

COMPAQ


Integrating Lotus Notes R4 on Compaq Servers

*Performance Management and
Capacity Planning of
Lotus Notes R4
for Microsoft Windows NT on
Compaq ProLiant Servers*

Compaq TechNote

Includes information on:

- Atomic Task profiling
- Hardware performance comparison
- Performance tuning
- Capacity planning



.....

NOTICE

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

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

This publication contains information protected by copyright. No part of this publication may be photocopied or reproduced in any form without prior written consent from Compaq Computer Corporation.

The software described in this guide is furnished under a license agreement or non-disclosure agreement. The software may be used or copied only in accordance with the terms of the agreement.

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

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

© 1996 Compaq Computer Corporation.
All rights reserved. Printed in the U.S.A.

Compaq, Compaq Insight Manager, Fastart, NetFlex, ProLiant, ProLinea,
QVision, SmartStart, Systempro, and Systempro/LT
Registered United States Patent and Trademark Office.

ProSignia and Systempro/XL are trademarks of Compaq Computer Corporation.

Microsoft is a registered trademark of Microsoft Corporation and Windows and
Windows NT are trademarks of Microsoft Corporation.

Performance Management and Capacity Planning of Lotus Notes R4 for Microsoft Windows NT on Compaq ProLiant Servers

First Edition (September, 1996)
Part Number 242835-001

Contents

Chapter 1

About This TechNote

Objective.....	1-1
Notation Conventions.....	1-3
Other Resources to Use	1-4

Chapter 2

Atomic Task Profiling

Profiling Methodology	2-1
Test Configuration.....	2-1
Data Collection	2-2
Charts and Graphs	2-2
Axis Scales.....	2-2
Profile Graphs	2-4
CPU Utilization.....	2-5
Disk Throughput	2-8
LAN Throughput.....	2-11
Atomic Task Profiling Conclusions	2-13
Research Conclusions.....	2-13
Administrative Conclusions.....	2-14

Chapter 3

Performance Management

Test Methodology	3-4
NotesBench	3-4
Test Configuration and Procedure	3-10
Subsystem Performance Comparison.....	3-13

System Processor (CPU).....	3-13
Memory.....	3-20
Disk.....	3-23
Bus Architecture (PCI versus EISA).....	3-33
File Systems (FAT versus NTFS).....	3-36
Networking.....	3-39
Performance Tuning.....	3-39
Hard Disk Controller Tuning.....	3-39
Notes Server Tuning.....	3-41
NT Server Tuning.....	3-41
NT System Tuning.....	3-42
Networking Subsystem.....	3-43
Performance Conclusions.....	3-47
System Processor.....	3-47
Memory.....	3-47
Disk Subsystem.....	3-47
System Tuning.....	3-48

Chapter 4

Capacity Planning

Definition of Capacity Planning.....	4-1
Importance of Capacity Planning.....	4-2
Assessing Your Server Requirements.....	4-3
Evaluating Server Resource Alternatives.....	4-4
The Planning Process.....	4-4
Capacity Planning Methodology.....	4-5
Capacity Characteristics Observations.....	4-6
Relate Real World Users to NotesBench Workload Users.....	4-8
Select the System.....	4-9
Determine the Memory Configuration.....	4-11
Determine the Disk Subsystem Configuration.....	4-12

Capacity Planning Case Study.....	4-13
Step 1: Relate Real Word Users to NotesBench Workload Users.....	4-13
Step 2: Select the System	4-14
Step 3: Determine the Memory Configuration	4-14
Step 4: Determine the Disk Subsystem Configuration.....	4-16
Step 5: Determine the Network Configuration	4-18

Appendix A Multivendor Integration From Compaq

Index

Figures

Figure 2-1.	CPU Utilization - Part 1	2-5
Figure 2-2.	CPU Utilization - Part 2	2-6
Figure 2-3.	Disk Throughput - Part 1	2-8
Figure 2-4.	Performance Monitor—Index chart	2-9
Figure 2-5.	Disk Throughput - Part 2	2-10
Figure 2-6.	LAN Throughput - Part 1	2-11
Figure 2-7.	LAN Throughput - Part 2	2-12
Figure 3-1.	NotesBench test configuration	3-10
Figure 3-2.	Performance Monitor chart showing system resource utilization	3-12
Figure 3-3.	CPU performance comparison	3-15
Figure 3-4.	CPU response time comparison	3-16
Figure 3-5.	CPU performance comparison with SMP for the P5/133	3-17
Figure 3-6.	CPU performance comparison with SMP for the P5/166	3-18
Figure 3-7.	CPU performance comparison with SMP for the P6/166	3-19
Figure 3-8.	Memory performance	3-21
Figure 3-9.	Drive spindle performance comparison	3-24
Figure 3-10.	Response time comparison of hardware/software striping	3-26
Figure 3-11.	CPU Utilization comparison of hardware/software striping	3-27
Figure 3-12.	Fault tolerance type comparison of response time	3-31
Figure 3-13.	Fault tolerance type comparison of CPU utilization	3-32
Figure 3-14.	EISA test system - resource utilization rates	3-34
Figure 3-15.	PCI test system - resource utilization rates	3-35
Figure 3-16.	FAT versus NTFS comparison of response time	3-37
Figure 3-17.	FAT versus NTFS comparison of CPU utilization	3-38
Figure 3-18.	Smart-2P Disk Array Controller Array Accelerator Configured with various Read/Write Ratios	3-40
Figure 3-19.	NetBench 4.0 result of total throughput	3-46
Figure 4-1.	CPU utilization versus workload performance	4-6
Figure 4-2.	CPU utilization versus response time performance	4-7

Tables

Table 1-1 Notation Conventions	1-3
Table 2-1 Atomic Task Abbreviations	2-4
Table 2-2 CPU Utilization by Server Tasks	2-6
Table 2-3 CPU Utilization by Client Tasks	2-7
Table 2-4 Disk Throughput by Server Tasks	2-8
Table 2-5 Disk Throughput by Client Tasks	2-10
Table 2-6 LAN Throughput by Server Tasks	2-12
Table 2-7 LAN Throughput by Client Tasks	2-13
Table 3-1 Typical Benchmark Results	3-6
Table 3-2 Drive Spindle Performance Comparison Mixed Load	3-24
Table 3-3 Single Segment LAN Throughput	3-45

ACKNOWLEDGMENTS

Compaq and Lotus have worked cooperatively on the testing, analysis, and development of this TechNote as part of an aggressive effort to provide our ProLiant and Notes customers with R4 performance information. Our joint efforts will increase as we move forward with future Notes releases. Special effort was taken toward the generation of capacity and performance information for all R4 Notes versions which Compaq platforms will support – especially the multi-CPU configurations which enable large-scale Notes deployments.

James Grigsby
Notes Product Manager
Lotus Development Corporation

Compaq is very pleased with our relationship with Lotus Development Corporation and the work we have performed together in developing this TechNote. We look forward to our increasing efforts in delivering this type of performance information for future Lotus Notes releases. We at Compaq feel that strong cooperative efforts, such as this, benefit our mutual customers by enabling us to deliver useful information such as this White Paper and better integrated solutions.

Steve Siegel
Manager, Groupware/Intranet Engineering
Compaq Enterprise Solutions

Chapter 1

About This TechNote

This Compaq TechNote provides the results of a performance analysis conducted by Compaq Engineers on the Lotus Notes 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 Notes, and Windows NT. It is a supplement to the *Compaq Hardware Reference* document and the *Lotus Notes 4.0* documentation. The results and conclusions of this TechNote will give you:

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

Objective

One objective of this TechNote is to provide current customers who are running Lotus Notes R4 for Windows NT on Compaq servers information that will assist them in optimally configuring their server(s) for their given operating environment to achieve the highest possible performance from their hardware and software. Information is also provided that will assist customers to make configuration upgrade decisions that can be based on an anticipated return in performance gains.

.....

1-2 About This TechNote

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



Notation Conventions

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

Convention	Use
<i>Screen selections, variables, and new terms</i>	These items always appear in italics.
<i>FILENAMES</i>	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.

Other Resources to Use

You may want to consult the following resources for additional information on obtaining the best possible performance and throughput with Lotus Notes Server for NetWare:

- Compaq TechNote: *Configuring Compaq RAID Technology for Database Servers*, P/N 184206-001.

This TechNote provides general to help determine the best configuration of the Compaq SMART SCSI Array Controller for database server needs. It also contains information on Compaq drive subsystem terminology and RAID.

- Lotus Notes R4 documentation.

The Lotus Notes R4 documentation set provides a comprehensive set of documents covering installation, reference, and an administrator's guide with detailed information on Lotus Notes 4.0.

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

This book can help to determine bottlenecks in networks, servers. Gain 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.

- *Compaq Hardware Reference* documentation.

The *Compaq Hardware Reference* documentation can help isolate a problem to a certain computer component or cable segment. Some network problems occur because of malfunctioning hardware; check the documentation that came with the hardware involved.

- World Wide Web on Internet.

<http://www.compaq.com/support/techpubs/>

<http://www.lotus.com>

NOTE: This list is not intended to be all comprehensive of the materials available but access to these materials will be of benefit to the reader.

Chapter 2

Atomic Task Profiling

The goal of Atomic Task profiling is to identify and measure the impact of User and Server Atomic Tasks on a *Lotus Notes Server for Windows NT*. This information can be used as follows:

- Identification of the stressful Atomic Tasks highlights certain basic administrator options to improve overall system performance.
- Identify the pattern of Lotus Notes applications, find the bottleneck, and recommend optimum performance management.
- These stressful tasks can be used to create a system workload so that hardware can be varied to determine the optimum hardware configuration of a Lotus Notes Server for Windows NT.
- Those identified tasks can be used in combination to stress the server to determine boundary conditions. The boundary condition information can then be used for performance tuning and capacity planning.

Profiling Methodology

The profiling methodology consists of data collection, analysis and recommendation.

Test Configuration

- Lotus Notes Version: 4.0
 - Processor: one P5/133
 - Memory: 256 megabytes
 - Detached doc size: 100K
 - Database size for non-indexed search: 100 megabytes
-

Data Collection

Data from each test script was collected in two ways. The response time measured by Microsoft Test was recorded in a centralized log file. Performance Monitor provided the CPU, disk, and LAN utilization on the Windows NT server.

The utilization numbers were then imported into a spreadsheet and were graphed as raw data to reveal the greatest subsystem stress. The raw data graphs of the disk and LAN subsystems were excluded from further analysis because they did not reveal sufficient stress on the server to be of interest. The CPU utility data are discussed in the following sections.

Charts and Graphs

The data for each subsystem is summarized in the following figures. The entire set of bar charts is sorted in descending order based on CPU utilization. This is because the CPU is the most stressed subsystem in the server.

In the subsystem graphs that follow, the bars appear in pairs. The first bar of each pair reflects the average of the subsystem utilization resulting from an Atomic Task. The second bar in each pair reflects the standard deviation of these averages which, in turn, illustrates the choppy character of the data. A tall standard deviation bar implies a choppy data set when graphed, and a short standard deviation bar implies a flat data set when graphed. In contrast, choppy data sets contain intervals of lower CPU utilization which means that the CPU is free during that time to perform other tasks.

Axis Scales

The axis scales of the raw data graphs have been adjusted to visually reflect the relative impact on the three measured subsystems: CPU, disk, and LAN. In each of the profile bar graphs that follow, tall bars indicate a strong impact on the particular subsystem and short bars indicate a weak impact. The graphs are displayed as follows:

- The CPU utilization was recorded in percentages and varied between 0% and 100%. Hence, the y-axis upper limit is set to 100% for each of the CPU graphs.
- The disk throughput was recorded in bytes per second. Unlike a percentage measurement that has an obvious upper limit of 100%, a realistic disk upper limit is difficult to determine. However, this upper limit directly effects the visual information that is derived from the graphs. Previous disk throughput measurements at Compaq have shown 1,700 kilobytes per second (KB/s) to be a reasonable upper limit under similar conditions.

NOTE: Disk throughput depends upon a variety of factors. The upper limit can vary depending upon the hardware and the operating system environment.

- The LAN throughput was also recorded in bytes per second. A realistic upper limit of 70% of the Ethernet theoretical maximum has been selected for display purposes. Hence, 70% of a 10 megabits per second (Mb/s) maximum is 7 Mb/s which equates to 875 KB/s (kilobytes per second). An upper limit of 900 KB/s was selected.
-

Profile Graphs

Three sets of graphs (CPU, Disk, LAN) are displayed in this section to show the relative impact of each of the Atomic Tasks on a specific subsystem. For display purposes, each subsystem graph has been split into two separate graphs which should be viewed as a single chart.

