation to determine whether a memory upgrade is necessary. The memory aspect of capacity planning is discussed in additional detail in Chapter 4, "Capacity Planning."

Table 3-2
Memory Recommendation - Mail

Number of Users	Minimum Memory Required (MB)	Recommend Memory Configuration (MB)
150 or less	64	128
300	128	192
500	192	256
600	256	320
800	384	448
1000	512	576
1200	640	704
1400	768	832
1600	896	960
1800	1024	1088
2000	1152	1216
2200	1280	1344
2400	1408	1472
2500	1472	1536

The memory recommendations in Table 3-2 are based many NotesBench tests that were run, including the 1200 user mail workload test previously mentioned. Compaq engineers used the mail workload run optimal memory configuration findings and added 64-MB to derive the recommended memory configuration. The optimal memory findings are listed as minimal memory required for varying number of users. The recommended memory configuration amount includes 64-MB added for other background tasks.

Table 3-3
Memory Recommendation – Groupware

Number of Users	Minimum Memory Required (MB)	Recommend Memory Configuration (MB)
150 or less	192	256
300	320	384
500	576	640
600	704	768
800	960	1024
1000	1216	1280

The memory recommendations in Table 3-3 are based on several NotesBench test that were run, including the 300 user Groupware_A workload test previously mentioned. Compaq engineers used the Groupware_A workload for NotesBench optimal memory configuration findings and added 64MB to derive the recommended memory configuration. The optimal memory findings are listed as minimal memory required for varying number of users. The recommended memory configuration amount includes 64-MB added as a cushion to support other background tasks.

Disk

The disk subsystem has an impact on performance for all applications. The amount of I/O required by your application determines the degree of impact on the disk subsystem performance. Since Lotus Domino is a very I/O oriented application, the disk subsystem is an important contributor to overall system performance. Determining the impact of the disk subsystem involved analyzing the following options:

- File System Type – FAT vs. NTFS
- Single Drive Spindle vs. Striping over 4, 7, and 14 Drives
- Hardware Striping versus Software Striping
- Fault tolerance: RAID 0, RAID 1, RAID 5
- Array Controller Accelerator Read/Write Ratio

File Systems (FAT versus NTFS)

The file systems supported by Windows NT include the File Allocation Table (FAT) File System and NT File System (NTFS). The file system decision for your Lotus Domino Server depends on a number of factors related to the features of the file system. A brief description of the differences between the file systems is included for your information.

FAT File System

The FAT file system is also supported by DOS and therefore is a very common file system for the PC arena. The FAT file system is limited to a 8.3 naming convention and thus does not support long file names. No NT File or Directory Security can be set for a partition formatted as a FAT file system under NT Server. An NT Server that has the NT system partition formatted as FAT can be easily accessed when booting from a DOS diskette even if NT system configuration changes are necessary.

NT File System (NTFS)

The NT File System (NTFS) is supported by NT Server. NTFS can handle DOS names in the typical 8.3 format as well as long file names. Additionally, NTFS supports both file and directory security so that the data on the server can be protected to the degree appropriate for your application and business operations. A partition formatted as NTFS cannot be accessed from DOS. Therefore if configuration problems are encountered which prevent your server from starting NT successfully, you cannot access that drive's files when booting up from a DOS diskette.

NTFS is typically thought to be the faster operating file system when running under NT Server. The testing performed by Compaq engineers shows that this depends on the exact disk configuration. The performance impact of FAT versus NT file systems is presented with the following response time, CPU utilization, and disk queue maximum length data. These tests were run using the MailDB workload because of the amount of disk I/O for this workload. The tests were designed to support 1,000 users and were run on a ProLiant 5000 1xP6/200-512-KB L2 cache system configured with 512-MB of memory. The system tested was configured with a Smart/2P array controller and 7 drives were used in a RAID 0 array.

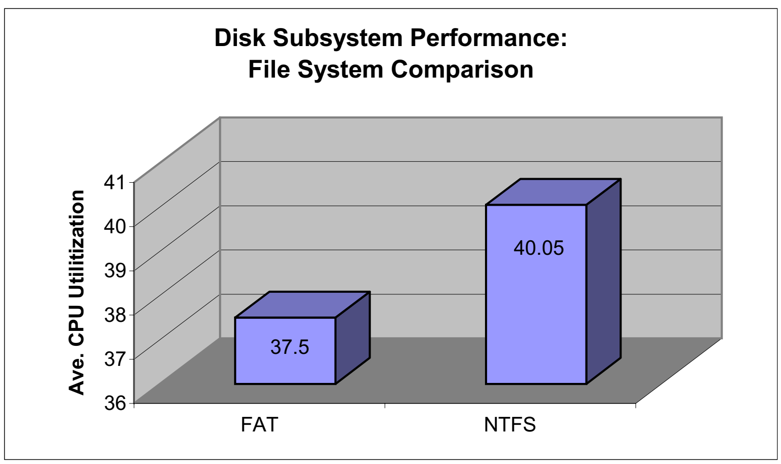

Figure 3-9. FAT versus NTFS comparison of response time for RAID 0 seven drive array running the MailDB workload.

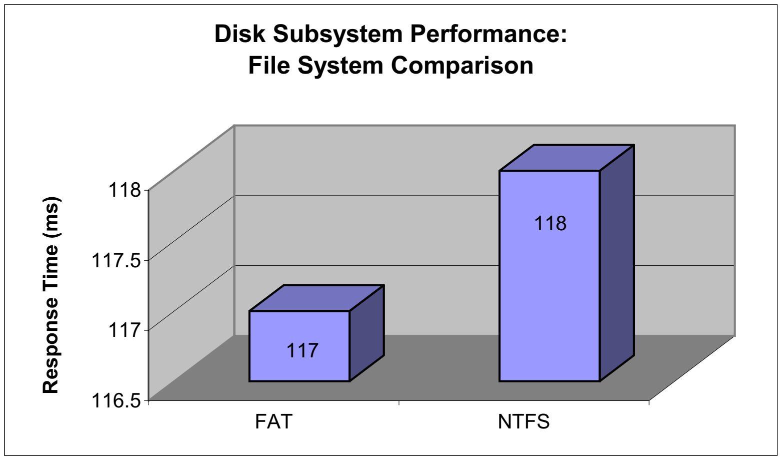

Figure 3-10. FAT versus NTFS comparison of CPU utilization for RAID 0 seven drive array running the MailDB workload.

FAT and NTFS testing resulted in very similar performance. As illustrated in Figure 3-10, the CPU utilization of the FAT system was slightly lower than the NTFS system. As illustrated in Figure 3-11, the FAT disk queue length was 33% longer than the disk queue length for the NTFS system. Due to this combination of better response time performance and lower CPU requirements, Compaq engineers recommend using NTFS as the file system for the Lotus Domino Server. While the difference in the performance of this test comparing NTFS and FAT was not tremendous, Compaq engineers have seen NTFS perform as much as three times faster than FAT in previous testing while using very little CPU bandwidth for support. Since disk I/O is a resource which is often a bottleneck in a Domino installation, NTFS is recommended as the file system.

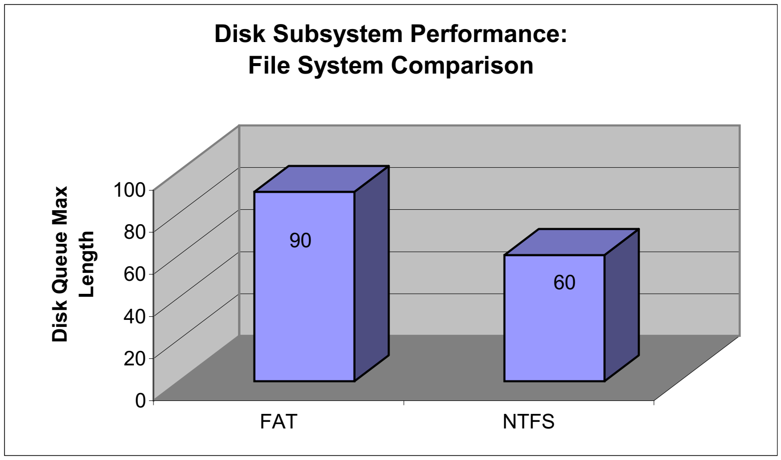

Figure 3-11. FAT versus NTFS comparison of Disk Queue Maximum Length for RAID 0 seven drive array running the MailDB workload.

Drive Spindles/Striping

If your applications generate significant disk I/O, there is probably more concurrent use of system services. You can improve the performance of your disk subsystem under load conditions by having your hardware logical drive span multiple physical drives using “striping”. Striping allows the data to be written “across” a series of physical drives which is viewed by the system to be one logical drive. This data distribution across drives makes it possible to access data concurrently from multiple physical drives that have been defined to be one logical drive array.

You achieve performance gains when you read from or write to the drive after the series of physical drives is united into one or more logical drive arrays. By distributing the data or “striping” the data evenly across the drives, it is then possible to access data concurrently from multiple drives in the series or “array”. The concurrent access of the data leads to higher I/O rates for the drive arrays than the spindles, thus improving your total system performance.

Table 3-4
Drive Spindle Performance Comparison
Mixed Load

Drive Configuration	Response Time (Seconds)	Average CPU Utilization Rate (%)
One Drive	631	39.2
4 Drives Hardware Striping (No fault tolerance)	137	43.9
7 Drives Hardware Striping (No Fault Tolerance)	129	39.2
14 Drives Hardware Striping (No fault tolerance)	124	40.9

Table 3-4 illustrates how multiple drives in a logical array can improve the response time by up to 500% over the response time when accessing a single drive. The response time increases by over 10% when comparing a logical drive consisting of an array of 4 drives to a logical drive consisting of an array of 7drives. If comparing striping across 7 drives to striping across 14 drives, the performance difference is slightly over a 4% increase in response time.

A graphical representation of the information in Table 3-4 illustrates the performance variance in response time when comparing the single drive spindle to a logical drive consisting of four, seven, and 14 physical drives.

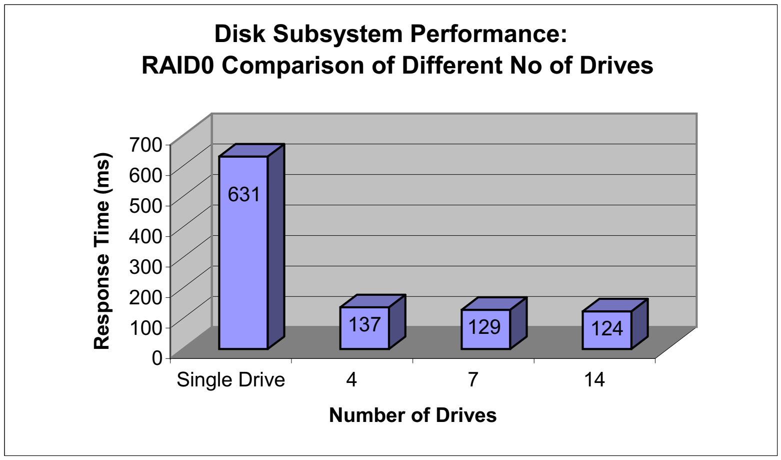

Figure 3-12. Performance Comparison of Single Drive vs. Drive Array of 4, 7, and 14 Drives

Figure 3-12 illustrates how striping across multiple drives in a logical array can improve the response time by up to 500% over accessing a single drive. The response time increases by over 10% when comparing a striping across 4 drives to striping across 7 drives. If comparing a 7 to a 14 drive striping scenario, the difference is slightly over a 4% increase in response time.

