Microsoft Windows Server System

Exchange Server 2003 Performance and Scalability Guide

Valid Until:	October 1, 2004
Product Version:	Exchange Server 2003
Reviewed By:	Exchange Product Development
Latest Content:	www.microsoft.com/exchange/library
Author:	Pablo Ang

Exchange Server 2003 Performance and Scalability Guide

Pablo Ang

Published: May 2004
Applies to: Exchange Server 2003

Acknowledgments

Project Editor: Cathy Anderson

Contributing Writers: Robert Quimbey, Nicole Bonilla, Paul Bowden, Don Mace

Contributing Editors: Cathy Anderson, Alison Hirsch, Elizabeth Inghram (Linda Werner & Associates)

Technical Reviewers: Matt Gossage, Kevin Chase, Jeff Chow, Dave Lalor, Jeff Mealiffe, Gideon Wout

Graphic Design: Kristie Smith

Production: Sean Pohtilla

Table of Contents

Introduction

This guide provides information about Microsoft® Exchange Server 2003 performance and scalability. You will learn about the various factors that affect performance, and it provides recommendations about how to optimize your Exchange 2003 environment. Finally, the guide reviews how Exchange performs under various user loads and provides a method to determine your hardware needs.

This guide is a companion to the *Exchange Server 2003 High Availability Guide*. As you plan your Exchange 2003 deployment, review both guides to help you design and optimize your environment.

> **Note**
> Exchange 2003 is supported on both Microsoft Windows® 2000 Server with Service Pack 3 (SP3) or later and Windows Server™ 2003. Exchange 2003 is optimized on Windows Server 2003; however, this guide discusses its performance on both operating systems.

Table I.1 provides an overview of related documentation that you might find helpful to read in addition to this guide.

Table I.1 Exchange documentation related to planning an Exchange environment

Exchange 2003 guide	Material covered
Exchange Server 2003 High Availability Guide http://go.microsoft.com/fwlink/?LinkId=21277	• Understanding high availability concepts • Planning your high availability strategy
Planning an Exchange Server 2003 Messaging System http://go.microsoft.com/fwlink/?LinkId=21766	• Complete planning guide for Exchange 2003
Server Consolidation Using Exchange Server 2003 http://go.microsoft.com/fwlink/?LinkId=25209	• Strategies for server consolidation • Understanding how to deal with continued growth
Troubleshooting Exchange Server 2003 Performance http://go.microsoft.com/fwlink/?LinkId=22811	• How to isolate performance degradations and what tools to use

Who Should Read This Guide?

This guide is designed for information technology (IT) professionals who are responsible for deploying Exchange messaging systems for their organizations. Such professionals may be in the following roles:

Systems Architects
Those individuals responsible for planning and crafting overall business strategies and solutions

Enterprise Exchange Administrators
Those individuals responsible for installation, maintenance, and administration of software in the enterprise

What Will You Learn from This Guide?

This guide provides detailed answers to the following questions:

- What components affect Exchange 2003 performance?
- How do I determine my hardware requirements?
- How does Exchange 2003 perform under various user loads?
- What should I tune?
- How do I tune my Exchange 2003 server?
- How can I measure the performance of my server?
- How can I improve my client performance?

Terminology

Before reading this guide, familiarize yourself with the following terms.

dynamic buffer allocation (DBA)
A cache mechanism used by the Store.exe process to govern how much memory it uses and balance the memory use with other applications running on the server.

peak averages
The averages during peak load, such as during the morning when people log on to the network.

redundant array of independent disks (RAID)
A technology for implementing fault tolerance on a disk subsystem by using data redundancy, by using either software or a separate hardware RAID storage unit.

storage area network (SAN)
Specialized storage hardware that incorporates RAID technology to help ensure high availability and performance.

transient state
The time it takes for the memory structures to populate after you bring an Exchange server online. The duration of the transient state varies depending on the size of the server (for example, number of processors and size of memory) and the load applied to it. Transient state can range from two hours to three days.

For more information about terminology, see the *Exchange Server 2003 Glossary* (http://go.microsoft.com/fwlink/?LinkId=24625).

How Is This Guide Structured?

This guide has four chapters and five appendixes. For best results, review these chapters in order because each chapter builds upon the concepts revealed in preceding chapters:

Chapter 1, "Understanding Exchange Performance"
This chapter describes the components that affect Exchange 2003 performance.

Chapter 2, "Scaling Exchange Server 2003"
This chapter presents information on how Exchange 2003 performs under specific conditions, and it provides general hardware recommendations.

Chapter 3, "Maximizing Client Performance"
This chapter introduces new client software that is optimized for Exchange 2003.

Chapter 4, "Tuning Exchange Performance"
This chapter provides recommendations and tuning tips that you can use when you configure Exchange 2003.

Appendix A, "Exchange Performance Tools"
This appendix describes specific tools that are used to measure Exchange 2003 performance.

Appendix B, "Performance Counter Definitions"
This appendix contains the definitions of the performance counters that are used in this guide.

Appendix C, "Calculate Your Server Size"
This appendix provides a method that shows how to determine your server sizing requirements, specifically focusing on the hardware needed to support a group of users.

Appendix D, "Resources"
This appendix contains links to additional resources that will help you maximize your understanding of Exchange performance and scalability.

Appendix E, "Accessibility for People with Disabilities"
This appendix contains information about features, products, and services that make Windows 2000, Windows Server 2003, and Exchange 2003 more accessible for people with disabilities.

CHAPTER 1

Understanding Exchange Performance

Tuning a system for optimum performance is an iterative process. You must analyze, test, and adjust your system as many times as needed, and this iterative process includes Microsoft® Exchange Server 2003. You must take the time to understand all the variables that affect your system, including user profile, architecture, and hardware.

Generally, the performance of a server is determined by the component that has the lowest performance—the bottleneck in the system. The key to improving performance is being able to identify bottlenecks, determine their cause, and apply the appropriate corrective action. As you plan your Exchange Server 2003 deployment, use this guide to help design and optimize your environment. Later chapters provide metrics and tuning tips to help you reach optimum performance from your Exchange server.

The concept of performance is closely related to the concept of scalability. When you have a solid understanding of the factors influencing the performance of system components, you can deploy components in a way that scales to support periods of high demand. Later chapters in this guide cover scaling front-end and back-end servers and provide detailed metrics about how Exchange 2003 scales under different conditions.

> **Note**
> For Exchange 2000 Server users, many of the concepts are the same. Later chapters discuss the differences between the two versions. Exchange 5.5 users should review this guide along with the other guides that are recommended in the Introduction.

Measuring Performance

Several tools for measuring performance are included with Exchange 2003, including Exchange Server Stress and Performance (ESP) 2003, Jetstress, and Load Simulator 2003 (LoadSim). Microsoft Windows Server™ 2003 also includes some general performance tools including Network Monitor and System Monitor. For more information about tools, see Appendix A, "Exchange Performance Tools."

In addition to these tools, analyze your current user loads to establish a minimum server requirements baseline. Understanding how your users use the system is one of your biggest challenges. Later chapters in this guide provide a method on how to measure specific CPU, memory, and storage loads in relation to your current user loads. After you determine your hardware requirements, you should conduct a pilot test to make sure performance levels are acceptable. For more information about pilot testing, see "Laboratory Testing and Pilot Deployments" in the *Exchange Server 2003 High Availability Guide* (http://go.microsoft.com/fwlink/?LinkId=21277).

Hardware Performance

The hardware that you select for your Exchange deployment has the greatest effect on performance. Because of the large number of variables that affect performance, it is difficult to predict the effects on performance of any particular hardware component. The following sections provide general information about what components affect Exchange 2003 performance, including processors, memory, network, and storage.

Processor Performance

The processor usage on a server should maintain a load of about 60 percent during peak working hours. This percentage level allows room for periods of extreme load. If the processor usage is consistently greater than 75 percent, processor performance is considered a bottleneck.

There are several factors by which the CPU in a server affects performance. These include:

- The processor clock speed, measured in megahertz (MHz) or gigahertz (GHz).
- The number of processors.
- The type of processor.

For performance, selecting the fastest processor yields the best results. However, budget cost dictates most companies' choices.

Besides the clock speed, the technology used in a processor can affect performance. For example, some processors use Hyper-Threading Technology, which enables a single processor to act as two virtual processors. Such processors typically incorporate advanced cache management and increased bus speed features.

Exchange can fully use multiple processors and, in many cases, using servers with more processors improves performance. However, the relationship between the number of processors and performance is complex. If the server has too many processors, the overhead associated with context switching can be greater than the benefit from the additional processors. The optimum number of processors is partly determined by the role that a server plays. For example, a back-end mailbox server hosting many MAPI connections may make efficient use of an eight-processor computer. By contrast, a server used to host Microsoft Outlook® Web Access users makes better use of a four-processor computer.

For information about how different processors perform, see Chapter 2, "Scaling Exchange Server 2003."

Memory Performance

Exchange services typically consume no more than 3 gigabytes (GB) of physical memory. After you add operating system requirements and antivirus, backup, and management software, the total physical memory that is used can approach 4 GB. On servers that are dedicated to Exchange, you do not need more than 4 GB of memory.

The biggest individual consumer of memory in Exchange Server 2003 is the Store.exe process, which manages mailboxes and public information storage

Besides the Store.exe process, other processes that consume memory (and may affect performance) include:

- **Inetinfo.exe** Process that handles Internet protocols
- **Emsmta.exe** Microsoft Exchange MTA Stacks service
- **Mad.exe** Microsoft Exchange System Attendant

For more information about memory optimization, see Chapter 4, "Tuning Exchange Performance."

Network Performance

Much of the network interface subsystem is tuned automatically. Server-based network adapters are capable of detecting the type and level of traffic passing through the network interface, and they self-tune to reflect this information. Beyond making sure that you have the latest device driver on the server, there is not much to do here.

For mailbox servers, a full duplex 100 megabits per second (Mbps) network connection is typically sufficient. However, if you plan to back up and restore across the network, consider using gigabit Ethernet (1,000 Mbps or 1 gigabits per second [Gbps]).

Generally, the greatest bottleneck in a front-end and back-end server configuration is the network that separates the two sets of servers. Front-end servers can consume a 100 Mbps LAN connection. Therefore, consider multiple switched fast Ethernet networks of gigabit Ethernet connections.

Performance-related issues may be because your hardware, firmware, or software drivers are not designed to work in your configuration. For more information, see the Products Designed for Microsoft Windows Web site (http://go.microsoft.com/fwlink/?LinkId=26219).

For more information about network components, see "Network Components" in the *Exchange Server 2003 High Availability Guide (*http://go.microsoft.com/fwlink/?LinkId=21277).

Storage Performance

As storage requirements increase and companies consolidate servers, you must balance the cost, availability, and performance when you design a storage system. Take time to invest in good storage design before you implement it; unlike processors and memory, which you can scale while the network is active, storage redesign requires network downtime to implement. Tuning your Exchange storage becomes the most critical component.

There are many storage solutions available, including locally attached storage and storage area networks (SANs). The storage requirement of an Exchange server depends on the role of the server. For example, a back-end server would benefit from a SAN because of the large amount of critical data the server must store and present. SANs are specialized storage hardware that incorporates redundant array of inexpensive disks (RAID) technology to ensure high availability and performance. In contrast, a front-end server is more processor intensive, and it does not require an advanced storage solution.

With the advancements in data capacity, adding a larger capacity hard disk drive does not solve performance issues related to increased user loads. You must consider the ability of each hard disk drive to respond sufficiently to various user loads. This ability can be measured by analyzing your current user loads. Chapter 2 discusses a method you can use to analyze your current database usage. With this data, you can estimate your storage requirements.

For more information about storage strategy, see the *Exchange Server 2003 High Availability Guide* (http://go.microsoft.com/fwlink/?LinkId=21277).

General Architecture Considerations

Whether you deploy a small (single server) or large (multiple front-end and back-end servers) scale environments, you must consider challenges that affect your overall performance.

A front-end server is a server that receives requests from clients and relays them to the appropriate back-end server. A back-end server is a server that hosts at least one database to which front-end servers connect when relaying requests from clients.

Regardless of architecture, many factors affect the performance of an Exchange server. These factors include the protocols being used, number of installed processors, available memory, network traffic anticipated, use of secure authentication, and use of Secure Sockets Layer (SSL) to encrypt network traffic. You must take such factors into account before you select your hardware for a specific Exchange 2003 configuration.

For more information about front-end and back-end server performance, see Chapter 2, "Scaling Exchange Server 2003."

Troubleshooting Performance

With the troubleshooting tools listed in Appendix A, you can start to diagnose where Exchange 2003 performance has decreased. Poor server performance is frequently a result of an underperforming subsystem. For an Exchange server, performance degradation has the symptoms of increasing mail queues or poor client response.

For more information about troubleshooting performance, see *Troubleshooting Exchange Server 2003 Performance* (http://go.microsoft.com/fwlink/?LinkId=22811).

Summary

Understanding the different factors that affect Exchange performance is the first step in achieving optimum performance. You must continue to test, analyze, and adjust your system. Additionally, understanding your current user loads can help you determine what your scaling requirements are. With this information, you can better predict what your hardware requirements will be. The following chapters provide specific examples and techniques for determining hardware needs and tuning specific components.

C H A P T E R 2

Scaling Exchange Server 2003

This chapter provides information about how Microsoft® Exchange Server 2003 performs under different configurations and user loads. Using this data, you can build a highly scalable messaging system, customized to the requirements of your organization.

What You Will Learn in Chapter 2

- What is scalability?
- How do I determine what hardware I require?
- How does Exchange 2003 perform under different configurations and user loads?

What Is Scalability?

Scalability is the ability for a system to grow to meet increasing performance demands. When applied to clustering, scalability is the ability to incrementally add systems to an existing cluster when the overall load of the cluster exceeds the cluster's capabilities—either by scaling up or scaling out. Scaling up involves increasing system resources (such as processors, memory, disks, and network adapters) to your existing hardware or replacing existing hardware with greater system resources (for example, faster CPU and network adapter, more memory, and more storage). Scaling out involves adding servers to meet demand. For more information about scalability strategies, see the *Exchange Server 2003 High Availability Guide* (http://go.microsoft.com/fwlink/?LinkId=21277).

This chapter discusses how Exchange 2003 scales under different configurations and user loads.

Front-End and Back-End Server Architectures

Implementing a front-end and back-end server environment presents different challenges that affect your overall performance.

Front-end servers, such as those that serve Microsoft Outlook® Web Access, Outlook Mobile Access, Exchange ActiveSync®, RPC over HTTP, authentication, IP address checking, Secure Sockets Layer (SSL) protocol, and encryption schemes, have security features that require significant processing. For these servers, you are likely to see an increase in processor activity, both in privileged and user mode, and an increase in the rate of context switches and interrupts. If the processors in the server cannot handle this increased load, queues are likely to develop.

Factors that affect front-end servers include:

- The protocols being used.
- The number of processors installed.
- The available memory.

- The network traffic.
- The authentication methods.
- The use of SSL to encrypt network traffic.

Because front-end servers forward all requests to the back-end servers, back-end servers have the same processor and processing issues that front-end servers have. Back-end servers may also experience storage issues because of the read and write activity when retrieving and storing data. With public folders, replication traffic between public folders (if there is more than one public folder in the topology) can affect all the servers involved.

Factors that affect back-end servers include:

- The protocols being used.
- The number of processors installed.
- The available memory
- The type of storage used.
- The storage available.
- The replication of public folder information.

For more information about how these factors affect front-end and back-end servers, see "Baseline Data" later in this chapter.

Front-End Server Licensing Guide

Microsoft Windows® 2000 Server and Windows Server™ 2003 support two types of licensing models: per seat and per server. The Windows 2000 License Logging Service maintains a list (on disk and in memory) of all users who authenticate against a server using the per-seat licensing model. The per-server model does not keep a list of users. The list of authenticated users in the per-seat configuration does not consume too much memory on servers with fewer than 50,000 users. However, the memory footprint of the License Logging Service can grow too large when it runs with a per-seat licensing model in a front-end and back-end topology, with hundreds of thousands of users.

In the per-seat scenario, with the front-end servers load balancing the client requests, the front-end server's License Logging Service builds a list of all users in the site (including all users on all back-end servers). Depending on the size of the site, the server can consume hundreds of megabytes (MB) of memory in the License Logging Service on the front-end server. Therefore, it is recommended that the per-server licensing model be used in large front-end and back-end topologies. For more information about the License Logging Service, see the Windows 2000 Resource Kits Web site (http://go.microsoft.com/fwlink/?LinkId=6545).

Determining Your Server Sizing

Because of the wide variety of Exchange configurations and user profiles, it is difficult to accurately determine the number of users supported by a server. You must consider the different types of clients, how active the users are, the capacity of the storage subsystem, and how the Exchange server is configured to use the disk resources. For more information about how to calculate your server sizing requirements, see Appendix C, "Calculate Your Server Size."

Baseline Data

This section describes how Exchange 2003 performs under different configurations and user loads. This information can help you establish a baseline when determining what minimum hardware you require. The following scenarios include:

- Mailbox Server (using MAPI)
- Outlook Web Access
- Post Office Protocol version 3 (POP3)
- Internet Message Access Protocol version 4rev1 (IMAP4)
- Simple Mail Transfer Protocol (SMTP)

Each scenario analyzes the following areas:

- Processor
- Memory
- Disk usage
- Network usage

Mailbox Server (Using MAPI)

This section provides baseline performance data on Exchange 2003 mailbox server under specific MAPI client loads. The key hardware criteria for a mailbox server are processor type and speed, memory size, network speed, and disk performance and configuration.

> **Note**
> Exchange 2000 scalability data is used as a basis for comparison. Load Simulator 2000 and Load Simulator 2003 are used to simulate Outlook 2000 and Outlook 2003 MAPI clients respectively. For more information about Load Simulator (LoadSim), see Appendix A, "Exchange Performance Tools."

Hardware

Table 2.1 shows the hardware specifications used in the following scenarios:

- Scenario 1: Exchange 2000 and Outlook 2000 Online versus Exchange 2003 and Outlook 2003 Online
- Scenario 2: Exchange 2000 and Outlook 2003 Cached Exchange Mode versus Exchange 2003 and Outlook 2003 Cached Exchange Mode

Table 2.1 Mailbox hardware configuration

Server type	Processor type	RAM	Storage
Back-end server	Intel P4 Xeon 4 processors, 1.4 GHz (Hyper-Threading disabled)	4 GB	• SAN with RAID0+1 for database volumes • 2 spindles of RAID1 for each transaction logs for each storage group • 8,000 users were spread evenly across 12 databases, in three storage groups with an average user mailbox size of 25 MB

Scenario 1: Exchange 2000 and Outlook 2000 Online vs. Exchange 2003 and Outlook 2003 Online

This scenario compares the load characteristics of Exchange 2000 and Outlook 2000 with Exchange 2003 and Outlook 2003. Load Simulator 2003 is used to generate the load. The following LoadSim configuration was used:

- **LoadSim Client Configuration**
  - Intel P4 600 megahertz (MHz)
  - 512 MB RAM
  - Single Integrated Drive Electronics (IDE) disk
  - 1,000 LoadSim clients per computer
- **LoadSim Configuration**   Tables 2.2, 2.3, and 2.4 show the values that were used.

Table 2.2   Mailbox configuration settings

Mailbox configuration	Value
Messages in Inbox	100
Message in Deleted Items	1
Number of new folders	10
Messages per new folder	10
Calendar Appointments	25
Number of Contacts	128
Average Message Size	75 KB

Table 2.3 User action settings

User actions per 8-hour day	Value
Send Mail	12
Messages Received/day	208
Messages Sent/day	52
Process Inbox	12
Browse Mail	20
Check Free/Busy	1
Request Meeting	1.4
Make Appointment	2.8
Browse Calendar	6
Journal Applications	0
Logoff	0
Browse Contacts	0
Create Contact	0

Table 2.4 Distribution list settings

Distribution lists	Value
Distribution lists per site	100
Minimum distribution list size	2
Average distribution list size	10
Maximum distribution list size	20

Table 2.5 shows the results produced by the mailbox server in this scenario.

Table 2.5 Mailbox server performance comparison

Server platform	Windows 2000 SP3, Exchange 2000 SP3	Windows 2000 SP3, Exchange 2003	Windows Server 2003, Exchange 2003	Windows Server 2003, Exchange 2003
Client platform	**Windows XP, Outlook 2000 Online**	**Windows XP, Outlook 2003 Online**	**Windows XP, Outlook 2003 Online**	**Windows XP, Outlook 2003 Online**
Hyper-Threading enabled?	No	No	No	Yes
Users	8,000	8,000	8,000	8,000
% Processor Time	75%	76%	75%	57%
Context Switches/sec	8,300	7,833	10,183	13,487
Local Delivery Rate	30	30	30	30

Server platform	Windows 2000 SP3, Exchange 2000 SP3	Windows 2000 SP3, Exchange 2003	Windows Server 2003, Exchange 2003	Windows Server 2003, Exchange 2003
Client platform	**Windows XP, Outlook 2000 Online**	**Windows XP, Outlook 2003 Online**	**Windows XP, Outlook 2003 Online**	**Windows XP, Outlook 2003 Online**
Network Usage (in Kbps)	3,154	1,594	1,604	1,613
DB Disk Transfers/sec	2,088	1,980	1,929	1,955
Log Writes/sec	377	314	353	398
Disk Bytes/sec (MB)	19.9	20.4	18.9	17.9
Database Cache Size	864	896	896	896
RPC Operations/sec	950	955	940	925
RPC Requests	7	8	8	7
Client Latency (ms)	96	100	95	105
Store Virtual Bytes	2,047	2,075	2,080	2,082

For more information about the performance counters used in this scenario, see Appendix B, "Performance Counter Definitions."