NOTE: The graphs reflect only a single instance of each Atomic Task. Actual loads will vary because multiple instances of User Atomic Tasks can be run simultaneously.

Table 2-1
Atomic Task Abbreviations

Abbreviations	Atomic Task
Rebld View	Rebuilding a View Index
Create Indx	Creating a Full Text Index
N Indx Srch	Executing a Non-Indexed Text Search
Indx Srch	Executing an Indexed Full Text Search
Expand	Expanding or Collapsing a Category
Scroll	Scrolling in a View
Mail	Sending a Mail Message
Edit	Editing a Document
Detach	Detaching a File
Read	Reading a Document

CPU Utilization

Of the three subsystems analyzed, the CPU utilization was most affected. The three server tasks with the highest impact on the CPU are: Rebuilding a View Index, Creating a Full Text Index, and Executing a Non-Indexed Text Search. The four most CPU-intensive client tasks are Expand the view, Edit documents, Detach documents and Scrolling through a view. The two illustrations that follow show the averages for CPU utilization. Tables 2-2 and 2-3 show the percentage for each task.

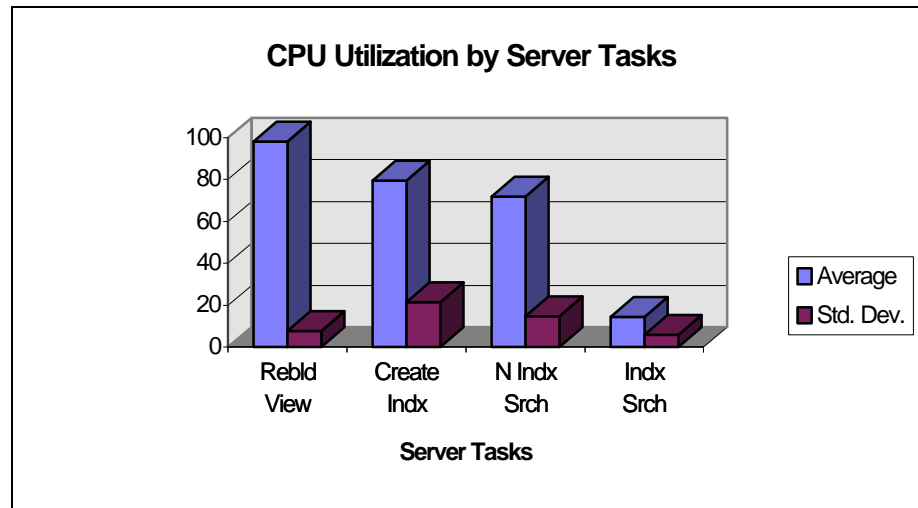


Figure 2-1. CPU Utilization - Part 1

2-6 Atomic Task Profiling

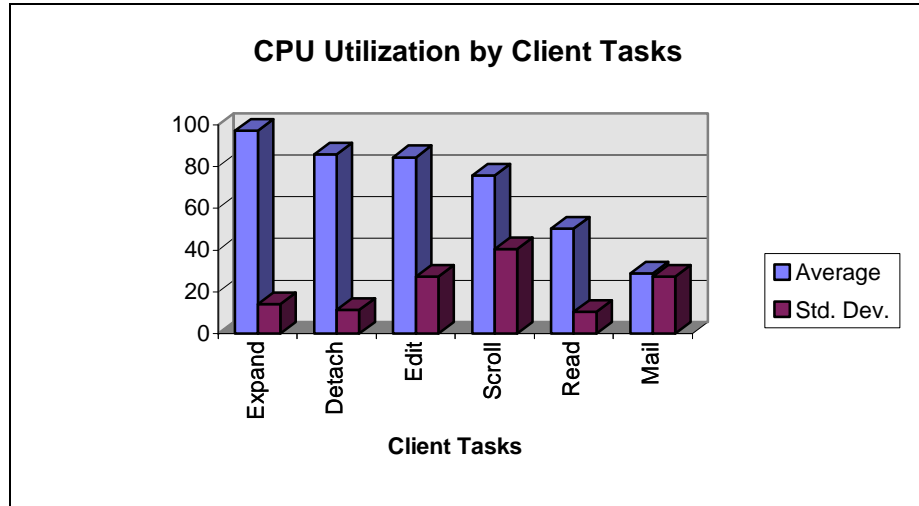


Figure 2-2. CPU Utilization - Part 2

Table 2-2
CPU Utilization by Server Tasks

	Average (%)	Standard Deviation (%)
Rebld view	98.24	7.55
Create Indx	79.45	21.36
N Indx Srch	71.92	14.59
Indx Srch	14.32	5.79

Table 2-3
CPU Utilization by Client Tasks

	Average (%)	Standard Deviation (%)
Expand	97.23	13.96
Detach	85.89	11.39
Edit	84.57	27.24
Scroll	75.82	40.45
Read	50.34	10.31
Mail	28.76	27.26

2-8 Atomic Task Profiling

Disk Throughput

Disk Throughput is only moderately impacted by a Non-Indexed Full Text Search or replication. : All other Atomic Tasks have very little impact on the server disk subsystem. Figures 2-3 and 2-5 show the high averages for disk throughput. Tables 2-4 and 2-5 shows the KB/s for each task.

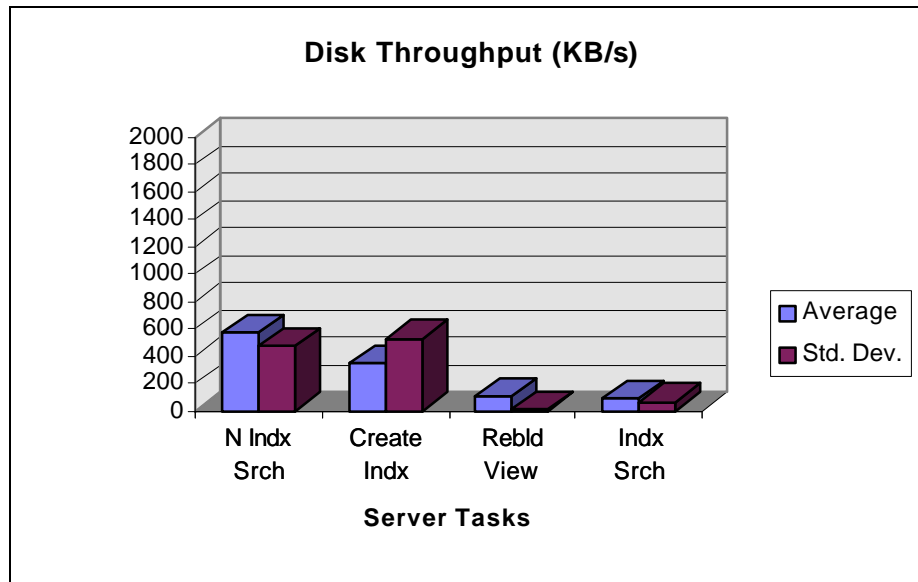


Figure 2-3. Disk Throughput - Part 1

Table 2-4
Disk Throughput by Server Tasks

	Average (KB/s)	Standard Deviation (KB/s)
N Indx Srch	562.84	465.53
Create Indx <i>(see Note following)</i>	337.77	528.62
Reblid view	98.24	7.47
Indx Srch	86.93	64.0

NOTE: Creating Index generates a very choppy utilization pattern as shown in the Performance Monitor data below.

In Figure 2-4, the top line of the graph represents CPU performance and the bottom line represents total bytes per second.




Figure 2-4. Performance Monitor—Index chart




Figure 2-5. Disk Throughput - Part 2

Table 2-5
Disk Throughput by Client Tasks

	Average (KB/s)	Standard Deviation (KB/s)
Detach	214.98	51.21
Read	184.54	94.94
Edit	166.52	120.63
Mail	121.42	140.22
Scroll	4.64	10.37
Expand	3.15	15.63

LAN Throughput

The server task has very little impact on the LAN throughput due to the client/server architecture. For example, the client issues a create index command, the server performs the task and only acknowledges the client when the job is done, generating no excess network traffic. The client task generates some stress, but those tasks are highly intensified. For example, our “Edit” task actually modified 2000 documents (average size from four to ten pages) in less than 200 seconds—an unusually heavy load in the real world. Even so, the impact on the network wire is minor relative to the full capacity of the network. The two illustrations that follow show the high averages for LAN throughput. Tables 2-6 and 2-7 shows the KB/s for each task.

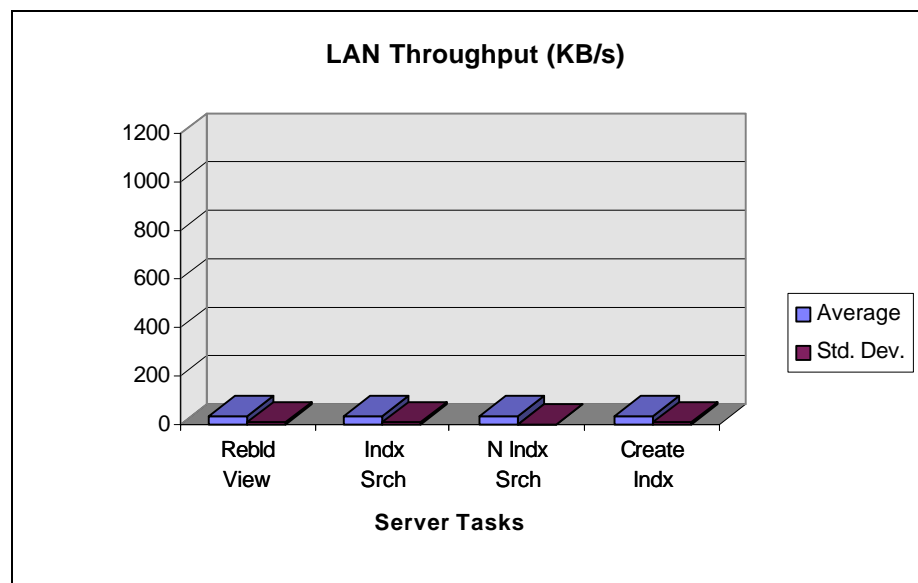


Figure 2-6. LAN Throughput - Part 1

2-12 Atomic Task Profiling




Figure 2-7. LAN Throughput - Part 2

Table 2-6
LAN Throughput by Server Tasks

	Average (KB/s)	Standard Deviation (KB/s)
Rebld view	35.4	12.12
Indx Srch	32.72	9.41
N Indx Srch	32.14	3.14
Create Indx	31.95	9.41

Table 2-7
LAN Throughput by Client Tasks

	Average (KB/s)	Standard Deviation (KB/s)
Detach	484.56	57.41
Read	465.55	103.53
Edit	202.46	68.94
Scroll	68.89	21.82
Expand	46.99	6.78
Mail	42.16	9.76

Atomic Task Profiling Conclusions

The following sections describe the research and administrator conclusions for Atomic Task Profiling.

Research Conclusions

The following list describes the research conclusions for Atomic Task Profiling.

- As is true of most client/server applications, Lotus Notes Server is a CPU intensive program. This suggests that the CPU capacity, multi-processor systems, clock speed, and memory caches should be the focal points of further research. The Atomic Tasks that create the greatest stress on the server are:
 - Executing a Non-Indexed Text Search
 - Creating a Full Text Index

.....

2-14 Atomic Task Profiling

- Lotus Notes Server does not significantly stress the disk subsystem in many cases. However, performing a non-indexed full text search or replication of a large database might place undue stress on the disk subsystem. Those Atomic Tasks are useful for an analysis to optimize the disk subsystem.
- Lotus Notes Server does not significantly stress the LAN subsystem. Among the server tasks, only the Replicating Server to Server (Read and Write) moderately impacts the LAN subsystem. However, network design as well as Notes domain/logical layout is still very important for those WAN users and enterprise network users.
- Replicating workstation to server is significantly slower than replicating server to server. Hence, this Atomic Task is useful in investigating workstation hardware configurations and the client impact on system performance. Generally the slower the client, the greater capacity of the server for workstation-oriented Atomic Tasks. Further research on the client performance such as Win95 or protocol performance will be performed in the future.

Administrative Conclusions

The following are the administrative conclusions for Atomic Task Profiling.

- Notes Server 4.0 has significant improvement from Notes Server 3.x. The Compaq Groupware Engineering team found 5% - 35% performance improvement with 4.0. The capacity of the server showed even greater improvement.
- Non-indexed full text search is the most CPU-intensive task, therefore, users should always use indexed databases to avoid this costly operation.
- Create full text indexes during non-peak hours to prevent system delays.

Creating a full text index is the second most stressful Atomic Task. It utilizes 79.45% of the CPU capacity.