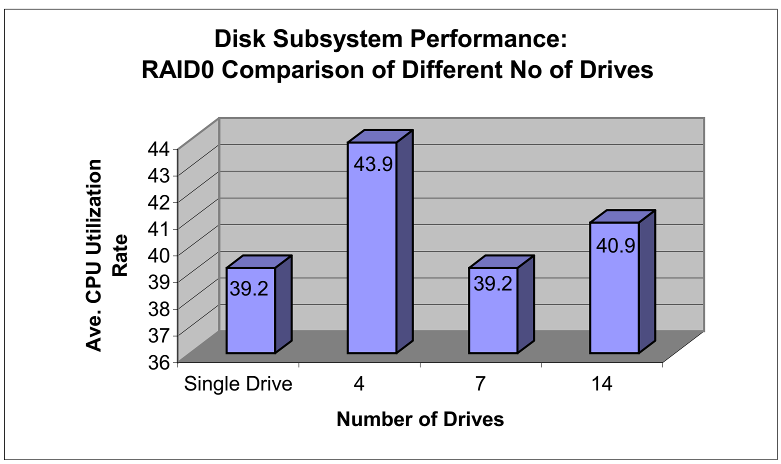

Figure 3-13. CPU Utilization Rate Comparison of Single Drive vs. Hardware Array of 4, 7, and 14 Drives

Figure 3-13 illustrates that using striping across 4, 7 or 14 drives, rather than using a single drive spindle, can achieve better performance with very minimal cost from a resource utilization standpoint. Compaq strongly recommends striping numerous smaller drives rather than a few large drives to achieve better performance with comparable overall storage capacity.

Hardware versus Software Striping

When referring to the action of *striping* data across logical drive arrays or volumes as they are sometimes called, hardware or software can control the striping process. Compaq recommends hardware striping over software striping due to a number of advantages. The main advantage that users appreciate from hardware striping is a performance gain as well as additional protection of their data. The Compaq option that supports hardware striping is the Smart-2 Array Controller which has a number of data protection features built-in to the controller as well as 4-MB of read/write cache. The new Smart-2DH controller has 16-MB of read/write cache. The performance tests included in this paper include the Smart-2 Array Controller. Testing has begun using the Smart-2DH Array Controller and performance results will be included in future papers. All testing for this paper was performed using the Smart-2/P Array Controller except for the Partitioned Server testing.

Software striping is controlled by the operating system. In these tests, NT Disk Administrator offers the capability to do software striping by specifying Create Stripe Set from the Partition menu option. The disadvantage of software striping is that the system must bear the entire burden of managing this software striping for all disk I/O. The additional system work of managing the software striping decreases the overall system performance.

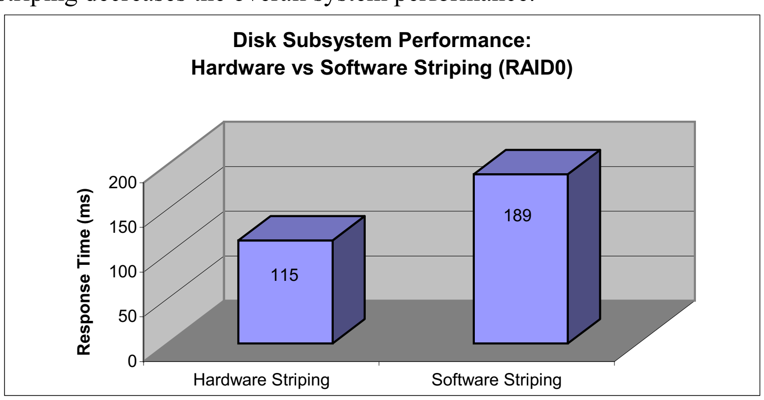

Figure 3-14. Response time comparison of hardware vs. software striping

Figure 3-14 indicates a comparable response time between hardware and software striping. Notice the 64% increase in response time when software striping is used versus software hardware striping. Hardware striping provides better performance (shorter response time) than software striping. This can be attributed to the overhead that is required of the operating system to manage the software striping.

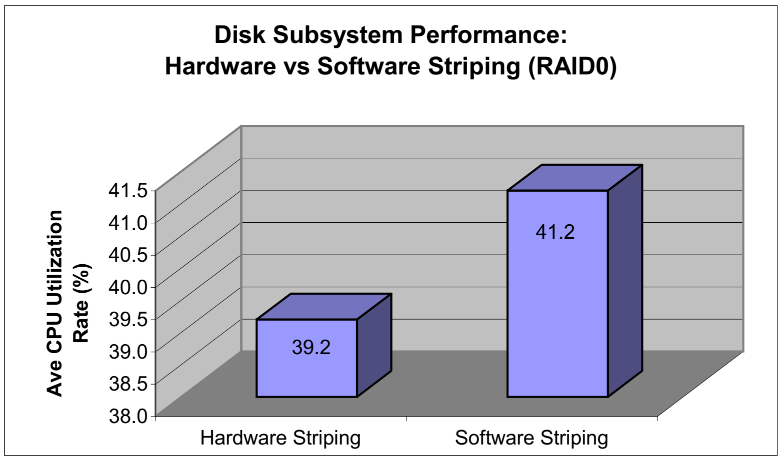

Figure 3-15. CPU Utilization comparison of hardware vs. software striping

Note, however, that hardware striping uses slightly less CPU bandwidth. With more users, it is reasonable to assume that software striping performance will decrease. Testing showed an improvement exceeding 60% in performance and about five percent lower CPU utilization rates using hardware striping. We strongly recommend the use of hardware striping for better performance and more efficient use of system and operating system resources.

Fault Tolerance

You have several available options when configuring the Lotus Domino Server and making a decision about the level of fault tolerance the system requires. Redundant Arrays of Inexpensive Disks (RAID) level is a term used to refer to an array technology that provides data redundancy to increase the overall system reliability and performance. The fault tolerance method that you select affects the amount of available disk storage and the performance of the drive array. The following levels of fault tolerance support are available:

- RAID 5 - Distributed Data Guarding

- RAID 4 - Data Guarding
- RAID 1 - Disk Mirroring
- RAID 0 - No Fault Tolerance Support

The Compaq Smart-2 Array Controller is needed to support hardware striping and all levels of fault tolerance support. Features offered by the Compaq Smart-2 Controllers follow:

- Support for RAID 0, RAID 1, RAID 4 and RAID 5 Hardware Striping and Fault Tolerance
- Dual Fast-Wide SCSI-2 channels on a single board support up to 14 drives (7 per channel)
- Support for multiple logical drives per drive array
- Removable Array Accelerator - battery-backed 4MB Read/Write cache with Error Checking and Correcting (ECC)
- Read-ahead caching
- Online Capacity Expansion and Disk Drive Upgrades
- Fault Management Features

RAID 0 - No Fault Tolerance

No fault tolerance is provided in a disk configuration that utilizes RAID 0. The data is still striped across the drives in the array, but it does not include a method to create redundant data. If one of the logical drives fails, data on that drive is lost. None of the logical drive capacity is used for redundant data, so RAID 0 typically offers the best processing speed as well as capacity. RAID 0 is appropriate for applications that deal with non-critical data requiring high speed access.

RAID 1 - Drive Mirroring

RAID 1 is also referred to as drive mirroring. This is typically the highest performance fault tolerance method. RAID 1 is the only option for fault tolerance if no more than two drives are selected. Drive mirroring works as its name implies, storing two sets of duplicate data on a pair of disk drives. Therefore RAID 1 always requires an even number of disk drives. From a cost standpoint, RAID 1 is the most expensive because 50 percent of the drive capacity is used for fault tolerance.

If a drive fails, the mirror drive provides a backup copy of the data and normal system operation is not interrupted. A system with more than two drives may be able to withstand multiple drive failures as long as the failed drives are not mirrored to one another.

RAID 4 - Data Guarding

RAID 4 is referred to as data guarding because it uses parity data to guard against the loss of data similar to RAID 5, which is called distributed data guarding. The difference is that RAID 4 writes all of the parity data to a dedicated single drive in the array. If a drive fails, the parity data and the remaining functioning drives with data use the parity information to reconstruct data from the failed drive. A problem would arise if the drive(s) that failed contained the parity information. The parity drive data could not be reconstructed.

The usable disk space for a Raid 4 configuration is the same as for RAID 5 mentioned below, but different in that the space required for parity is on a single drive. For example if you have a three-drive system, then 33 percent of the total drive space, one drive, would be used for fault tolerance. In this case, two drives would store data, and one drive would store parity data. There is a maximum of 14 drives supported. Of these 14 drives, 7 percent of total space available, or one drive, would be used for parity data. Writing all the parity data to a single drive also introduces a degradation in performance, since parity data is not being striped across all drives as when using RAID 5. Therefore disk configuration using RAID 4 for fault tolerance has historically proven to yield performance results at a level below RAID 5 configurations. For this reason, the testing performed and results included in this paper do not include Raid 4 configurations.

RAID 5 - Distributed Data Guarding

RAID 5 is referred to as *distributed* data guarding because it uses parity data to guard against the loss of data. The parity data is distributed or striped across all the drives in the array. RAID 5 therefore provides very good data protection because if a single drive fails, the parity data and the data on the remaining drives is used to reconstruct the data on the failed drive. With Compaq Smart-2 controller technology this reconstruction process allows the failed drive to be replaced while the system continues to operate at a slightly reduced performance. RAID 5 also offers good performance because spreading the parity across all the drives allows more simultaneous read operations.

The usable disk space when using RAID 5 depends on the total number of drives in the array. If there are three drives, 67 percent of the disk space is usable for data with the remainder being used to support fault tolerance. If there are fourteen drives, then 93 percent of the disk would be available. The tests that follow used seven drives.

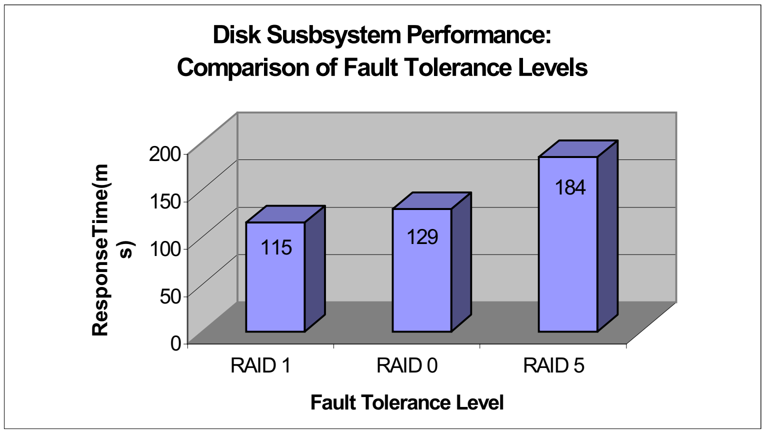

Figure 3-16. Fault tolerance level response time comparison

Compaq tests were performed to enable the fault tolerance levels of RAID 0, RAID 1, and RAID 5 to be compared. The tests were run on a ProLiant 5000, 1P6/200 – 512-KB L2 cache, 512-MB of RAM. The test was run to simulate 1000 MailDB users. The RAID level configured for the test drive array applies only to the Notes data volume. The Domino and system volume including paging file space was configured the same for each test run using two internal SCSI drives.