Processor

Exchange 2000 SP3 and Exchange 2003 show comparable levels of processor performance. The 8,000-user Exchange 2000 SP3 test reaches 75 percent processor usage, and the Exchange Server 2003 test reaches 76 percent processor usage. Both tests show similar workloads (950 **RPC Operations/sec** to 955, with the same **Local Delivery Rate**).

Table 2.5 shows that Windows 2000 Enterprise Server with SP3 and Windows Server 2003 perform comparably in terms of processor usage with Exchange 2003. However, Windows Server 2003 has key memory manager optimizations that significantly reduce virtual memory fragmentation, which is not shown in Table 2.5.

Exchange 2003 takes advantage of Intel Hyper-Threading Technology to increase server scalability by 25 percent. Table 2.5 shows that processor use is reduced by 25 percent (from 76 percent to 57 percent) when the same test is run with Hyper-Threading enabled. The benefit of Hyper-Threading is the same for Cached Exchange Mode client scenarios.

Memory

The Exchange store memory footprint of Exchange 2003 is roughly the same as that of Exchange 2000. With the increase in **Database Cache Size** (896 MB compared to 864 MB) in Exchange 2003 and performance improvements in Jet, random Exchange store disk I/O operations can be reduced by up to 10 percent.

Disk Usage

Exchange 2003 produces 10 percent fewer Exchange Database disk I/O operations than Exchange 2000 SP3. Because of the increase in **Database Cache Size** (896 MB compared to 864 MB) in Exchange 2003, random Exchange database disk I/O operations can be reduced by up to 10 percent.

Network Usage

Exchange 2003 and Outlook 2003 can reduce network bandwidth when compared to Exchange 2000 and Outlook 2000 by up to 50 percent. This 8,000-user test shows that Exchange 2000 and Outlook 2000 use 3,154 Kbps, and Exchange 2003 and Outlook 2003 use approximately half the bandwidth at 1,594 Kbps. Because of the compression technology implemented in Exchange 2003 and Outlook 2003, the network bandwidth required to host a particular number of users is significantly less than that of Exchange 2000 and Outlook 2000. Many performance optimizations were implemented in Exchange 2003 so this compression would not detract from the scalability of the server. In effect, the reduction of bytes over the wire is free in terms of server CPU usage.

Scenario 2: Exchange 2000 and Outlook 2003 Cached Exchange Mode vs. Exchange 2003 and Outlook 2003 Cached Exchange Mode

This test compares the load characteristics of Exchange 2000 and Outlook 2003 with Exchange 2003 and Outlook 2003 in a Cached Exchange Mode client scenario. LoadSim 2003 is used to generate the load.

- **LoadSim Client Configuration**
  - Intel P4 600 MHz
  - 512 MB RAM
  - Single IDE disk
  - 1,000 LoadSim clients per computer
- **LoadSim Configuration** Tables 2.6, 2.7, and 2.8 show the values that were used.

Table 2.6 Mailbox configuration settings

Mailbox configuration	Value
Messages in Inbox	100
Message in Deleted Items	1
Number of new folders	10
Messages per new folder	10
Calendar Appointments	25
Number of Contacts	128
Average Message Size	75 KB

Table 2.7 User action settings

User actions per eight-hour day	Value
Send Mail	7
Messages Received/day	161.90
Messages Sent/day	44.12
Process Inbox	20

User actions per eight-hour day	Value
Browse Mail	0
Check Free/Busy	4
Request Meeting	2
Make Appointment	4
Browse Calendar	1
Journal Applications	0
Logoff	3
Browse Contacts	0
Create Contact	0

Table 2.8 User action settings

Distribution lists	Value
Distribution lists per site	100
Minimum distribution lists size	2
Average distribution lists size	10
Maximum distribution lists size	20

Table 2.9 shows the results produced by the mailbox server in this scenario.

Table 2.9 Exchange 2000/Exchange 2003 Cached Mode comparison

Server platform	Windows 2000 SP3, Exchange 2000 SP3	Windows 2000 SP3, Exchange 2003
Client platform	**Windows XP, Outlook 2003 Cached Exchange Mode**	**Windows XP, Outlook 2003 Cached Exchange Mode**
Hyper-Threading enabled?	No	No
Users	8,000	8,000
% Processor Time	64%	57%
Context Switches/sec	7,981	6,488
Local Delivery Rate	23	23
Network Usage (in Kbps)	3,613	1,765
DB Disk Transfers/sec	1,496	1,355
Log Writes/sec	304	315
Disk Bytes/sec (MB)	17.7	16.2
Database Cache Size	896	864
RPC Operations/sec	1,205	1,193

Server platform	**Windows 2000 SP3, Exchange 2000 SP3**	**Windows 2000 SP3, Exchange 2003**
Client platform	**Windows XP, Outlook 2003 Cached Exchange Mode**	**Windows XP, Outlook 2003 Cached Exchange Mode**
RPC Requests	8	7
Store Virtual Bytes	2,016	2,076

For more information about the performance counters used in this scenario, see Appendix B, "Performance Counter Definitions."

> **Note**
> Results of the Outlook 2003 online tests and the Outlook 2003 Cached Exchange Mode tests are not comparable. The user profile and user actions are quite different. For specific performance data about Outlook 2003, see *Client Network Traffic with Exchange 2003* (http://go.microsoft.com/fwlink/?LinkId=27020).

Processor

In this test, processor performance of Cached Exchange Mode MAPI clients on Exchange 2003 is approximately 10 percent better than Exchange 2000 SP3. The 8,000-user Exchange 2000 SP3 test reaches 64 percent processor usage, and the Exchange 2003 test reaches 57 percent processor usage. Both tests show similar workloads (1,193 **RPC Operations/sec** compared to 1,205, with the same **Local Delivery Rate**). Exchange 2003 Server is optimized for Outlook 2003 Cached Exchange Mode clients.

Memory

The Exchange store memory footprint of Exchange 2003 is roughly the same as that of Exchange 2000. With the increase in **Database Cache Size** (896 MB compared to 864 MB) on Exchange 2003 and performance improvements in Jet, the random Exchange database disk I/O operations can be reduced by up to 10 percent.

Disk Usage

Exchange 2003 produces 10 percent fewer Exchange database disk I/O operations than Exchange 2000 SP3. Because of the increase in the size of the database cache (896 MB compared to 864 MB) in Exchange 2003, the random Exchange database disk I/O operations can be reduced by up to 10 percent.

Network Usage

Exchange 2003 and Outlook 2003 can reduce the network bandwidth when compared to Exchange 2000 and Outlook 2003 by up to 50 percent. This 8,000-user test shows Exchange 2000 and Outlook 2003 uses 3,613 Kbps, and the Exchange 2003 and Outlook 2003 test uses approximately half the bandwidth at 1,765 Kbps. Because of the compression technology implemented in Exchange 2003 and Outlook 2003, the network bandwidth required to host a particular number of users is significantly less than that of Exchange 2000 and Outlook 2003. Many performance optimizations were implemented in Exchange 2003 so this compression would not detract from the scalability of the server. In effect, the reduction of bytes over the wire is free in terms of server CPU usage.

Mailbox Scalability Guidelines

When you design a mailbox server, consider the following recommendations.

- A back-end mailbox server hosting many MAPI connections scales well on 4-processor servers.
- Generally, one processor for every 1,000 users is a good guideline for all but the heaviest of mail users. For more information on processor requirements, see Appendix C, "Calculate Your Server Size."
- Exchange uses a maximum of 3 GB of memory. To increase performance, increase memory up to 4 GB to reduce the paging to disk.

Outlook Web Access

This section provides baseline performance data on Outlook Web Access. Outlook Web Access is a Web interface that uses HTTP to access an Exchange server. With any Web browser, you can access most of the feature that are available to a client using Outlook. This section describes the following scenarios:

- **Scenario 1** Compares 10,000 Outlook Web Access users on the front-end server with similar mail flow on two different configurations: Exchange 2003 and Windows 2000, and Exchange 2003 and Windows Server 2003.
- **Scenario 2** Compares the performance when additional features are enabled while keeping the mail flow constant.
- **Scenario 3** Compares the performance when additional features are enabled while under various loads, using LoadSim.

Scenario 1

This scenario compares 10,000 Outlook Web Access users on the front-end server with similar mail flow on two different configurations: Exchange 2003 and Windows 2000, and Exchange 2003 and Windows Server 2003. The following configuration is used in this scenario:

- There were four storage groups with three private mailbox stores in each storage group.
- Clients send messages with an average size of 20 KB.
- Each user's Inbox is populated with 31 IMAP4 messages before the test begins.
- Transport traffic occurs as each user sends several e-mail messages to the Internet.
- Each user connection lasts approximately 10 minutes.

Table 2.10 shows the actions performed by each user.

Table 2.10 User test script

Action	Times performed
Log on	1
Check mail	2
Send messages	2
Recipients of message sent	1
Receive messages	4
Read messages	4

Action	Times performed
Move messages	1
Delete messages	1

Hardware

Table 2.11 shows the specifications of the three servers used in this scenario.

Table 2.11 Outlook Web Access scenario 1 hardware

Server type	Processor type	RAM	Storage
Front-end server	Intel P4 Xeon 2 processors, 2.6 GHz (Hyper-Threading)	1 GB	Not applicable
Back-end server 1	Intel P4 Xeon 4 processors, 1.4 GHz	4 GB	• 20 spindles of RAID0+1 for database volumes • 4 spindles of RAID0+1 for the transaction logs for each storage group
Back-end server 2	AMD Opteron 4 processors, 1.6 GHz	2 GB	• 20 spindles of RAID0+1 for database volumes • 4 spindles of RAID0+1 for the transaction logs for each storage group

Outlook Web Access Front-End Server

Client access to Outlook Web Access is provided through HTTP. Table 2.12 gives an overview of processor usage, context switches each second, network traffic, and memory usage on the front-end server with both Windows 2000 and Windows Server 2003.

Table 2.12 Outlook Web Access front-end comparison

Front-end server	Exchange 2003 Windows 2000	Exchange 2003 Windows Server 2003
Network Usage (in Kbps)	4,679	6,313
Inetinfo Private Bytes	518 MB	38 MB
W3WP Private Bytes	Not applicable	79 MB
Available Mbytes	276 MB	484 MB
% Processor Time	52%	21%
Context Switches/sec	11,795	13,791
Web Bytes Total/sec	2,720 KB	3,706 KB
Web ISAPI Extension Requests/sec	98	135

For more information about the performance counters used in this scenario, see Appendix B, "Performance Counter Definitions."

Processor

In this scenario, Outlook Web Access uses significantly less CPU on the front-end server when running on Windows Server 2003. Context switching is stable and fairly low for 10,000 users.

Memory

The ASP.NET process (W3WP) in Windows Server 2003 is more memory-efficient than Internet Information Services (IIS) in Windows 2000 at servicing Outlook Web Access requests. In this scenario, Exchange 2003 running on Windows 2000 Server consumes 756 MB of RAM, and the server running Windows Server 2003 consumes only 540 MB. The core Inetinfo process in Windows Server 2003 pushes most of the work onto the ASP.NET process. The store process on both front-end servers consumed a negligible amount of RAM.

Disk Usage

For information about disk usage for a dedicated Outlook Web Access front-end server, see "Disk Usage" in "POP3 Front-End Server" later in this chapter.

Network Usage

For more information about network usage for a dedicated Outlook Web Access front-end server, see "Network Usage" in "POP3 Front-End Server" later in this chapter.

Outlook Web Access Back-End Server

Table 2.13 shows the results produced by the Outlook Web Access back-end server.

Table 2.13 Outlook Web Access back-end server comparison

Back-end server	Exchange 2003 Windows 2000	Exchange 2003 Windows Server 2003
Database Cache Size	896 MB	896 MB
Available Mbytes	409 MB	93 MB
Local Delivery Rate	7	6
Network Usage (in Kbps)	4,705	6,313
Disk Bytes/Sec	14,739 KB	15,365 KB
DB Disk Transfers/sec	1,959	1,723
Inetinfo Private Bytes	83 MB	124 MB
Store Virtual Bytes	1,788 MB	1,764 MB
% Processor Time	27%	38%
Context Switches/sec	15,769	14,363
Web ISAPI Extension Requests/sec	59	81

For more information about the performance counters used in this scenario, see Appendix B, "Performance Counter Definitions."

Processor

Although CPU usage is reduced on the front-end server under Windows Server 2003, the back-end server shows increased CPU usage. Higher throughput on the Windows Server 2003 test consumes more CPU but the cost per operation is similar in both tests.

Memory

The four-processor Outlook Web Access back-end servers require at least 500 MB of RAM. During most of the four-processor scenarios described in this chapter, the Store.exe process consumed more than 1 GB of RAM. Exchange 2003 uses a maximum of 3 GB of memory. To increase performance, increase memory up to 3 GB to reduce the paging to disk.

Disk Usage

When using a back-end server dedicated to Outlook Web Access clients, put each storage group's database files on a dedicated RAID0+1 array and use cache-backed controllers.

Network Usage

A 100 megabits per second (Mbps), full duplex, network connection is sufficient for all Outlook Web Access mailbox server applications. The Outlook Web Access tests conducted to support this guide do not exceed 7 Mbps network usage even under the most severe of test conditions. This level is well below the network saturation point of a 100-Mbps, full duplex network connection.

Scenario 2

This scenario compares the performance of additional Outlook Web Access features to the baseline Outlook Web Access test (Scenario 1). In this scenario, the following features are enabled:

- **Spelling check** The Outlook Web Access baseline test is modified to check the spelling on every send, reply, and forward action to simulate a user using the spelling checker feature on sent messages.
- **SSL** The Outlook Web Access baseline test is modified to use SSL encryption providing the security and an example of an Outlook Web Access SSL baseline test.
- **Gzip compression** The Outlook Web Access SSL baseline test is modified to use gzip compression, reducing the size of the files sent over the wire.
- **S/MIME** The Outlook Web Access SSL baseline test is modified to use Secure/Multipurpose Internet Mail Extensions (S/MIME) to encode every message sent, replied to, or forwarded.

Additionally, the following configuration is used in this scenario:

- Each storage group contains four storage groups with three private mailbox stores.
- 2,000 Outlook Web Access users are used with 3.5 messages per second sent through SMTP inbound from the Internet.
- Clients sent messages with an average size of 20 KB.
- Each user's Inbox is populated with 31 IMAP4 messages before the test begins.
- Transport traffic occurs as each user sends several e-mail messages to the Internet.
- Each user connection lasts approximately 10 minutes.

Table 2.10 shows the actions performed by each user. (These are the same actions described in Scenario 1.)

Hardware

This scenario uses the same hardware configuration used in Outlook Web Access Scenario 1. For information about specific hardware, see Table 2.11.

Outlook Web Access Features - Front-End Server

Based on the user action listed in Table 2.10, the following data was observed.

Table 2.14 Outlook Web Access front-end feature comparison

Front-end server	Outlook Web Access baseline	Spelling checker	SSL	SSL gzip	SSL S/MIME
Inetinfo Private Bytes	36 MB	36 MB	29 MB	29 MB	29 MB
Lsass Private Bytes	50 MB	58 MB	211 MB	212 MB	222 MB
Available MB	517	484	335	338	289
% Processor Time	12%	23%	27%	29%	33%
Context Switches/sec	3,303	3,685	3,462	3,461	3,794
Web Bytes Total/Sec	332 KB	494 KB	377 KB	384 KB	508 KB
Web ISAPI Extension Requests/sec	34	46	34	35	45

For more information about the performance counters used in this scenario, see Appendix B, "Performance Counter Definitions."

Processor

When the spelling checker is enabled, processor usage from the Outlook Web Access baseline almost doubles. The spelling checker is CPU intensive on the front-end server, and it can be manually or automatically enabled on every sent message. SSL uses 125 percent more CPU than the Outlook Web Access baseline. Enabling gzip compression increases CPU usage by 7 percent over the SSL baseline. Enabling S/MIME was the most CPU intensive, increasing CPU usage 22 percent over the SSL baseline. Context switches are low and similar among all the tests.

Memory

All tests show similar memory consumption. The SSL enabled tests show a decrease in available megabytes because the Lsass process consumes more memory. The Lsass process handles the credential and encryption of an SSL connection.

Disk Usage

An Outlook Web Access front-end server rarely uses its hard disk. For information about disk usage for a dedicated Outlook Web Access front-end server, see "Disk Usage" in "POP3 Front-End Server" later in this chapter.

Network Usage

All tests show similar network consumption. The spelling checker consumes 5 percent more network consumption on the front-end server because of the additional action of having to send the message to the front-end server to check spelling before finally sending the message for delivery. S/MIME is the notable exception with 27 percent more network consumption on the front-end server.

Outlook Web Access Features - Back-End Server

Most of the features tested affect the front-end server and not the back-end server.

Table 2.15 Outlook Web Access back-end feature comparison

Back-end mailbox server	Outlook Web Access	Spelling checker	SSL	SSL GZIP	SSL S/MIME
Local Delivery Rate	24	22	23	23	8
DB Disk Transfers/sec	935	745	1,021	1,031	711
Network Usage (in Kbps)	687	723	675	695	876
Inetinfo Private Bytes	120 MB	107 MB	72 MB	75 MB	112 MB
% Processor Time	40%	41%	42%	40%	37%
Context Switches/sec	11,842	12,486	11,878	11,997	10,931
Web Bytes Total/sec	271 KB	312 KB	275 KB	291 KB	400 KB
Web ISAPI Extension Requests/sec	34	41	34	35	45

For more information about the performance counters used in this scenario, see Appendix B, "Performance Counter Definitions."

Processor

Processor usage and context switching is similar among the tests with approximately 40 percent CPU usage and 12,000 context switches.

Memory

Memory consumption is similar among all the tests. The SSL enabled tests use less memory in the Inetinfo process.

Disk Usage

An Outlook Web Access back-end server is very disk intensive. Use RAID0+1 for the database volumes and at least RAID1 for each transaction log in a storage group.

Network Usage

Network consumption is similar among all the tests and well within the requirements for a 100 Mbps network adapter.

Scenario 3

This scenario compares the performance when additional features are enabled under various MAPI user loads. In this scenario, 8,000 MAPI users are used with 3.5 messages a second sent through SMTP inbound from the Internet. The following configurations are tested:

- **MAPI baseline** In this baseline test, 8,000 MAPI users directly access the back-end server and 3.5 messages are sent to the back-end server each second through SMTP inbound from the Internet.
- **MAPI, Exchange ActiveSync, and AUTD** In this baseline test, 8,000 MAPI users directly access the back-end server and 3.5 messages are sent to the back-end server each second through SMTP inbound from the Internet. Two thousand, four hundred out of the 8,000 users use Exchange ActiveSync to synchronize their mobile devices with their mailboxes through the front-end server. Half of the Exchange ActiveSync users (1,200) have Always Up-To-Date (AUTD) set up for notifications. Each Exchange ActiveSync connection lasts approximately 3 minutes. Table 2.16 shows the actions each user performs.
- **RPC over HTTP** This baseline test is routed through a front-end RPC proxy server.
- **MAPI and Outlook Web Access** In this baseline test, 8,000 MAPI users directly access the back-end server and 3.5 messages are sent to the back-end server each second through SMTP inbound from the Internet. Five hundred Outlook Web Access users access their mailboxes through the front-end server. Each user connection lasted approximately 10 minutes, and each user performed the same actions described in Table 2.10.

The average client message size is 20 KB. Each user's Inbox is populated with approximately 31 IMAP4 messages before the test begins. Transport traffic occurs as each user sends several e-mail messages to the Internet.

Table 2.16 Exchange ActiveSync user test script

Action	Times performed
GETHIERARCHY	2
sync calendar	5
sync contacts	5
sync inbox	2
getestimate calendar	1
getestimate contacts	1
getestimate inbox	1
addmailrecipient	0.6
sendmail	1
syncadd contacts	1
syncchange FOLDER("contacts").rnd	2
syncdel FOLDER("contacts").rnd	2
syncadd calendar	1

Hardware

Table 2.17 shows the specifications of the four servers used in this scenario.

Table 2.17 Outlook Web Access scenario 3 hardware configuration

Server type	Processor type	RAM	Storage
Front-end server 1 (MAPI, AUTD, Exchange ActiveSync)	Intel P4 Xeon 2 processors, 2.6 GHz	1 GB	Not applicable
Front-end server 2 (MAPI, RPC over HTTP)	AMD Athlon 2 processor 1.2 GHz	2 GB	Not applicable
Front-end server 3 (MAPI, Outlook Web Access)	Intel P2 Xeon 2 processors, 450 MHz	1 GB	Not applicable
Back-end server	Intel P4 Xeon MP 4 processors, 1.4 GHz (Hyper-Threading)	4 GB	• 40 spindles of RAID0+1 for database volumes • 2 spindles of RAID0+1 for the transaction logs for each storage group

Outlook Web Access Features Using MAPI - Front-End Server

Table 2.18 shows the performance of the three front-end servers under similar MAPI loads.

Table 2.18 Outlook Web Access features using MAPI - front-end

	Front-end server 1 (MAPI, Exchange ActiveSync, AUTD)	Front-end server 2 (RPC over HTTP)	Front-end server 3 (MAPI, Outlook Web Access)
% Processor Time	3%	23%	2%
Context Switches/sec	610	10,447	776
Web Bytes Total/sec	6 KB	253 KB	123 KB
Web ISAPI Extension Requests/sec	4	61	3

For more information about the performance counters used in this scenario, see Appendix B, "Performance Counter Definitions."