Indexed full text searches are very fast compared to non-indexed full text searches. Also, a non-indexed full text search utilizes 71.92% of the CPU capacity. Thus, a 79.45% CPU hit during non-peak hours (Create Index) is much better than a 71.92% CPU hit during peak hours (Non-Indexed Search). If CPU utilization remains above 80%, we recommend adding a CPU to handle CPU-intensive tasks such as view indexing.

NOTE: The administrator has the option of adjusting the frequency of the full text index updates. The list box under *File*→*Full Text Search*→*Information* contains various options of frequency of index updates. By reducing the frequency, the impact on the CPU during peak hours can be delayed. However, the index will not be current until it is updated.

Chapter 3

Performance Management

Performance Management can only be successfully achieved by fully understanding the performance impact that system resources such as the system processor, memory, and the disk subsystems components have on the overall operation of your entire system. By changing the configuration of these components, performance will be affected in some way. The goal of this chapter is to help the customer better understand the relationship between system resources and Lotus Notes Server performance so decisions can be made regarding changes to an existing server configuration as well as complete configuration of a new installation.

Therefore the 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 network administrator, performance means effective management of system resources. A system administrator's concerns are 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.

The overall purpose of this chapter is to provide the customer with an understanding of how Lotus Notes 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 a Lotus Notes running on a Compaq server.

.....

3-2 Performance Management

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 that is 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 Notes server system resource areas:

- System Processor (CPU) Performance
- Memory
- Disk Subsystem
- Bus Architecture (PCI versus EISA)
- File Systems
- Networking

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 affect 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 would allow 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

- Workload Generator (NotesBench)
- Performance Monitor Tool

NotesBench

Lotus NotesBench is an implementation of a standard benchmark tool. Lotus NotesBench for Lotus Notes R4 is a collection of benchmarks and documentation for evaluating the performance of Notes R4 servers. The benchmarks (usually called tests in this user guide) model the behavior of Notes workstation-to-server or server-to-server operations. They return measurements to evaluate server performance in relation to the server system's cost of ownership.

This chapter includes these topics:

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

Some 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
 - Different systems running the same software
-

Notes Bench benchmarks are in the last category of this list.

■ **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.

■ **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/second. 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 five-year 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 least cost of ownership.

■ **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.

3-6 Performance Management

■ 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 (k\$/tps). These benchmark results apply to different systems running the same software application.

Table 3-1
Typical Benchmark Results

System Under Test	tps	k\$/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.

What is NotesBench?

NotesBench gives you an objective way of evaluating the performance of different platforms running Notes. It is intended for testing engineers at hardware and software vendors. NotesBench requires that vendors run the NotesBench tests and publish their performance results. Each vendor runs the same tests in the same manner. The published tests must be audited.

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—Notes R4. You can also use it to compare different versions (or code builds) of Notes 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. There are NotesBench workloads for the following:

- A replication hub—a server that exists to propagate changes among a collection of other servers.
 - A 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.
 - A server for mail users—a workload that models sites that rely only on mail for communication.
 - A server for mail and shared database users—a workload that models active users who are only performing mail and simple shared database operations.
-

.....

3-8 Performance Management

- 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.
- 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.
- An idle usage workload that establishes an upper bound on the number of sessions (which do nothing) that a Notes server can support. You can use this metric to aid in setting up the other NotesBench tests.

You specify a workload by running the NotesBench commands NOTEBNCH and NOTESNUM 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.

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 in the test and for the average response time.

Results Metrics discusses NotesBench performance results in more detail.

NotesBench Pricing

To calculate the weighted performance of a Notes server (the price/performance ratio) you must calculate the cost of the system under test. This includes the cost of the system and the network components required to achieve the reported workload performance.

NotesBench Price/Performance

Price/performance for NotesBench consists of the price/NotesMark and the price/user. These are the performance values for a NotesBench test divided into the price of the system under test.

NotesBench Scaling

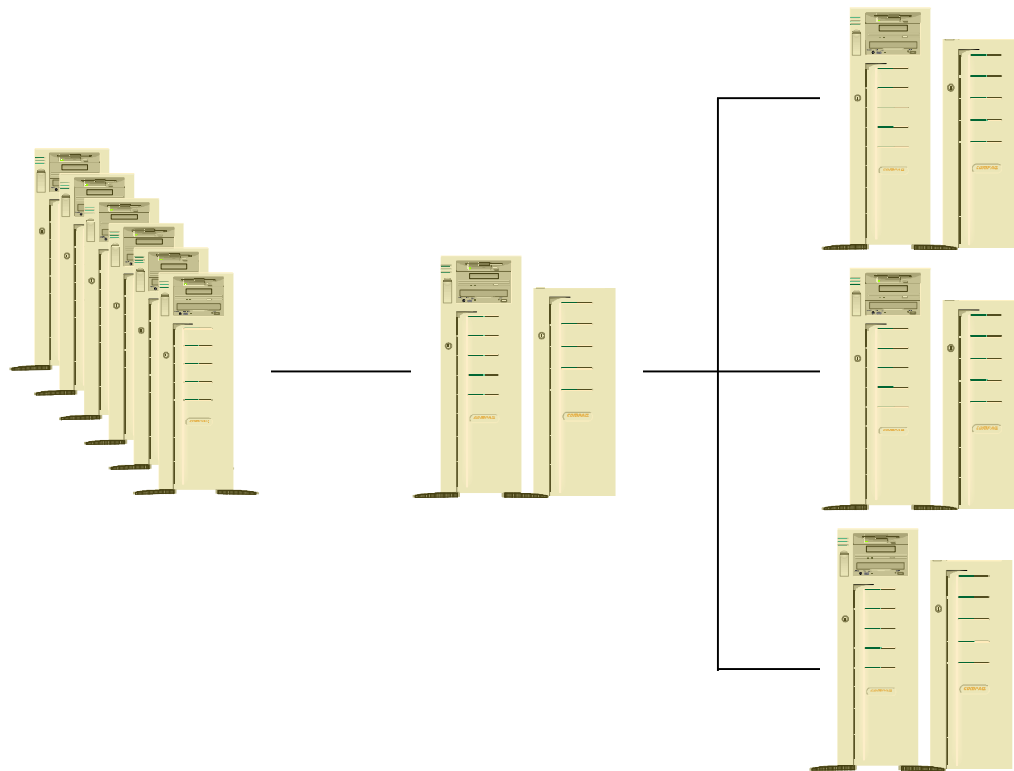
NotesBench tests maintain a proportional relationship between the capacity of the system under test and the workloads presented to the system. To accomplish scalability for the workloads, you set Notes environmental variables to values outlined in the NotesBench manual.

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



6 Driving Systems (Clients)

ProLiant 2000
1xP5/66MHz
32MB Memory
2x1GB Disk

System Under Test

ProLiant 5000
2xP6/166MHz
512MB Memory
1x2GB System Volume
7x2GB Data Volume
RAID 0

3 Destination Servers

ProLiant 4000
4xP5/133MHz
192MB Memory
1x1GB Systems Volume
5x1GB Data Volume

Figure 3-1. NotesBench test configuration

Software Versions

- Client Systems: Windows NT Server with Notes Server 4.0 and NotesBench (Notes version 138 - Windows/32)
- System Under Test: Windows NT Server 3.51 with Notes Server 4.0 and NotesBench (Notes version 138 - Windows/32)
- Destination Server: Windows NT Server with Notes Server 4.0 and NotesBench (Notes version 138 - Windows/32)

Test Procedure

Compaq Groupware Engineers first performed several trial runs to determine the best test duration and confirmation of Steady State. Both test duration and steady state were determined using real time monitor utilities from Windows NT. During the trial runs, we monitored the Windows NT Performance Monitor and also logged the whole test process. The results were compared with the NotesBench Specification for conformity.

Actual testing began with three clients. The system under test was allowed to reach a steady state for about 10 minutes, then client systems were added gradually, about 3-10 minutes apart. This was done because the system under test usually showed tremendous stress during the workload initialization stage. After the last client was started, we turned off Performance Monitor and only recorded the Notes related activities for the best performance results.

The tests ran continuously run for 4-5 hours after reaching steady state with all client systems connected.

The test data collected includes the following files:

- All client's *NOTES.INI* files
 - The SUT *NOTES.INI* file and *LOG.NSF* file
 - All destination server's *NOTES.INI* file and *LOG.NSF* files
 - The NOTESNUM utility results file
-

.....

3-12 Performance Management

Performance Monitor Tool

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 typical resource usage of the system under test during a test workload of 1200 mail users. The system was configured with 384MB memory and four P6/166MHz CPU's.

The following is an example of the Performance Monitor results:

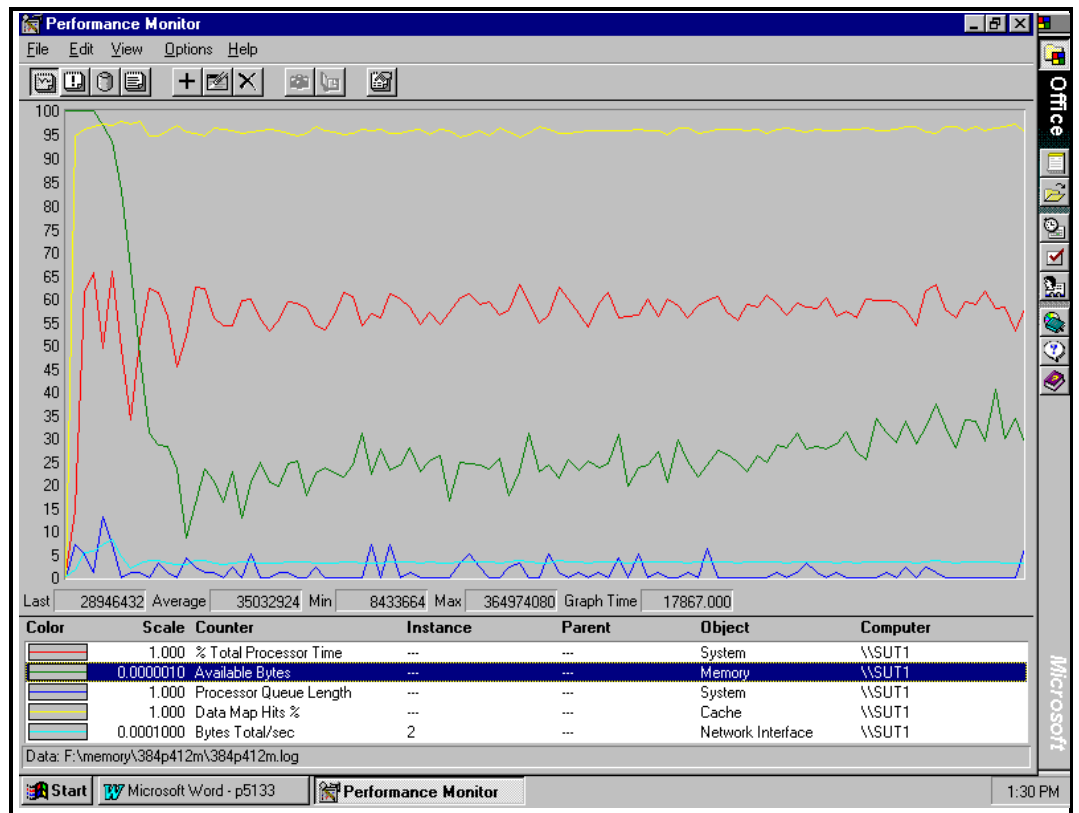


Figure 3-2. Performance Monitor chart showing system resource utilization

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 engineers. 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
- Bus Architecture (PCI versus EISA)
- File Systems (FAT versus NTFS)
- Networking

System Processor (CPU)

In contrast to a resource sharing (file server) environment, a faster processor in a Lotus Notes 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 Lotus Notes, 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 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:

.....

3-14 *Performance Management*

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

NOTE: If you use the 133MHz Pentium processor in your Lotus Notes Server, be sure to verify that the system ROM in your Compaq ProLiant Server is dated 08/16/95 or later for optimal performance. You can observe the ROM date on your screen when you power up the server. You must also upgrade the System Configuration for EISA and EISA/PCI to V2.28 Rev. A or later.