The RAID 0 and RAID5 tests were performed using seven 2.1-GB drives in an array. The RAID 1 test was run using fourteen total drives, with seven drives being a mirror of the other seven. The RAID 1 14 drive configuration required two Smart-2P Array controllers to be used in the test system. Each Smart-2 was controlling seven drives which explains why the RAID 1 configuration appears to have outperformed the RAID 0 configuration. The RAID 0 and RAID 5 configurations used seven total drives all under the control of a single Smart-2P Array controller.

Compaq test results show that there is a slight but measurable difference in response time rates between RAID 5 and RAID 1 or RAID 0. RAID 1 achieved the best performance, outperforming RAID 0 by 13% in response time and performing about 60% better than RAID 5. Keep in mind that while RAID 0 does utilize available disk space most efficiently, this level offers no fault tolerance protection. Based solely on response time, the recommendation is to use RAID 1 over RAID 5 for systems with critical data because of the performance gains expected combined with the hardware fault tolerance protection. RAID 5 is frequently the choice of customers for economic reasons because it offers better usage of disk capacity than RAID 1. The CPU utilization between RAID 0, RAID 1, and RAID 5 should also be considered.

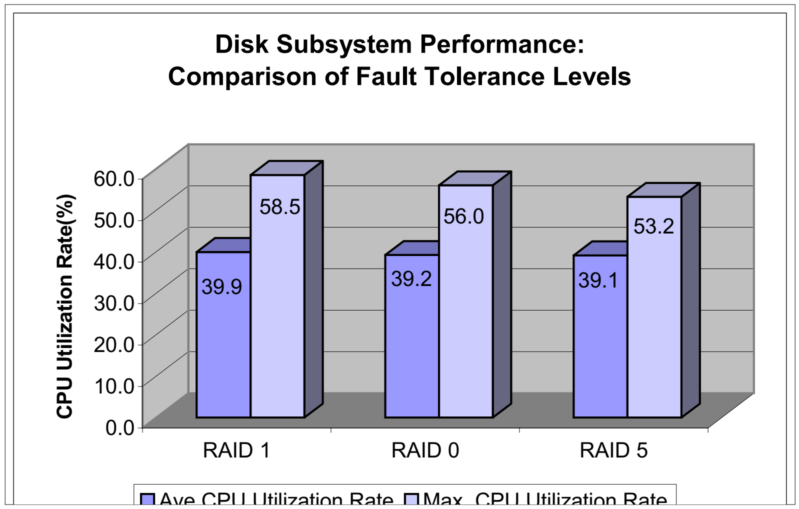

Figure 3-17. Fault tolerance level comparison of CPU utilization

Compaq tests results quantified these performance differences as follows. RAID 1, RAID 0 and RAID 5 tests illustrated almost identical average CPU utilization rates. RAID 1 maximum CPU utilization rate peaked slightly higher than RAID 0 with RAID 5 having the lowest peak in CPU usage. Overall the variance in maximum CPU utilization was only 5.3 percentage points or 9% of the total utilization.

For systems with data that is not critical, RAID 5 is recommended because of its fault tolerance support with distributed parity data performing very well at the lowest system utilization cost. For systems with mission critical data, RAID 1 is recommended because it provides the system with proper protection at the cost of slightly more system overhead combined with excellent performance. As previously mentioned, the main cost involved with RAID 1 is the usable disk storage capacity being 50% of the overall total disk space due to mirroring. For systems with frequent backups, non-critical data and replication, RAID 0 might be considered because of its good performance as well as its total usable data storage equating total disk storage.

Bus Architecture (PCI versus EISA)

The Compaq Pentium Pro processor-based servers have PCI bus implementations to meet the needs of different classes of applications as well as an EISA bus for backward compatibility. As mentioned in the controller discussion, the Smart-2 Array Controller is available for the PCI or EISA bus architecture. For very demanding critical business applications, the Compaq ProLiant 5000 uses dual, peer PCI buses that provide a total throughput capability of 267-MB/s to ensure the balanced performance of four processor systems. A bridged PCI bus implementation provides a total throughput capability of 133-MB/s and is typically suited more for two-processor systems.

The impact of the Smart-2 Controller bus type selected for a configuration could have a measurable difference in performance, particularly in the Compaq ProLiant 5000 using the dual, peer PCI bus architecture. Compaq engineers have seen the use of PCI disk controllers versus EISA disk controllers boost overall system performance from 10-15% in previous tests. On previous Compaq ProLiant servers, there was no marked difference in performance of EISA versus PCI because they did not use the dual, peer PCI architecture of the ProLiant 5000. Due to performance differences being quantified in the past, this performance test was not repeated.

Domino Advanced Server Features Performance Comparison

This section offers guidelines for obtaining the optimum value and performance from your Compaq server when using Domino Advanced Server features. These guidelines are based on tests designed by Compaq engineers. The tests are based on the analysis of the data gathered from NotesBench testing. This section contains a description of the advanced server feature tested, the data collected from testing, and recommendations for the configuration of your Compaq server when using this feature. The advanced server feature discussed in this paper is Domino Partitioned Servers.

The Domino Server Clusters advanced server feature will be discussed in a future paper.

Domino Partitioned Servers

The Domino partitioned server is an advanced server feature that enables multiple Domino server sessions to run on a single computer. In other words, one physical computer may have two or more instances of Domino Server running simultaneously. Running a partitioned server requires a Lotus Domino Advanced Server license. Lotus Domino Release 4.5 supports partitioned servers on Windows NT and UNIX operating systems.

There are several advantages to partitioning a single computer into separate Domino servers. Benefits to partitioned server features include:

- Full Domino security for users of partitioned servers
- Reduced number of computers to own and administer in order to support independent groups of users
- Easy migration from partitioned servers to individual servers

Partitioned servers may provide the scalability and security needed for an enterprise Notes system. The use of partitioned servers provides administrators a way to give independent groups of users their own private Notes domain on a shared computer. At the same time, the administrator can reduce the number of computers relative to the number of supported users. As the Notes system grows, users can be migrated from partitioned servers to individual servers if necessary.

Domino 4.5 supports up to six partitioned servers on a single computer. The number of partitioned servers installed depends on user requirements and on the ability of a computer to run more than one Domino server. The *Lotus Domino Install Guide for Servers* describes the requirements for a Domino partitioned server.

Partitioned servers on the same computer can belong to different Notes domains. If user requirements include high availability of Notes databases, a server can also be a member of a cluster. Domino server clustering will be discussed in detail in the next section.

Lotus recommendations concerning partitioned server include carefully planning the naming scheme you use and if possible, avoid changing a server name after you install. When you use the TCP/IP protocol and choose to assign a unique IP address to each partitioned server, Lotus recommends that you use the computer host name as the name of each partitioned server. In general, choose names that are easy to associate with specific servers, for example, names of departments or customers. Anyone performing an administrative task should be able to immediately distinguish one partitioned server from another.

When installing a partitioned server for the first time, a name must be specified for the program directory, which defaults to *Notes*, and the data directory, which defaults to *Data*. Domino partitioned servers use different Notes data directories and NOTES.INI files. The second installation of partitioned server on the same computer requires naming the data directory only. The same program directory is used by both partitions. Domino increments the data directory name specified during the initial install. For example, if the default data directory name data was accepted during the initial installation of partitioned server, the second install of partitioned server defaults to *data2*. The data directory entry can be edited and changed to another name if the default is not desired.

During Notes Setup, a unique partition number process is automatically assigned to each partitioned server. The partition number is between 1 and 99. This partition number guarantees separate system resources for each partitioned server installed. Each partitioned server has its own administrative client that uses the same partition number used by its corresponding partitioned server and shares its system resources and Notes data directory. For instance Client1 is used for administration for Server1. On Windows NT systems, one administrative client can be run at a time on a single computer. The Show Server command displays a server's partition number.

For more information, see the *Lotus Domino Install Guide for Servers*.

Partitioned servers operated completely independently of one another even though residing on the same computer. If one partitioned server on a computer shuts down or encounters a fatal error, the other partitioned servers running on the same computer are not affected. If one server shuts down, the others continue to run. If a partitioned server encounters a fatal error, an automatic cleanup procedure allows the server to restart without restarting the computer.

Test Setup

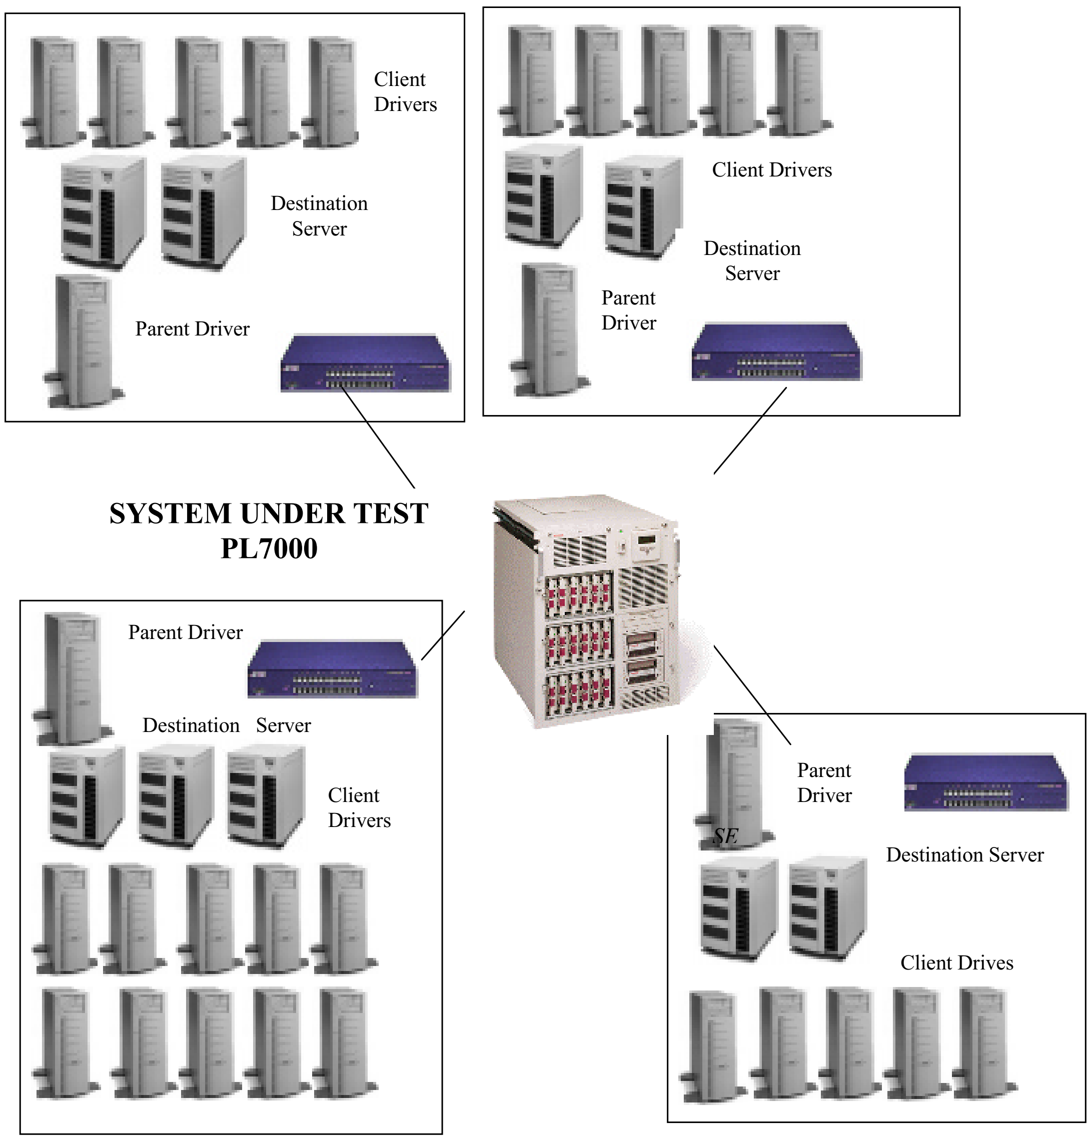
Client
Drivers
Destination
Server
Parent Driver
Client Drivers
Destination
Server
Parent
Driver
SYSTEM UNDER TEST
PL7000
Parent Driver
Destination Server
Client
Drivers
Parent
Driver
Destination Server
Client Drives

System Under Test

ProLiant 7000
4xP6/200MHz
1536-MB Memory

1x2.1-GB System Volume, & paging File NTFS (C:)

7x2.1-GB Data Volume, NTFS RAID 0 (E:)

6x2.1-GB Data Volume, NTFS RAID 0 (F:)

6x2.1-GB Data Volume, NTFS RAID 0 (G:)

5x2.1-GB Data Volume, NTFS RAID 0 (H:)

Testing was performed to determine the performance impact on the underlying computer resources of this advanced server feature and to verify proper operation of the partitioned server feature on Compaq hardware.