When you use a front-end server for RPC over HTTP proxy server, it consumes 23 percent of the two-processor 1.2-GHz server for 8,000 users with 10,000 context switches. The Outlook Web Access and ActiveSync stress is minimal on the front-end server. Memory and disk consumption is minimal and the network usage is 2 MB, well within the recommended 100-Mbps network adapter.

Outlook Web Access Features Using MAPI - Back-End Server

Table 2.19 shows the performance of the back-end server when specific Outlook Web Access features are enabled.

Table 2.19 Outlook Web Access features using MAPI - back-end

	MAPI baseline	RPC over HTTP	MAPI and Outlook Web Access	MAPI, ActiveSync, and AUTD
Database Cache Size	896	896	896	896
Log Writes/sec	265	270	229	220
DB Disk Transfers/sec	1,609	1,494	1,767	1,530
Disk Bytes/sec	18.7 MB	16.9 MB	20.5 MB	18.6 MB
Local Delivery Rate	26	25	26	22
RPC Operations/sec	1,208	1,136	1,226	1,188
RPC Requests	9	6	22	7
Network Usage (in Kbps)	2,452	2,031	2,237	2,008
Inetinfo Private Bytes	190	51	187	45
Store Virtual Bytes	2,127	2,063	2,104	2,124
% Processor Time	50	44	73	91
Context Switches/sec	13,875	13,292	18,824	12,365
Web Bytes Total/Sec	Not applicable	Not applicable	109 KB	108 KB
SMTP Messages Del/Sec	7.6	7.6	8	7
Web ISAPI Extension Requests/sec	Not applicable	Not applicable	3	17

For more information about the performance counters used in this scenario, see Appendix B, "Performance Counter Definitions."

Based on this test, the following conclusions were drawn:

- **MAPI, Exchange ActiveSync, and AUTD** versus **MAPI**
  - Four AUTD and Exchange ActiveSync notifications are sent per second.
  - On the back-end server, Exchange ActiveSync and AUTD notifications push the CPU usage over 90 percent to service the Exchange ActiveSync feature. The load on the front-end server is negligible.
- **RPC over HTTP** versus **MAPI**
  - When MAPI traffic is routed through a front-end proxy server, 8,000 MAPI users consume 23 percent of the front-end processor. The back-end server load is similar to the MAPI baseline test.
- **MAPI and Outlook Web Access** versus **MAPI**
  - With both Outlook Web Access and MAPI running, server CPU usage increases 23 percent in the back-end server. Context switching increases at a similar rate.

Outlook Web Access Scalability Guidelines

When you design an Outlook Web Access server, consider the following guidelines:

- Outlook Web Access scales well on four-processor servers.
- Use a ratio of one front-end server to four back-end servers.
- Outlook Web Access requires 30 KB of RAM per active connection.
- Outlook Web Access uses virtually no disk resources, unless the front-end server is paging or HTTP-DAV protocol logging is turned on.
- Outlook Web Access requires a second 100-Mbps network adapter if it is servicing more than 5,000 connections.
- Use Network Load Balancing to load balance an Outlook Web Access front-end server.
- Outlook Web Access SSL connections require up to three times more processing power and 60 percent more memory on their front-end server.

POP3

This section provides baseline performance data on Post Office Protocol version 3 (POP3) servers. POP3 is an Internet protocol that allows a POP3 client to download e-mail from a server. This protocol works well for computers that cannot maintain a continuous connection to a server.

POP3 Front-End Server

This scenario measures how well a dedicated POP3 front-end server scales under different amounts of client requests. The front-end server uses Pentium 4 Xeon processors with Hyper-Threading.

Hardware

Table 2.20 shows the specifications of the five servers used in this scenario.

Table 2.20 POP3 front-end server hardware configuration

Server type	Processor type	RAM	Storage
Front-end server	Intel P4 Xeon 2 processors, 2.6 GHz (Hyper-Threading)	1 GB	Not applicable
Four back-end servers	Intel P4 Xeon 8 processors, 550 MHz	4 GB	• 12 spindles of RAID0+1 for database volumes • 2 spindles of RAID0+1 for the transaction logs for each storage group

Scenario

An average message size of 26 KB is sent. Each user's Inbox is populated with 31 IMAP4 messages before the test begins. Transport traffic occurs when the POP3 front-end server receives incoming e-mail from the Internet during the session targeted to one recipient. Activity on the POP3 front-end server includes users logging on to the server, retrieving all mail, and then deleting all mail from the server. The number of SMTP messages being delivered to mailboxes each second equals the number of messages retrieved and deleted through POP3 each second. The POP3 commands associated with message retrieval and deletion are RETR and DELE. The POP3 command issued by a POP3 client to determine the number of messages in a mailbox is STAT. Statistical counters for RETR, DELE, and STAT can be used to determine the number of POP3 transactions occurring on the front-end server each second. These counters can be accessed through System Monitor, which is included with the Microsoft Windows Server 2003 operating system.

In this scenario, the number of STAT commands received by the POP3 server is almost double the number of RETR and DELE commands. The Exchange Stress and Performance (ESP) tool was used to generate this server load. For more information about this tool, see Appendix A, "Exchange Performance Tools."

Table 2.21 shows the variance in hardware usage when using a range of back-end servers with the particular profile. Note that context switching on the front-end server does not grow significantly as processor usage increases.

Table 2.21 POP3 front-end performance

	Back-end server 1	Back-end server 2	Back-end server 3	Back-end server 4
Front-End Processor %	13.0	27.9	55.1	81.0
Context Switches/sec	11,423	20,495	24,872	23,217
POP3 DELE/sec	48	94	102	174
POP3 STAT/sec	124	215	356	380
Network Usage (in Kbps)	4,245	6,580	8,227	12,846
Inetinfo Working Set	127 MB	172 MB	172 MB	172 MB

For more information about the performance counters used in this scenario, see Appendix B, "Performance Counter Definitions."

Processor

POP3 scales well on a two-processor server, based on a low number of context switches while under heavy processor load. Figure 2.1 shows how context switching begins to level off on a two-processor front-end server as processor usage increases.

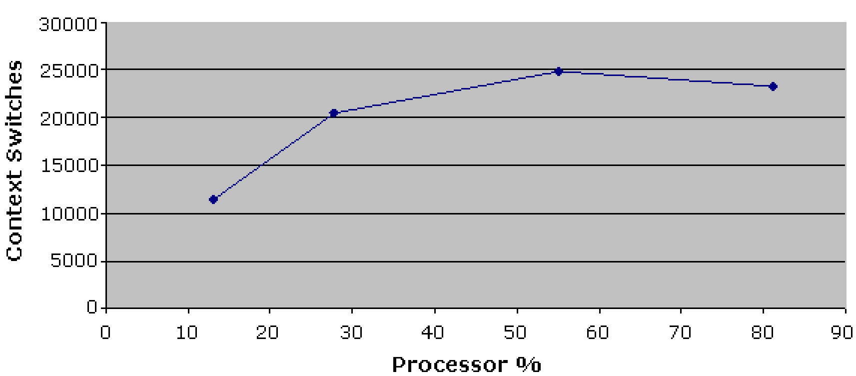

Figure 2.1 POP3 front-end server performance

Memory

POP3 front-end servers require almost no memory to operate efficiently. As the number of simultaneous POP3 sessions on the POP3 front-end server increases, memory usage does not increase significantly. Memory does not increase significantly because POP3 clients do not maintain long connections to the front-end server, which makes the use of memory relatively small. The MSExchangeIS (Store.exe) service can be disabled on POP3 front-end servers, which yields additional memory savings. If this service is disabled, a POP3 front-end server can run efficiently with 256 MB of RAM.

Disk Usage

When determining your hardware requirements for a dedicated POP3 front-end server, consider the disk space you require. A POP3 front-end server rarely uses its hard disk, because it acts as a proxy server, passing each protocol session to the appropriate back-end server. If protocol logging is enabled in Exchange System Manager for a POP3 virtual server, the hard disk is used on the front-end server to store the requested protocol log. Cache Manager in Windows Server 2003 also uses the disk to page information to and from the paging file. The cache manager uses the paging file to temporarily store information from RAM that has not been accessed recently when additional memory is required by active system processes. You can minimize paging activity by increasing the RAM on the server.

A POP3 front-end server with 256 MB or more of physical memory rarely pages. One disk spindle for a POP3 front-end server is sufficient for most applications. If you run large servers with protocol logging enabled, consider adding a second spindle.

Network Usage

On POP3 front-end servers, you must consider network traffic when you try to determine the type of hardware you require. Because a POP3 front-end server can service multiple back-end servers, the network traffic that occurs on a front-end server is frequently very high. The minimum network requirement for any high-end front-end server is a single 100-Mbps network adapter running in full duplex mode (meaning that data can be transmitted and received at the same time). Using a ratio of one front-end server to four back-end servers, a two-processor 2.6-GHz front-end server can transfer around 13 Mbps of data to a back-end server, requiring a gigabit network card or multiple 100 Mbps network cards. This example creates extremely heavy network traffic as the saturation point of a 100-Mbps full duplex network connection is considered to be around 7 to 8 Mbps.

On high-end, front-end servers with two or more 2.6-GHz or faster processors, it is recommended that you use two 100-Mbps full duplex network connections, or a single gigabit Ethernet connection. Servers of this class can easily exceed the capacity of a single 100-Mbps full duplex connection.

To balance the client load across multiple POP3 front-end servers, you can use Network Load Balancing. Network Load Balancing enables multiple front-end servers to appear as one server, with incoming connections intelligently spread among the pool of available front-end servers.

POP3 Back-End Mailbox Server

This scenario measures how well a POP3 back-end server scales under different amounts of client requests that pass through a dedicated front-end server. The back-end server uses the Pentium 4 Xeon processor with Hyper-Threading.

Hardware

Table 2.22 shows the specifications of the five servers used in this scenario.

Table 2.22 POP3 back-end hardware configuration

Server type	Processor type	RAM	Storage
Front-end server	Intel P4 Xeon 2 processors, 2.6 GHz (Hyper-Threading)	1 GB	Not applicable
Four back-end servers	Intel P4 Xeon 8 processors, 550 MHz	4 GB	• 12 spindles of RAID0+1 for database volumes • 2 spindles of RAID0+1 for the transaction logs for each storage group

Scenario

An average message size of 26 KB is used. Each user's Inbox is populated with 31 IMAP4 messages before the test begins. Transport traffic occurs when the POP3 front-end server receives incoming e-mail from the Internet during the session targeted to one recipient. Activity on the POP3 front-end server includes users logging on to the server, retrieving all mail, and then deleting all mail from the server. The number of SMTP messages being delivered to mailboxes each second equals the number of messages retrieved and deleted through POP3 each second. The POP3 commands associated with message retrieval and deletion are RETR and DELE. The POP3 command issued by a POP3 client to determine the number of messages in a mailbox is STAT. Statistical counters for RETR, DELE, and STAT can be used to determine the number of POP3 transactions occurring on the front-end server each second. These counters can be accessed through System Monitor, which is included with the Microsoft Windows Server 2003 operating system.

In this scenario, the number of STAT commands received by the POP3 server is close to double the number of RETR and DELE commands. The Exchange Stress and Performance (ESP) tool was used to generate this server load. For more information about this tool, see Appendix A, "Exchange Performance Tools."

Table 2.23 shows how an Exchange 2003 back-end server servicing POP3 and inbound SMTP requests performs under different user loads.

Table 2.23 POP3 back-end performance

	100 POP3 users	150 POP3 users	175 POP3 users	200 POP3 users
% Processor Time	28.8%	51.7%	67.3%	84.7%
Context Switches/sec	16,201	21,436	23,286	24,537
SMTP Messages Del/sec	23.5	34.9	40.3	44.6
SMTP Local Queue	3.2	7.4	12.9	24.2
POP3 STAT/sec	173	257	296	332
POP3 DELE/sec	23.3	34.8	40.2	44.2
Disk Transfers/sec	631	891	989	1,053
Network Usage (in Kbps)	1,926	2,760	3,136	3,459

For more information about the performance counters used in this scenario, see Appendix B, "Performance Counter Definitions."

Processor

POP3 scales well on four processor servers. When the processor is at 84.7 percent, the server is stable and only 24,537 context switches occur. Figure 2.2 shows the relatively linear scalability of running a POP3 back-end server on a four-processor server.

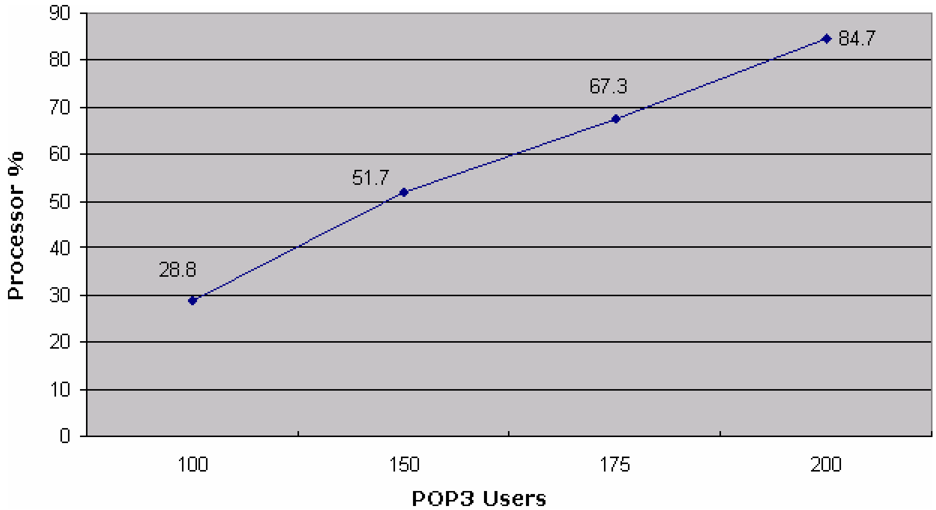

Figure 2.2 POP3 4-processor CPU usage

Figure 2.3 shows the context switches used during the tests. This figure demonstrates how the context switches grow linearly as more load is applied.

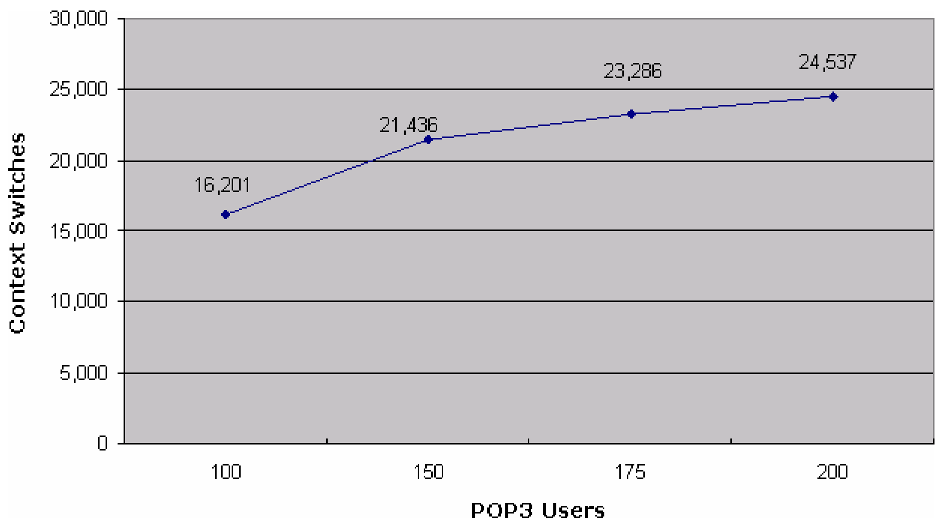

Figure 2.3 POP3 4-processor context switches

Memory

POP3 clients do not remain logged on to the server for extended periods of time. The typical behavior is to log on, retrieve all e-mail, delete all e-mail, and then log off the server. The four-processor POP3 back-end servers require at least 500 MB of RAM; otherwise, you can experience high latencies and your queues can increase greatly. The Inetinfo and Store.exe processes consume a combined total of 350 MB of memory during the tests. Exchange uses a maximum of 3 GB of memory. To increase performance by reducing paging to disk, increase memory to 3 GB.

Disk Usage

It is recommended that you have at least two spindles for the log drive and at least two spindles for the database files. It is a good idea to add a spindle for each additional 100 disk I/O operations expected.

On a production four-processor Exchange 2003 back-end server servicing POP3 requests, it is recommended that you have a minimum of 10 hard disks in addition to the operating system requirements for correct performance:

- Two mirrored disks for the log files
- Two mirrored disks for the SMTP queue
- At least six disks striped in a RAID0+1 configuration for the Exchange database files

Each spindle can handle approximately 100 random disk I/O operations per second. As the disks become saturated, add more spindles to the database.

Network Usage

A single 100-Mbps, full duplex, network connection is sufficient for almost all POP3 back-end applications.

POP3 Scalability Guideline

When you design a POP3 server, consider the following recommendations:

- A POP3 front-end server scales well on two-processor servers.
- A POP3 back-end server scales well on four-processor servers.
- Use a ratio of one front-end server to four back-end servers.
- 256 MB of RAM is sufficient for front-end servers.
- A POP3 front-end server uses virtually no disk resources, unless the server is paging or POP3 protocol logging is turned on.
- A POP3 front-end server requires a second 100-Mbps network adapter, or a gigabit Ethernet connection, if run on a high-end, 800-MHz, two-processor server.
- A POP3 front-end server can be load balanced using Network Load Balancing.
- Double the processor capacity of a POP3 front-end server if all the connections are performed over SSL.

IMAP4

This section provides baseline performance data on Internet Message Access Protocol version 4rev1 (IMAP4) servers. IMAP4 is an Internet protocol that allows an IMAP4 client to download e-mail from a server. This protocol works well for computers that can maintain a continuous connection to a server. IMAP4 offers more features than POP3, including folder hierarchies, flags, and searching.

IMAP4 Front-End Server

This scenario measures how well IMAP4 scales on a four-processor Xeon MP with Hyper-Threading server. Different levels of load are run against the server.

Hardware

Table 2.24 shows the specifications of the five servers used in this scenario.

Table 2.24 IMAP4 front-end hardware configuration

Server type	Processor type	RAM	Storage
Front-end server	Intel P4 Xeon 2 processors, 2.6 GHz (Hyper-Threading)	1 GB	Not applicable
Four back-end servers	Intel P4 Xeon 8 processors, 550 MHz	4 GB	• 12 spindles of RAID0+1 for database volumes • 2 spindles of RAID0+1 for the transaction logs for each storage group

Scenario

An average message size of 26 KB is sent. Each user's Inbox is populated with 31 IMAP4 messages before the test begins. Transport traffic occurs as incoming SMTP mail for local delivery at a rate of 24 messages per second.

Table 2.25 provides an overview of the IMAP4 traffic that can be handled by a single two-processor 2.6-GHz front-end server with a range of back-end servers. The **IMAP4 UID/sec** rate counter measures the number of unique ID (UID) commands received by the IMAP server per second. This counter is a good measure of the overall number of IMAP4 transactions occurring per second.

Table 2.25 IMAP4 front-end server performance

IMAP4 front-end server	Back-end server 1	Back-end server 2
% Processor Time	25.4%	47.3%
Context Switches /sec	41,945	39,781
IMAP4 UID/sec	436	2,084
IMAP4 Connections	15,000	30,000
Network Usage (in Kbps)	4,628	5,108
Inetinfo Working Set	221 MB	387 MB

For more information about the performance counters used in this scenario, see Appendix B, "Performance Counter Definitions."

Processor

An Exchange 2003 front-end server that only services IMAP4 clients does not use much CPU. Table 2.25 shows how 15,000 concurrent IMAP4 users consume only 25.4 percent of available CPU.

IMAP4 is most efficient when running on a two-processor front-end server. Like POP3, IMAP4 begins to context switch too much under moderate load on four-processor front-end server. Subsequently, it is not recommended that you run an IMAP4 front-end server on hardware that has more than four processors.

Memory

Like POP3, IMAP4 does not require much physical memory; 256 MB of RAM is sufficient for IMAP4 front-end servers that are servicing fewer than five back-end servers. If your front-end server services more than five back-end servers, you must have 512 MB of RAM installed on your front-end server.

Disk Usage

For information about disk usage for a dedicated IMAP4 front-end server, see "Disk Usage" in "POP3 Front-End Server" earlier in this chapter.

Network Usage

An IMAP4 front-end server uses many network resources when servicing multiple back-end servers. The minimum network requirement of a high-end IMAP4 front-end server is a single 100-Mbps network adapter running full duplex. Depending on the type of users accessing the system and the number of back-end servers being serviced, it may be necessary to add an additional 100 Mbps to each network adapter, or move to a gigabit-based network.

Network Load Balancing can also be used to balance the client load across multiple IMAP4 front-end servers.

IMAP4 Back-End Mailbox Server

This scenario compares how well IMAP4 scales on a four-processor Hyper-Threading Xeon MP server. Different levels of load are run against the server.

Hardware

Table 2.26 shows the specifications used in this scenario.

Table 2.26 IMAP4 back-end server hardware

Server type	Processor type	RAM	Storage
Back-end	Intel P4 Xeon 4 processors, 1.4 GHz (Hyper-Threading)	4 GB	• 12 spindles of RAID0+1 for database volumes • 2 spindles of RAID0+1 for the transaction logs for each storage group

Scenario

An average message size of 26 KB is sent. Each user's Inbox is populated with 31 IMAP4 messages before the test begins. Transport traffic occurs as incoming SMTP mail for local delivery at a rate of 24 messages per second.

Table 2.27 shows how an IMAP4 back-end server performs under various user loads.