Performance Results

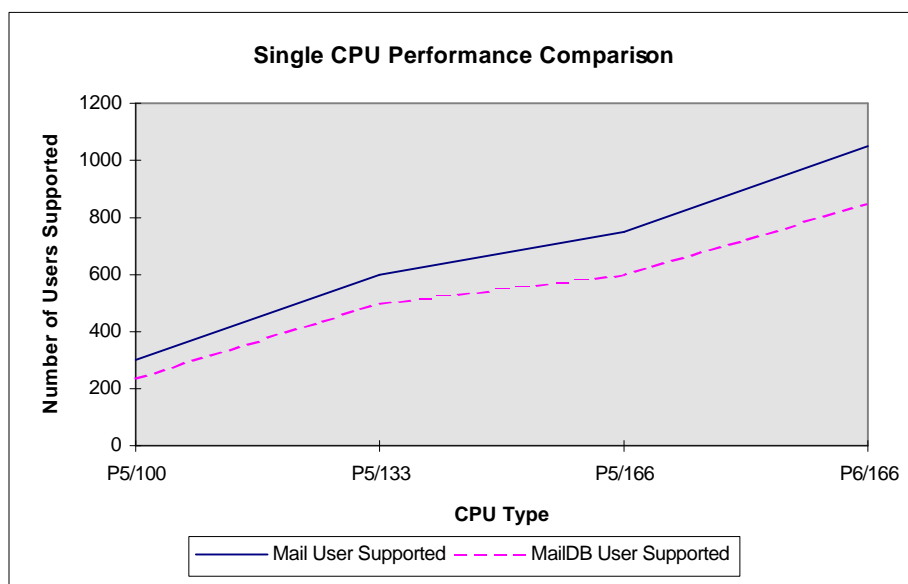


Figure 3-3. CPU performance comparison

The data charted above validates the conclusion regarding the importance of the processor. A linear relationship exists between number of users supported and the clock speed of particular processor chip, illustrating that the more processing power that the server has the higher the number of users supported. The results also show that 40-42 percent additional users can be supported by utilizing the Pentium Pro Processor over a P5 processor rated at the same clock speed. The results also quantify the 150 percent performance gain you might expect when you upgrade your server from a P5/100 MHz processor to a P6/166 MHz processor.

NOTE: The measurement of number of users was obtained by running tests that pushed the number of users as much as possible while complying with Lotus' recommendation of keeping response time under 5 seconds.

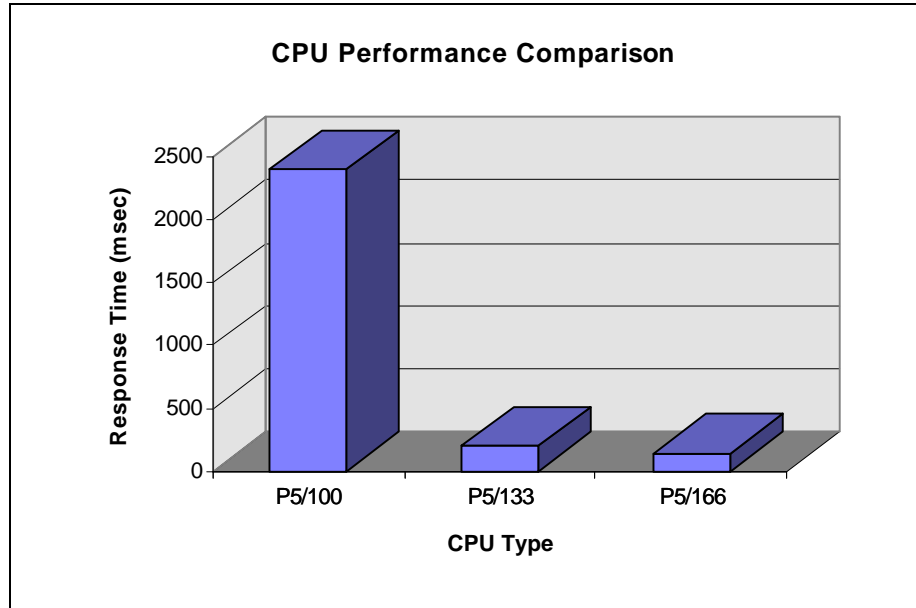


Figure 3-4. CPU response time comparison

The great performance gain from P5/100 to P5/133 is due not only to the CPU clock speed, but also to the large L2 cache. The P5/133 processor has 2MB of L2 cache not present on the P5/100 processor.

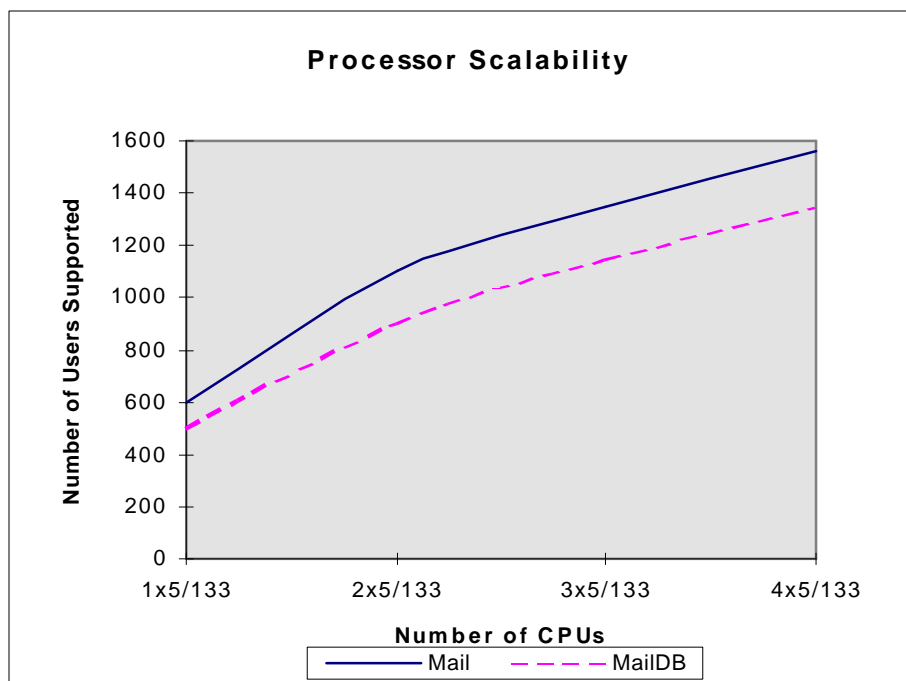


Figure 3-5. CPU performance comparison with SMP for the P5/133

The data charted above introduces the twist of multiple processors (Pentium with 133 MHz) 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 of 83.3 percent (Mail users) and 80 percent (MailDB users) increase in number of supported users can be seen moving from 1P to 2P. The performance gain from 2P to 3P is an increase of 22.73 percent (Mail users) and 27.78 percent (MailDB users) in supported users. An additional 15.56 percent and 21.74 percent of Mail and MailDB users are supported when moving from 3P to 4P.




Figure 3-6. CPU performance comparison with SMP for the P5/166

While the chart above shows the number of users supported increasing as you move from 1P through 4P, the most significant increase in the number of users—60 percent for Mail and 83.3 percent for MailDB—can be seen moving from 1P to 2P. The performance gain from 2P to 3P is an increase of 18.2 percent (Mail users) and 25 percent (MailDB users) in supported users. An additional 20 percent and 15.4 percent of respective Mail and MailDB users are supported when moving from 3P to 4P.




Figure 3-7. CPU performance comparison with SMP for the P6/166

The P6/166 chart above shows the number of users supported increasing as you move from 1P through 2P is 55.23 percent and 56.67 percent respectively for Mail and MailDB applications.

Due to limitations within the current test software, we could only test 1P and 2P scenarios. Corrected software versions are expected to be available for future testing.

The number of Groupware_A users supported on our test systems were:

- 4xP5/133 - 400 users
- 4xP5/166 - 500 users
- 2xP6/166 - 500 users

Memory

Memory is a parameter that is going to be dependent upon the total functional responsibilities of your Lotus Notes server. At all times you must consider what background tasks your Notes 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 Lotus Notes, the system and other applications you are running will probably 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 with a workload equivalent of 1200 NotesBench Mail users. All other parameters were fixed with only the memory configuration changing. Tests show that more memory is not necessarily always better. However, testing showed that a system without proper memory configuration will not only experience performance degradation, but will also begin 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.




Figure 3-8. Memory performance

Notice that the 1200 user workload test shows an optimal memory configuration of 384MB. More than 384MB shows a slight increase in performance time that could be attributed to the additional work the processor had to do to manage the larger memory pool. Adding 128MB of memory to the optimal 384MB configuration leads to a degradation of 25.8% in response time for the 512MB system. A much more pronounced performance degradation of 157 percent is seen when reducing the 384MB configuration by 128MB in the 256MB system. This could be attributed to lack of memory resources causing more paging to virtual memory which adds to the processor utilization as well as the involving the disk subsystem.

.....

3-22 **Performance Management**

Memory Recommendation		
Number of Users	Minimum Memory Required (MB)	Recommend Memory Configuration (MB)
500 or less	128	192
600	192	256
800	256	320
1000	320	384
1200	384	448
1400	448	512
1600	512	576
1800	576	768

The memory recommendations in the chart above are based on the 1200 user mail workload tests that were run. Compaq engineers used mail for test run 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 64MB added for background tasks, mail and shared db, groupware or other 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 Notes 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:

- Drive Spindles/Striping
- Hardware Striping versus Software Striping
- Fault tolerance: RAID 0, RAID 1, RAID 4, RAID 5
- Controller type

Drive Spindles/Striping

If your applications generate significant disk I/O, there will likely be a lot 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 will 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 3-2
Drive Spindle Performance Comparison
Mixed Load

	Response Time (Seconds)	Average CPU Utilization Rate (%)
One Drive	178	26.9
7 Drives Hardware Striping (No Fault Tolerance)	94	27.5

The table above illustrates how multiple drives in a logical array can improve the response time over spindles by up to 89.4%.




Figure 3-9. Drive spindle performance comparison

Figure 3-9 illustrates that using striping across 7 drives, rather than spindles, can achieve better performance with very minimal cost from a resource utilization standpoint. Compaq engineers strongly recommend 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 engineers recommend 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 4MB of read/write cache.

Software striping is under control of 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 will decrease the overall system performance.

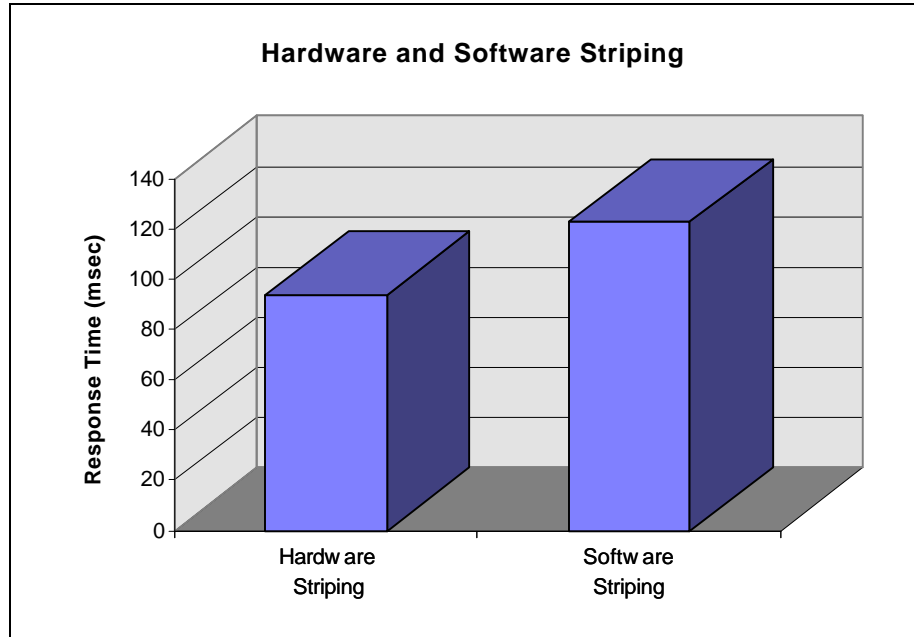


Figure 3-10. Response time comparison of hardware/software striping

The chart above indicates a comparable response time between hardware and software striping. Notice the 30.85% decrease in response time when hardware striping is used versus software striping. This can be attributed to the higher utilization of system resources required for managing the software striping.