NotesBench was the tool used to help determine the performance impact of running four partitioned server instances on a single ProLiant 5000 system. The ProLiant 5000 was configured with four Pentium Pro 200 MHz 1MB cache processors, 1536MB RAM, 4 Smart-2DH controllers, 18 internal drives and 7 external drives, configured in four RAID 0 data arrays, and four network controllers. One NIC connected the first Domino Server partition to segment one Parent, Child, and Destination servers.

Partitioned Server Performance Results

By using the Domino Advanced Server Partitioned Server feature, Compaq engineers were able to obtain the highest number of NotesBench Mail workload users supported on a single system to date - 5150. The Partitioned Server feature also demonstrated the most effective use of four processors in a single system to date. Previous processor scalability testing showed that the customer received good scaling ranging from 30-50 percent when upgrading from one processor to two. Satisfactory scaling ranging from 20-30 percent was observed when upgrading from two processors to three, and then scaling flattened ranging from 5-15 percent when upgrading from three processors to four. Using the Partitioned Server feature, the overall number of simulated users supported could be scaled higher than for a system running a single instance of Domino Server.

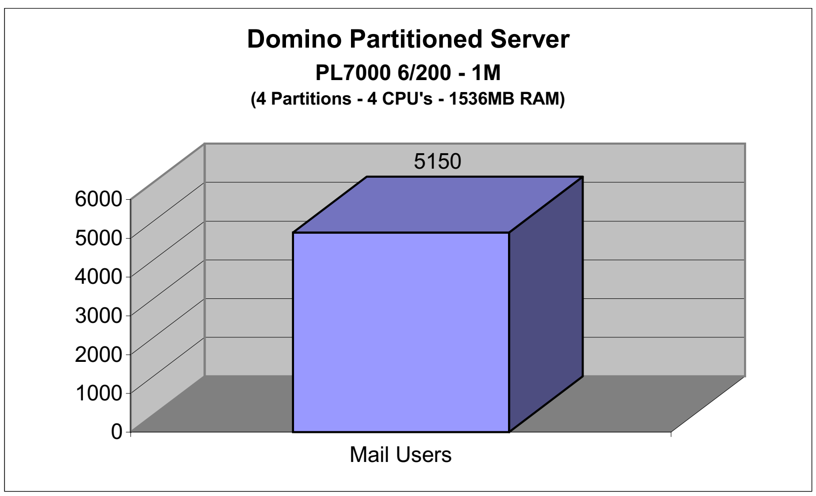

Figure 3-18. Lotus NotesBench Mail Workload – Number of Simulated Users

Additional NotesBench performance testing is planned for the ProLiant 7000. This 5,150 simulated Mail user partitioned server result can be attributed to several hardware changes made when running this partitioned server test. The hardware enhancements that were used for this test include the following:

- Compaq Processor Upgrade Kit providing the Intel Pentium Pro 200 MHz processor which utilizes a 1M L2 cache

- Compaq Smart-2DH Array Controller that supports Wide-Ultra SCSI transfer rates
- Compaq Smart-2DH Controller Array Accelerator provides 16MB of onboard cache (as compared to 4MB of onboard cache available with the Smart2/P Controller Array Accelerator)
- Wide-Ultra SCSI drives

This partitioned server NotesBench performance test result can be attributed to the optimal combination of performance tuned hardware components and the Domino Partitioned Server software.

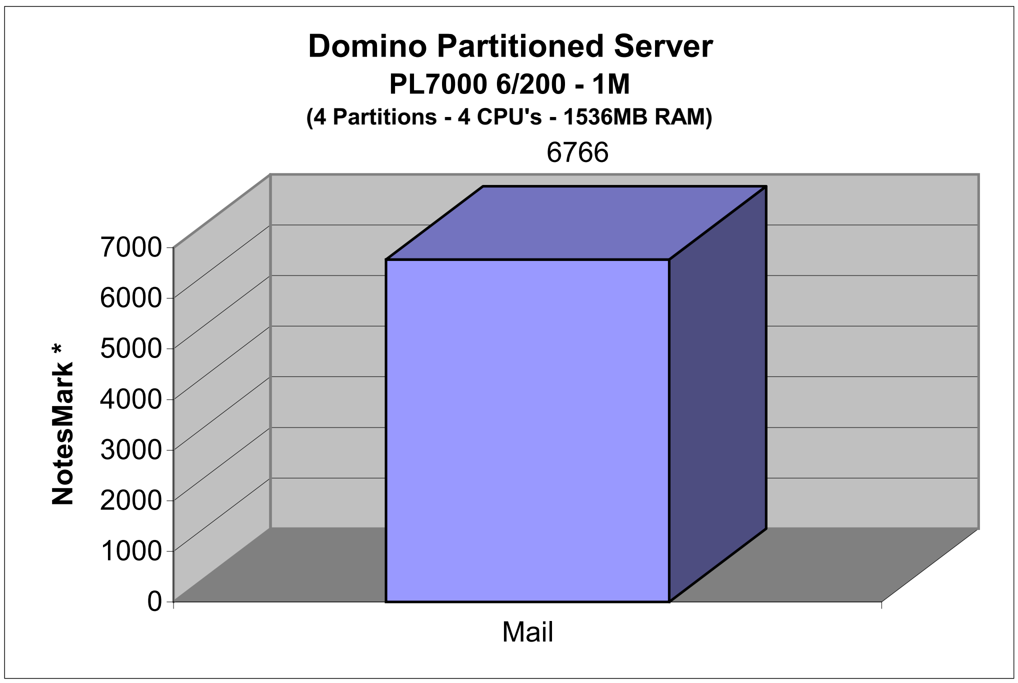

Figure 3-19. Lotus NotesMark Values for Benchmarked Configurations -
***** NotesMark – transactions per minute or tpm

The chart above illustrates the total transactions per minute (tpm) for the partitioned server test to be 6,766. The price performance ratio information for the ProLiant 7000 supporting 5,150 users is found below:

Mail Workload	6,766 NotesMark (tpm)	Price/Performance for System Throughput:	$13.22/NotesMark
	5150 Users	Price/Performance for System Capacity:	$17.37/user

The two following Performance Monitor charts shows the resource usage of the system under test during ramp up and after achieving a steady state supporting a NotesBench workload of 5,150 Mail Users. The system was configured with 1536-MB memory and 4xP6/200MHz CPU with 1-MB L2 Cache.

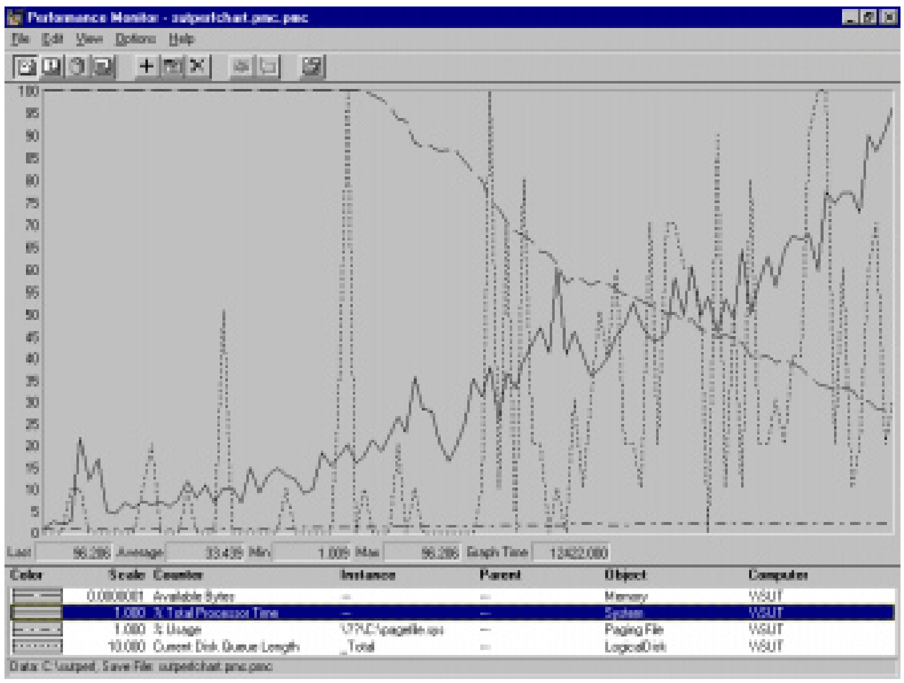

Figure 3-20. Resource Usage of System Under Test During Ramp Up

During the ramp up time, as shown in Figure 3-21, the system's CPU went as high as 96% utilization and it continued to be stressed until all the users were connected to the server. The CPU utilization rate (% Total Processor Time) is represented by the solid line in the following Performance Monitor chart. CPU utilization is shown climbing to 96% as all simulated user connections are established. The CPU subsystem utilization then dropped off and settled into a lower, steady rate of utilization averaging 89% during the remainder of the test. The available memory (Available bytes) is represented by the line made up of long dashes. Available memory is shown to be decreasing steadily as the system ramped up.

The CPU utilization rate (% Total Processor Time) is represented by the solid line at the top of the following Performance Monitor chart. The CPU utilization rate averaged 89% after all users were connected to the system under test. The NT Performance Monitor parameter, DiskPerf, provides useful information about the disk usage, but it also has a significant impact on the system under test, so we usually turned it off when we ran the actual test. The following Performance Monitor chart shows the resource usage of the system under test during a test workload of 5,150 Mail Users. The system was configured with 1536MB memory and 4xP6/200MHz CPU with 1MB L2 Cache. Available memory remained at a satisfactory level as represented by the line made up of long dashes in Figure 3-21.