Table 2.27 IMAP4 back-end server performance

	5,000 IMAP4 users	15,000 IMAP4 users	20,000 IMAP4 users
% Processor Time	35.3%	57.2%	71.7%
Context Switches/sec	22,929	50,565	46,978
SMTP Messages Del/sec	23.7	23.9	23.5
SMTP Local Queue	8.6	38.3	32.9
IMAP4 UID/sec	542	961.2	1,055
DB Disk Transfers/sec	870	1,304	1,921
Network Usage (in Kbps)	1,844,838	1,986,082	2,318,340

For more information about the performance counters used in this scenario, see Appendix B, "Performance Counter Definitions."

Processor

IMAP4 scales well on a four-processor back-end server. When determining the appropriate hardware for your IMAP4 back-end server, realize that the disk I/O operations increase as the average user's number of Inbox messages increases. Figure 2.4 shows as the user load increases, processor usage scales within an acceptable range.

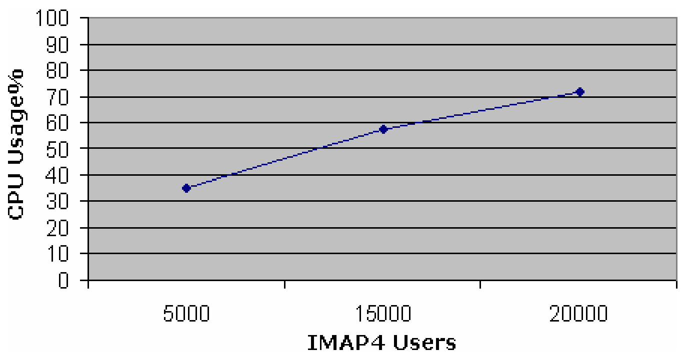

Figure 2.4 IMAP4 back-end server CPU usage

Memory

The four-processor IMAP4 back-end servers require at least 500 MB of RAM. Exchange uses a maximum of 3 GB of memory. To increase performance, increase memory to 3 GB to reduce the paging to disk.

Disk Usage

It is recommended that you have at least 2 RAID1 spindles for the log drive and at least 10 RAID0+1 spindles for the database files to support about 8,000 users. You can add a spindle for each additional 100 disk I/O operations expected. The more users you add to your server, the less likely a user will be cached and the more likely disk usage will increase. Because your disk subsystem is the first bottleneck on an IMAP4 back-end mailbox server, make sure that the **Current Disk Queue Length** performance counter stays below 10 (or *n*, where *n* equals the number of spindles in the array).

Network Usage

A single 100-Mbps, full duplex, network connection is sufficient for almost all IMAP4 back-end applications.

IMAP4 Scalability Guidelines

When you design an IMAP4 server, consider the following recommendations:

- An IMAP4 front-end server scales well to a two-processor front-end server.
- Use a ratio of one IMAP4 front-end server to eight back-end servers.
- An IMAP4 front-end server requires a minimum of 256 MB of RAM. However, if you are using a front-end server that services more than five back-end servers, use 512 MB of RAM.
- An IMAP4 front-end server uses virtually no disk resources, unless the front-end server is paging or IMAP4 protocol logging is turned on.
- A single 100-Mbps full duplex network connection is sufficient for all but the most demanding front-end server applications such as environments where large message attachments are common. Depending on the type of users you plan to service and the number of back-end servers being serviced, it may be necessary to add an additional 100-Mbps full duplex network adapter or move to a gigabit-based network.
- An IMAP4 front-end server can be load balanced using Network Load Balancing.
- SSL connections generate a 50 percent increase in CPU activity, and they require an additional 10 percent of physical memory.

SMTP

The section provides baseline performance data on Simple Mail Transfer Protocol (SMTP) gateways. SMTP is a network protocol that is designed to move electronic mail messages across a network to a destination server. In large data centers, it is recommended that you dedicate specific Exchange 2003 servers to handle only inbound and outbound SMTP traffic. These servers are typically named SMTP gateways or SMTP hubs. They are responsible for moving SMTP mail between clients and Exchange 2003 mailbox servers. This section describes only SMTP characteristics; message transfer agent (MTA) and message conversion are not covered here.

Hardware

Table 2.28 shows the hardware specifications used in this scenario.

Table 2.28 SMTP hardware

Server type	Processor type	RAM	Purpose
SMTP gateway	Intel 4 processors, 1.4 GHz	8 GB	• Two 1-Gbps network adapters • SMTP Queue directory was placed on a series of inexpensive disks with RAID0+1 disk array.

Scenario

The average message size transferred between clients and the SMTP gateway is 50 KB. Anonymous authentication is used with no SSL Transport Layer Security (TLS) encryption of the network traffic.

Table 2.29 provides an overview of performance on a four-processor 1.4-GHz SMTP gateway serving a range of back-end servers (which are the final destination for messages sent through the SMTP gateway).

Table 2.29 SMTP gateway server performance

	Windows 2000 Server and Exchange 2000	Windows 2000 Server and Exchange 2003	Windows Server 2003 and Exchange 2003
% Processor Time	56.0%	66.1%	66.7%
Context Switches/sec	34,782	31,854	18,850
SMTP Messages Sent/Sec	142.3	141.2	142.1
Network Usage (in Mbps)	12.8	12.8	12.8
Disk Writes/sec	810	920	842

For more information about the performance counters used in this scenario, see Appendix B, "Performance Counter Definitions."

Processor

The service scales well on a four-processor server. Although the CPU usage increases 17 percent after an upgrade to Exchange 2003, context switching drops 40 percent after an upgrade to Windows Server 2003.

Memory

SMTP gateway servers use memory primarily for maintaining connections and keeping track of vital information about messages in the queues. SMTP servers generally do not have thousands of connections; therefore, the memory required for this purpose is not significant. However, the memory used to store message properties on queued e-mail can be significant. SMTP stores messages in the queue in two states: opened (that it, it keeps a handle open) or closed (that is, it closes the handle). SMTP can keep 1,000 messages in the queue open at any particular time, closing old messages as new ones arrive. An open message in the queue consumes approximately 10 KB of memory in the Inetinfo process, and a closed message consumes approximately 4 KB of memory in the Inetinfo process.

The memory required by an SMTP gateway server is a function of how many messages the server can be expected to have queued at a particular time. The following examples demonstrate how to calculate the approximate memory that will be consumed by SMTP queues:

- 1,000 open messages = 10 MB of Inetinfo memory
- 1,000 open messages + 20,000 closed messages = 80 MB of Inetinfo memory
- 1,000 open messages + 89,000 closed messages = 366 MB of Inetinfo memory

> **Note**
> By default, the maximum number of messages that an SMTP gateway queues before refusing new messages is 90,000.

SMTP gateways are frequently used for distribution list expansion. Distribution lists enable messages to be sent to a single address and then distributed to multiple recipients. Distribution list expansion refers to the process of getting the list of recipients in the distribution list. Distribution list expansion also affects the memory used on an SMTP gateway. Each recipient expanded by a distribution list uses 1 KB of Inetinfo memory. Servers that expand very large distribution lists, or frequently expand medium sized distribution lists, require more memory.

Low-traffic SMTP gateway servers can perform adequately with 256 MB of RAM. In high-traffic data centers, where large queues are common and large distribution lists are expanded, at least 512 MB of RAM is preferred. Generally, an SMTP gateway server does not benefit by having more than 1 GB of memory. If a large SMTP queue is generated on a server that does not have sufficient memory, excessive paging occurs. This paging dramatically increases the time it takes for the server to handle the queues.

> **Note**
> The maximum number of messages that are kept open and the maximum number of messages allowed in the queue can be configured manually. For more information about this topic, see the following Microsoft Knowledge Base articles:
>
> - 271084, "XGEN: Exchange 2000 Server SMTP Optimized with Maximum Handle Threshold Registry Key" (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=271084)
> - 258748, "XGEN: Setting a Limit on the Number of SMTP Messages in Queues" (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=258748)

Disk Usage

An SMTP gateway is disk intensive. Every message that is received by an SMTP gateway is saved to a disk. With a default message size of 50 KB, about 7 to 8 disk writes occur for each message processed by the SMTP gateway. This behavior is expected. Generally, SMTP performs 7 disk writes for every message queued that is smaller than 32 KB. An SMTP gateway's write buffer is 32 KB. Therefore, messages that are larger than 32 KB require an additional disk write for every 32 KB. For example, a 100-KB message requires 10 disk writes to save the message in the queue.

For these reasons, an SMTP gateway server requires a high performance disk-subsystem. It is recommended that you use RAID0+1 array with multiple disk spindles for the SMTP queue drive. A RAID0+1 array is configured with half of the spindles in the array striped together to create one large volume, and this volume is mirrored to the other half of the available spindles for data redundancy. The number of disk spindles and the size of the write cache can be based on the expected SMTP message throughput of the server.

Network Usage

An SMTP gateway can process lots of SMTP data. As the load increases, the server's network resources are put under pressure. The minimum network capacity should be a single 100-Mbps with full duplex network connection. Full duplex connections can receive and transmit data at the same time. Depending on the expected server load, a second 100-Mbps connection may be required, or moving to a gigabit Ethernet-based network. Sending around 140 messages per second, when the average message size is 50 KB, generates almost 13 Mbps of network traffic. Servers expected to handle a load greater than 7 Mbps should have either two 100-Mbps, full duplex network connections, or a single gigabit Ethernet network connection.

You can also use Network Load Balancing to balance SMTP traffic across multiple SMTP gateway servers.

SMTP Scalability Guidelines

When you design an SMTP server, consider the following recommendations:

- An SMTP gateway scales well on four-processor servers.
- An SMTP gateway requires a minimum of 256 MB of RAM to manage large queues and expand distribution lists. A minimum of 512 MB of RAM is recommended.
- SMTP gateways require lots of disk resources; at least seven disk writes are required for each message queued. The drive containing the queue on high-end servers should be configured as RAID0+1 with multiple disk spindles that use a write-caching array controller.
- Two-processor gateways require a single 100-Mbps connection. High-end, four-processor servers should be fitted with two 100-Mbps connections or a single gigabit Ethernet connection.
- SMTP traffic can be load balanced using Network Load Balancing, and encrypting SMTP traffic using TLS does not have a significant effect on the processor or memory requirements.

CHAPTER 3

Maximizing Client Performance

In addition to Microsoft® Exchange Server 2003 improvements, several client performance factors can be achieved only through the new version of Microsoft Office Outlook® 2003 and Outlook Web Access. This chapter highlights the specific improvements over earlier client versions.

What You Will Learn in Chapter 3

- What are the benefits of using the new version of Outlook?
- What are the benefits of using the new version of Outlook Web Access?
- How can I monitor client performance?

Outlook 2003

Outlook 2003 contains several new features that can optimize client performance and can improve the overall user experience when used with Exchange Server 2003. These features include:

- **Cached Exchange Mode** This feature enables you to work in a messaging environment with a perceived connection between the Outlook 2003 client and the Exchange server. Typically, users do not notice any difference in messaging performance when using Exchange cached mode. For servers, processor performance was 10 percent better in comparison to Exchange 2000 Server with Service Pack 3 (SP3). For more information about server data, see "Mailbox Server (Using MAPI)" in Chapter 2, "Scaling Exchange Server 2003."

  Although both Exchange 2000 and Exchange 2003 support Cached Exchange Mode, Exchange 2003 includes several performance improvements implemented specifically to improve the performance of Outlook 2003 clients.

- **RPC over HTTP** This feature enables you to access your Exchange 2003 account from the Internet when you are working outside your organization's firewall without any special connections or hardware, such as smart cards and security tokens. For more information about configuring Exchange 2003 to use RPC over HTTP, see *Exchange Server 2003 RPC over HTTP Deployment Scenarios* (http://go.microsoft.com/fwlink/?LinkId=24823).

For more information about:

- How Exchange 2003 improves performance for Outlook 2003, see "Outlook Improvements" in *What's New in Exchange Server 2003* (http://go.microsoft.com/fwlink/?LinkId=21765).
- These features, see the *Exchange Server 2003 Client Access Guide* (http://go.microsoft.com/fwlink/?LinkId=27739).
- Performance comparisons of different versions of Outlook, see *Client Network Traffic with Exchange 2003* (http://go.microsoft.com/fwlink/?LinkId=27020).

Outlook Web Access

Outlook Web Access in Exchange 2003 provides several new features and enhancements. The new version uses gzip compression to compress data, which reduces overall network traffic by as much as 50 percent. In addition, underlying components were redesigned to take up less memory, resulting in lower bandwidth requirements. For a performance comparison of different versions of Outlook Web Access, see "Outlook Web Access" in Chapter 2, "Scaling Exchange Server 2003."

Additionally, Secure Sockets Layer (SSL) is a prerequisite to enabling compression and it does not add much overhead traffic, as shown in "Scenario 2" in "Outlook Web Access" in Chapter 2, "Scaling Exchange Server 2003." The only effect of SSL activation is the additional processor resource that you require on your Exchange servers. If you start SSL with Outlook Web Access, your messaging traffic decreases and becomes more secure at the same time.

For best connection performance, users should select the Outlook Web Access front-end server physically located closest to their mailbox server, not the Outlook Web Access front-end server closest to their present location.

For more information about:

- The new features for Outlook Web Access, see *What's New in Exchange Server 2003* (http://go.microsoft.com/fwlink/?LinkId=21765).
- How to set up Outlook Web Access with Exchange 2003, see the *Exchange Server 2003 Client Access Guide* (http://go.microsoft.com/fwlink/?LinkId=27739).
- Detailed performance on Outlook Web Access, see *Client Network Traffic with Exchange 2003* (http://go.microsoft.com/fwlink/?LinkId=27020).

Other Mobile Services with Exchange

Exchange 2003 includes Outlook Mobile Access and Exchange ActiveSync® features previously included with Microsoft Mobile Information Server. Previously, you had to install Mobile Information Server in every network domain where these services were required. Because Exchange includes built-in mobile services, installation on network domains is no longer necessary.

- **Outlook Mobile Access** This feature enables users to use mobile devices to access their Inbox, Contacts, Calendar, and Tasks folders. Previous support through Exchange 2000 and Mobile Information Server posed limits when mobile devices were used outside users' home domains. Exchange 2003 eliminates this domain boundary limitation.
- **Exchange ActiveSync** This feature enables users to synchronize their mobile devices directly with the Exchange Server 2003. Users can use their mobile carrier connection to synchronize their Exchange information to their Pocket PC Phone Edition or Smartphone device and then access this information when they are offline. Like Outlook Mobile Access, Exchange 2003 eliminates the domain boundary limitation in Mobile Information Server.

For more information about Outlook Mobile Access and Exchange ActiveSync, see the *Exchange Server 2003 Client Access Guide* (http://go.microsoft.com/fwlink/?LinkId=27739).

Monitoring Client Performance

Earlier versions of Exchange do not monitor the performance experience for Outlook client users. However, with Exchange 2003 and Outlook 2003, administrators can analyze performance for these users.

Exchange 2003 servers record both RPC latency and errors on client computers running Outlook 2003. An administrator can use this information to determine the overall client performance and to monitor the Exchange server for errors.

Outlook clients send RPC data (for example, latency data or error code) to the Exchange 2003 server on subsequent successful RPC calls.

> **Note**
> RPC data sent from the client computers to the Exchange server are not the primary method for detecting individual real time errors.

The client-side performance counters include:

- Client: RPCs attempted
- Client: RPCs succeeded
- Client: RPCs failed
- Client: RPCs failed: Server unavailable
- Client: RPCs failed: Server too busy
- Client: RPCs failed: all other errors
- Client: RPCs attempted / sec
- Client: RPCs succeeded / sec
- Client: RPCs failed / sec
- Client: RPCs failed / sec: Server unavailable
- Client: RPCs failed / sec: Server too busy
- Client: RPCs failed / sec: all other errors
- Client: Total reported latency
- Client: Latency > 2 sec RPCs / sec
- Client: Latency > 5 sec RPCs / sec
- Client: Latency > 10 sec RPCs / sec

You can monitor these counters with System Monitor or Microsoft Operations Manager. Additionally, the Exchange 2003 Management Pack for Microsoft Operations Manager provides an Outlook Client RPC Performance and Failure report. The Management Pack also has rules that are triggered by client monitoring events. These rules are located in Microsoft Exchange Server 2003\Exchange Event Monitoring\Information Store service\ of the Exchange 2003 Management Pack:

- **Client: Percentage of successful client RPCs is less than the specified threshold.**

  This rule is triggered by event MSExchangeIS 9640.

- **Client: Averaged client RPC latency is greater than the specified threshold.**

  This rule is triggered by event MSExchangeIS 9641.

- **Client: The number of client RPC errors and latency warnings issued is greater than the specified threshold.**

  This rule is triggered by event MSExchangeIS 9642.

For more information about the RPC-related operations that you can monitor using Microsoft Operations Manager, see *What's New in Exchange Server 2003* (http://go.microsoft.com/fwlink/?LinkId=21765).

For more information about Microsoft Operations Manager, see the Microsoft Operation Manager: Be Accountable Web site (http://go.microsoft.com/fwlink/?LinkId=16198) and the Better Together: Microsoft Operations Monitor and Exchange 2003 Web site (http://go.microsoft.com/fwlink/?LinkId=18176).

CHAPTER 4

Tuning Exchange Performance

This chapter provides information about how to tune and improve the performance of Microsoft® Exchange Server 2003. This chapter includes a collection of tuning tips and recommendations that may not all apply to your current deployment. Review the chapter to determine what settings apply to your situation.

For Exchange 2000 Server customers, some of the recommendations in this chapter are the same recommendations from the previous performance guide. Additionally, some previous changes are no longer required. For more information about settings that you may want to remove, see "Obsolete Settings" later in this chapter.

Performance Improvements

Exchange Server 2003 includes a number of new tuning enhancements such as link state traffic and virtual address space management. For a detailed description of these performance improvements, see *What's New in Exchange Server 2003* (http://go.microsoft.com/fwlink/?LinkId=21765).

Disk Subsystem

The following sections provide recommended storage configuration for the various types of Exchange 2003 servers. For Exchange 2000 customers, this section contains the same recommendations as previous Exchange 2000 performance documentation.

SMTP Bridgehead Server

If the server is a Simple Mail Transfer Protocol (SMTP) bridgehead server, generally the best disk layout is to create one partition. Messages arrive on the SMTP interface, are written to the Mailroot directory (which should be on an NTFS file system-formatted system), and then are passed to the next computer. To get the best performance, span the Mailroot directory over as many disks as possible. A single spindle should yield a performance rate of approximately 30 small messages per second. Therefore, when you increase the number of spindles, relay performance increases.

Your hardware redundant array of inexpensive disks (RAID) configuration is largely dictated by the number of disk spindles available, as shown in Table 4.1.

Table 4.1 RAID recommendations

Number of disks	Recommended hardware RAID level and partitioning
2	RAID1, mirroring, one partition
4 or more	RAID0+1, data striping and mirroring, one partition

If you have more than one partition on an SMTP bridgehead server, relocate the SMTP Mailroot directory to the fastest partition.

If your hardware RAID controller has a mirrored, battery-backed, write-back cache, and it allows you to tune the read/write cache ratio, set the ratio to 100 percent write because the server can only acknowledge receipt of a message after it has been written to the disk. Therefore, the faster the server can write to the disk, the more responsive the server is to other computers. Although the message contents must be reread before being relayed to the next computer, the SMTP service has an open file handle to the data, and it can retrieve the contents from the NTFS cache.

X.400 Bridgehead Server or Other Connector Server

If your Exchange 2003 server contains many X.400 connectors, or if it connects to other messaging systems (such as Lotus cc:Mail, Lotus Notes, Novell GroupWise, or Microsoft Mail), create a separate disk partition for transaction logs if you have sufficient spindles. Messages that use these Exchange connectors move data through the Store.exe process and, if the server is under a heavy load, additional performance is gained by splitting the transaction logs from the database.

When a message is received by the message transfer agent (MTA), the data is written to the Mtadata directory on an NTFS partition and passed to the Store.exe process through an MTS-IN virtual queue. Because the message goes to the Store.exe, data is written to the transaction log files. The message then goes through the categorization and routing process. If routing has determined that the message should be sent out through an X.400 connector, the data is moved to the MTS-OUT virtual queue. Other messaging connectors use a similar process.

The number of disk spindles available affects your hardware RAID configuration (see Table 4.2).

Table 4.2 RAID recommendations

Number of disks	Recommended hardware RAID level and partitioning
2 - 4	RAID1, mirroring, two partitions (operating system and pagefile on partition 1 and Exchange on partition 2)
5	RAID5 (drive C) for three disks; binaries and database RAID1 (drive D) for two disks; log files
6	RAID5 (drive C) for three disks; binaries and database RAID1 (drive D) for two disks; log files No RAID (drive E) for one disk; pagefile

Although the 6-disk configuration does not use a RAID partition for the pagefile, if high availability of the server is critical, locate the pagefile on a protected partition.

If your hardware RAID controller has a mirrored, battery-backed, write-back cache, and it allows you to tune the read/write cache ratio, set the ratio to 100 percent write for configurations that use a single partition, and 100 percent write on the transaction log partition when you use five or more disks. You will see a better performance gain by using write cache on the transaction logs instead of read cache on the database drive.

Mailbox and Public Folder Servers

If you configure a single database server, your disk design is similar to that of Exchange Server 5.5 and Exchange 2000 Server. Your first priority is to split the transaction logs and database onto separate disk spindles. Not only should the logs be on a different spindle set, but that spindle set should be dedicated only to logs. For example, you do not receive the performance benefits of the split if you put the pagefile or the Windows® 2000 Server or Windows Server™ 2003 operating system on the same spindle set as the logs.

If you intend to create multiple storage groups on the server, your number one priority, again, is to split the logs on to dedicated spindles. However, to gain any performance, you must have a dedicated spindle set for each storage group. After you have a dedicated spindle set, determine how many spindles you have left. On the majority of servers, not many disk slots will be available. Therefore, your best option may be to put all databases on the same partition.