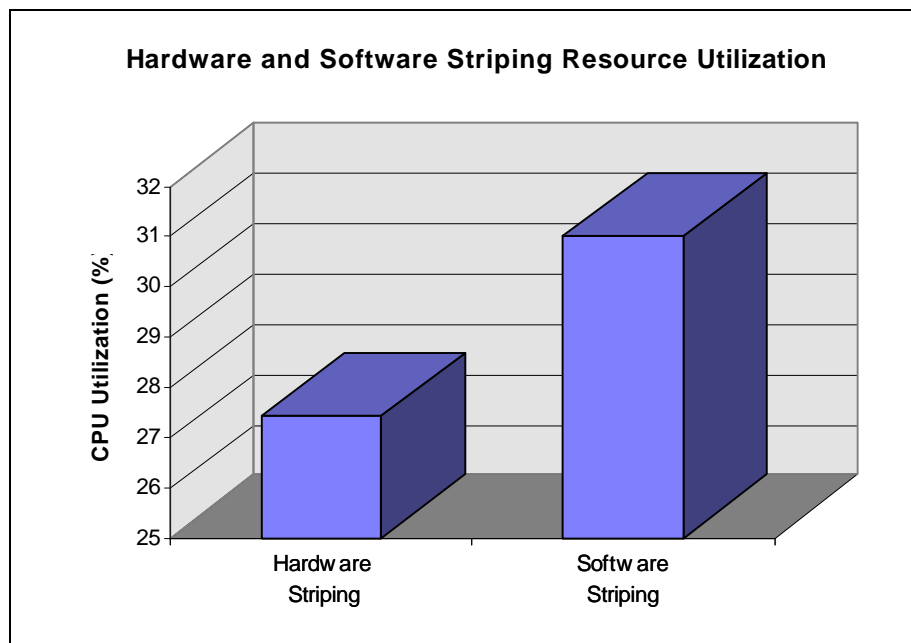


Figure 3-11. CPU Utilization comparison of hardware/software striping

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 of about 30-40% in performance and between 10-15% lower CPU utilization rates using hardware striping. We strongly recommend the use of hardware striping for better performance and more system capacity.

Fault Tolerance

The customer has several available options when configuring the Lotus Notes 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 the customer selects effects 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 that are not found with Fast-Wide SCSI-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 5 - Distributed Data Guarding

RAID 5 is also 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 which follow used seven drives.

RAID 4 Data Guarding

RAID 4 is also referred to as data guarding because it assures data reliability similar to RAID 5. 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. The 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 is the same as for RAID 5, 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. In most tests, RAID 4 performs below RAID 5.

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 0 - No Fault Tolerance

No fault tolerance is provided. 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 will be lost. None of the logical drive capacity is used for redundant data, so RAID 0 offers the best processing speed as well as capacity. RAID 0 is appropriate for applications that deals with non-critical data requiring high speed access.




Figure 3-12. Fault tolerance type comparison of response time

Compaq test results show that there is a measurable difference in response time rates between RAID 4 and RAID 1, RAID 5, or RAID 0. RAID 0 achieved the best performance, outperforming RAID 5 by 10-20% in response time and performing about 5 times better than RAID 4. RAID 5 response times are more than 4 times faster than tests run using RAID 4. You should 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 would be to use RAID 1 over RAID 5 or RAID 4 for systems with critical data because of the performance gains expected combined with the hardware fault tolerance protection. RAID 5 would be fine for data that is not mission critical, and would offer better usage of disk capacity than RAID 1. You should also take into consideration the CPU utilization between RAID 0, RAID 1, RAID 4, and RAID 5.




Figure 3-13. Fault tolerance type comparison of CPU utilization

Compaq tests results quantified these performance differences as follows. While RAID 1 and RAID 5 used 2-3% more CPU resources than RAID 0, RAID 4 used the least CPU resources. Even though RAID 4 used less CPU resources than RAID 5 or RAID 1, its performance is too poor to recommend for a Lotus Notes Server. For systems with data that is not critical, RAID 5 would be recommended because of its fault tolerance support with distributed parity data performing very well at a very low system utilization cost. For systems with mission critical data, and those used for online applications (which is the most common scenario), RAID 1 is recommended because it provides the system with proper protection at the cost of slightly more system overhead combined with excellent performance. 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 performance advantage combined with the total usable disk 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 267MB/s to ensure the balanced performance of four processor systems. A bridged PCI bus implementation provides a total throughput capability of 133MB/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 mentionable difference in performance of EISA versus PCI because they did not use the dual, peer PCI architecture of the ProLiant 5000.

3-34 Performance Management

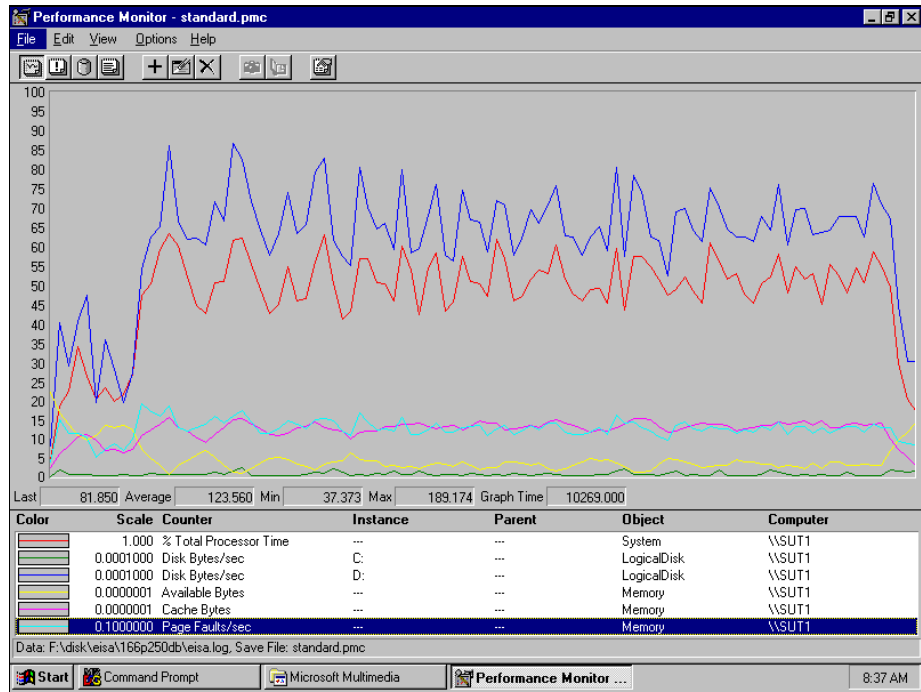


Figure 3-14. EISA test system - resource utilization rates

Based on the tests performed with the Compaq ProLiant 1500, the Performance log was used to graphically represent system resource utilization during a test run with the 1500 configured with a PCI bus disk controller and a workload simulating 500 users. As is apparent when comparing the next chart (Figure 3-15) to the previous chart (Figure 3-14) which showed the system resource utilization during a test run with the 1500 configured with an EISA bus disk controller, there is virtually no mentionable difference between the system resource utilization of the PCI and EISA configurations with the 1500.

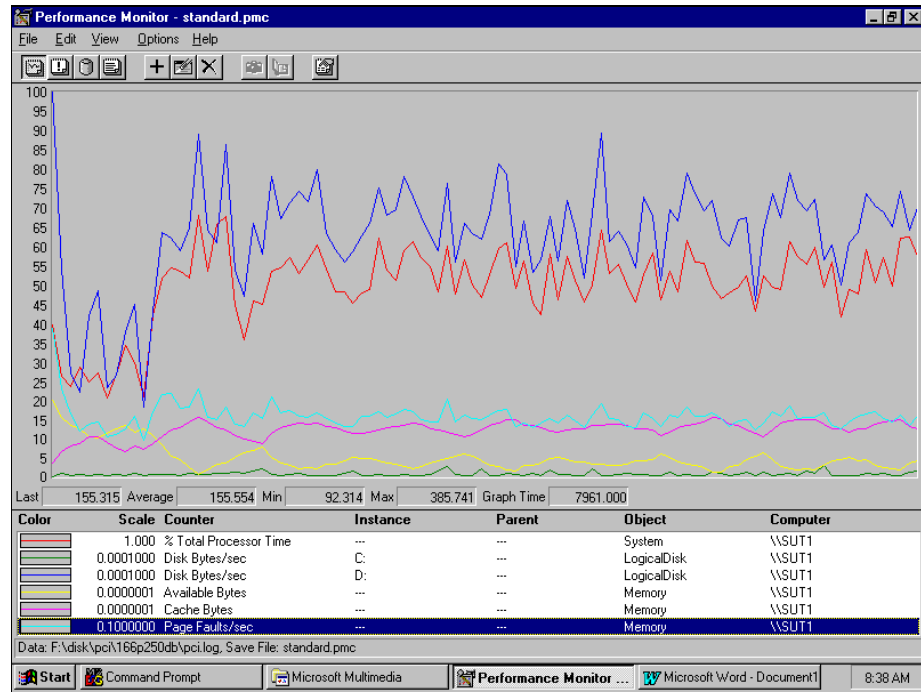


Figure 3-15. PCI test system - resource utilization rates

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 Notes server will probably depend 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 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 the faster operating file system when running under NT Server. The testing performed by Compaq engineers supports this claim. The performance impact of FAT versus NT file systems is presented with the following data and recommendations.

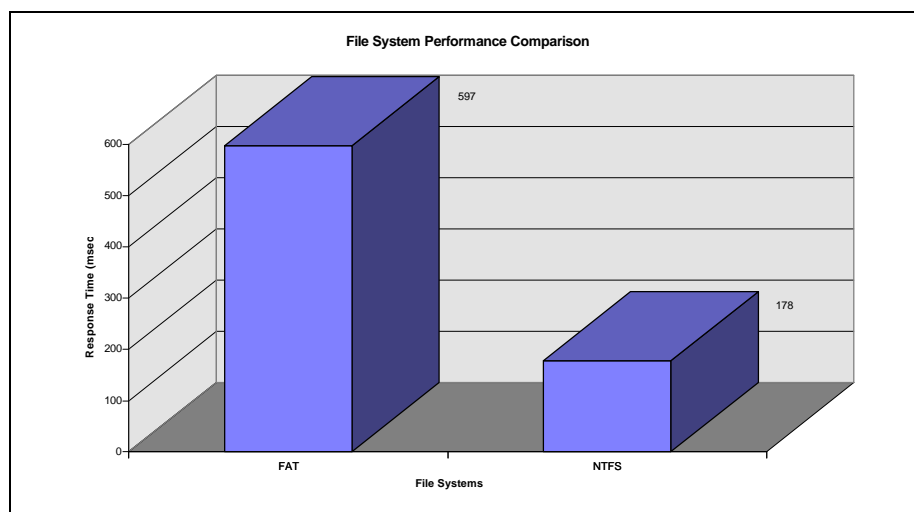


Figure 3-16. FAT versus NTFS comparison of response time

The response time rate for tests run on a system using NTFS as the file system showed a performance advantage almost three times better than when using a FAT file system.

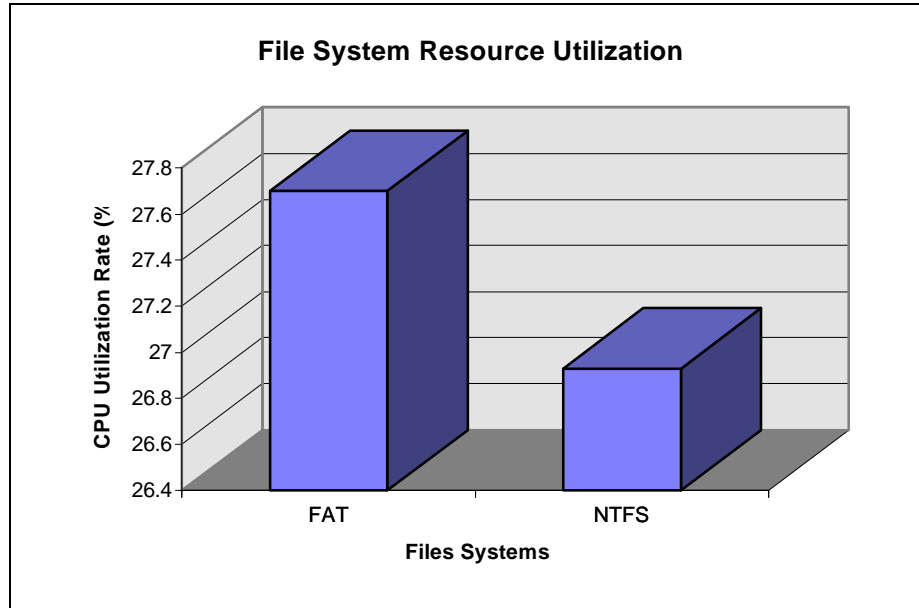


Figure 3-17. FAT versus NTFS comparison of CPU utilization

NTFS is much better in NT/Notes environment than FAT. As illustrated in the above chart, the CPU Utilization of NTFS versus FAT is lower. Due to the combination of better response time performance and lower CPU requirements, Compaq engineers recommend using NTFS as the file system for the Lotus Notes server. NTFS response time performance is about three times faster and about 1 percent less CPU bandwidth is required.