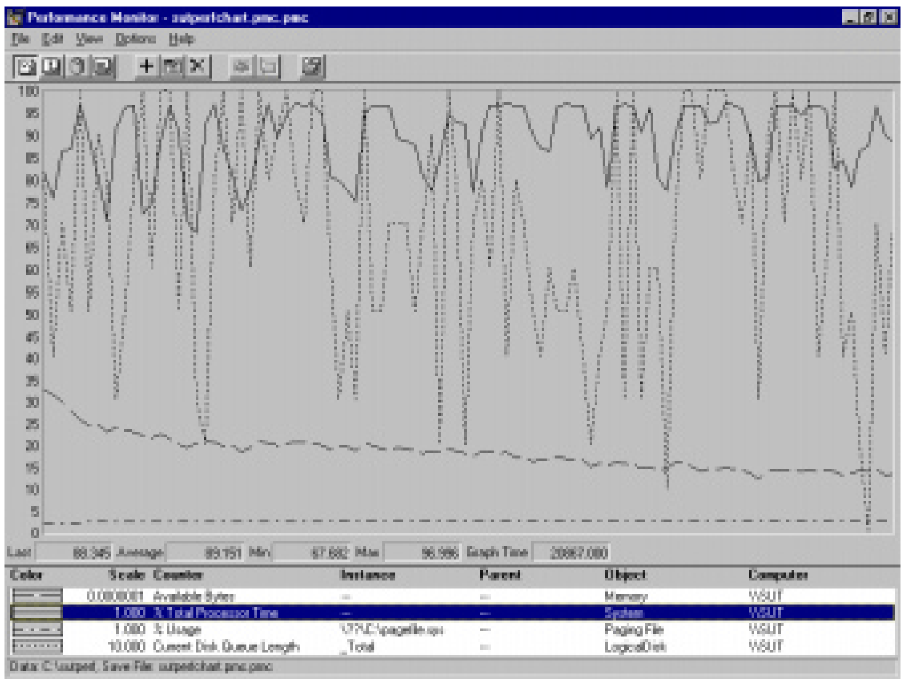

Figure 3-21. CPU Utilization During Steady State

Performance Tuning

Hard Disk Controller Tuning

Some of these features offer performance and fault tolerance advantages as discussed in the section on hardware versus software striping. Also discussed were the number of drives supported in an array. Now, the performance impact of the Smart-2 Controller Array Accelerator is examined.

The Smart-2 Controller Array Accelerator serves as a read-ahead and write cache which dramatically improves the performance of read and write commands. The Array Accelerator performance gains are best seen in database and fault tolerant configurations. The Smart-2 Controller writes data to 4-MB of cache memory on the Array Accelerator rather than directly to the drives, allowing the system to access this cache more than 100 times faster than accessing the disk. The data in the Array Accelerator is written to the drive array later by the Smart-2 Controller when the controller is otherwise idle. The Smart-2DH Controller writes data to 16MB of cache memory on the Array Accelerator rather than directly to the drives even further improving performance.

The Array Accelerator also anticipates requests as another method of increasing performance. A multi-threaded algorithm is used to predict the read operation most likely for the array. That prediction is used to pre-read data into the Array Accelerator so that data may be there before you access it. If the Smart-2 Controller receives a request for cached data, it can be burst into system memory at PCI or EISA bus speeds.

The Array Accelerator has a read/write cache ratio that can be customized to fit your Lotus Domino Server activity using the Compaq Array Controller Configuration Utility. The default setting is 50% Read /50% Write, but several other ratios are possible.

Testing to obtain a performance comparison with the Array Accelerator read/write cache configured with various read/write ratios was completed in Compaq labs. Test results indicated that the 25% Read / 75% Write ratio yields the best response time for both RAID 1 and RAID 5. This improvement in performance can be explained by the additional write related work that the controller has to perform when writing data and parity data with RAID 5, and when writing the data through two channels to both drives when mirroring with RAID 1.

Domino Server Tuning

NSF Buffer Size

The NSF buffer is the most important Domino Server tuning parameter that impacts performance. Other parameters do not appear to significantly impact performance.

The NSF buffer size specifies the maximum size in bytes of the NSF buffer pool. The NSF buffer pool is the section of memory dedicated to buffering I/O transfers between the NIF indexing functions and disk storage. Unless specified, the Domino Server automatically determines this value. By default, 25% of available memory is allocated. The maximum default value that the Server will allocate is 160Mb. The SHOW STAT command can be typed at the Server console to determine how much memory is available. The output includes a listing for Memory Available.

Compaq engineers recommend that the NSF buffer size be allowed to take the default value. Should you, however, have a need to set a specific value, the syntax for this parameter is:

NSF_BUFFER_POOL_SIZE=*value*

NT Server Tuning

When running under NT Server, consider changing the following operating system parameter values:

- Foreground and Background Applications set to “Equally Responsive” - this is set under *Control Panel→System→Tasking*

- NT Registry - Hkey_Local_Machine/System/CurrentControlSet/Control/PriorityControl/Win32PrioritySeparation:REG_DWORD:0x0
- NT Registry - Hkey_Local_Machine/System/CurrentControlSet/Control/SessionManager/MemoryManager/LargeCacheSystemCache:REG_DWORD:0x0

NT System Tuning

Smart-2 Controller's Array Accelerator Read/Write Ratio

The Smart-2 Array Controller's Array Accelerator provides a configuration utility which assigns 4-MB of cache memory to read/write operations. Ratios of 0% Read/100% Write, 25% Read/ 75% Write, 50% Read/ 50% Write, 75%Read/ 25% Write and 100% Read/ 0% Write are possible.

The 25% Read / 75% Write ratio, yielding the best response time for both RAID 1 and RAID 5, is recommended by Compaq engineers. This improvement in performance can be explained by the additional write related work that the controller has to perform when writing data and parity data with RAID 5, and when writing the data through two channels to both drives when mirroring with RAID 1.

Domino Server NSF Buffer Size

A Domino Server tunable parameter that impacts the system performance is the NSF Buffer Size which is the amount of memory allocated to the Domino Server NSF buffer specified in bytes.

Compaq engineers recommend that the NSF buffer size be allowed to take the default value. Should you, however, have a need to set a specific value, the syntax for this parameter is:

NSF_BUFFER_POOL_SIZE=*value*

NT Server Tuning

When running under NT Server, consider changing the following operating system parameter values:

- Foreground and Background Applications set to “Equally Responsive” - this is set under *Control Panel→System→Tasking*
- NT Registry - Hkey_Local_Machine/System/CurrentControlSet/Control/PriorityControl/Win32PrioritySeparation:REG_DWORD:0x0
- NT Registry - Hkey_Local_Machine/System/CurrentControlSet/Control/SessionManager/MemoryManager/LargeCacheSystemCache:REG_DWORD:0x0

Performance Summary and Recommendations

Based on the performance tests and data analysis carried out by Compaq engineers, the conclusions and recommendations for performance management are as follows:

System Processor

Research clearly showed that the CPU was found to be the most important server subsystem to affect overall system performance of the Lotus Domino Server. The conclusion is that the faster the processor the better the performance gains for the system. Therefore Compaq engineers recommend the fastest processor that can be purchased within the budgetary limitations of your project. Furthermore, the performance of the Pentium Pro Processor clearly showed that its superior features help contribute to the improvement in performance over the Pentium Processor rated at the same clock speed.

The results also showed that adding processors to the server helped to support additional users. The most marked increase in capacity of users was seen upgrading from 1P to 2P. Thus far performance testing has indicated that using Domino Server partitioning advanced features provides the best scalability story.

Memory

Memory was found to be a resource that depended upon the type of Domino activity that was taking place. The optimal memory configuration recommendation varies for the same number of users as the workload itself varies. The amount of memory needs to be properly balanced with the system's need for the resource. Too much or too little memory can have a negative impact on performance, depending upon the specific server activity involved.

The memory recommendation tables that follow can be used by the system administrator as a rule of thumb guideline for initially determining the system memory requirements. Once in production, the administrator can use a tool such as NT Performance Monitor to follow the memory resource utilization during operation to determine whether a memory upgrade is necessary. The memory aspect of capacity planning will be discussed in additional detail in Chapter 4 - Capacity Planning.

Table 3-5
Memory Recommendation - Mail

Number of Real World Users	Minimum Memory Required (MB)	Recommend Memory Configuration (MB)
150 or less	64	128
300	128	192
500	192	256
600	256	320
800	384	448
1000	512	576
1200	640	704
1400	768	832
1600	896	960
1800	1024	1088
2000	1152	1216
2200	1280	1344
2400	1408	1472
2500	1472	1536

The memory recommendations in Table 3-5 are based on many NotesBench tests that were run. Compaq engineers used the mail workload run optimal memory configuration findings and added 64-MB to derive the recommended memory configuration. The optimal memory findings are listed as minimal memory required for varying number of users. The recommended memory configuration amount includes 64 MB added for other background tasks.

Table 3-6
Memory Recommendation – Groupware

Number of Real World Users	Minimum Memory Required (MB)	Recommend Memory Configuration (MB)
150 or less	192	256
300	320	384
500	576	640
600	704	768
800	960	1024
1000	1216	1280

The memory recommendations Table 3-6 are based on several NotesBench Groupware_A tests that were run. Compaq engineers used the NotesBench run optimal memory configuration findings and added 64-MB to derive the recommended memory configuration. The optimal memory findings are listed as minimal memory required for varying number of users. The recommended memory configuration amount includes 64-MB added for background tasks.

Disk Subsystem

Compaq engineers recommend that disk striping be implemented to benefit from the gain in I/O performance. The recommendation is to use numerous smaller drives in an array rather than a few larger drives to achieve the best overall system performance providing comparable storage capacities.

Hardware striping is recommended due to performance gains as well as more system resource efficiencies than when using software striping. Hardware striping is achieved by Compaq's Smart-2 Array Controller which also has built-in data protection features, adding another benefit over software striping.

Fault Tolerance is strongly recommended by Compaq engineers. RAID 1 is the preferred level of fault tolerance for systems that have mission critical data, while RAID 5 is recommended for systems storing non-critical data. RAID 1 is the preference due to a combination of a high level of performance and protection of data. RAID 1 uses disk mirroring, providing very good data protection at the cost of low utilization of the actual disk capacity. Disk mirroring uses 50% of available disk space for fault tolerance support. RAID 5 uses distributed data guarding, striping data and parity data across all drives in the array. The more drives in the array, the lower the portion of each drive reserved for fault tolerance support.

System Tuning

Smart-2 Controller's Array Accelerator Read/Write Ratio

The 25% Read / 75% Write ratio, yielding the best response time for both RAID 1 and RAID 5, is recommended by Compaq engineers. This improvement in performance can be explained by the additional write related work that the controller has to perform when writing data and parity data with RAID 5, and when writing the data through two channels to both logical drives when mirroring with RAID 1.

Domino Server NSF Buffer Size

A Domino Server tunable parameter that impacts the system performance is the NSF Buffer Size which is the amount of memory allocated to the Domino Server NSF buffer specified in bytes.

Compaq engineers recommend that the NSF buffer size be allowed to take the default value of 25 percent of available memory.

NT Server Tuning

When running under NT Server , consider changing the following operating system parameter values:

- Foreground and Background Applications set to “Equally Responsive” - this is set under *Control Panel→System→Tasking*
- NT Registry - Hkey_Local_Machine/System/CurrentControlSet/Control/PriorityControl /Win32PrioritySeparation:REG_DWORD:0x0
- NT Registry - Hkey_Local_Machine/System/CurrentControlSet/Control/SessionManager/MemoryManager/LargeCacheSystemCache:REG_DWORD:0x0

Chapter 4
Capacity Planning

Definition of Capacity Planning

Capacity planning is a method of determining the balance between your Lotus Domino Server workload and its configuration at minimum cost, while meeting necessary user response time objectives. The goal of capacity planning is finding the best server and equipment to cost-effectively meet network workload demands and performance requirements. Capacity planning allows you to balance demand and supply—the demand for present and anticipated workload and the supply of present and future computer resources. A basic objective is consistent and acceptable user response times.