If you spread databases from multiple storage groups over the same array and that array becomes unavailable, all databases from those storage groups (even those databases located on other arrays) temporarily become unavailable. The bad database is marked as down, and the good databases reconnect. The outage should last only a few minutes, and then MAPI clients such as Microsoft Outlook® recover their connections.

With sufficient disks, you can split the properties database (.edb) files from the streaming database (.stm) files. The types of clients that your server accommodates dictate the performance benefit of splitting these files. For example, if you expect to service only Internet Message Access Protocol version 4rev1 (IMAP4) and Post Office Protocol version 3 (POP3) users with large attachments, the read/write size might be larger than the 4 KB that is typically seen with MAPI users on an .edb database. In this case, splitting these files is recommended. You can optimize your disks for .stm files to support the large attachments.

MAPI and Microsoft Outlook Web Access clients read and write data to and from the .edb file. Technically, clients communicate to the Store.exe process. Therefore, they have no understanding of .edb files versus .stm files. All other clients use the .stm file. Table 4.3 shows the preferred location that is used by different clients.

Table 4.3 Preferred client storage locations

Client type	Submitting data	Retrieving data
Outlook in MAPI mode	.edb files	.edb files
POP3	.stm files (through SMTP)	.stm files
IMAP4	.stm files (through SMTP)	.stm files
Outlook Web Access	.stm files (but full promotion to .edb files)	.edb files
Installable File System (IFS)	.stm files	.stm files
SMTP	.stm files	Not applicable
SMTP (with MAPI data)	.stm files (but full promotion to .edb files)	Not applicable
Microsoft ActiveX® Data Objects (ADO) or OLE DB	.stm files	.stm files
Collaboration Data Objects (CDO) for Exchange 2000	.stm files	.stm files
CDO 1.21	.edb files	.edb files

Although Table 4.3 indicates the preferred file type, cross conversions may occur. For example, if a POP3 client submits a message, the message data is physically stored in the .stm file. There is some automatic property promotion to the .edb file of the message header and other data. However, now assume that a MAPI client attempts to read the message. In this scenario, the message data in the .stm file is fully promoted to the .edb file (attachments are left in the .stm file and converted in memory). Promotion only occurs from the .stm file to the .edb file. There are no circumstances where data is promoted to the .stm file; all client conversions are performed in memory in these scenarios (with the aid of temporary files).

In the majority of scenarios, there is no real benefit to splitting the .edb files and .stm files on separate spindles. For example, if you have six disks, you must decide whether creating two RAID5 partitions over three disks each would give better performance than a single RAID0+1 (stripe and mirror) partition. In this case, RAID0+1 provides better write performance.

As disks become larger, it is tempting to purchase fewer larger disks to handle user data. Unfortunately, the speed of disks is not relative to the size. Therefore, you may have sufficient physical disk space, but the performance may not be sufficient for the user load. Allocate sufficient physical disk spindles for mailbox servers. It should be noted that specialized storage subsystems that use gigabytes of cache and many racks of hard disks typically have their own partition management and fault-tolerant operating system. When you work with these devices, seek specialist assistance from the hardware vendor.

In summary, the following are your best guidelines for configuring mailbox servers:

- Create a RAID1 partition for Windows and Exchange binary files.
- Put the pagefile on a separate RAID1 spindle. For mailbox servers, you should never put the pagefile on a non-RAID partition because the loss of the volume causes the server to stop.
- Create one dedicated fault-tolerant partition for each storage group for the transaction logs (for example, RAID1 or RAID0+1). A two-disk RAID1 partition should yield approximately 300 sequential write I/O operations a second. This capacity should be more than sufficient for a busy five-database storage group.
- Create at least one fault-tolerant partition for your databases. If you have only one array, put all databases on this array. If you have multiple arrays, use one array for each storage group (databases only).

Databases, SMTP queues, and transaction log files can be moved to different partitions by using Exchange System Manager.

For more information about moving databases and logs, see Microsoft Knowledge Base article 257184, "XADM: How to Move Exchange Databases and Logs in Exchange 2000 Server" (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=257184).

Aligning Exchange I/O Operations with Storage Track Boundaries

With a physical disk that maintains 64 sectors per track, Windows always creates the partition starting at the sixty-forth sector, therefore misaligning it with the underlying physical disk. To be certain of disk alignment, use Diskpar.exe, a disk partition tool. Diskpar.exe is a utility provided by Microsoft in the Windows 2000 Resource Kit. Diskpar.exe is a command-line utility that can explicitly set the starting offset in the master boot record (MBR). By setting the starting offset, you can track alignment and improve disk performance. Exchange 2003 writes data in multiples of 4 KB I/O operations (4 KB for the databases and up to 32 KB for streaming files). Therefore, make sure that the starting offset is a multiple of 4 KB. Failure to do so may cause a single I/O operation spanning two tracks, causing performance degradation.

Optimizing Memory Usage

This section contains information about monitoring and optimizing memory usage on your servers.

Monitoring Memory Usage

To monitor the Event Viewer application log and Performance Logs and Alerts for virtual memory problems, in Administrative Tools, click **Event Viewer** or **Performance**. In the application log, an Event ID 9582 warning appears when the largest free block of virtual memory decreases to 32 MB. If you see this warning, you should restart the Exchange store process at the next opportunity. If the largest block decreases to 16 MB, an Event ID 9582 error appears again; this error means the server might fail, and you should restart the server at the earliest opportunity. Failure to act on these events can cause sporadic mail delivery and IMAIL conversion failures (Event ID 12800).

In Performance Logs and Alerts, monitor the following counters:

- **VM Largest Block Size counter in the MSExchangeIS object** A healthy server has more than 200,000,000 bytes (200 MB) as the largest free block. If the value is lower, carefully monitor the server.
- **Pool Pages Bytes in the Memory object** When the **/3GB** switch is used, amounts larger than 200 MB indicate a problem except when backups are running. During backups, each page in the cache manager is copied into the pool page, which causes an increase in pool page size.
- **Pool Nonpaged Bytes in the Memory object** When the **/3GB** switch is used, amounts larger than 100 MB indicate a problem.
- **Free System Page Table Entries in the Memory object** Amounts less than 3000 indicate a problem.
- **Working Set in the Process object** An upward trend indicates a potential memory leak.

If a server shows signs of low virtual address space, you should adjust the following settings. If these settings are not optimized for Exchange, event 9665 appears in Event Viewer.

- If the server is running Windows 2000 Advanced Server or Windows Server 2003 and has 1 GB or more of physical memory, set the **/3GB** switch in the Boot.ini file.
- If the server is running Windows Server 2003 (any edition), configure the /USERVA switch registry key.
- If the server is running Windows 2000, make sure that Windows 2000 SP3 or later is installed.
- If the server is running Windows Server 2000 Advanced Server, configure the SystemPages registry key.
- If the server has 1 GB or more of physical memory, set the HeapDeCommitFreeBlockThreshold registry parameter.
- If necessary, tune the store database cache size.

Event 9665

Exchange performs an optimal memory configuration check when the store process starts. If the memory settings are not optimal, you receive an event 9665 in the Event Viewer. This message appears in the following instances:

- The server is running Windows 2000 and the SystemPages value in the registry is set outside the range of 24,000 to 31,000.
- The server has 1 GB of memory or more and does not have the **/3GB** switch.
- The server is running Windows Server 2003, has 1 GB of memory or more and the **/3GB** switch is set, but the /USERVA setting is not present or is outside the range of 3,030 to 2,970.

If you see this event, check the SystemPages and HeapDeCommitFreeBlockThreshold settings in the registry, in addition to the **/3GB** switch and the /USERVA setting in the Boot.ini file. SystemPages applies only to Windows 2000. The following sections contain recommendations for each of these settings.

If you want to turn off the memory configuration check, create the registry key shown in Table 4.4.

> **Warning**
> Incorrectly editing the registry can cause serious problems that may require you to reinstall your operating system. Problems resulting from editing the registry incorrectly may not be able to be resolved. Before editing the registry, back up any valuable data.

Table 4.4 Registry key to turn off memory configuration check

Path	HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\ Services\MSExchangeIS\ParametersSystem\
Parameter	Suppress Memory Configuration Notification
Type	REG_DWORD
Setting	1

> **Note**
> The memory configuration check does not occur on servers running Microsoft Small Business Server.

Setting the /3GB Switch

By default, Windows 2000 Advanced Server and Windows Server 2003 allocate 2 GB of virtual address space to user mode processes such as Store.exe. If a server has 1 GB or more of physical memory, set the **/3GB** switch in the Boot.ini file to increase virtual address space. For more information about the **/3GB** switch, see Microsoft Knowledge Base article 266096, "XGEN: Exchange 2000 Requires /3GB Switch with More Than 1 Gigabyte of Physical RAM" (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=266096).

> **Important**
> The /3GB switch is designed for Windows 2000 Advanced Server and all editions of Windows Server 2003. Do not set the /3GB switch in Windows 2000 Standard Edition.