Networking

In a test environment that is purely Lotus Notes, the networking subsystem is less likely to cause performance problems than the subsystem areas previously discussed. However, in an enterprise network environment, the network subsystem becomes a performance factor due to the replication that occurs between servers. The “Networking Subsystem” segment under “Performance Tuning” offers guidelines for identifying performance problems that are network related. Also presented are network management guidelines as well as strategies for increasing network throughput—should this subsystem be the source of performance issues.

Performance Tuning

Hard Disk Controller Tuning

Some of these features offer performance and fault tolerance advantages which were discussed in an earlier section discussing 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 will be 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 4MB 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 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.

3-40 Performance Management

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

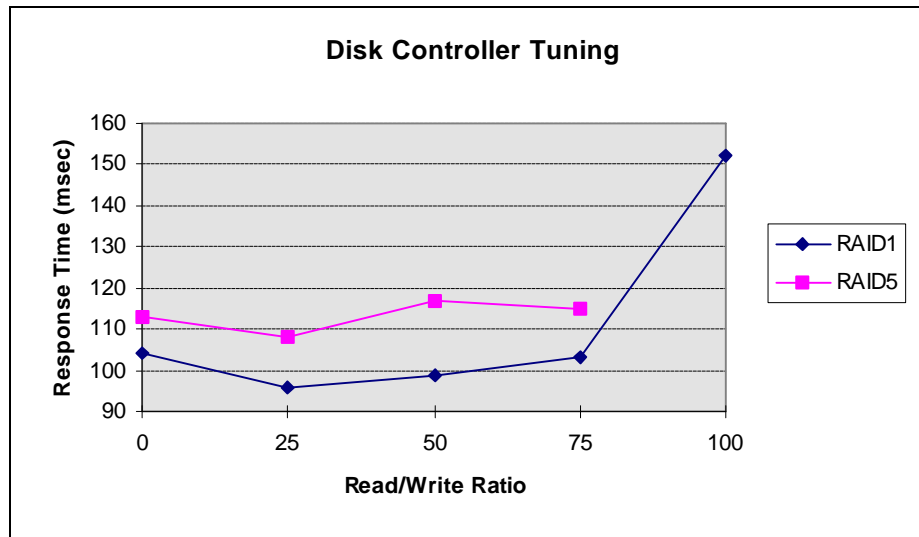


Figure 3-18. Smart-2P Disk Array Controller Array Accelerator Configured with various Read/Write Ratios

The chart above shows a performance comparison with the Array Accelerator read/write cache configured with various read/write ratios. As the chart illustrates, 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.

Notes Server Tuning

NSF Buffer Size

The NSF buffer is the most important Notes 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 Notes Server automatically determines this value. By default, 25 percent 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 4MB 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.

Notes Server NSF Buffer Size

A Notes Server tunable parameter that impacts the system performance is the NSF Buffer Size which is the amount of memory allocated to the Notes 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

Networking Subsystem

This segment deals with two performance enhancement strategies: segmenting the LAN and increasing the LAN bandwidth.

Segmenting the LAN

A key strategy that can increase networking subsystem performance is to divide a single Ethernet segment into multiple network segments. If you determine the networking subsystem is not reaching optimum throughput, there are two network implementations that can improve the overall throughput and general performance gain of a network.

- Physical segmentation

To physically segment a network, you first add more NICs to the server and then balance the network load among the multiple NICs. Segmenting a network by adding additional NICs and hubs has the added benefit of creating separate collision domains. Creating additional collision domains minimizes packet collisions by decreasing the number of workstations on the same physical network.

- Network switching technology (microsegmenting)

Switching-hubs, similar to routers and bridges, also provide LAN segmentation capabilities. LAN switches provide dedicated, packet-switched connections between their ports. The packet-switched connection provides simultaneous switching of packets between the hub ports which increases the available bandwidth.

Migrating to 100-Mb/s Technology

Migrating a network Ethernet implementation from 10Base-T to 100Base-TX or 100VG-AnyLAN provides 100 Mb/s of shared bandwidth for the LAN

.....

3-44 *Performance Management*

clients. Implementing this type of change can substantially improve network throughput and overall performance.

A gradual migration to the faster Ethernet technology does not have to be expensive and time consuming. Partially converting your LAN is a viable alternative to converting all clients on the LAN simultaneously.

The advantages of upgrading a server to a 100-Mb/s NIC while accommodating existing LAN clients with a bandwidth of 10 Mb/s are as follows:

- Cost effectiveness: it's not as expensive as converting all clients at the same time
- Better throughput: aggregate network throughput is improved because the transmission speed from the server to the hub is faster
- Ease of upgrade: replacing your 10-Mb/s NIC with a 100-Mb/s NIC is not a difficult process
- No complex cable requirements: you can use your existing 10-Mb/s cable

The disadvantages of upgrading the server NIC to 100 Mb/s while leaving clients at 10 Mb/s are:

- Replacing existing 10-Mb/s NICs with the more expensive 100-Mb/s NICs might be cost prohibitive depending on the number of NICs being replaced.
- Re-routing all existing clients to a switching hub is required. Depending on the amount of clients, this can be an inconvenience to an administrator.

Reviewing Migration Results

NOTE: The testing tool, NetBench, used for network subsystem analysis is not the same as and should not be confused with the NotesBench tool mentioned previously in this paper.

To compare and evaluate 10-Mb/s Ethernet and 100-Mb/s Ethernet, parallel test environments were set up in the integration testing labs at Compaq. The results, shown in Table 3-3, compare a 10-Mb/s Ethernet LAN with that of a 100-Mb/s Ethernet LAN. Windows NT Performance Monitor indicates the total throughput for the NetFlex-3/P controller installed in the server. Compare the throughput of the 10-Mb/s NIC to that of the 100-Mb/s NIC. Theoretically, the maximum data transmission rate should increase by a factor of 10 when migrating from the 10-Mb/s NIC to the 100-Mb/s NIC. Table 3-3 shows the NetBench 4.0 throughput results for a maximum of 10 clients, running at 10-Mb/s and 100-Mb/s Ethernet.

Table 3-3
Single Segment LAN Throughput

Number of Clients	Ethernet Bandwidth	Total Throughput (Mb/s)
4	10 Mb/s	9.3
10	10 Mb/s	9.4
4	100 Mb/s	69.3
10	100 Mb/s	90.1

Figure 3-19 is a graphical representation of one Ethernet segment of 10-Mb/s and 100-Mb/s clients (IPX protocol). The results are from a NetBench 4.0 monitoring session where Total Throughput was captured for a ProLiant 1500 server equipped with a NetFlex-3/P Controller (100-Mb/s TX Module). The graph illustrates that the Total Throughput increased as the number of clients increased, then saturated at 4 clients.

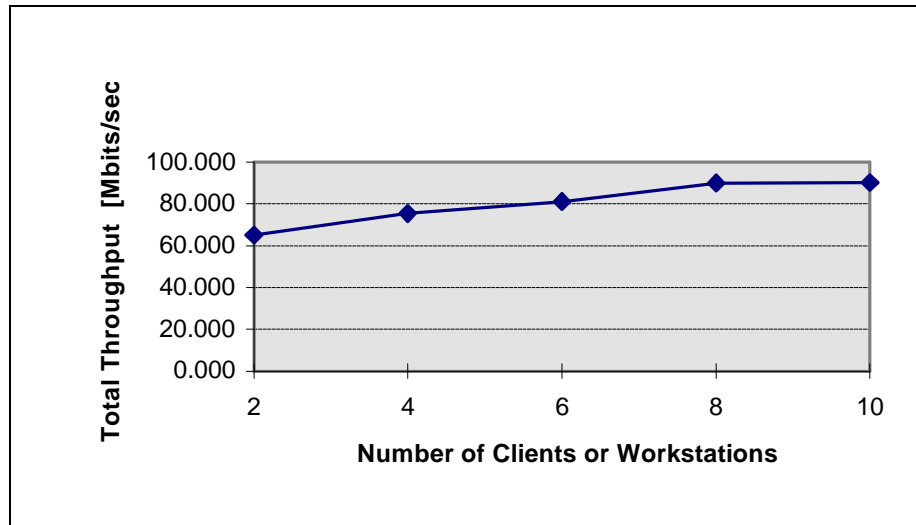


Figure 3-19. NetBench 4.0 result of total throughput

In general, if total throughput stays around 50 percent or above on a consistent basis, then your LAN is approaching network saturation or it may be bottlenecked. In both of our test cases, the LAN saturated at 4 clients. Reaching saturation level with such a low number of clients indicates a definite need for segmenting the LAN to distribute the work load.

As noted, our testing results indicate wire saturation for a very low number of users because NetBench creates a test environment that simulates network demand placed on a file server; every client reads the same data from a datafile. The use of a synthetic network measuring program (NetBench) and the even distribution of work caused the low saturation point. Thus, our testing does not represent a typical LAN environment of several hundreds or thousands of users that arbitrarily broadcast over the entire LAN via router(s), bridge(s), and gateways. In a real world environment, network clients should not reach wire saturation for so few users, as indicated in Table 3-3 and Figure 3-19. The table and graphical data show the wire bandwidth difference between 10 Mb/s and 100 Mb/s as well as the effect of increasing the user load.

Performance Conclusions

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 Notes 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.

Memory

Memory was found to be a resource such that more is not necessarily always better. The amount of memory needs to be properly balanced with the system's need for the resource. Too much or too little memory had a negative impact on performance.

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 was 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 the protection of the 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 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 drives when mirroring with RAID 1.

Notes Server NSF Buffer Size

A Notes Server tunable parameter that impacts the system performance is the NSF Buffer Size which is the amount of memory allocated to the Notes 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 Notes 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 Notes administrator or integrator. Capacity planning is closely tied to performance management. Notes 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 err on the side of caution by not planning 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 ongoing 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 Notes server implementation will be 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 Notes server requirements including accurately capturing the number of users, application profile, and so on. When recording this information, it would be 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 Notes usage. For example, perhaps Company XYZ is preparing to implement Lotus Notes for a group of 1800 mail users, 1500 mail and shared database users, and 300 groupware users. From a capacity planning standpoint, we need to know if there is any overlap between these user counts and their needs. For instance we 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 Notes 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 will be most critical to your environment.
-

Evaluating Server Resource Alternatives

Upon completion of your Lotus Notes 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 Notes 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.
- 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:

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

Capacity Planning Methodology

- Capacity Characteristics Observations
 - Relate Real World Users to NotesBench Workload Users
 - Select the System
 - Determine the Memory Configuration
 - Determine the Disk Subsystem Configuration
 - Determine the Network 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 engineers also saw a relationship develop between CPU utilization and response time as shown in the following chart. 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 percent utilization would not be desired.




Figure 4-2. CPU utilization versus response time performance

Compaq engineers consider 70 percent CPU utilization to be a rule of thumb threshold that leaves a 30 percent 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 percent, scaling utilization back to about 50 percent. The additional 20 percent 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.

Relate Real World Users to NotesBench Workload Users

To begin our model, we must create a formula that relates real users to a corresponding number of NotesBench workload users. We know that given a P5/166MHz four-processor system, the maximum number of supported NotesBench mail workload users is 1800. Using the 50 percent utilization discount discussed previously, we could say that 900 NotesBench mail workload users can be supported providing very good performance while allowing some resources for other activities. Therefore we could assume the following model equating NotesBench mail workload users to planned mail users:

Maximum NotesBench mail users 50% = # of planned mail users

Thus if we are provided the number of planned users and want to calculate the number of NotesBench users, we could assume the following equation is also true:

NotesBench mail users = # of planned mail users 50%

The above formula is mathematically equivalent to the equation below that we will use as part of our methodology for converting planned mail users to NotesBench mail users:

NotesBench mail users = # of planned mail users 2

Based on the previous explanation, the number of required mail users or mail users that also share a discussion database (MailDB) will also be multiplied by a factor of two to derive an equivalent representation of NotesBench mail and MailDB workload number of users that would fall within the comfort zone of the CPU utilization range.

Because Groupware_A workload includes a lot of background tasks a weight of 100 percent will be used, thus basically equating NotesBench Groupware_A workload users to planned groupware users.

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

Mail: *CP NotesBench Mail Users = # of planned Mail users 2*

Mail and shared database: *CP NotesBench MailDB Users = # of planned Mail/Shared DB users 2*

Groupware: *CP NotesBench Groupware_A Users = # of planned groupware users*

Select the System

The selection of the system type is going to be derived from the NotesBench users calculated in the above section. For instance, we can take the calculated number of NotesBench users and refer back to Figure 3-4, CPU Performance Comparison, Figure 3-5, P5/133 Processor Scalability, Figure 3-6, P5/166 Processor Scalability, and Figure 3-7, P6/166 Processor Scalability. From these figures, we have the maximum number of users supported for NotesBench Mail and MailDB workloads.