Capacity planning may be one of many responsibilities of the Lotus Domino administrator or integrator. Capacity planning is closely tied to performance management. Domino Server performance depends on the number of users on the system, the operating environment of the server and workstations, and the bandwidth and speed that are available to the physical network. The type of server, NICs, and cabling systems play an important role in how the network operates under heavy traffic conditions.

In capacity planning, the planner must balance complex, vague, and sometimes confusing data about workload, user needs, and computer resources, devise a coherent plan, and make these needs known to others. Although capacity planning requires the use of statistical data and mathematical techniques, it also requires a planner with practical experience and expert-level knowledge of the computer industry. It is not an exact science.

This Compaq TechNote offers data from Compaq integration labs testing to help you in these efforts.

Importance of Capacity Planning

Planning the appropriate hardware server platform to meet the needs of users is one of the most important strategic planning tasks in the Lotus Notes administrator's responsibilities. Poor planning, whether it be over estimating or under estimating your computer resource needs, affects the corporate “bottom line.” Over estimating results in a network server that costs more and has more capacity than users will ever need. This is a waste of resources and corporate funds. On the other hand, most planners do not plan far enough ahead into the future, resulting in insufficient computer capacity. This can create unhappy users, affect group productivity, negatively impact a company's bottom line, and place the company at a competitive disadvantage. Planning for sufficient computer capacity is an on-going process that allows you to avoid both overspending and insufficient capacity.

In the current drive to reduce corporate spending, allowance for the planning function is sometimes trimmed or even overlooked. However, reducing or neglecting this task exposes the danger of a poorly planned Lotus Notes implementation with insufficient capacity. In this case, the Notes administrator spends a large amount of time, effort, and cost reacting to user and management complaints and creating short-term fixes, rather than providing support, development, and strategic planning functions which are a critical part of the administrative responsibilities.

The amount of time, effort and cost spent properly planning the Lotus Domino Server implementation is worth the investment when the system adequately meets user response-time expectations and optimally utilizes system resources.

Assessing Your Server Requirements

Assessing network requirements includes both current and future Lotus Domino Server requirements including accurately capturing the number of users, application profile, and so on. When recording this information, it is wise to record both current user counts and application utilization as well as accurately projecting future growth in user counts and estimating potential changes to the application profiles.

The server requirements assessment should be carried out paying particularly close attention to accurately capturing the number of users and properly classifying their Domino usage. For example, perhaps Company XYZ is preparing to implement Lotus Domino Server for a group of 1800 mail users, 1500 mail and shared database users, and 300 groupware users. From a capacity planning standpoint, you need to know if there is any overlap between these user counts and their needs. For instance you need to ask if any of the 300 groupware users are included in the count of 1800 mail users and/or 1500 mail and shared database users.

The steps for assessing your Lotus Domino Server requirements include the following:

- List the Lotus Notes applications planned.
- Estimate the number of users for each application.
- Categorize the applications based on application type.
- Identify key subsystems that are most critical to your environment.

Evaluating Server Resource Alternatives

Upon completion of your Lotus Domino Server needs assessment, you must analyze the current hardware platform offerings and identify which products meet the needs of your specific configuration.

Evaluating server resource alternatives is a very important part of your planning for the implementation of your Domino Server. When evaluating your resource alternatives, you should complete the following steps:

- List available server features.
- Evaluate server features based on maximum capacity.
- Analyze server platform capabilities, both subsystem and overall server.
- Establish an upgrade plan if you are planning to upgrade rather than purchase a system.
- Establish an upgrade plan if you are planning to upgrade your system as your user population grows.
- Select a server configuration based on your requirements and the capabilities of the available server resources.

The Planning Process

The capacity planning process includes the following tasks:

- Define the current or planned Lotus Domino environment including the number of users and application mix.
- Define the current and future workload by monitoring the existing environment and basing new installations on experience with existing workloads.
- Evaluate computer resource alternatives based on available technology, relative cost, and capacity.

Capacity Planning Methodology

Compaq developed a capacity planning methodology to provide customers with a five step approach to follow when trying to answer capacity planning questions. The methodology was developed after careful study of general capacity characteristics observed during the analysis of numerous collections of NotesBench performance data. These capacity characteristics are described in this section. The five steps of the methodology are described and an example case study is also included. The remainder of this chapter contains the following information:

- General Capacity Characteristics Observations
- Determine the Application Mix or User Profile
- Convert Real World Users Requirements to NotesBench Workload User Results
- Select the System
- Determine the Memory Configuration
- Determine the Disk Subsystem Configuration

Capacity Characteristics Observations

NotesBench performance data revealed the peak workloads that could be supported by a configuration using a particular processor and amount of memory. For example, the following charts illustrate the CPU utilization and workload performance relationship that Compaq engineers saw when running real world applications. CPU utilization stayed relatively consistent with the curve, having only a very slight incline until the workload was pushed to a point which began to really stress the server. Then you can see that the CPU utilization curve began a much steeper incline as the workload increased.

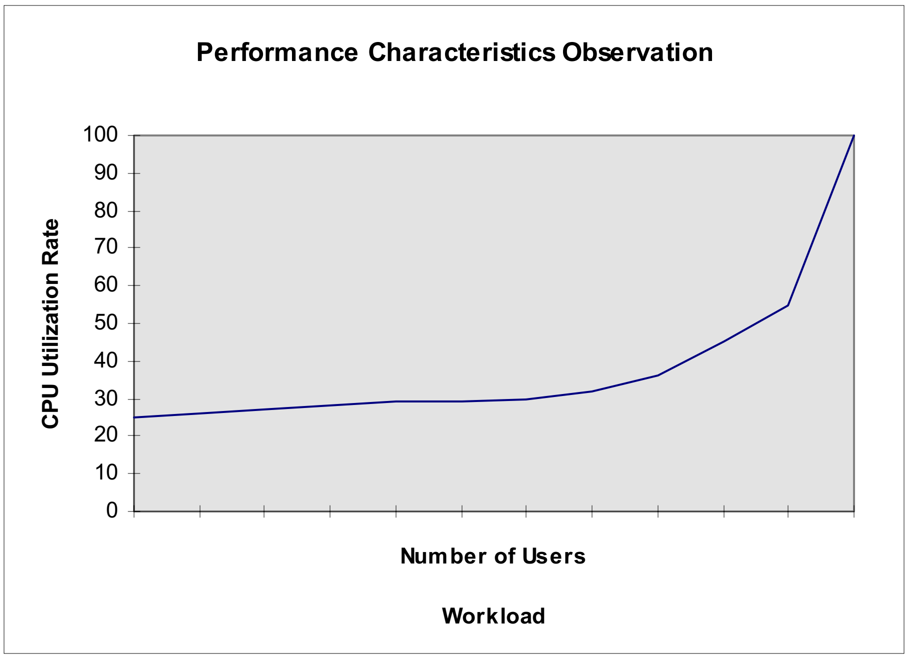

Figure 4-1. CPU utilization versus workload performance

Compaq also saw a relationship develop between CPU utilization and response time as shown in Figure 4-2. Response time is represented by a gently sloping curve as CPU utilization increased until the server began to be significantly stressed. The curve began a very steep incline from that point as the CPU utilization increased to the point that the server became stressed. When considering both CPU utilization and response time, a configuration that was running at constant rate greater than 70% utilization would not be desired.

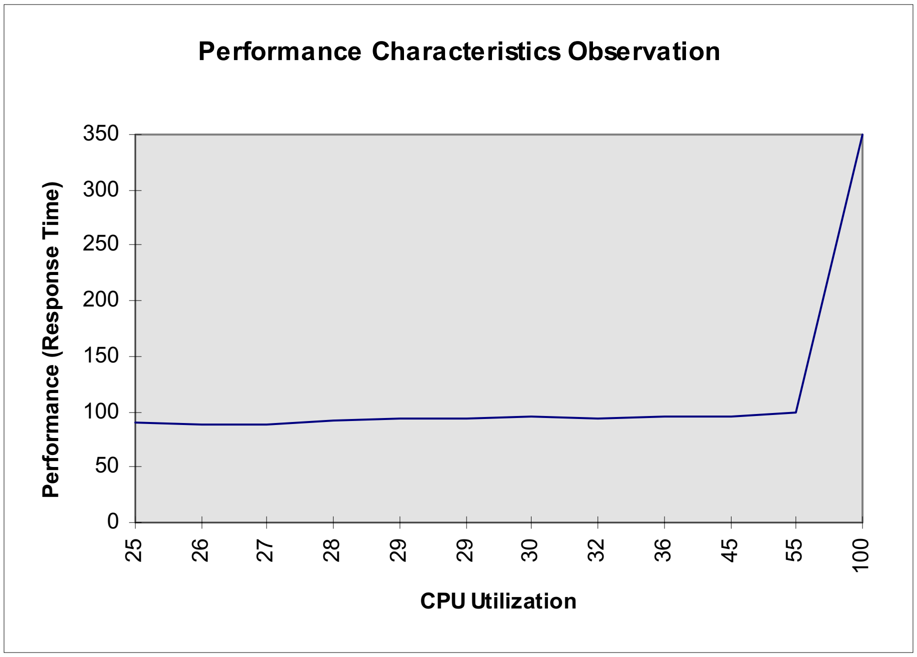

Figure 4-2. CPU utilization versus response time performance

Compaq considers 70% CPU utilization to be a rule of thumb threshold that leaves a 30% buffer of additional resources to be used for other background tasks that the server may be running. Ideally, Compaq engineers would like to see the CPU utilization buffer include a comfort zone of an additional 20%, scaling utilization back to about 50%. The additional 20% allowance is included to account for the fact that during test runs there was no full blown name and address book, no agents were running, nor was replication impacting performance.

Five Steps to Compaq Capacity Planning

The capacity planning methodology recommended by Compaq engineers involves a step-by-step process that begins with defining the application mix or user profile requirements. Customers are walked through the task of relating the real world user to a discounted NotesBench simulated user, selecting the system, and determining the memory and disk subsystem configurations.

- Determine the Application Mix or User Profile
- Convert Real World Users to NotesBench Workload Users
- Select the System
- Determine the Memory Configuration
- Determine the Disk Subsystem Configuration

Determine the Application Mix or User Profile

The first step in the capacity planning methodology is to determine the user requirements. This is a simple task of determining the user requirements and classifying these requirements into profiles that can be related to NotesBench performance results.

For example, if Domino is used for mail alone, then all users would be categorized as Mail users. When determining that users are mail users it is also important to classify their mail usage as light, medium, or heavy depending on the number of messages that are sent with attachments as well as the size of these messages. Light mail users are defined as sending and receiving mail messages averaging between 1 and 10-KB. Medium mail users are defined here as average mail messages between 10 and 20-KB. Heavy mail users are defined here as average mail messages from 20 to 30-KB.

Note: If users' average message size is considerably larger than 40-KB, then the NotesBench results for the GroupWare workload should be considered for capacity planning purposes rather than Mail workload results.

For Mail and Shared Discussion Database users, determine the level of users' discussion database activity. For example if the activity is related more to information retrieval (read activity) than update (or write) activity, consider users to be shared discussion database readers. Classify the shared discussion database user as an active participant if the user activity includes main topic and response creation which involves write file activity.