Configuring /USERVA and SystemPages

If the server is running Windows 2000, you should set the SystemPages registry key to a value between 24,000 and 31,000. The SystemPages registry key is located in the following path:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\
Session Manager\Memory Management\SystemPages
```

If the server is running Windows Server 2003, set the SystemPages value to zero, and set the /USERVA=3030 parameter in the Boot.ini file. These settings let you enter more system page table entries on the server, which is critical for systems that scale-up.

For more information, see Microsoft Knowledge Base article 810371, "XADM: Using the /Userva Switch on Windows Server 2003-Based Exchange Servers" (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=810371).

Setting the HeapDeCommitFreeBlockThreshold Registry Key

The HeapDeCommitFreeBlockThreshold registry key controls the free space required before the heap manager decommits (or frees up) memory. The default is zero, which means that the heap manager decommits each 4-KB page that becomes available. Over time, virtual address space can become fragmented. On servers that have 1 GB or more of physical memory, you should set the registry key to a higher value to reduce fragmentation. Set the registry key as shown in Table 4.5, and then restart the server. For more information about the HeapDeCommitFreeBlockThreshold registry key, see Microsoft Knowledge Base article 315407, "XADM: The 'HeapDecommitFreeBlockThreshold' Registry Key" (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=315407).

Table 4.5 HeapDeCommitFreeBlockThreshold settings

Path	HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\ Control\Session Manager
Parameter	HeapDeCommitFreeBlockThreshold
Type	REG_DWORD
Default	Zero
Recommended setting	262144 **Note** This value is the number of blocks in decimal. The recommended value is 262144. This value corresponds with a hexadecimal value of 0x00040000.

Exchange Store and Extensible Storage Engine Tuning

Exchange 2003 can handle up to 25 databases per server. The total storage space is made up of a maximum of four production storage groups with five databases per storage group and a Restore Storage Group that handle up to five databases for recovery purposes. Wherever possible, you should fill out your storage groups to the maximum number of databases before creating a new group. The advantages to filling out your storage group are:

- Reduced memory consumption
- Reduced disk overhead

However, there are some tradeoffs to make:

- You can control logging only at the storage group level.
- If one database in the group goes offline because of an error, the other databases in that group temporarily go offline and then reconnect (note that this scenario does not apply to out-of-disk-space errors).
- Only one backup or restore process can occur in a single storage group at any one time.
- Backing up one database in a storage group stops the online maintenance of all other databases in the storage group.

The most efficient solution uses a combination of storage groups and databases. For example, a server with two storage groups and three databases in each provides a good balance between performance, resource consumption, and manageability. If this server has sufficient disk spindles, consider using three storage groups with two databases in each. This design provides you with a smaller recovery window if parallel streaming is used. Only on the largest of servers should you create and populate the fourth storage group.

Single Instance Support

Exchange 2003 fully supports single instance messaging. If a single message is addressed to multiple recipients and those recipients are located in the same database, only one message is sent to the database. The recipients receive pointers to the single message.

If message recipients are located on multiple databases in a storage group, a single copy of the message is delivered to each database where recipients are present. Additionally, multiple copies of the message are held in the transaction log set for the storage group.

There are many theories about how to make the best use of single instance support. Some people prefer to locate whole departments of users in the same database, therefore potentially saving disk space; others do not like this approach as a database failure can make an entire department become unavailable.

Single instance was designed to improve the delivery performance of the server when delivering to many people, and not necessarily to save space on the disk.

Single instance is also retained if you move mailboxes between Exchange databases, storage groups, and servers.

Increasing Page Table Entries on Servers Running Windows 2000 Server

On large scale-up mailbox servers that have the **/3GB** switch set, the Store.exe process has more virtual address space in user mode and the kernel address space is reduced to 1 GB. Unfortunately, this issue can cause an imbalance if a heavily-loaded server can run out of page table entries. If the issue continues, your server can experience resource failures when writing data to the disk or network connection failures.

For more information about what to do if your server is running out of page table entries, see Microsoft Knowledge Base article 313707, "XADM: An Exchange 2000 Server with the '/3GB' Switch in the Boot.ini File May Lose Network Connectivity Under a Heavy Messaging Load" (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=313707).

> **Note**
> This workaround is not required when you use Exchange Server 2003 on Windows Server 2003 operating systems where the /USERVA=3030 option has been set.

Online Database Maintenance

> **Note**
> For Exchange 2000 customers, this topic contains the same recommendation as previous Exchange 2000 performance documentation.

Each mailbox and public folder store requires periodic online maintenance to be run. By default, each database is set to run online maintenance between the times of 01:00. and 05:00. Online maintenance performs a variety of tasks necessary to keep the store in good health. These tasks include, but are not limited to:

- **Task 1** Check Microsoft Active Directory® directory service to determine whether there are any deleted mailboxes.
- **Task 2** Permanently remove any messages or mailboxes that are older than the configured retention policy.
- **Task 3** Perform online defragmentation of the data in the database.

These operations, performed by online maintenance, have specific performance costs, and you should understand these costs before you implement an online maintenance strategy.

Task 1 performs an Active Directory lookup for each user in the database. The more users you have in each database, the more Active Directory (LDAP) searches are made. These searches are used to keep the mailbox store synchronized with any Active Directory changes (specifically, look for deleted mailboxes). The performance cost of this task is negligible on the Exchange server, but it can be significant on Active Directory depending on the number of users, the number of databases, and the online maintenance times of each database. In a corporate scenario, the online maintenance typically occurs at night, when very few users are logged on and the load on Active Directory servers is very low. The extra domain controller load created by online maintenance should not be a problem in this scenario.

If Exchange 2003 is installed in a global data center and serves customers from multiple time zones, the default online maintenance time may become an issue. The effect online maintenance has on Active Directory is proportional to the number of users in each of the server's databases. A check for a deleted mailbox is performed against each user in each database. Therefore, if you have 10,000 users spread over multiple databases on a server, 10,000 Lightweight Directory Access Protocol (LDAP) searches against Active Directory will occur at the default time of 01:00. If Active Directory servers are always (around the clock) under moderate load, stagger the online maintenance (set each database to start maintenance at a different time). Staggering the online maintenance is especially critical if you have hundreds of thousands of users spread across dozens of servers and hundreds of databases.

Tasks 2 and 3 are very disk intensive and only affect the server where the maintenance is being run. During this part of online maintenance, the server may be perceived by users as sluggish if many databases are set to perform online maintenance at the same time. Again, in corporate scenarios, this maintenance would occur at night when the server can easily handle the extra load. In a global data center, it may make sense to stagger the database schedule (in respect to each other on a single server) to spread disk-intensive tasks over a greater period of time.

Online backup requires additional consideration. Backing up an Exchange 2003 database stops the maintenance of any database in that storage group (maintenance restarts if the backup finishes before the maintenance interval has passed). If you have two databases in a single storage group and one is running online maintenance, if a backup is started against either database, the online defragmentation on the database (which is running online maintenance) is stopped. It is critical that the backup time for any database in a storage group not conflict with the maintenance times of any database in the same storage group. If it does, backup stops the online defragmenting part of the online maintenance and the database may never finish defragmenting.

The correct online maintenance strategy can be devised by examining the user profile (such as times of low activity); knowing how many users, databases, and servers are in the site; and coordinating this information with the online backup strategy.

Here is an example of an online store maintenance schedule for a corporate Exchange 2003 mailbox server, which hosts users for a single time zone:

First Storage Group

Database One: Online maintenance runs between 21:00 and 01:00.

Database Two: Online maintenance runs between 21:30 and 01:30.

Database Three: Online maintenance runs between 22:00 and 02:00.

Online backup begins at 02:00 and backs up all databases in the first storage group when all databases have finished online maintenance.

Second Storage Group

Database Four: Online maintenance runs between 22:30 and 02:30.

Database Five: Online maintenance runs between 23:00 and 03:00.

Database Six: Online maintenance runs between 23:30 and 03:30.

Online backup begins at 03:30 and backs up all databases in the second storage group when all databases have finished online maintenance.

This configuration staggers the Active Directory LDAP queries generated by online maintenance, which are performed in the first minutes of the procedure, and makes sure that online backup does not interfere with online defragmentation.

The third task of defragmenting the database is made up of 18 separate subtasks. After a subtask has started, it must complete. Therefore, if subtask 12 is still running at the end of the online maintenance window, this task completes fully before the process exits. Therefore, online maintenance can run over the time window. Subtask 13 runs during the next online maintenance window. Depending on the run window and backup schedule, it may take more than one day for a full defragmentation to finish.

Message Promotion on Move Mailbox

Although Internet protocol clients such as IMAP4 and POP3 use the streaming store (.stm file) for reading and writing data, you should understand that the move mailbox function in Active Directory Users and Computers moves data into the .edb file. Therefore, clients who use POP3 and IMAP4 to access their mailboxes may see decreased performance after their mailbox has been moved between databases or servers. When logging on, message size is calculated and a MAPI to MIME conversion occurs in the memory and on the disk of the server. In extreme cases this can cause very large temporary files being created on the Exchange 2003 server.

If you move hundreds or thousands of mailboxes, you can mitigate potential issues by using the following recommendations. On the destination server, do both of the following:

- Make sure that the TMP/TEMP environment variables point to a very fast RAID0+1 spindle set (up to 12 disks for large mailbox servers). For stand-alone mailbox servers, the system environment TMP/TEMP variables should be adjusted. For clustered servers, the variables should be configured for the service account that the Cluster service is running under.
- Set the following registry parameters to inform the store that it should use approximate instead of exact calculations for message sizes.

Table 4.6 Compatibility setting (POP3svc)

Location	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\POP3svc\Parameters
Parameter	Compatibility
Type	REG_DWORD
Default setting	Not present
When to change	Change when you want the store to use approximate message size calculations. **Note** Changing this setting causes some older mail clients (such as Fetchmail) to break. Setting this registry key also breaks RFC-compliance for the POP3 protocol.
Recommended setting	0xfffffffe

Table 4.7 Compatibility setting (IMAP4svc)

Location	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\IMAP4svc\Parameters
Parameter	Compatibility
Type	REG_DWORD
Default setting	Not present
When to change	Change when you want the store to use approximate message size calculations. **Note** Changing this setting causes some older mail clients (such as Fetchmail) to break. Setting this registry key also breaks RFC-compliance for the IMAP4 protocol.
Recommended setting	0xfffffffe

After you set these registry parameters, you must restart the Store.exe and Inetinfo processes.

For more information, see Microsoft Knowledge Base Article 317722, "XADM: Client Latencies Occur When Exchange 2000 Converts Mail from MAPI to MIME Format" (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=317722).

Virtual Address Space and Store Database Cache Size

The Store.exe process in Exchange 2003 has a finite amount of memory that it can address. This amount is known as the virtual address space. For larger servers, you may want to manually adjust the virtual address space for optimal memory usage.

The Store Database Cache, also known as the ESE buffer, provides a large caching area for database transactions before they are committed to the store. If a server is heavily loaded or if disk performance is not optimal, increasing the ESE buffer improves overall system performance. Depending on your configuration, you may have to increase or reduce the size of this buffer to obtain the best overall performance.

For more information about how to adjust these settings, see Microsoft Knowledge Base article 815372, "How to Optimize Memory Usage in Exchange Server 2003" (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=815372).

Backup Tuning

You may experience performance problems during backup operations on computers running Windows Server 2003. This problem occurs because of the way that the Jet database cache size increases or decreases depending on the Transition Fault/sec counter from the operating system. For example, currently backup is performed by using the cache, which affects the Transition Fault/sec counter. This example may cause inconsistencies in how the Jet database memory management feature optimizes its memory usage and in how it tries to free more memory for the operating system.

You can modify the ExchESEParamCacheSizeMin parameter to improve backup performance. This setting also resolves performance issues related to context indexing.

For more information, see Microsoft Knowledge Base Article 822894, "Jet Database Does Not Work Correctly During a Backup or During a Content Indexing Operation" (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=822894).

Message Transfer Agent Tuning

Exchange 2003 does not include a Performance Optimization wizard; the majority of Exchange 2003 components are self-tuning. For Exchange 2000 customers, this section contains the same recommendations as previous Exchange 2000 performance documentation.

In scenarios in which only Exchange 2003 servers exist in an organization, the MTA does almost no processing and, therefore, does not need performance tuning. However, when you use both Exchange 5.5 and Exchange 2003, the MTA is used for all communications within a site, and possibly between sites, if site connectors using remote procedure calls (RPCs) or X.400 connectors are deployed. The MTA also processes messages coming from or going to other messaging systems such as Lotus cc:Mail, Lotus Notes, Novell GroupWise, and Microsoft Mail.

In smaller organizations, the MTA process does not require manual tuning. For larger companies that encompass many servers in the same site, many sites in the organization, or earlier messaging connectors, consider tuning the MTA registry parameters. In Exchange 2003, the MTA places a much greater load on system resources (CPU, memory, and disk) than in Exchange 5.5, because all messages now go through the Store.exe process. You should consider this increased load when planning the size of your hardware. In a large mixed Exchange 5.5 and Exchange 2000 and/or Exchange 2003 site, you may want to divide your servers into different routing groups, and dedicate one of the Exchange 2003 servers to handling the Exchange 2000 and/or Exchange 2003 to Exchange 5.5 MTA communication. The other Exchange 2003 servers in the site can use this dedicated bridgehead server to send messages to the Exchange 5.5 servers, which reduces the MTA processing overhead.

The rest of this section details the registry parameters that may require adjustment and gives examples as to where tuning may be appropriate. Increasing the number of threads for a process does not always increase performance (it may actually decrease performance). For example, consider the relationship between Reliable Transfer Service threads and Kernel threads in the MTA. Reliable Transfer Service (RTS) threads are responsible for putting data in the queues; Kernel threads take the data out and moves if off of the server. If the Reliable Transfer Service threads setting is too high, local queues fill quickly and may prevent the Kernel threads from reading the queue data, causing message transfer to slow.

MSExchangeMTA\Parameters Registry Key Settings

The registry parameters in this section can be found in the following registry key:

HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\
MSExchangeMTA\Parameters

Table 4.8 DB data buffers per object setting

Parameter	DB data buffers per object
Type	REG_DWORD
Explanation	This value is the number of database server buffers that are configured for each database object. More buffers require more memory, but they make it less likely for a database object to be rolled out to disk because of a lack of buffer space.
Default setting	0x00000003
When to change	Adjust if this MTA communicates with multiple Exchange 5.5 servers either in a site, or between sites. Additionally, you should tune if another messaging connector is homed on this server.
Recommend setting	0x00000006

Table 4.9 Dispatcher threads setting

Parameter	Dispatcher threads
Type	REG_DWORD
Explanation	This value is the number of MTA Dispatcher threads, which are responsible for the processing of messages. This is multiplied by three for the three subtypes (Router, Fanout, Result) of Dispatcher thread.
Default setting	0x00000001
When to change	Adjust if this MTA communicates with other Exchange 5.5 servers. Additionally, consider tuning if more than five Exchange 5.5 servers exist in the site, or if groups (that is, distribution lists) are heavily used.
Recommend setting	0x00000003

Table 4.10 Kernel threads setting

Parameter	Kernel threads
Type	REG_DWORD
Explanation	This value is the number of platform threads that handle the Presentation and Session level of the Open Systems Interconnection (OSI) stack. These threads are at the heart of MTA message processing.
Default setting	0x00000001
When to change	Adjust if this MTA communicates with other Exchange 5.5 servers using RPC over slow or highly latent network connections.
Recommend setting	0x00000003 — Standard recommendation. 0x00000008 — If the Exchange 2003 MTA belongs to a site containing more than 15 Exchange 5.5 servers. 0x0000000C (12) — If the Exchange 2003 MTA belongs to a site containing more than 30 Exchange 5.5 servers.

Table 4.11 Max RPC Calls Outstanding setting

Parameter	Max RPC Calls Outstanding
Type	REG_DWORD
Explanation	This value is the maximum number of RPC threads. This setting limits the maximum number of RPCs that are guaranteed to be processed at the same time.
Default setting	0x00000032 (50)
When to change	Adjust if you increase the number of Gateway In/Out threads in the Store.exe process (which is recommended in Exchange 5.5 and Exchange 2000 co-existence scenarios).
Recommended setting	0x00000080 (128)

Table 4.12 MDB users setting

Parameter	MDB users
Type	REG_DWORD
Explanation	Defines the number of distinguished names (also known as DNs) to cache from the directory.
Default setting	0x000001F4 (500)
When to change	Adjust when your organization contains more than 1,500 users. Change the value to one-third the size of the global address list (GAL), to a maximum of 5000.
Recommended setting	0x00001388 (5000)

Table 4.13 RTS threads setting

Parameter	RTS threads
Type	REG_DWORD
Explanation	Reliable Transfer Service threads. This value is the number of platform threads that handle the reliable transfer service element (RTSE) level of the OSI stack.
Default setting	0x00000001
When to change	Adjust if this MTA communicates with multiple Exchange 5.5 servers, either in a site or between sites.
Recommended setting	0x00000003

Table 4.14 TCP/IP control blocks setting

Parameter	TCP/IP control blocks
Type	REG_DWORD
Explanation	This value is the maximum number of concurrent RFC 1006 (TCP/IP) connections that are supported. This setting controls the number of buffers available for X.400 connections.
Default setting	0x00000014 (20)
When to change	Adjust if hosting more than one X.400 connector on the server.
Recommended setting	10 control blocks for each hosted X.400 connector, plus 1 control block for incoming connections.

Table 4.15 Transfer threads setting

Parameter	Transfer threads
Type	REG_DWORD
Explanation	This value is the number of MTA Transfer threads. It is multiplied by two for the two subtypes (Transfer In, Transfer Out) of the Transfer thread.
Default setting	0x00000001
When to change	Adjust if this MTA communicates with multiple Exchange 5.5 servers, either in a site or between sites.
Recommended setting	0x00000003

You must restart the MTA before registry adjustments take effect.

> **Note**
> After you upgrade your last Exchange 5.5 server, restore the original registry parameters for the MTA. Using the MTA tunings described in Table 4.10 through Table 4.17 in a native mode Exchange organization decreases the performance of your Exchange installation.

MSExchangeIS Registry Key Settings

When messages are received in the MTA from an Exchange 5.5 server or a previous gateway, they are handed off to the Store.exe process, and then they go to the advanced queuing engine. In environments in which many messages reach the MTA at the same time, it is a good idea to increase the number of processing threads between the Store.exe process and MTA. You can monitor the MTA to Store.exe process queue build-ups using either System Monitor or Exchange System Manager.

The registry parameters in this section can be found in the following key:

HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\
MSExchangeIS*server_name*\Private-*database_guid*

Table 4.16 Gateway In Threads setting

Parameter	Gateway In Threads
Type	REG_DWORD
Explanation	Defines the number of threads available for retrieving messages from the MTA process in the Store.exe process.
Default setting	Not present, but defaults to 0x00000001.
When to change	Adjust if this MTA communicates with multiple Exchange 5.5 servers or other messaging connectors.
Recommended setting	0x00000003
Note	Each thread consumes about 1 MB of virtual memory. Additionally, the actual number of threads created is this value multiplied by the number of databases. This may be an issue on servers with many private databases. For example, if you have 10 private databases, and you increase this parameter and the following parameter from 1 to 3 (a total increase of 4 threads), you actually creates 4 x 10 = 40 threads, which together consume 40 MB of virtual memory.

Table 4.17 Gateway Out Threads setting

Parameter	Gateway Out Threads
Type	REG_DWORD
Explanation	Defines the number of threads available for sending messages from the Store to the MTA process.
Default setting	Not present, but defaults to 0x00000001.
When to change	Adjust if this MTA communicates with multiple Exchange 5.5 servers or other messaging connectors.
Recommended setting	0x00000003

You should add these parameters to all private databases configured on the server. After you make this change, you must increase the **Max RPC Calls Outstanding** registry parameter for the MTA process as discussed in the "MSExchangeMTA\Parameters Registry Key Settings" section earlier in this chapter. For more information, see the Microsoft Knowledge Base article 264075, "Performance Tuning for Microsoft Exchange" (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=264075).

MTA File Directories

By default, the Exchange MTA database and run directories are located under the path where Exchange 2003 is installed (*drive*:\Program Files\Exchsrvr\mtadata). On some servers, especially where Exchange is acting as a bridgehead server, you improve performance by relocating the MTA database on a fast disk array, such as a RAID0+1 partition. Do not attempt to relocate the MTA run directory because doing so can cause unwanted side effects.

To relocate the MTA database, use the Exchange System Manager administration interface.

Routing Tuning

This section contains information about message routing on your servers.

Suppressing Link-State Changes

Exchange 2003 includes functionality to automatically detect changes in the state of a link. This information can be used to inform other servers running Exchange 2003 that an alternative route should be used instead of the lowest-cost primary route. Link state information is broken down into major and minor changes. A major change occurs when the administrator changes the routing topology, such as the addition of a new connector or a cost change. Minor updates occur when the system automatically detects the failure or restoration of a link.

This feature works well in small to medium sized organizations. However, in large multisite environments, mass network fluctuation can cause link update floods for the minor version. To be truly effective, link state data must be broadcast to all the servers in the organization. Additionally, when state changes, the whole link state table is rebroadcast, which can cause significant data to be transmitted over the network. In these scenarios, it may be useful to suppress minor link state changes. To disable minor link state changes, implement the registry parameter shown in Table 4.20.

Table 4.18 SuppressStateChanges setting

Location	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\RESvc\Parameters
Parameter	SuppressStateChanges
Type	REG_DWORD
Default setting	Not present, but defaults to 0x00000000.
When to change	Change to 0x00000001 to disable broadcast of minor link state changes.

In hub-and-spoke environments in which there is no alternative path between the spoke bridgehead and the hub, Exchange 2003 servers automatically suppresses link state changes for that connector. In this scenario, you do not have to enforce the registry parameter described in Table 4.20.

Glitch Retry

If Exchange 2003 tries to route a message to a heavily loaded external SMTP system such as a virus wall, it can receive a Server Busy error. In these situations, the Exchange transport goes into a state known as 'glitch retry'. In this state, Exchange 2003 waits 60 seconds before attempting to resend the message, and it repeats this process three times before resorting to other actions. If external SMTP servers are consistently busy, reduce the glitch retry wait time to prevent mass message queuing. Table 4.21 shows the registry parameter that controls this time.

Table 4.19 GlitchRetrySeconds setting

Location	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\SMTPSvc\Queuing
Parameter	GlitchRetrySeconds
Type	REG_DWORD
Default setting	Not present, but defaults to 0x0000003C (60)

Creating Routing Groups

After you create a routing group in Exchange System Manager, information about that group, contained servers, and any connectors associated with the group is broadcast through link state updates to all the other Exchange 2000 and Exchange 2003 servers in the organization. If the routing group is removed, the object becomes orphaned in the link state table; however, the data continues to be broadcast as part of the link state. Removing the group does not cause routing problems, but the link state table will be larger than it should be. The only way to permanently remove all orphaned routing groups from the link state table is to shut down the Exchange 2000 and Exchange 2003 servers in the organization at the same time.

Therefore, you should keep the creation and deletion of routing groups in a production Exchange environment to a minimum.

SMTP Transport Tuning

This section contains information about transport tuning on your servers. For Exchange 2000 customers, this section contains the same recommendations as previous Exchange 2000 performance documentation.

Mailroot Directory Location

In Exchange 2003, when messages arrive through SMTP, the data is written to a disk in the form of an NTFS file system file (specifically, an .eml file). By default, these files are written to a directory (<*drive*>:\Program Files\Exchsrvr\Mailroot) on the same disk partition where the Exchange 2003 binary files are installed.

In some scenarios, such as configuring a bridgehead or relay server, relocating the SMTP Mailroot directory to a faster disk partition may improve performance.

In Exchange 2003, you can use Exchange System Manager to move the Mailroot directory. To move the Mailroot directory, use the **Messages** tab in the **SMTP Virtual Server Properties** dialog box. For more information about how to move the Mailroot directory, see Exchange Server 2003 Help.

SMTP MaxMessageObjects

The MaxMessageObjects setting correlates to the number of messages that can be queued up at a particular time by SMTP. Each mail message resident in the SMTP queue uses at least 4 KB of memory; therefore, you can experience low memory with a very large queue. Lowering this setting reduces the maximum number of messages that can reside in the queue, therefore decreasing the maximum memory footprint for SMTP. After this limit is reached, each SMTP connection made to the server returns with an out-of-memory error. For example, if this value is reduced to 10,000, SMTP refuses inbound mail after the queue reaches 10,000 messages.

Table 4.20 MaxMessageObjects setting

Location	HKEY_LOCAL_MACHINE\Software\Microsoft\Exchange\Mailmsg
Parameter	MaxMessageObjects
Type	REG_DWORD
Default setting	Not present, but defaults to 0x000186a0 (100000).
When to change	Adjust if the Exchange 2003 server is running out of memory because the number of incoming messages is too great for the server to process.

Active Directory Connector Tuning

The Active Directory Connector (ADC) requires almost no tuning during typical operation. There are two scenarios in which you might want to consider manually tuning the ADC process: during sleep time and block searching. For Exchange 2000 customers, this section contains the same recommendations as previous Exchange 2000 performance documentation.

Sleep Time

After the ADC has fully replicated Exchange and Active Directory data, it performs replication on the changes made to those directories. In most circumstances, those changes are small. During a connection agreement's activation time, the ADC is permitted to work continuously for 5 minutes. After that, the ADC sleeps for 5 minutes to allow other applications, such as replication processing time on domain controller or global catalog servers, to run. However, if a connection agreement is running for the first time or if many changes are made to one of the directories, you may want to permit the ADC to perform the replication without sleeping (and therefore speed up the replication cycle). You can configure the maximum time that the ADC is permitted to work without sleeping and the maximum time the ADC should sleep.

> **Note**
> These changes affect all the connection agreements running on the ADC server, and they may adversely affect other Active Directory applications.

Table 4.21 Max Continuous Sync (secs) setting

Location	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\MSADC\Parameters
Parameter	Max Continuous Sync (secs)
Type	REG_DWORD
Default setting	Not present, but defaults to 0x0000012c (300).
When to change	Adjust if you want the ADC to continue processing even after replicating solidly for 5 minutes.
Recommended setting	Set to no more than 20 minutes (1,200 seconds). Setting this value too high can have an adverse effect on Active Directory and other applications.

Table 4.22 Sync Sleep Delay (secs) setting

Location	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\MSADC\Parameters
Parameter	Sync Sleep Delay (secs)
Type	REG_DWORD
Default setting	Not present, but defaults to 0x0000012c (300).
When to change	Adjust if you want to change the default sleep time of the ADC.
Recommended setting	Set to a minimum of 1 minute (60 seconds). Whenever possible, leave this setting at the default. Be careful when raising this value. For example, if you set it to 3,600 seconds (1 hour) and set the connection agreement replication schedule for individual 15-minute segments, the ADC may never replicate. It is recommended that you only change the sleep delay when setting the replication schedule to **Always**.

For more information, see the Microsoft Knowledge Base article 253825, "XADM: How the Active Directory Connector Polling Period Works" (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=253825)

Block Searching

By default, the ADC requests changes from Exchange and Active Directory in blocks of 10,000. If more than 10,000 objects are to be replicated, the ADC requests the first 10,000 entries, processes them, and then prompts you for the next 10,000. If the ADC communicates with directory servers over an error-prone network, it may be useful to reduce the block size. If the ADC receives a partial block caused by a connection failure, the complete block must be replicated again. Reducing the block size reduces the number repeat replications caused by connection failures.

Table 4.23 Export Block Size setting

Location	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\MSADC\Parameters
Parameter	Export Block Size
Type	REG_DWORD
Default setting	Not present, but defaults to 0x00002710 (10000).
When to change	Adjust if you want the ADC to commit before 10,000 new or changed objects have been received. You would generally lower this value in difficult wide area network (WAN) situations.
Recommended setting	Set as appropriate. If you set the value too low (never go below 100), you receive poor performance from the ADC.

For more information, see the following Microsoft Knowledge Base articles:

- 253665 "XADM: How the Active Directory Connector Uses Block Search to Replicate Changes" (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=253665)
- 253840 "XADM: When the Active Directory Connector Commits Changes to Active Directory" (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=253840)

Active Directory Integration Tuning

This section contains information about tuning Active Directory components to optimize Exchange 2003.

Global Catalog-to-Exchange Ratios

Planning the number of global catalog servers required for your Exchange 2003 installation can be time-consuming. You should implement one global catalog processor for every four Exchange 2003 processors and then fine-tune as necessary. With this rule, it is assumed that the processors are of the same class and speed. For example:

- A single-processor global catalog server can handle the load of a single 4-processor Exchange 2003 server.
- Two single-processor global catalog servers can handle the load of a single 8-processor Exchange 2003 server.
- Four 4-processor global catalog servers can handle the load of eight 8-processor Exchange 2003 servers.

Exchange 2003 can use global catalog servers that have up to 8 processors installed. However, you must take into account the 75 percent scaling factor between 4 and 8 processors for Active Directory servers. Therefore, an 8-processor global catalog should be considered as a 7-processor server when performing global catalog to Exchange ratio calculations.

Dedicated Active Directory Servers for Exchange

If you have a high concentration of Exchange 2003 servers, you must dedicate a set of global catalog servers for Exchange. Create a dedicated Active Directory site that contains both the Exchange 2003 servers and any dedicated global catalog servers. There are several positive effects here:

- Traffic from systems other than Exchange 2003 is distributed to other Active Directory servers in the organization.
- Performance analysis and management of Active Directory is easier.
- The Exchange administrator has better control over the Active Directory servers that are dedicated for Exchange.

Set PDC Avoidance

Applications other than Exchange can make heavy use of the primary domain controller (PDC) emulator computer. If Exchange 2003 also tries to use the PDC emulator computer for its requests, the performance of the PDC emulator and Exchange can decrease. By default, DSAccess picks up the PDC emulator computer for request together with other servers in the local Active Directory site. However, you can edit the registry to change this behavior. Table 4.24 shows the registry parameter to set.

Table 4.24 MinUserDc setting

Location	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\MSExchangeDSAccess\Profiles\Default
Parameter	MinUserDc
Type	REG_DWORD
Default setting	Not present, but defaults to 0xffffffff (-1)
When to change	Set this value if the PDC emulator is located in the same Active Directory site as your Exchange 2003 servers. The value that you set relates to the minimum number of domain controllers that must be detected in the same site and same domain before the PDC emulator is excluded from the server list. For example, if you set this value to 3, DSAccess does not use the PDC for LDAP requests if it detects three or more domain controllers in the Active Directory site. If fewer than three domain controllers are detected, Exchange continues to use the PDC. To force PDC avoidance in all scenarios, set this registry value to 1.

Use of the /3GB Switch on Active Directory Servers

The Active Directory process uses the Extensible Storage Engine (ESE) for its database. By default, the ESE cache size is 512 MB. However, if you use Windows 2000 Advanced Server or Windows Server 2003 (any edition) for your global catalog servers, and you have more than 2 GB of physical RAM installed, you should set the **/3GB** switch in the Boot.ini file of these servers. This action automatically increases the ESE cache to 1024 MB. In most circumstances, having a larger cache reduces the number of disk reads by 20-40 percent and dramatically reduces LDAP response times.

Using Exchange 2003 on Active Directory Servers

For the best scalability and administrative flexibility, Exchange 2003 should be installed on a Windows member server instead of a domain controller or global catalog server. The latter scenario is supported, but you should be aware of the consequences:

Threads in the Lsass.exe (Active Directory) process run at a higher priority than Exchange threads. An increase in the Lsass.exe process can adversely affect Exchange processing time.

If the Exchange server also acts as a domain controller, the server spends resources on other non-Exchange requests, such as user authentication and directory lookups for other applications. This additional activity affects performance on the Exchange 2003 server.

Although DSAccess detects all domain controllers and global catalog servers in the local Active Directory site, it does not use them. All directory requests are sent to the local directory service. Load-balancing and fail-over do not occur in this scenario.

To manage Exchange 2003 services, the administrator must be defined as a local administrator. If Exchange 2003 is running on a domain controller, the Exchange administrator must belong to the Administrators group in the domain. Membership implicitly gives the Exchange administrator additional access to other computers in the domain.

> **Important**
> Administrators who have logon access to domain controllers must be trusted in the Active Directory forest, because it is possible for these users to elevate their permissions in Active Directory. Administrators who can only log on to member servers cannot elevate their permissions in Active Directory.

Tuning DSAccess on Mailbox Servers

By default, DSAccess caches objects. The cache is broken into two sections: one section for user objects (that is, domain naming context), and the other section for configuration data such as store and routing objects. User objects are cached for 5 minutes, and configuration data is cached for 15 minutes.

In Exchange 2000, each cache pool was 25 MB. For better performance, the default settings have changed. By default, the configuration data cache in Exchange 2003 is 5 MB and the user object is 140 MB.

In very large topologies containing more than 100 administrative or routing groups, increased performance can be achieved by manually tuning the DSAccess cache sections. Table 4.25 shows the registry parameter to set.

Table 4.25 MaxMemoryConfig setting

Location	HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\MSExchangeDSAccess\Instance0
Parameter	MaxMemoryConfig
Type	REG_DWORD
Default setting	Not present, but defaults to 0x00001400 (5 MB).
When to change	Adjust when there are more than 100 administrative or routing groups in the organization.
Recommended setting	0x0000028f5 (approximately10 MB)

Tuning DSAccess on Branch Office Servers

If you run Exchange 2003 in a large branch office environment, some manual performance adjustments may be necessary to get the best efficiency:

- You must always have LAN-speed access to a global catalog server.
- If there are hundreds of routing groups in the organization, you may notice a mass spike of activity on the local global catalog every 15 minutes. This spike occurs because DSAccess rereads the routing group configuration. In severe cases, the local global catalog may become consumed with this activity. To minimize the burden, you can index the Routing Group Back-Link property (msExchRoutingGroupMembersBL). Use Active Directory Schema Manager to enable the **Index this attribute in Active Directory** option for this property. For more information about Active Directory Schema Manager, see Windows 2000 Administration Tools, which is included on the Windows 2000 Server and Windows 2000 Advanced Server CD.

Increasing the Maximum Active LDAP Queries

If there are many Exchange 2003 servers in a Windows Active Directory site, a very large LDAP load can be put on the Active Directory servers. By default, an Active Directory server is configured to support a maximum of 20 active LDAP queries. If this limit is reached, Active Directory returns the LDAP_ADMIN_LIMIT_EXCEEDED error and stops processing any more LDAP queries. A setting of 20 is generally sufficient for most Active Directory servers, but it is necessary to increase this value when you are running Exchange 2003 on eight-processor servers, or if the LDAP_ADMIN_LIMIT_EXCEEDED error message is logged.

The maximum LDAP queries can be configured through the **MaxActiveQueries** attribute. This setting can be adjusted using the Ntdsutil.exe tool. Increasing this setting uses more memory in the Lsass.exe process on the Active Directory server. Therefore, do not increase this value any higher than is necessary.

To change the MaxActiveQueries setting using NTDSUTIL

1. Click **Start**, and then click **Run**.
2. Type **NTDSUTIL**, and then click **OK**.
3. Type **LDAP POLICIES**.
4. Type **CONNECTIONS**.
5. Type **CONNECT TO SERVER** *domain_controllor_or_global_catalog_name*.
6. Type **Q**.
7. Type **SHOW VALUES**.
8. Type **SET MAXACTIVEQUERIES TO 40**.
9. Type **COMMIT CHANGES**.
10. Type **SHOW VALUES**.
11. Verify that the new setting is shown.
12. Type **Q**.
13. Type **Q**.

> **Note**
> This setting replicates to all Active Directory servers in the forest. You do not have to restart the domain controller or global catalog servers for this change to take effect.

Disabling Unused Services

For additional performance and security, you may decide to disable certain services from running. For example, although Exchange 2003 requires the Network News Transfer Protocol (NNTP) stack to be installed before Exchange 2003 is installed, if you have no requirements for NNTP access, you can shut down and disable this service after Exchange 2003 is installed. The same is true for POP3 and IMAP4 services.

Do not disable services that Exchange 2003 requires for normal operation. For example, the MTA service is required for internal operations even when working in a native mode Exchange organization. Therefore, you should never disable the MTA or stop it for any long period of time.

Obsolete Settings

If a server previously ran Exchange 2000, you may have performed the manual tuning changes that were recommended for Exchange 2000. If you later upgraded that server to run Exchange 2003, it no longer requires these manual tuning changes. You should manually remove them from the server. For more information about the settings that must be removed, see "Removing Exchange 2000 Tuning Parameters" in the *Exchange Server 2003 Deployment Guide* (http://go.microsoft.com/fwlink/?LinkId=21768).

Appendixes

APPENDIX A

Exchange Performance Tools

This appendix describes the tools that can help you verify the performance of your Microsoft® Exchange Server 2003 environment. Some of these tools are installed with Microsoft Windows® 2000 Server or Windows Server™ 2003, some with Exchange, and others can be found at the "Downloads for Exchange 2003" Web site (http://go.microsoft.com/fwlink/?LinkId=25097). Table A.1 lists specific performance-related tools.

> **Warning**
> Some tools can cause serious, sometimes irreversible, problems if they are used incorrectly. Before you use tools in your production environment, always familiarize yourself with them on test servers first. Be sure to read the documentation associated with any tool, and familiarize yourself with the risks involved.

Table A.1 Exchange performance tools

Tool name	Description	Run from	Install from
Exchange Stress and Performance (ESP)	Use to test stress and performance. This tool simulates large numbers of client sessions, by concurrently accessing one or more protocol servers.	Command prompt	2003 version: http://go.microsoft.com/fwlink/?LinkId=27881 2000 version: http://go.microsoft.com/fwlink/?LinkId=1709
Jetstress	Use to test the performance and stability of the disk subsystem.	Command prompt	http://go.microsoft.com/fwlink/?LinkId=27883
Load Simulator (LoadSim)	Use as a benchmarking tool to test the response of servers to mail loads.	For setup and installation instructions, see http://go.microsoft.com/fwlink/?LinkID=1710.	2003 version: http://go.microsoft.com/fwlink/?LinkId=27882 2000 version: http://go.microsoft.com/fwlink/?LinkId=1710
Network Monitor	Use to diagnose issues with server connectivity.	Start ▶ All Programs Administrative Tools ▶ Network Monitor	Add/Remove Programs ▶ Add/Remove Windows Components

Tool name	Description	Run from	Install from
System Monitor (also known as Performance Monitor)	Use for establishing a baseline of performance and for troubleshooting performance issues.	Start ▸ All Programs ▸ Administrative Tools ▸ Performance	Installed during Windows setup.

LoadSim and ESP are helpful when you test systems to ensure the health of the systems before going into production. You can use the other tools to help diagnose bottlenecks in production servers.

Exchange Server Stress and Performance 2003

You can use Exchange Server Stress and Performance (ESP) 2003 to simulate arbitrary several client sessions that are concurrently accessing one or more Exchange 2003 servers.

ESP provides modules that simulate client sessions over the following Internet protocols and APIs:

- WebDAV (for Microsoft Office Outlook® Web Access)
- Internet Message Access Protocol version 4rev1 (IMAP4)
- Lightweight Directory Access Protocol (LDAP)
- OLE DB
- Network News Transfer Protocol (NNTP)
- Post Office Protocol version 3 (POP3)
- Simple Mail Transfer Protocol (SMTP)
- Exchange ActiveSync®
- Outlook Mobile Access

ESP is similar to LoadSim; however, use ESP when you are validating deployments that use mobility features and Internet protocols that LoadSim does not cover.

For more information about ESP, see the documentation that comes with the tool.

Jetstress

Jetstress helps administrators verify the performance and stability of the disk subsystem prior to putting their Exchange server into production.

Jetstress helps verify disk performance by simulating Exchange disk input/output (I/O) load. Jetstress simulates the Exchange database and log file loads produced by a specific number of users. You use System Monitor, Event Viewer, and Exchange Server Database Utilities together with Jetstress to verify that your disk subsystem meets or exceeds the performance criteria you establish.

With Jetstress, you can perform two types of tests:

- The Jetstress Disk Performance Test runs for two hours. You can verify the performance and sizing of your storage solution.
- The Jetstress Disk Subsystem Stress Test runs for 24 hours. You can test your server load using a much larger load over a more significant amount of time.

Running both tests is the best way to verify the integrity performance of your disk subsystem. After a successful completion of the Jetstress Disk Performance Test and Jetstress Disk Subsystem Stress Test in a nonproduction environment, you are ready for the next step in your Exchange 2003 deployment process. By running the tests, you help ensure that your Exchange 2003 disk subsystem is adequately sized (in terms of performance criteria that you establish) for the user count and user profiles you have established.

For more information about Jetstress, see the documentation that comes with the tool.

Load Simulator 2003

Load Simulator 2003 (LoadSim) simulates the performance load of MAPI clients. LoadSim helps you determine if each of your servers can handle the load that you intend for them to carry. Another use for LoadSim is to help you validate your deployment plan.

However, LoadSim does not account for all of the factors involved in sizing servers. LoadSim does not simulate the following factors that can affect your server capacity planning:

- Incoming unsolicited commercial e-mail (also known as spam) from the Internet
- Incoming SMTP mail flow from the Internet or other sites within your organization
- Use of non-MAPI protocols for account access, such as POP3 and IMAP4
- Use of mobile devices
- Public folder usage

In addition, LoadSim does not give a complete picture with regards to user experience, and its results should not be interpreted in that aspect.

For more information about LoadSim 2003, see the documentation that comes with the tool.

Network Monitor

With Network Monitor, you can detect and troubleshoot problems on LANs. Using Network Monitor, you can:

- Identify network traffic patterns and network problems. For example, you can locate client-to-server connection problems, find a computer that makes a disproportionate number of work requests, and identify unauthorized users on your network.
- Capture frames (packets) directly from the network.
- Display, filter, save, and print the captured frames.

For more information about Network Monitor, see the following Microsoft Knowledge Base articles:

- 294818, "Frequently Asked Questions About Network Monitor" (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=294818)
- 148942, "How to Capture Network Traffic with Network Monitor" (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=148942)

System Monitor

System Monitor is a Microsoft Management Console (MMC) snap-in. You can use it to monitor a wide range of subsystems and software. It provides a common infrastructure for reporting data based upon performance counters. These counters are organized hierarchically by object, counter, and (optionally) instance, as follows:

- **Performance Object** The part of the computer you can monitor. Some of the most commonly used objects are **Processor**, **Memory**, and **PhysicalDisk**. When you install Exchange 2003, new objects, such as the Microsoft Exchange Information Store (MSExchangeIS), are added to the performance object list.
- **Counters** The parts of the object you can monitor. For example, you can monitor the available bytes, kilobytes, and megabytes of memory in addition to the page faults per second or total pages per second for a Memory object.
- **Instances (optional)** Multiple objects or counters to monitor on the computer. For example, when you look at counters under the **Processo**r object on a multiple processor computer, you see as many instances as there are processors on that computer. You can choose to monitor only a specific processor or all processors.

For more information about System Monitor, see the Windows Help.

A P P E N D I X B

Performance Counter Definitions

Available Mbytes
Displays the amount of physical memory, in bytes, available to processes running on the computer.

Bytes Total/sec
The total rate of bytes transferred by the Web service. This counter is the sum of Bytes Sent/sec and Bytes Received/sec.

Client Latency
The latency of MAPI/remote procedure call (RPC) actions measured at the LoadSim/Microsoft® Office Outlook® client. This counter measures the time it takes for the server to fulfill the client request. It can be used to estimate the time a user would have to wait between initiating individual Outlook actions.

Database\Database Cache Size
The average amount of system memory used by the database cache manager to hold commonly used information from the database files to prevent file operations. If the database cache size seems too small for optimal performance and there is very little available memory on the system (see Memory/Available Mbytes), adding more memory to the system may increase performance. If there is a lot of available memory on the system and the database cache size is not growing beyond a certain point, the database cache size may be restricted to an artificially low limit. Increasing this limit may increase performance.

DB Disk Transfers/sec
The average sum of all random read/write input/output (I/O) operations to the Microsoft Exchange Database disk volumes (both .edb and .stm files).

Disk Bytes/sec
The average number of disk bytes written or read per second across all disk volumes.

IMAP4 Connections
The number of current Internet Message Access Protocol version 4rev1 (IMAP4) client connections.

IMAP4 UID/sec
The number of unique identifier (UID) commands per second.

ISAPI Extension Requests/sec
The number of requests per second for Outlook Web Access transactions.

Log Writes/sec
The average sum of all sequential write I/O operations to the Exchange log file disk volumes (.log files).

MSExchangeIS Mailbox\Local Delivery Rate
The average rate at which messages are delivered locally to the Exchange store.

MSExchangeIS\RPC Operations/sec
The rate at which RPC operations occur. This counter is a good rate counter to measure Exchange workload because all MAPI-based actions use the RPC protocol.

MSExchangeIS\RPC Requests
The number of client requests that are currently being processed by the Exchange store.

Network Interface\Bytes Total/sec
The average rate at which bytes are sent and received over each network adapter, including framing characters. Network Interface\Bytes Total/sec is the sum of Network Interface\Bytes Received/sec and Network Interface\Bytes Sent/sec.

Network Usage
Measures network traffic on the server going to and from the server's network adapter.

POP3 DELE/sec
The number of message delete commands per second.

POP3 STAT/sec
The number of STAT commands per second. A STAT command is issued once per each user's connection.

Private Bytes
Displays the current number of bytes this process has allocated that cannot be shared with other processes.

Processor\% Processor Time
The average percentage of elapsed time that the processor spends to execute a non-idle thread. It is calculated by measuring the duration of time the idle thread is active in the sample interval, and subtracting that time from the interval duration. (Each processor has an idle thread that consumes cycles when no other threads are ready to run.) This counter is the primary indicator of processor activity, and it displays the average percentage of busy time observed during the sample interval.

SMTP Local Queue
The number of messages in the local queue waiting delivery to local users.

SMTP Messages Del/sec
The number of messages being delivered each second to local users.

SMTP Messages Sent/sec
The number of messages being sent each second to a remote server.

Store Virtual Bytes
The average size, in bytes, of the virtual address space that the Store.exe process is using. Use of virtual address space does not necessarily imply corresponding use of either disk or main memory pages. Virtual space is finite, and the process can limit its ability to load libraries.

System\Context Switches/sec
The combined average rate at which all processors in the computer are switched from one thread to another. Context switches occur when a running thread voluntarily relinquishes the processor, is preempted by a higher priority ready thread, or switches between user-mode and privileged (kernel) mode to use an Executive or subsystem service. This counter is the sum of Thread\Context Switches/sec for all threads running on all processors in the computer, and it is measured in numbers of switches. There are context switch counters on the **System** and **Thread** objects. This counter displays the difference between the values observed in the last two samples, divided by the duration of the sample interval.

Web ISAPI Extension Requests/sec
The rate at which Internet Server Application Programming Interface (ISAPI) extension requests are received by the Web service. Internet server API requests are used by Outlook Web Access to access the Exchange server.

Working Set
The set of memory pages (areas of memory allocated to a process) recently used by the threads in a process. If available memory on the server is above a specified threshold, pages remain in the Working Set of a process even if they are not in use. When available memory falls below a specified threshold, pages are removed from the Working Set. If these pages are needed, they will be returned back to the Working Set before they leave main memory and are made available for other processes to use.

APPENDIX C

Calculate Your Server Size

This appendix provides you with a method that shows how to determine your server sizing requirements, specifically focusing on the hardware that is needed to support a group of users. Because of the wide variety of Microsoft® Exchange configurations and user profiles, it is difficult to determine accurately the number of users supported by a server. You must consider the different types of clients, how active the users are, the capacity of the storage subsystem, and how the Exchange server is configured to use the disk resources.

Use the following steps to help you assess these issues and determine what hardware you require:

1. Determine your usage profile.
2. Select a server based on your usage profile.
3. Validate your disk subsystem capacity.

> **Note**
> The method described in this section also applies to Exchange 2000 Server. With Exchange Server 2003, you will see slightly lower user loads and significantly better use of memory than with Exchange 2000. Exchange 2003 uses approximately 10 percent fewer disk resources than Exchange 2000 for the same user profile. If you upgrade and select a new server, include this adjustment in your estimate.

Determine Your Usage Profile

To calculate how many users a server can support, you must first determine your current usage profile. You can calculate a usage profile by using the following two key metrics together:

- **Megacycles/mailbox** Megacycles per second, per mailbox. The raw processor usage required per mailbox that is measured during the peak two-hour period on a production server. For example, if a user uses one megacycle/second during peak operation and there are 1,000 users on the server (1,000 megacycles/second), a single 2,000-MHz processor operates at 50 percent CPU usage.

  > **Note**
  > The actual units used in this measurement are megacycles per second, per mailbox. For brevity, "per second" is omitted in this section.

- **IOPS/mailbox** Input/output per second, per mailbox. The raw database (DB) disk usage (input/output per second) required per user that is measured during the peak two-hour period on a production server. This metric does not include transaction log input/output (I/O) operations. For example, if each mailbox uses 0.5 DB IOPS during peak activity and there are 1,000 users homed on the server, there are 500 DB IOPS. IOPS/mailbox metrics are based on random read/write Exchange database I/O operations.

  > **Note**
  > The actual units on this measurement is IOPS per second, per mailbox. For brevity, "per second" is omitted in this section.

Usage profiles are based on production data that may include third-party applications (for example, BlackBerry and Trend) in addition to Microsoft Outlook®. The recommendations in this section are not specific to any particular client or client version. When you calculate megacycles/mailbox and IOPS/mailbox, use the current number of mailboxes on that server. If the server contains many unused mailboxes or runs other applications that do not add much load during the peak two hours, your results will not represent a typical user load.

Choose a server that has typical user mailboxes for your measurements, or do not include the unused mailboxes in your calculation.

Be aware that different days of the week have slightly different usage loads. For example, in many companies, Mondays have a heavier load than other days of the week. A good time to measure typical peak activity is between 08:00 and 10:00 locally on a Monday.

To calculate megacycles/mailbox

1. Select a production server with a typical user load.
2. Use the System Monitor tool to monitor the **Processor\% Processor Time** counter over the peak two hours of server activity.
3. Calculate the average CPU usage (percent) from the data you obtain in Step 2.
4. Calculate your current megacycles/mailbox as described in the following formula:

   megacycles/mailbox = ((average CPU usage) × (speed of processors in megacycles) × (number of processors)) ÷ (number of mailboxes)

If the server runs other processes that consume significant server resources, consider using the **Process\% Processor Time** counter for the Store.exe process instead of total CPU usage. Because there are many factors that affect CPU usage that are nonlinear (such as the effect of memory caches and how servers scale with the number of CPUs), use this calculation as a guideline to determine your processing needs. The actual amount of processing needs depends on how much your final hardware differs from the hardware that you use currently for the measurement.

To measure IOPS/mailbox

1. Select a production server with a typical user load.
2. Use the System Monitor tool to monitor **Physical Disk\Disk Transfers/sec** counter over the peak 2 hours of server activity.
3. Calculate your current IOPS/mailbox as described in the following formula:

   IOPS/mailbox = (*average disk transfer/sec*) ÷ (*number of mailboxes*)

> **Note**
> If the users in a company are diverse usage requirements, you may have to measure usage profiles separately for different groups of users. For example, the sales engineers may have a different usage profile than the local marketing group. Having separate measurements is helpful only if the groups of users are significantly different.

Select a Server Based on Your Usage Profile

After you have determined your usage profile (megacycles/mailbox and IOPS/mailbox), you can calculate your CPU and disk subsystem requirements.

The following sections provide four example usage profiles, and a sample server hardware recommendation is also given. You can compare your user's profile with the sample profiles, determine which profile best matches the needs of your company, and use the recommended hardware as a guideline. For example, if you have both heavy users and light users, use the guidelines for the heavy users.

Each guideline below is specific to a usage profile and server/storage area network configuration. The Hewlett Packard StorageWorks Enterprise Virtual Array or CLARiion FC-4500 storage area network (SAN) is used in these examples, but any SAN that provides the same disk throughput will work. After you select appropriate hardware, validate that the disk subsystem meets your needs. For more information, see "Validate Your Disk Subsystem Capacity" later in this appendix.

It is recommended that you use a 4-processor (2.8 GHz) server for high-end server configurations. The recommended hardware does not take into account other performance factors, such as network capacity, server memory, and cache sizes. However, by using the example usage profiles, you can estimate if a server has sufficient CPU and disk capacity.

To calculate CPU requirements

- To calculate the maximum number of users your processor can support, multiply the total processor megacycles by 0.80 and divide by your current megacycles per mailbox.

  $\mathbf{N_{max\ users}}$ = 0.80 × (*number of processors*) × (*speed of processors in megacycles*) ÷ (*current megacycles per mailbox*)

- Similarly, to determine if your server has sufficient megacycles for the number of users you want to put on the server, multiply the number of users by your current megacycles/mailbox, and then divide that result by the megacycles of the server (the number of processors multiplied by the speed in megacycles), shown below. If the value is less than 0.80, the server can support the number of users. Do not exceed 80 percent usage, or the server will be overloaded.

  CPU Usage = ((*number of users*) × (*current megacycles per mailbox*)) ÷ ((*number of processors*) × (*speed of processors in megacycles*))

To calculate disk subsystem requirements

- To calculate if the disk subsystem is adequate, multiply the number of users that will be on the server by your current IOPS/mailbox, and then divide that result by the total IOPS/second capacity that the disks provide. Your disk usage should be under 0.80, or the server will be overloaded.

  Disk Usage = ((*number of users*) × (*current IOPS per mailbox*)) ÷ (*total IOPS/sec*)

- It is recommended that you calculate disk usage separately for each database. If all your databases in a storage group share a single logical unit number (LUN), you can calculate per storage group instead of per database, as shown below.

  Disk Usage (per database) = ((*number of users in the database*) × (*current IOPS per mailbox*)) ÷ (*total IOPS/sec*)

Example Usage Profiles

This section provides example usage profiles and recommended hardware for each profile. Use the information you collected in the previous section to determine the example profile that most closely matches your current requirements.

Depending on your mailbox size and user activity, your IOPS/mailbox measurement may be significantly higher or lower than the examples listed in the following sections. For example, one company had user profiles with 4 IOPS/mailbox. The large IOPS/mailbox were mostly because the users had no mailbox quotas (typical mailbox sizes were 1 to 10 GB). The users also sent messages with large attachments (attachment limits had been raised to 25 MB).

Heavy Knowledge Worker Profile

A Heavy Knowledge Worker (HKW) is a very intense knowledge worker profile. Users who fit this profile have jobs that depend heavily on email. Users may have Cached Exchange Mode clients. With this profile, you can expect the following usage load:

- Megacycles/mailbox: approximately 2.5
- IOPS/mailbox: approximately 0.75

Table C.1 Sample large capacity server for an HKW profile

Server hardware	4 processor, 1,996 MHz (Hyper-Threading), 4 GB RAM
Storage area network hardware	Hewlett Packard StorageWorks Enterprise Virtual Array 4 storage groups, 5 databases per storage group, spread across 48 disk spindles using RAID0+1
Mailboxes per storage group	1,150
Mailboxes per server	4,600
Peak processor usage	80%
Peak disk usage	84%

In this sample configuration, the server can support 5,100 HKW users. With 5,100 HKW users on the server, the peak processor usage is 80 percent, which leaves sufficient overhead for periods of extremely high load.

It is estimated that 30 spindles can handle 4,800 IOPS/sec (assuming disks support 100 IOPS/spindle). Therefore, with 4,600 users requiring 0.75 IOPS/mailbox, the peak disk usage for the database drives is 72 percent. If you consider that a RAID1 configuration requires two I/O operations for every write, the estimated throughput is reduced to 3,840 IOPS/sec (for an explanation of how this value was calculated, see "Estimating Disk Capacity" later in this section). The actual peak disk usage shown in Table C.1 is slightly higher because it is based on a measurement of actual disk capacity instead of an estimate.

Medium Knowledge Worker Profile

A Medium Knowledge Worker (MKW) is an intense knowledge worker profile. Clients may be using BlackBerry or other roaming devices. Users who fit this profile have jobs that depend heavily on e-mail. With this profile, you can expect the following usage load:

- Megacycles/mailbox: approximately 1.9
- IOPS/mailbox: approximately 0.4

Table C.2 Sample large capacity server for an MKW profile

Server hardware	4 processor, 2,800 MHz, 4 GB RAM
Storage area network hardware	Hewlett Packard StorageWorks Enterprise Virtual Array 3 storage groups, 1 database per storage group, spread across 30 disk spindles using RAID0+1
Mailboxes per storage group	1,575
Mailboxes per server	4,725
Peak processor usage	80%
Peak disk usage	67%

In this sample configuration, the server can support 4,725 MKW users. With 4,725 MKW users on the server, the peak processor usage is 80 percent, which leaves sufficient overhead for periods of extremely high load.

It is estimated that 30 spindles can handle 3,000 IOPS/sec (assuming disks support 100 IOPS/spindle). Therefore, with 4,725 users requiring 0.4 IOPS/mailbox, the peak disk usage for the database drives is 63 percent. The actual peak disk usage shown in Table C.2 is slightly higher because it is based on a measurement of actual disk capacity instead of an estimate.

Light Knowledge Worker Profile

A Light Knowledge Worker (LKW) is a light knowledge worker profile. Users who fit this profile typically have small mailbox quotas. With this profile, you can expect the following usage load:

- Megacycles/mailbox: approximately 0.75
- IOPS/mailbox: approximately 0.18

Table C.3 Sample large capacity server for an LKW profile

Server hardware	4 processor, 2,800 MHz, 4 GB RAM
Storage area network hardware	Hewlett Packard StorageWorks Enterprise Virtual Array 3 storage groups, 1 database per storage group, spread across 30 disk spindles using RAID0+1
Mailboxes per storage group	3,000
Mailboxes per server	9,000
Peak processor usage	76%
Peak disk usage	46%

Very Light Knowledge Worker Profile

A Very Light Knowledge Worker (VLKW) is a very light e-mail user. Users who fit this profile probably use Post Office Protocol version 3 (POP3) and have small mailbox quotas. With this profile, you can expect the following usage load:

- Megacycles/mailbox: approximately .33
- IOPS/mailbox: approximately .078

Table C.4 Sample large capacity server for a VLKW profile

Server hardware	4 processor, 2,000 MHz, 4 GB RAM
Storage area network hardware	CLARiion FC-4500 5 storage groups, 1 database per storage group, spread across 18 disk spindles using RAID0+1
Mailboxes/Storage Group	6,700
Mailboxes/Server	20,100
Peak Processor Usage	76%
Peak Disk Usage	46%

Validate Your Disk Subsystem Capacity

The final step in determining server sizing is to validate your disk subsystem capacity. After you select a disk subsystem, you should test the throughput of the hardware to make sure it meets your requirements. You can use the Jetstress tool, which Microsoft provides, to measure the performance of your disk subsystem. Jetstress generates stress loads that simulate Exchange Database read/write load. When you run the tool, load each SAN with the maximum number of IOPS (I/O operations per second) without exceeding 20 ms of read or write latencies. For more information about Jetstress, see Appendix A, "Exchange Performance Tools."

Many different types of disk subsystems can service an Exchange messaging deployment. The sample subsystem discussed in the following section is an example, and it is not intended to be a recommendation of a particular storage subsystem. Whatever disk subsystem you select, you should first verify through testing that the subsystem meets your requirements.

Sample Test Results on a Fibre Channel SAN

The data in Table C.5 shows the test results obtained when testing the maximum sustainable throughput on a Fibre Channel SAN. The test was conducted in a lab environment.

Table C.5 Jetstress SAN test

Function	Log	Database
Storage group configuration	6 disk with RAID0+1	6 disk with RAID0+1
Disk write latency (ms)	3	10
Disk read latency (ms)	0	20
Disk transfers/sec	135	430
Disk reads/sec	0	285
Disk writes/sec	135	145
IOPS/spindle	Not applicable	71.7

In this example, Jetstress testing sustained a maximum rate of 430 IOPS per storage group database. The Jetstress test was run with the following parameters:

```
jetstress -l L:\logfile_location -Z -A -I 50 -D 50 -R 0 -N 0
```

> **Note**
> If Exchange logical units share spindles with other nonmessaging applications or servers, performance may decline. Exchange performs best when disks are dedicated to the Exchange server. If Exchange is sharing spindles, the actual performance may be worse than the performance observed during lab tests.

Based on the measured throughput that the Jetstress test reveals, you can determine how many users your disk subsystem can support. For example, in this scenario, the storage area network can support 1,075 HKW mailboxes.

Estimating Disk Capacity

If you want to estimate disk capacity, a good guideline is to expect about 100 IOPS/sec for each spindle (this assumes 10,000 rpm). Depending on your disk configuration, you may need to make adjustments. For the Exchange database disks, a reasonable ratio of disk reads to disk writes is 3:1. However, you may want to measure the ratio yourself for your users. Assuming a ratio of 3:1, Table C.6 shows the throughput estimates for the RAID0, RAID1, RAID0+1, and RAID5 configurations.

Table C.6 Estimated RAID throughput per spindle

Raid configuration	Estimated IOPS/second per spindle
Raid0	100
Raid1	80
Raid0+1	80
Raid5	57

These calculations estimate that 48 disks striped together are capable of 3,840 IOPS/sec. Similarly, 5 disks in a RAID5 configuration are capable of 285 IOPS/sec.

In a RAID0 configuration, each read and each write generates one I/O operation. In the RAID1 and RAID0+1 configurations, each read generates one I/O operation, but each write requires two I/O operations (a write to each mirrored disk). In RAID5, each write requires four I/O operations: two reads to calculate parity and two writes (one for data, one to write parity). Therefore, the initial number of reads and writes is expanded for RAID1, RAID0+1, and RAID5. For an example of the increase in I/O operations, see "Sample Calculation" later in this appendix. In terms of the initial number of reads and writes, the apparent throughput is diminished.

Sample Calculation

Table C.7 shows how many I/O operations are required for 300 read I/O operations and 100 write I/O operations for each RAID configuration.

Table C.7 Sample RAID I/O operation performance

RAID configuration	Number reads and writes	Total I/O operations
RAID0	1 read + 1 write	400
RAID1	1 read + (2 × write)	500
RAID0+1	1 read + (2 × write)	500
RAID5	1 read + (4 × write)	700

In this example, you can see that 400 transactions (300 reads, 100 writes) produces 500 I/O operations in a RAID1 configuration. The apparent throughput is reduced by the ratio of 400/500, or 0.8. Therefore, instead of estimating 100 IOPS per spindle for RAID0, a better estimate is 80 IOPS per spindle.

Summary

The three steps to take when sizing a server are:

- Determine the usage profile.
- Select hardware and calculate if the hardware CPU and disk is adequate for the usage profile.
- Validate the performance of the disk subsystem.

Usage profiles can vary over time; therefore, you must monitor your servers regularly to maintain overall good performance and proper load.

A P P E N D I X D

Resources

Resources Cited in This Guide

The following resources have been cited in this guide.

Exchange Server 2003 Guides

- Exchange Server 2003 High Availability Guide (http://go.microsoft.com/fwlink/?LinkId=21277)
- Troubleshooting Exchange Server 2003 Performance (http://go.microsoft.com/fwlink/?LinkId=22811)
- Planning an Exchange Server 2003 Messaging System (http://go.microsoft.com/fwlink/?LinkId=21766)
- Exchange Server 2003 RPC over HTTP Deployment Scenarios (http://go.microsoft.com/fwlink/?LinkId=24823)
- Exchange Server 2003 Client Access Guide (http://go.microsoft.com/fwlink/?LinkId=27739)
- What's New in Exchange Server 2003 (http://go.microsoft.com/fwlink/?LinkId=21765)
- Server Consolidation Using Exchange Server 2003 (http://go.microsoft.com/fwlink/?LinkId=25209)
- Exchange Server 2003 Deployment Guide (http://go.microsoft.com/fwlink/?LinkId=21768)

Microsoft Knowledge Base Articles

The following Microsoft® Knowledge Base articles are available on the Web at http://go.microsoft.com/fwlink/?LinkId=14898

- XGEN: Exchange 2000 Server SMTP Optimized with Maximum Handle Threshold Registry Key (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=271084)
- XGEN: Setting a Limit on the Number of SMTP Messages in Queues (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=258748)
- XADM: How to Move Exchange Databases and Logs in Exchange 2000 Server (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=257184)
- XGEN: Exchange 2000 Requires /3GB Switch with More Than 1 Gigabyte of Physical RAM (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=266096)
- XADM: Using the /Userva Switch on Windows Server 2003-Based Exchange Servers (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=810371)

- XADM: The "HeapDecommitFreeBlockThreshold" Registry Key (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=315407)
- XADM: An Exchange 2000 Server with the '/3GB' Switch in the Boot.ini File May Lose Network Connectivity Under a Heavy Messaging Load (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=313707)
- XADM: Client Latencies Occur When Exchange 2000 Converts Mail from MAPI to MIME Format (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=317722)
- How to Optimize Memory Usage in Exchange Server 2003 (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=815372)
- Jet Database Does Not Work Correctly During a Backup or During a Content Indexing Operation (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=822894)
- Performance Tuning for Microsoft Exchange (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=264075)
- XADM: How the Active Directory Connector Polling Period Works (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=253825)
- XADM: How the Active Directory Connector Uses Block Search to Replicate Changes (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=253665)
- XADM: When the Active Directory Connector Commits Changes to Active Directory (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=253840)
- Frequently Asked Questions About Network Monitor (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=294818)
- How to Capture Network Traffic with Network Monitor (http://go.microsoft.com/fwlink/?LinkId=3052&kbid=148942)

Other Web Sites

- Products Designed for Microsoft Windows (http://go.microsoft.com/fwlink/?LinkId=26219)
- Windows 2000 Resource Kits (http://go.microsoft.com/fwlink/?LinkId=6545)
- Client Network Traffic with Exchange 2003 (http://go.microsoft.com/fwlink/?LinkId=27020)
- Microsoft Operation Manager: Be Accountable (http://go.microsoft.com/fwlink/?LinkId=16198)
- Better Together: Microsoft Operations Monitor and Exchange 2003 (http://go.microsoft.com/fwlink/?LinkId=18176)
- Downloads for Exchange 2003 (http://go.microsoft.com/fwlink/?LinkId=25097)
- Exchange 2003: Exchange Server Stress and Performance 2003 (http://go.microsoft.com/fwlink/?LinkId=27881)
- Exchange Stress and Performance Tool (ESP) (http://go.microsoft.com/fwlink/?LinkId=1709)
- Exchange 2003: Load Simulator 2003 (LoadSim) (http://go.microsoft.com/fwlink/?LinkId=27882)
- Exchange 2003: Jetstress (http://go.microsoft.com/fwlink/?LinkId=27883)
- Load Simulator (http://go.microsoft.com/fwlink/?LinkId=1710)

Additional Resources

Besides the resources cited in this guide, you may find the following resources useful in your implementation of Exchange Server 2003.

Web Sites

- Exchange Server 2003 Technical Library (http://go.microsoft.com/fwlink/?LinkId=21277)
- Exchange Server 2003 Tools and Updates (http://go.microsoft.com/fwlink/?LinkId=25097)
- Microsoft Developer Network (MSDN®) (http://go.microsoft.com/fwlink/?LinkId=21574)
- Exchange Server 2003 MAPI Messaging Benchmark 3 (MMB3) (http://go.microsoft.com/fwlink/?LinkId=27675)

Exchange Server 2003 Guides

- Exchange Server 2003 Administration Guide (http://go.microsoft.com/fwlink/?LinkId=21769)
- Exchange Server 2003 Glossary (http://go.microsoft.com/fwlink/?LinkId=24625)

Resource Kits

- Microsoft Exchange 2000 Server Resource Kit (http://go.microsoft.com/fwlink/?LinkId=6543)

> **Note**
> You can order a copy of the *Microsoft Exchange 2000 Server Resource Kit* from Microsoft Press® at http://go.microsoft.com/fwlink/?LinkId=6544.

APPENDIX E

Accessibility for People with Disabilities

Microsoft is committed to making its products and services easy for everyone to use. This appendix provides information about features, products, and services that make the Microsoft Windows Server™ 2003 family, the Windows® 2000 Server family, Microsoft Exchange Server 2003, and Microsoft Office Outlook® Web Access 2003 more accessible for people with disabilities. The following topics are covered:

- Accessibility in Microsoft Windows
- Adjusting Microsoft products for people with accessibility needs
- Microsoft product documentation in alternative formats
- Microsoft services for people who are deaf or hard-of-hearing
- Specific information about Exchange 2003 and Outlook Web Access 2003
- Other information resources for people with disabilities

> **Note**
> The information in this appendix applies only if you acquired Microsoft products in the United States. If you acquired Windows outside the United States, your package contains a subsidiary information card listing Microsoft support services telephone numbers and addresses. Contact your subsidiary to find out whether the type of products and services described in this appendix are available in your area. See the International Microsoft Accessibility Site (http://go.microsoft.com/fwlink/?LinkId=22008) for information available in the following languages: Chinese, English, French, Italian, Japanese, Portuguese, Spanish (Latin America), and Spanish (Spain).

Accessibility in Microsoft Windows

Many accessibility features have been built into the Windows operating system, starting with the introduction of Windows 95. These features are useful for individuals who have difficulty typing or using a mouse, are blind or have low vision, or who are deaf or hard-of-hearing. The features can be installed during setup.

For more information about the accessibility features of the various Windows operating systems, go to the Microsoft Products Accessibility Web site (http://go.microsoft.com/fwlink/?LinkId=22010).

Accessibility Files to Download

If you have a modem or another type of network connection, you can download accessibility files from the following network services:

- The Microsoft Accessibility Web site at http://go.microsoft.com/fwlink/?LinkId=21487.
- The Microsoft Help and Support Web site at http://go.microsoft.com/fwlink/?LinkId=14898. Select the **Knowledge Base Article ID Number Search** option, type **165486**, and then click the arrow. The search displays the Knowledge Base article, "Customizing Windows for Individuals with Disabilities," which includes links to documents about customizing various versions of Microsoft Windows.

  For other accessibility articles, from the Microsoft Help and Support Web site, select the **Search the Knowledge Base** option, select **All Microsoft Products**, and in **Search for**, type **kbenable**, and then click **Go**.

- Microsoft Internet server at ftp://ftp.microsoft.com/, in softlib/MSLFILES.
- Microsoft Download Service (MSDL), which you can reach by dialing (425) 936-6735 in the United States or (905) 507-3022 in Canada. Direct modem access to MSDL is available 24 hours a day, 365 days a year. Outside the United States and Canada, contact your local Microsoft subsidiary for information.

> **Note**
> MSDL supports 1200, 2400, 9600, or 14,400 baud data transmission with no parity, 8 data bits, and 1 stop bit. MSDL does not support 28.8 Kbps, 56K, or Integrated Digital Network (ISDN) connections.

Adjusting Microsoft Products for People with Accessibility Needs

Accessibility options and features are built into many Microsoft products, including the Windows operating system. Accessibility options and features are useful for individuals who have difficulty typing or using a mouse, are blind or have low vision, or who are deaf or hard-of-hearing.

Free Step-by-Step Tutorials

Microsoft offers a series of step-by-step tutorials to help you learn how to adjust the accessibility options and settings on your computer. The free tutorials provide detailed procedures on how to adjust options, features, and settings to meet your accessibility needs. Information related to the use of the mouse, the keyboard, or a combination of both is presented in a side-by-side format to help you learn.

Visit the Microsoft Accessibility Step by Step Tutorials Overview Web site (http://go.microsoft.com/fwlink/?LinkId=14899) to find the latest step-by-step tutorials.

Assistive Technology Products for Windows

A wide variety of assistive technology products are available to make computers easier to use for people with disabilities.

Microsoft provides a searchable catalog of assistive technology products that run on the Windows operating systems at the Microsoft Overview of Assistive Technology page (http://go.microsoft.com/fwlink/?LinkId=14901).

As an example, products available for the MS-DOS®, Windows, and Windows NT® operating systems are:

- Programs that describe information on the screen in Braille, or that provide synthesized speech for people who are blind or have difficulty reading.
- Hardware and software tools that modify the behavior of the mouse and keyboard.
- Programs that enable people to type by using a mouse or their voice.
- Word or phrase prediction software that people can use to type more quickly and with fewer keystrokes.
- Alternative input devices, such as single switch or puff-and-sip devices, for people who cannot use a mouse or a keyboard.

Upgrading an Assistive Technology Product

If you use an assistive technology product, be sure to contact your assistive technology vendor to check compatibility with products on your computer before upgrading. Your assistive technology vendor can also help you learn how to adjust your settings to optimize compatibility with your version of Windows or other Microsoft products.

Microsoft Documentation in Alternative Formats

Documentation for many Microsoft products is available in several formats to make it more accessible. Exchange 2003 documents are available as Help on the CD included with the product and on the Exchange Web site at http://go.microsoft.com/fwlink/?LinkId=21573.

If you have difficulty reading or handling printed documentation, you can obtain many Microsoft publications from Recording for the Blind & Dyslexic, Inc. (RFB&D). RFB&D distributes these documents to registered, eligible members of their distribution service in a variety of formats, including audiocassettes and CDs. The RFB&D collection contains more than 90,000 titles, including Microsoft product documentation and books from Microsoft Press®. You can download many of the Microsoft books from the Accessible Documentation for Microsoft Products Web site (http://go.microsoft.com/fwlink/?LinkId=22007).

For more information, contact RFB&D at the following address or contact information:

Recording for the Blind & Dyslexic
20 Roszel Road
Princeton, NJ 08540
Phone from within the United States: (866) 732-3585
Phone from outside the United States and Canada: (609) 452-0606
Fax: (609) 987-8116
Web: http://www.rfbd.org/

Microsoft Services for People Who Are Deaf or Hard-of-Hearing

If you are deaf or hard-of-hearing, complete access to Microsoft product and customer services is available through a teletype/telecommunication device for the deaf(TTY/TDD) service.

Customer Service

Contact the Microsoft Sales Information Center on a TTY/TTD by dialing (800) 892-5234 between 06:30 and 17:30 Pacific Time [UTC-8, Coordinated Universal Time (Greenwich Mean Time)], Monday through Friday, excluding holidays.

Technical Assistance

For technical assistance in the United States, contact Microsoft Product Support Services on a TTY/TDD at (800) 892-5234 between 06:00 and 18:00 Pacific Time
(UTC-8), Monday through Friday, excluding holidays. In Canada, dial (905) 568-9641 between 8:00 and 20:00 Eastern Time (UTC-5), Monday through Friday, excluding holidays. Microsoft support services are subject to the prices, terms, and conditions in place at the time the service is used.

Exchange 2003

Section 508 of the Rehabilitation Act regulates how United States government agencies purchase electronic and information technology. It requires procurement officials to purchase only electronic and information technologies that are accessible to people with disabilities. Section 508 states that any "electronic and information technology" developed, procured, maintained, or used by federal agencies must be accessible to people with disabilities, including employees and members of the public, unless an undue burden would be imposed on the agency.

To view the Exchange 2003 VPAT (Voluntary Product Accessibility Template), which describes the accessibility features that address the Section 508 standards, go to http://go.microsoft.com/fwlink/?LinkId=22011.

Outlook Web Access

For customers who require assistive technology devices to interact with software applications, it is recommended that they use the Basic Outlook Web Access client. By default, the Basic client renders in all browsers except Microsoft Internet Explorer 5.01 to 6.*x*. However, an Exchange administrator can provide users of Internet Explorer 5.01 to 6.*x* with the option to choose the Basic client when logging on to Outlook Web Access. To do this, the administrator must use Exchange System Manager to enable forms-based authentication for Outlook Web Access. For details on enabling forms-based authentication, see the *Exchange Server 2003 Administration Guide* (http://go.microsoft.com/fwlink/?LinkId=21769).

Administrators also have the option of setting the Basic client as the default client for all browsers. For more information about this approach, see the Microsoft Knowledge Base Article 296232, "XCCC: Empty Inbox When Using Internet Explorer 5 and Later to Gain Access to OWA" (http://go.microsoft.com/fwlink/?LinkId=14919).

Getting More Accessibility Information

The Microsoft Accessibility Web site (http://go.microsoft.com/fwlink/?LinkId=21487) provides information for people with disabilities, their friends and family members, people in outreach organizations, educators, and advocates.

A free monthly electronic newsletter is available to help you keep up-to-date with accessibility topics about Microsoft products. To subscribe, visit the Accessibility Update subscription page (http://go.microsoft.com/fwlink/?LinkId=14920).

Does this book help you? Give us your feedback. On a scale of 1 (poor) to 5 (excellent), how do you rate this book?

Mail feedback to exchdocs@microsoft.com.

For the latest information about Exchange, see the following Web pages:

- Exchange Product Team technical articles and books
  http://go.microsoft.com/fwlink/?LinkId=21277
- Exchange Tools and Updates
  http://go.microsoft.com/fwlink/?LinkId=25097
- Exchange Server Community
  http://go.microsoft.com/fwlink/?LinkId=14927
- Self-extracting executable containing all Exchange Product Team technical articles and books
  http://go.microsoft.com/fwlink/?LinkId=10687