For example, referring to Figure 3-6 we see that the P5/166 4 Processor system can support 1,800 Mail workload users and 1500 MailDB workload users. The P5/133 2 Processor supports 1,100 Mail workload users and 900 MailDB workload users based on Figure 3-5.

The number of Groupware_A users supported on our test systems were:

- 4xP5/133 - 400 users
- 4xP5/166 - 500 users
- 2xP6/166 - 500 users

Based on the calculated CP NotesBench Mail, MailDB, and Groupware_A Users, we can derive the number of server(s) needed to meet all profile requirements by plugging numbers (obtained from the Figures) into the following formulas:

.....

4-10 Capacity Planning

$$\text{Mail:} \quad \frac{\text{CPNotesBenchMail Workload Users}}{\text{CPUMaxMail Workload}} \quad \# \text{ of Systems}$$

$$\text{MailDB:} \quad \frac{\text{CPNotesBenchMailDB Workload Users}}{\text{CPUMaxMailDB Workload}} \quad \# \text{ of Systems}$$

$$\text{Groupware_A:} \quad \frac{\text{CPNotesBenchGroupware A Users}}{\text{CPUMaxGroupware A Workload}} \quad \# \text{ of Systems}$$

We can then add the # of Server(s) for Mail, MailDB, and Groupware to determine the total server(s) required.

For example, let's select the P5/166 4P configuration and assume that we have 1800 CP NotesBench Mail Users, 700 CP NotesBench MailDB Users, and 200 CP Groupware_A Users. For this scenario, the formulas would look like the following:

$$\text{Mail:} \quad \frac{\text{CPNotesBenchMail Workload Usrs}}{\text{CPUMaxMail Workload}} \quad \# \text{ of Systems}$$

$$\text{MailDB:} \quad \frac{\text{CPNotesBenchMailDB Workload Users}}{\text{CPUMaxMailDB Workload}} \quad \# \text{ of Systems}$$

$$\text{Groupware_A:} \quad \frac{\text{CPNotesBenchGroupware A Workload Users}}{\text{CPUMaxGroupware A Workload}} \quad \# \text{ of Systems}$$

$$\text{Mail:} \quad \frac{1800}{1800} = 1 \text{ systems}$$

$$\text{MailDB:} \quad \frac{700}{1500} = .47 \text{ systems}$$

$$\text{Groupware_A:} \quad \frac{200}{500} = 0.4 \text{ systems}$$

Total = 1.87 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/XXX 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.

Determine the Memory Configuration

Memory requirements relate to the number of workload users as was discussed in the Performance Management chapter. The chart below is taken directly from that section of this TechNote.

Number of Users	Minimum Memory Required (MB)	Recommend Memory Configuration (MB)
500 or less	128	192
600	192	256
800	256	320
1000	320	384
1200	384	448
1400	448	512
1600	512	576
1800	576	768

The memory recommendations above is the same chart included previously in Chapter 3, "Performance Management".

Determine the Disk Subsystem Configuration

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

System Volume: The system volume should be a single drive attached to the Fast-SCSI 2 bus, formatted as NTFS. The following should be included on this volume:

- Windows NT Server
- Lotus Notes Server software, excluding data files
- Notes 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 servers. This drive array should be configured with a fault tolerance level 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 Notes data
 - NT Paging File (size dependent upon memory configuration)
-

Capacity Planning Case Study

A fictitious company, XYZ, wants to provide messaging capabilities to 900 employees. In addition, another 500 users must be provided with messaging and discussion database capabilities. Five group discussion databases will be used. One hundred additional users have groupware requirements. The future Notes Administrator of XYZ is asking for information that will tell him how many servers will be needed to meet these user requirements. The administrator also needs to know how the server should be configured regarding memory and disk storage and comments that good response time is a definite requirement for this implementation to be a success. Monetary resources will have to be justified, but availability of funds is not considered to be a problem. The company classifies this data as non-critical.

Step 1: Relate Real World Users to NotesBench Workload Users

As the first step in providing a solution, the real world user requirements must be plugged into the formulas mentioned in the methodology discussion which will convert planned real world user numbers into NotesBench workload users.

Mail: $CP\ NotesBench\ Mail\ Users = \#\ of\ planned\ Mail\ users \times 2$

$CP\ NotesBench\ Mail\ Users = 900 \times 2 = 1800$

Mail and shared database: $CP\ NotesBench\ MailDB\ Users = \#\ of\ planned\ Mail/Shared\ DB\ users \times 2$

$CP\ NotesBench\ MailDB\ Users = 500 \times 2 = 1000$

Groupware: $CP\ NotesBench\ Groupware_A\ Users = \#\ of\ planned\ groupware\ users$

$CP\ NotesBench\ Groupware_A\ Users = \#\ of\ Planned\ Groupware_A\ Users = 100$

Step 2: Select the System

For this scenario, the P5/166 4-processor system is assumed. The same model can be used to determine the number of other system types as well.

$$\text{Mail: } \frac{1800}{1800} = 1 \text{ Systems}$$

$$\text{MailDB: } \frac{1000}{1500} = .67 \text{ Systems}$$

$$\text{Groupware: } \frac{100}{500} = .20 \text{ Systems}$$

Total = 1.87 systems

Based on this Capacity Planning Model, two P5/166 4P 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—900 Mail Users
- Server 2—500 Mail & Shared DB Users + 100 groupware users

Step 3: Determine the Memory Configuration

Server 1

Server 1 should be configured with 384MB of memory. This recommendation is pulled directly from the memory recommendation table included in the methodology section. 384MB is derived from 320MB that would be recommended for the 900 real world Mail users plus 64MB added for additional server activities.

Server 2

Server 2 should be configured with 384MB of memory derived as follows:

- 500 Mail and Shared DB users + 100 groupware users = 600 Users
-

- 600 Users from the Memory Recommendation Chart = 256 MB
- + 1 MB × 100 groupware users = 100 MB
- Total Memory Calculated for Server 2 = 356 MB

NOTE: An assumption made is to provide an additional 1MB/groupware user

Total Memory Recommended for Server 2 = 384 MB

NOTE: The actual memory total of 356MB is not a possible configuration, so a slight adjustment upwards was made to the next memory configuration possible. This slight adjustment upwards does not necessarily buy better performance, but more memory provides security against lost connections.

Step 4: Determine the Disk Subsystem Configuration

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

- **System Volume:** The system volume should be a single drive attached to the Fast-SCSI 2 bus, formatted as NTFS. The following should be included on this volume:
 - ❑ Windows NT Server
 - ❑ Lotus Notes Server software, excluding data files
 - ❑ Notes 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 Notes data
 - ❑ NT Paging File (size dependent upon memory configuration)

Server 1

The size of the data volume allocated for Lotus Notes data on Server 1 should be at least 9GB, allowing approximately 10MB/user for Mail.
(900 users × 10MB = 9000MB or 9GB)

An additional 422MB 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 to be 384MB, there fore the NT Paging File size is calculated as 384MB + 38MB—or 422MB.

To support these storage requirements, six 2.1GB drives should provide over 9.4GB of adequate storage for the 900 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 2.1GB drives provides total disk space of 12.6GB. This value needs to be multiplied by 6/7 to calculate total usable disk space which is 10.8GB. This is the best fit for the drives, because five 2.1GB drives only allows 8.75GB 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 5GB, allowing approximately 10MB/user, for Mail.

$(500 \text{ users} \times 10\text{MB} = 5000\text{MB or } 5\text{GB})$

An additional 500MB 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 100MB for each database; $5 \text{ databases} \times 100\text{MB/database} = 500\text{MB}$).

A final 10GB of space is allowed for the 100 groupware users. (The rule of thumb used to calculate the storage requirements for the groupware users allows 100MB for each user; $100 \text{ groupware users} \times 100\text{MB/user} = 10000\text{MB or } 10\text{GB}$). The combined storage requirements are 15.9GB.

To support these storage requirements, five 4.3GB drives configured in an array using RAID 5 should provide 17.9GB** of usable storage for the 500 mail and discussion database users plus the 100 groupware users. Note that no other storage requirements were considered in calculating this minimum storage.

An additional 422MB 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 to be 384MB, therefore the NT Paging File size is calculated as $384\text{MB} + 38\text{MB}$ —or 422MB.

NOTE: To calculate the usable disk storage using RAID 5 hardware fault tolerance, the formula is $n/(n+1)$. For example, five 4.3GB drives provides total disk space of 21.5GB. This value needs to be multiplied by 5/6 to calculate total usable disk space which is 17.91GB. This is the best fit for the drives, because four 4.3GB drives only allows around 13.76GB of usable disk space which does not meet the calculated requirements.

Step 5: Determine the Network Configuration

Due to Notes client/server architecture, it does not impose any significant loads onto the Network. However, there are pragmatic steps you should take to ensure network constraints and overloads are minimized since end user response times can be significantly impacted by network-based bottlenecks. The following discussion is based on several assumptions:

Assumptions.

Two servers (Notes-1 and Notes-2) will host the Notes server application for active users.

One additional NT server will provide file and print servers as well as function as the NT Domain controller.

Mobile and remote users will attain access to Notes-1 and Notes-2 via:

- WindowsNT RAS (via the file and print server).
- User Notes server pass through capacity on Notes-1 (see note below)

Recommendations.

- **Disable NetBIOS.** Remote the NetBIOS protocol on the NT and Notes servers if this protocol is not used within the environment. This eliminates constantly NetBIOS broadcast messages on the LAN segment.
-

- **Schedule Notes indexing and compacting.** This task leverages available CPU resources and can increase response times experienced by active Notes users. Therefore, scheduled times should minimize overlap of peak remote use times (remote sessions use the X.PC protocol).
- **Server distribution.** When topology design accommodates server distribution, Notes-1 and Notes-2 should be on the same LAN segment. If the file server is on a separate segment, then non-Notes traffic (for example—print traffic, file I/O, and such) will not overlap completely with Notes based traffic.
- **Place common replicas on each server.** To minimize users maintaining servers sessions on both Notes servers, place common replicas on each server. This enable organizations to better retain their original user load distribution plans.
- **Use highspeed PCI-based NICs.** These NICs with speeds of 10/100Mbps can be used in servers to reduce network bandwidth utilization times for server to server functions (for example—mail routing and replication).
- **For remote locations with high user concentrations, consider placing a Notes server at the site.** While not necessary, this provides the benefit of (1) reducing WAN based delays for end users, (2) increasing the robustness of the architecture for disasters (i.e. loss of a key component does not immediately impact all users), and (3) simplifying deployment and upgrades of specific functional/geographical entities.

NOTE: Remote sessions (X.PC) stress CPU resources due to high frequency polling and therefore pass through should be placed on the Notes server supporting the least CPU intensive users.

Appendix A

Multivendor Integration From Compaq

To help you successfully integrate and optimize your network or multivendor system, Compaq has developed a library of systems integration TechNotes for the NetWare, Microsoft Windows NT, SCO UNIX, and OS/2 operating system environments. TechNotes provide you with important information on topics such as network performance management, server management, and operating system interconnectivity.

To get a copy of a particular TechNote in electronic format, you can access:

- Downloadable files from the Internet at:
[http:// www.compaq.com](http://www.compaq.com)
- Downloadable files from the Compaq Reference Library CD

Pricing

Compaq TechNotes U.S. \$15 (CND \$20)*

*Plus shipping & handling

How to Order

To order by PHONE (credit card orders only), call:

- 1-800-952-7689 (In U.S.: 7 a.m. - 7 p.m. C.S.T.)
- 1-800-263-5868 (In Canada: 24 hours)
- 1-317-364-7281 (From outside U.S. & Canada:
7 a.m. - 7 p.m. C.S.T.)

To order by MAIL or FAX, send complete order information to:

Worldwide: Compaq Fulfillment Center
(except Canada & P.O. Box 4100
France) Crawfordsville, Indiana 47933

FAX: 1-317-364-0787

Canada: Compaq TechNotes
Compaq Canada, Inc.
111 Granton Drive, Suite 101
Richmond Hill, ON L4B 1L5

FAX: 1-800-668-1835

France: Compaq TechNotes
Compaq Computer S.A.R.L.
5, Avenue de Norvege
91959 LES ULIS Cedex
France

FAX: 1-331-64-46-51-81

Current TechNotes

The current TechNotes available are listed in Table A-1.