These classifications of mail users are put to use during the second step of the capacity planning process when converting real world users to NotesBench simulated users.

Convert Real World Users to NotesBench Workload Users

To use a capacity planning model, a formula must be created that relates the required number of real world users to NotesBench workload user results. During NotesBench performance testing on the ProLiant 5000, it was demonstrated that given a P6/200MHz two-processor system, the maximum number of supported NotesBench mail workload users is 2850. Applying the 50 percent utilization discount discussed previously, this 2xP6/200 system should be able to support 1,425 real world users, providing very good performance while allowing some resources for other activities. Therefore the following model which discounts the maximum NotesBench mail users by at least 50% to relate the NotesBench mail workload users to planned real world mail users can be used for capacity planning purposes:

# of real world mail users = Maximum NotesBench mail users × 50%

or

# of planned real world mail users × 2 = Maximum NotesBench mail users

Thus, if you know how many users need to be supported, available NotesBench mail results can be used to help select the appropriate system and configuration. If you need to support 900 real world mail users for example, a system that had a NotesBench result of 1,800 mail users or greater should be selected. This would be the correct relationship for mail users classified as *light* mail users as discussed in the previous section.

If the users had average message sizes ranging from 10-KB to 20-KB, then they would be classified as *medium* mail users. Therefore the previous model which discounts the maximum NotesBench mail users by at least 50% to relate the NotesBench mail workload users to planned real world mail users can be discounted even more. The following formula uses a discount of 33% to take into account the additional work placed on the server by these more active users and can be used for capacity planning purposes:

# of real world mail users = Maximum NotesBench mail users × 33%

or

# of planned real world mail users × 3 = Maximum NotesBench mail users

If the users had average message sizes ranging from 20-KB to 30-KB, then they would be classified as *heavy* mail users. Therefore the previous model which discounts the maximum NotesBench mail users by 33% to relate the NotesBench mail workload users to planned real world mail users can be discounted even more. The following formula uses a discount of 25% to take into account the additional work placed on the server by these most active users and can be used for capacity planning purposes:

# of real world mail users = Maximum NotesBench mail users × 25%

or

# of planned real world mail users × 4 = Maximum NotesBench mail users

Based on the previous explanation, the number of required mail users that participate in shared discussion database (MailDB) activity will also be discounted by at least 50%. This enables relating NotesBench MailDB workload users to the number of planned real world mail and shared database users that would stay within the comfort zone of the CPU utilization range. If the users will mainly utilize the shared database for information retrieval purposes, then these users would be classified as *readers* and the following formula would apply:

# of real world mailDB users = Maximum NotesBench mailDB users × 25%

or

# of planned real world mailDB users × 4 = Maximum NotesBench mailDB users

If the users will utilize the shared database for information retrieval purposes and will create main topics and submit responses to existing topics, then these users would be classified as *active participants* and the following formula would apply:

# of real world mailDB users = Maximum NotesBench mailDB users × 33%

or

# of planned real world mailDB users × 3 = Maximum NotesBench mailDB users

Because Groupware_A workload includes a lot of background tasks a weight of 100 percent (or no discount) will be used, thus basically equating NotesBench Groupware_A workload users to planned real world groupware users.

The Shared Discussion Database (DiscDB) workload includes only shared database activity, but is much more intensive than the DB part of the MailDB workload. The MailDB workload includes only information retrieval activity. The DiscDB workload therefore will be assigned a weight of 100% (or no discount), thus equating NotesBench DiscDB users to planned real world users.

In summary, the following formulas are used to relate real world (planned) users to our Capacity Planning (CP) NotesBench users for these workloads:

Table 4-1
Mail Workload

Classification	Formula
Light	*# of planned Mail users = NotesBench Mail Users/2*
Medium	*# of planned Mail users = NotesBench Mail Users/3*
Heavy	*# of planned Mail users = NotesBench Mail Users/4*

Table 4-2
MailDB Workload

Classification	Formula
Reader	*# of planned MailDB users = NotesBench MailDB Users/2*
Active Participant	*# of planned MailDB users = NotesBench MailDB Users/3*

Table 4-3
GroupWare Workload

Classification	Formula
All	*# of planned GroupWare users = NotesBench GroupWare Users*

Table 4-4
Shared Discussion Database (DiscDB) Workload

Classification	Formula
All	*# of planned DiscDB users = NotesBench DiscDB Users*

Selecting the System

The selection of the system type depends upon the user requirements (NotesBench users) calculated during the previous step of capacity planning. Once NotesBench user levels have been calculated, NotesBench Report results can be used to determine the best system suited to support a given number of users. To obtain a complete list of certified (audited by independent auditing agency) NotesBench Audit Reports, refer to the NotesBench Reports section of the NotesBench Consortium web site, http://www.notesbench.org.

Appendices A through C contain summary charts of Compaq systems that have been used to run various NotesBench workloads. Based on the information on the web site and in Appendix A, "Compaq NotesBench Workload Results," the Compaq ProLiant 5000 1xP6/200-512 provides adequate support for this user load. Frequently more than one system provides the level of support required. In this case, the decision should be made based on other factors such as the potential growth in the number of users supported, the desired high availability and redundancy system features, and system pricing. The Compaq web site contains additional information about the features of each platform at:

http://www.compaq.com/products/servers/platforms.html

Based on the calculated Capacity Planning (CP) NotesBench Mail, MailDB, and Groupware_A Users, the number of server(s) needed to meet all profile requirements can be derived by plugging user count numbers (obtained from the NotesBench results tables) into the following formulas:

Mail:
$$\frac{CP\,NotesBench\,Mail\,Workload\,Users}{CPU\,Max\,Mail\,Workload} = \#of\ Systems$$

MailDB:
$$\frac{CP\,NotesBench\,MailDB\,Workload\,Users}{CPU\,Max\,MailDB\,Workload} = \#of\ Systems$$

Groupware_A:
$$\frac{CP\,NotesBench\,Groupware\,A\,Users}{CPU\,Max\,Groupware\,A\,Workload} = \#of\ Systems$$

DiscDB:
$$\frac{CP\,NotesBench\,DiscDB\,Users}{CPU\,Max\,DiscDB\,Workload} = \#of\ Systems$$

The number of Server(s) for Mail, MailDB, GroupWare and DiscDB can then be summed to determine the total server(s) required.

For example, select the ProLiant 5000 P6/200 single processor configuration and assume that the requirements include supporting 1800 CP NotesBench Mail users, 700 CP NotesBench MailDB users, and 200 CP Groupware_A users. For this scenario, the number of ProLiant 5000 single processor systems needed can be calculated by using the following formulas:

Mail:
$$\frac{CP\ NotesBench\ Mail\ Workload\ Users}{CPU\ \ Max\ Mail\ Workload} = \#of\ \ Systems$$

MailDB:
$$\frac{CP\,NotesBench\,MailDB\,Workload\,Users}{CPU\,Max\,MailDB\,Workload} = \#of\ Systems$$

Groupware_A:
$$\frac{CP\,NotesBench\,Groupware\,A\,Workload\,Users}{CPU\,Max\,Groupware\,A\,Workload} = \#of\ Systems$$

Mail: $$\frac{1800}{1920} = 0.94 \;\; systems$$

MailDB: $$\frac{700}{1440} = 0.49 \;\; systems$$

Groupware_A: $$\frac{200}{440} = 0.45 \; systems$$

Total = 1.88 systems - 2 servers

If the server total equates to a whole number, you have completed your server selection process.

If the server total has a fractional component such as 1.2, you have several options:

- Server 1, a P6/200 4P configuration, and Server 2, a less powerful server to handle the users that would overtax Server 1.
- Server 1 and Server 2 configured to evenly handle the projected total workloads.
- Server 1 and Server 2, etc. based upon geographic considerations, with the total power of the servers selected adequately handling its users.

The ProLiant 5000 was selected here because the customer wanted as much flexibility as possible to be able to upgrade the system to support future growth in the user population. The ProLiant 800, 850 or 2500 can be upgraded to two processors, but cannot be upgraded to the same extent as the ProLiant 5000. This customer is part of a rapidly growing company.

Determine the Memory Configuration

The appropriate memory configuration relates to the number of real world users determined during the requirement specifications. The following charts are taken directly from the Performance Management section of this TechNote and refer to Memory Recommendations for Mail and GroupWare users as previously discussed.

Table 4-5
Memory Recommendation - Mail

Number of Real-World Users	Minimum Memory Required (MB)	Recommend Memory Configuration (MB)
150 or less	64	128
300	128	192
500	192	256
600	256	320
800	384	448
1000	512	576
1200	640	704
1400	768	832
1600	896	960
1800	1024	1088
2000	1152	1216
2200	1280	1344
2400	1408	1472
2500	1472	1536

The memory recommendation chart above is included previously in Chapter 3, “Subsystem Performance Comparison.”

Table 4-6
Memory Recommendation – Groupware

Number of Real-World Users	Minimum Memory Required (MB)	Recommend Memory Configuration (MB)
150 or less	192	256
300	320	384
500	576	640
600	704	768
800	960	1024
1000	1216	1280

As mentioned in the previous performance management chapter, these memory recommendations are to serve as guidelines for customers. If the user activity involves a mixture of profiles, then a memory configuration needs to be derived. These charts are provided to assist customers with memory configurations for Compaq systems running Domino Server. If the user activity is a combination of mail and shared database activity, the memory configuration falls closer to the Mail user recommendation than the Groupware recommendation. If the user profile involves strictly shared database activity, the memory configuration recommendation falls closer to the Groupware user memory recommendation than the Mail user recommendation.

For example, if a customer needed to support 600 Shared Database (DiscDB type) real world users, then a memory configuration from 640 to 704-MB should provide adequate memory resources. Memory recommendations are as the name implies, guidelines that provide a good starting point for customers. Memory utilization should be monitored once the system is running to determine if adequate memory resources are available to the system. Response time experienced by users is in part affected by the server’s memory resource. Adding more memory often results in better response time for mail, shared database, and groupware users as discussed in the performance management chapter.

Determine the Disk Subsystem Configuration

The Lotus Domino Server(s) should be set up with two volumes:

System Volume: The system volume should be a single drive attached to the SCSI bus, formatted as NTFS. The following should be included on this volume:

- Windows NT Server
- Lotus Domino Server software, excluding data files
- Domino log file (*LOG.NSF*)

Data Volume: The data volume should be an array of drives controlled by a Smart-2 Array Controller. Hardware fault tolerance is recommended for all Domino Server systems. This drive array should be configured with a fault tolerance level of RAID 5 using hardware striping for non-critical data and RAID 1, mirroring, for mission critical data servers. The following should be included on this volume:

- Lotus Domino data
- NT Paging File (size dependent upon memory configuration)

When determining the disk subsystem configuration of a Domino server system, use Wide-Ultra drives connected to controllers that support the Wide-Ultra transfer rates when possible. Disk access is often a major bottleneck in the Domino server performance. Using the fastest available SCSI drives combined with disk controllers that support Wide-Ultra transfer rates, the system will be configured to keep this impact minimized as much as possible.

Capacity Planning Case Study

A fictitious company, XYZ, wants to provide simple messaging capabilities to 800 employees. In addition, another 300 users must be provided with messaging and discussion database capabilities. These 300 users will be using five group discussion databases that will be used actively and not just for information retrieval. Additionally, 100 users have very intensive power user requirements. The future Notes Administrator of XYZ is asking for information that will tell him how many servers are needed to meet these user requirements. The administrator also needs to know how the server should be configured regarding memory and disk storage, commenting that good response time is a definite requirement for this implementation to be a success. Monetary resources have to be justified, but availability of funds is not considered to be a problem. The company classifies this data as critical.

Step 1: Determine the Application Mix or User Profile

The 800 employees requiring simple messaging capabilities are classified as light NotesBench Mail users. The 300 additional mail users that also require shared database capabilities are profiled as NotesBench MailDB users and classified as active participants. The final 100 intensive power users are profiled as NotesBench GroupWare users.

Step 2: Convert Real World Users to NotesBench Workload Users

As the second step of capacity planning, the real world user requirements must be plugged into the formulas mentioned in the methodology discussion which converts planned real world user numbers into NotesBench workload users.

Light Mail:	*CP NotesBench Mail Users = # of planned Mail users × 2*
	CP NotesBench Mail Users = 800 × 2 = 1600
Mail and shared database:	*CP NotesBench MailDB Users = # of planned Mail/Shared DB users × 3*
	CP NotesBench MailDB Users = 300 × 3 = 900
Groupware:	*CP NotesBench GroupWare_A Users = # of planned groupware users*
	CP NotesBench GroupWare_A Users = 100

Step 3: Select the System

For this scenario, the ProLiant P6/200 single processor system is assumed. The same model can be used to determine the number of other system types as well.

Mail: $$\frac{1600}{1920} = 0.83\ Systems$$

MailDB: $$\frac{900}{1440} = 0.63\ Systems$$

Groupware: $$\frac{100}{440} = 0.23\ Systems$$

Total = 1.69 systems

Based on this Capacity Planning Model, two ProLiant 5000 single P6/200 processor systems are required. Multiple options are possible for allocating the users to the two servers. The assumption used for memory and disk subsystem configuration is:

- Server 1—800 Mail Users
- Server 2—300 Mail & Shared DB Users + 100 Groupware users

Step 4: Determine the Memory Configuration

Server 1

Server 1 should be configured with between 384 to 448-MB of memory. This recommendation is pulled directly from the memory guideline table included in the methodology section. The chart shows that 384-MB is required to support 800 real world users while 448-MB is recommended because 64-MB is included for additional server activities.

Server 2

Server 2 should be configured with at least 448-MB of memory derived as follows:

- 300 MailDB users can be derived by referring to the Mail and Groupware Memory Recommendation Charts
- 300 users from the Mail Memory Recommendation Chart = 192-MB
- 300 users from the Groupware Memory Recommendation Chart = 384-MB
- 300 MailDB users can be derived as 256-MB (this provides more memory than is recommended for 300 Mail users, not as much as recommended for 300 Groupware users, and more closely associated with the Mail recommendation.)
- 100 groupware users from the Groupware Memory Recommendation Chart = 192-MB
- Total Memory Calculated for Server 2 = 448-MB

Step 5: Determine the Disk Subsystem Configuration

Lotus Domino Server(s) should be set up with two volumes:

- **System Volume:** The system volume should be a single drive attached to the SCSI bus, formatted as NTFS. The following should be included on this volume:

    Windows NT Server

    Lotus Domino Server software, excluding data files

    Domino log file (*LOG.NSF*)

- **Data Volume:** The data volume should be an array of drives controlled by a Smart-2 Array Controller. This drive array should be configured with a fault tolerance level RAID 5 using hardware striping. The following should be included on this volume:

    Lotus Domino data

    NT Paging File (size dependent upon memory configuration)

Server 1

The size of the data volume allocated for Lotus Domino data on Server 1 should be at least 40-GB, allowing approximately 50-MB/user for Mail.
(800 users × 50-MB = 40000-MB or 40-GB)

An additional 500-MB should be allocated to the NT Paging File. The rule of thumb for this calculation is Amount of Memory + 10%. For our example, we determined the memory configuration for Server 1 and 2 to be 448-MB, therefore the NT Paging File size is calculated as 448-MB + 45-MB—or 493-MB. For capacity planning purposes, the 493-MB paging file was rounded to 500-MB.

To support these storage requirements of 40.5-GB, six 9.1-GB drives provide over 46-GB of adequate storage for the 800 Mail Users. Note that no other storage requirements were considered in calculating this minimum storage.

NOTE: To calculate the usable disk storage using RAID 5 hardware fault tolerance, the formula is n/(n+1). For example, six 9.1-GB drives provides total disk space of 54.6-GB. This value needs to be multiplied by 6/7 to calculate total usable disk space which is 46.8-GB. This is the best fit for the drives, because five 9.1-GB drives only allows 37.9-GB of usable disk space which does not meet the calculated requirements.

Server 2

The size of the data volume for Server 2 should include at least 15-GB, allowing approximately 50-MB/user, for Mail.

(300 users × 50-MB = 15000-MB or 15-GB)

An additional 2.5-GB of space should be allowed for the five discussion databases. (The rule of thumb used to calculate the storage requirements for the discussion databases allows at least 500-MB for each database;
5 databases × 500-MB/database = 2500-MB).

A final 10-GB of space is allowed for the 100 groupware users. (The rule of thumb used to calculate the storage requirements for the groupware users allows 100-MB for each user; 100 groupware users x 100-MB/ user = 10000-MB or 10-GB).

An additional 500-MB should be allocated to the NT Paging File. The rule of thumb for this calculation is Amount of Memory + 10%. For our example, we determined the memory configuration for Server 1 and 2 to be 448-MB, therefore the NT Paging File size is calculated as 448-MB + 45-MB—or 49-3MB. For capacity planning purposes, the 493-MB paging file was rounded to 500-MB.

The combined storage requirements are 15-GB for mail, 2.5-GB for discussion databases, 10-GB for the groupware users, and 500-MB for the paging file , totaling 28-GB.

To support these storage requirements, four 9.1-GB drives configured in an array using RAID 5 should provide 29-GB of usable storage for the 300 mail and discussion database users plus the 100 groupware users. Note that no other storage requirements were considered in calculating this minimum storage.

NOTE: To calculate the usable disk storage using RAID 5 hardware fault tolerance, the formula is n/(n+1). For example, four 9.1-GB drives provides total disk space of 36.4-GB. This value needs to be multiplied by 4/5 to calculate total usable disk space which is 29-GB. This is the best fit for the drives, because three 9.1-GB drives only allows around 20.5-GB of usable disk space which does not meet the calculated requirements.

Appendix A
Microsoft Windows NT NotesBench Results

Mail Workload – Audited Results

System Under Test	CPU	RAM	Number of Users	NotesMark (TPM)	$/User	$/TPM	Resp. Time	Domino Release
Compaq ProLiant 7000 P6/200-1M	4xP6/200 - 1M	1536-MB	5,150	6,766	$13.22	$17.37	.890	Domino 4.51 Server
Compaq ProLiant 6500	3xP6/200 - 1M	1280-MB	3,600	4765	$11.74	$15.54	.423	Domino 4.51 Server
Compaq ProLiant 6500	2xP6/200 - 1M	1024-MB	2,850	3734	$12.72	$16.66	.107	Domino 4.51 Server
Compaq ProLiant 3000	Pent. II 300 MHz	512-MB	2,200	2912	$7.10	$9.40	.496	Domino 4.51 Server
Compaq ProLiant 5000 P6/200 Model 1-X	P6/200 - 512	512-MB	1,920	2,546	$19.97	$15.06	.382	Domino 4.5 Server
Compaq ProLiant 5000 P6/200-512 Model 1-X 2P	2xP6/200 -512	768-MB	2,850	3,748	$13.77	$10.47	1.072	Domino 4.51 Server
Compaq ProLiant 5000 P6/200-512 Model 1-X 3P	3xP6/200 -512	1024-MB	3,400	4,484	$13.30	$10.09	.643	Domino 4.51 Server
Compaq ProLiant 800 P6/200 Model - 4300	P6/200 - 256	256-MB	1500	1,997	$13.48	$10.12	.237	Lotus Notes R4.5
Compaq ProSignia 200	Pent. II 233 MHz	128-MB	600	790	$7.59	$10.00	1.112	Lotus Intranet Starter Pack based on Domino

Shared Discussion Database (DiscDB) Workload – Audited Results

System Under Test	CPU	RAM	Number of Users	NotesMark (TPM)	$/User	$/TPM	Resp. Time	Domino Release
Compaq ProLiant 3000	Pent. II 300 MHz	512-MB	825	1,427	$25.13	$14.53	1.102	Domino Server 4.51
Compaq ProLiant 5000 P6/200 Model 1-X	P6/200 – 512	512-MB	720	1,248	$53.26	$30.73	1.025	Domino Server R4.5
Compaq ProLiant 5000 P6/200-512 Model 1-X 2P	2xP6/200 -512	768-MB	900	1,491	$43.60	$26.32	2.795	Domino Server R4.51
Compaq ProLiant 5000 P6/200-512 Model 1-X 3P	3xP6/200 -512	1024-MB	1,000	1,425	$45.24	$31.75	.147	Domino Server R4.51
Compaq ProSignia 200	Pent. II 233 MHz	256-MB	700	1,211	$6.80	$11.77	1.097	LISP based on Domino

MailDB Workload – Audited Results

System Under Test	CPU	RAM	Number of Users	NotesMark (TPM)	$/User	$/TPM	Resp. Time	Domino Release
Compaq ProLiant 5000 P6/200 Model 1-X	P6/200 – 512	512- MB	1,440	3,250	$26.63	$11.80	2.49	Domino Server R4.5

Groupware_A Workload – Audited Results

System Under Test	CPU	RAM	Number of Users	NotesMark (TPM)	$/User	$/TPM	Resp. Time	Domino Release
Compaq ProLiant 5000 P6/200 Model 1-X	P6/200 – 512	512-MB	440	1,466	$87.15	$26.16	.433	Domino Server R4.5

Appendix B
IBM OS/2 Warp Server Advanced NotesBench Results

Mail Workload – Audited Results

System Under Test	CPU	RAM	Number of Users	NotesMark (TPM)	$/User	$/TPM	Resp. Time	Domino Release
Compaq ProLiant 800 P6/200 Model 4300	P6/200 - 256	256-MB	1,300	1,722	$15.65	$11.81	.31	Domino Server R4.5
Compaq ProLiant 800 P6/200 Model 4300 2P	2xP6/200 -256	512-MB	1,900	2,521	$17.36	$13.09	.26	Domino Server R4.5

MailDB Workload – Audited Results

System Under Test	CPU	RAM	Number of Users	NotesMark (TPM)	$/User	$/TPM	Resp. Time	Domino Release
Compaq ProLiant 800 P6/200 Model 4300	P6/200 - 256	256-MB	1,050	2,434	$19.37	$8.36	.92	Domino Server R4.5

Appendix C
Novell NetWare 4.11 NotesBench Results

Mail Workload – Audited Results

System Under Test	CPU	RAM	Number of Users	NotesMark (TPM)	$/User	$/TPM	Resp. Time	Domino Release
Compaq ProLiant 800 P6/200 Model 4300	P6/200 - 256	256- MB	900	1,202	$22.46	$16.82	.098	Domino Server R4.5

MailDB Workload – Audited Results

System Under Test	CPU	RAM	Number of Users	NotesMark (TPM)	$/User	$/TPM	Resp. Time	Domino Release
Compaq ProLiant 800 P6/200 Model 4300	P6/200 - 256	256- MB	750	1,764	$26.96	$11.46	.172	Domino Server R4.5

Index

W