**Table A-1
Compaq TechNotes**

TechNote Name	Part Number
<i>Compaq Insight Management</i>	
Integrating Compaq Insight Manager with Enterprise Management Platforms (5/96)	145806-004
Integrating Compaq Insight Manager with Microsoft Systems Management Server (8/95)	184720-001
Integrating Compaq Insight Manager with ManageWise (8/95)	182074-001
<i>Microsoft</i>	
Sizing Microsoft BackOffice 1.5 on Compaq Servers (2/96)	185496-001
Implementing Microsoft Windows NT Server RAS, DHCP, and WINS on Compaq Servers (8/95)	184488-001
Migrating to Microsoft Windows NT Server 3.51 on Compaq Servers (8/95)	185026-001
Configuration and Tuning of Oracle7 for Windows NT on Compaq Servers (7/95)	184853-001
Configuration and Tuning of Microsoft SQL Server for Windows NT on Compaq Servers (7/94)	184207-001
Compaq Backup and Recovery for Microsoft SQL Server (3/95)	184489-001
Migrating from Microsoft Windows NT Advanced Server 3.1 to Microsoft Windows NT Server 3.5 (2/95)	184669-001

continued

.....

A-4 Multivendor Integration From Compaq

Compaq TechNotes Continued

TechNote Name	Part Number
Microsoft (continued)	
Drive Subsystem Performance with Microsoft LAN Manager 2.1 (11/92)	145517-001
Implementing Lotus Notes on Compaq Servers with Microsoft LAN Manager (9/92)	145201-001
Microsoft SQL Server and Novell NetWare Requester for OS/2 (9/92)	145376-001
NetWare	
SYBASE SQL Server 4.2.2 for NetWare: Tuning and Performance (2/95)	184429-001
NetWare 4 Performance Management (9/95)	137989-002
Compaq Answers to Technical Support Questions for NetWare Environments (2/95)	145710-002
Performance of Lotus Notes 3.1 on Compaq ProLiant Servers with NetWare 3.1x (2/95)	184430-001
NetWare Migration from 3.12 to 4 (11/94)	181152-001
Configuration and Tuning of Oracle7 for NetWare on Compaq Servers (11/94)	
NetWare Migration from 2.x to 3.12 (6/94)	181069-001
Tape Backup Using ARCserve v5.01 - Windows Edition from Compaq (1/94)	137672-001
NetWare SFT III v3.11 Installation (9/93)	145905-001
ORACLE Server Planning in a NetWare v3.11 Environment (5/93)	145143-001
Performance Management in a NetWare v3.1x Environment (4/93)	133399-002

continued

Compaq TechNotes Continued

TechNote Name	Part Number
<i>NetWare (continued)</i>	
NetWare Tape Backup Management Using ARCserve from Compaq (4/93)	137535-001
Compaq IDA-2 Configuration for NetWare (12/92)	145536-001
Compaq ProSignia Configuration Guidelines for NetWare v3.x (12/92)	145537-001
NetWare Server Management (9/92)	145231-001
<i>UNIX</i>	
Configuration and Tuning of Sybase System 10 for SCO UNIX Open Server 3.0 on Compaq Servers (9/95)	185111-001
Configuration and Tuning of Sybase System 10 for Novell UnixWare 2.0 on Compaq Servers (7/95)	184-942-001
Configuration and Tuning of Oracle7 for SCO UNIX on Compaq Servers (11/94)	184329-001
Disk Mirroring with VERITAS VxMirror (12/93)	145656-001
The Compaq, SCO UNIX, and ORACLE7 Database Server (9/93)	195818-001
Performance Tuning for SCO UNIX on Compaq Systems (5/93)	145730-001
SCO UNIX IP Routing Over X.25 Networks (4/93)	145711-001
Printing in SCO UNIX and NetWare Integrated Environments (3/93)	145538-001
SCO UNIX Connectivity to SNA Environments Over X.25 Networks (1/93)	145583-001

continued

Compaq TechNotes Continued

TechNote Name	Part Number
UNIX (continued)	
SCO UNIX Database Servers in Novell Networks (10/92)	145402-001
SCO UNIX in an SNA Environment Using Computone LYNX (7/92)	145116-001
SCO UNIX in an SNA Environment Using CLEO DataTalker U/X (5/92)	145081-001
IBM	
Compaq Answers to Technical Support Questions for OS/2 Environments (11/94)	
Performance of Lotus Notes 3.1.5 on Compaq ProLiant Servers with OS/2 2.11 (3/95)	
Technology Overview	
Configuring Compaq RAID Technology for Database Servers (7/94)	184206-001
Integration Server	
Updating Compaq System Software from an Integration Server (5/96)	219467-001




NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO. 44 CRAWFORDSVILLE, INDIANA

POSTAGE WILL BE PAID BY ADDRESSEE

COMPAQ TECHNOTES
COMPAQ COMPUTER CORPORATION
PO BOX 4100
CRAWFORDSVILLE IN 47933-9940



(Fold Here)

Please Tape Here

Please Tape Here

USER SURVEY

Please help us provide quality documentation by completing and returning this postage-paid form.

TechNote Title: _____

1. Overall rating of the document:
 Excellent Satisfactory Fair Poor
2. Does this document contain all the information you need? Yes No
If not, what information is missing? _____
3. Is the information appropriate to your experience? Yes No
Explain: _____
4. Are the checklists, charts, etc., provided useful? Yes No
What additional items would be helpful? _____
5. What part of this document do you use the most? _____
6. How often do you anticipate referring to this TechNote?
 Never Once Occasionally Often
7. How did you first learn about Compaq TechNotes?
 Compaq advertising Compaq Reseller/Distributor
 Compaq salesperson Compaq Customer support
 Other: *(explain)* _____
8. Does your company own any Compaq products?
 Server Desktops Portables Network laser printers
9. What operating system(s) does your company use?
 Microsoft Windows NT NetWare SCO UNIX Oracle7
 IBM LanServer Banyan VINES UnixWare
 Other *(specify)* _____
10. What is your role in the acquisition of computer products at your company?
 VP Finance Manager/Information Systems Network Administrator
 Purchaser User Other _____

(Optional) If you would like to receive future information about Compaq TechNotes, complete the following:

Name: _____ Title: _____

Company: _____

Address: _____

City: _____ State: _____ Zip: _____

If mailed outside of the U.S.A. please affix proper postage to ensure delivery.

Please fold card over and tape closed. Do not staple.




NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS MAIL PERMIT NO. 400 HOUSTON, TEXAS

POSTAGE WILL BE PAID BY ADDRESSEE

MAILCODE 090307
COMPAQ TECHNOTES
COMPAQ COMPUTER CORPORATION
PO BOX 692000
HOUSTON TX 77269-9976



(Fold Here)

Please Tape Here

Please Tape Here

Index

A

- abbreviations for atomic tasks 2-4
- atomic task
 - testing conclusions 2-13
- atomic task abbreviations 2-4
- atomic task profiling, uses for 2-1
- atomic tasks
 - LAN subsystem results 2-12, 2-13
 - performed on CPU 2-5
 - performed on disk subsystem 2-8
 - results of CPU testing 2-6, 2-7
 - results of disk subsystem testing 2-8, 2-10

B

- bandwidth, increasing 3-43
- bar charts, general description of 2-2
- benchmark, definition of 3-4
- benchmark, uses of 3-4

C

- cabling requirements
 - general 3-44
- capacity planning
 - assessing disk subsystem requirements 4-12
 - case study 4-13
 - computing memory requirements 4-11
 - computing server requirements 4-11

- defined 4-1
- formulas for 4-9, 4-10
- importance of 4-2
- process overview 4-4
- recommendations 4-18
- system selection process 4-9
- charts, bar, general description of 2-2
- client systems, software used on 3-11
- collision domains 3-43
- collisions, minimizing 3-43
- CPU
 - atomic tasks performed on 2-5
 - Pentium Pro, advantages of 3-13, 3-15
 - summary of testing 3-47
- CPU utilization
 - effects on response time 4-6
 - planning recommendations 4-7

D

- data, collection process for testing 2-2
- destination server, software used on 3-11
- disk array controller
 - Compaq Smart-2 bus compatibility 3-33
 - Compaq Smart-2 features 3-28, 3-29, 3-39
 - tuning recommendations 3-42
- disk subsystem
 - atomic tasks performed on 2-8
 - bus architecture performance comparison 3-34
 - configuring 4-12
 - drive spindle comparison chart 3-24
 - FAT versus NTFS 3-36

- fault tolerance 3-28
- fault tolerance recommendations 3-31, 3-48
- performance impact variables 3-23
- RAID 0 fault tolerance 3-30
- RAID 1 fault tolerance 3-30
- RAID 4 fault tolerance 3-29
- RAID 5 fault tolerance 3-29
- striping recommendations 3-25, 3-47
- summary of testing 3-47
- disk throughput
 - how measured 2-3

F

- FAT performance comparison 3-37
- fault tolerance
 - data guarding 3-29
 - distributed data guarding 3-29
 - drive mirroring 3-30
 - RAID 1 3-48
 - RAID 5 3-48
 - recommendations 3-31, 3-48

G

- graphs
 - axis scales described 2-2
 - general description of 2-2
 - upper limit setting 2-3

H

- hardware configuration for testing, illustration of 3-10

L

- LAN
 - segmenting guidelines 3-46
 - segmenting options 3-43
- LAN subsystem
 - overview of testing 2-11
 - performance impact overview 3-39

M

- mail workload planning 4-8
- measurements
 - overview of disk throughput 2-3
 - overview of LAN throughput 2-3
- memory
 - summary of testing 3-47
 - memory allocated to NSF buffer 3-42
 - memory recommendation chart 4-11
 - memory subsystem
 - recommendation chart 3-21
 - test configuration 3-20
 - memory, computing requirements 4-11
 - Microsoft Test 2-2

N

- networking subsystem
 - performance enhancement 3-43
 - performance impact overview 3-39
- Notes Server
 - evaluating resource alternatives 4-4
 - planning requirements 4-3

NotesBench
 commands for workloads 3-8
 overview of 3-6
 test execution described 3-9
 uses of 3-7
 workloads available 3-7
 NotesMark described 3-8
 NSF buffer size
 tuning recommendations 3-41,
 3-48
 NSF buffer size parameter 3-42
 NT Disk Administrator 3-25
 NT Server
 tuning recommendations 3-41
 NT Server parameters 3-42
 NTFS performance comparison 3-37

O

optimizing performance, guidelines
 for 3-13

P

parameters, for NT Server tuning 3-
 42
 performance
 conclusions and
 recommendations 3-47
 defined 3-1
 managing, overview of 3-1
 obtaining optimum 3-13
 performance analysis
 requirements 3-2
 performance enhancement
 CPUs, using multiple 3-17
 networking subsystem 3-43
 performance management
 Notes Server considerations 4-1

Performance Monitor 2-2
 DiskPerf parameter 3-12
 sample chart 3-12
 purpose of this TechNote 1-1

R

RAID 0 3-30
 RAID 1 3-30, 3-48
 RAID 4 3-29
 RAID 5 3-29, 3-48
 recommendations
 for capacity planning 4-18
 NT Server tuning 3-42
 summary of testing 3-47
 requirements for Notes Server
 planning 4-3
 requirements of performance analysis
 3-2
 resources, additional 1-4

S

scalability
 CPUs, gains with multiple 3-17
 workload overview 3-9
 servers, computing required 4-11
 striping recommendations 3-25
 striping, disk subsystem 3-23
 subsystem tests
 CPU 3-13
 disk 3-23
 memory 3-20
 supplemental reading 1-4
 switching-hubs 3-43
 for increasing network
 bandwidth 3-44
 system under test, software used on
 3-11

T

- Technote, purpose of this 1-1
- test configuration
 - memory subsystem 3-20
 - software versions used 3-11
- test data, collection process 2-2
- test execution, as performed by
 - NotesBench 3-9
- test results
 - CPU subsystem 2-6, 2-7, 3-13
 - disk subsystem 2-8, 2-10, 3-23
 - fault tolerance 3-31
 - LAN subsystem 2-12, 2-13
 - memory 3-20
 - research conclusions 2-13
- test tools
 - capacity planning, overview of 3-3
 - customized, overview of 3-3
 - Lotus NotesBench 3-1
 - Microsoft Test 2-2
 - NetBench 3-44, 3-46
 - NotesBench described 3-4
 - Performance Monitor 2-2, 3-12
 - standard, overview of 3-2
- throughput
 - increasing 3-44
- tuning recommendations
 - disk array controller 3-42
 - NSF buffer size 3-41
 - NT Server 3-41

U

- upgrading to 100Mb/s NIC
 - advantages 3-44
 - disadvantages 3-44
- upper limits for graphs 2-3

W

- workload
 - setting scalability, overview of 3-9
 - specifying a 3-8