

**Microsoft®**

# Troubleshooting Microsoft® Exchange 2000 Server Performance



Product Version:  
Latest Content:  
Author:

Exchange 2000 Server SP3  
[www.microsoft.com/exchange/library](http://www.microsoft.com/exchange/library)  
Dale Koetke

 Microsoft  
**Windows Server System**





# Troubleshooting Microsoft® Exchange 2000 Server Performance

**Published:** September 2002

**Updated:** May 2003

**Applies To:** Exchange 2000 Server SP3

**Contributing Writers:** Patricia Anderson, Teresa Appelgate, Susan Hill,  
Jon Hoerlein, Aaron Knopf, Jyoti Kulkarni, Michele Martin, Joey Masterson, John Speare,  
Randy Treit, Christopher Budd, Tammy Treit

**Project Editors:** Diane Forsyth, Susan Bradley

**Technical Reviewers:** KC Lemson, Jim Lucey, Nick Rosenfeld, Jason Hill, Michael Palermiti, Charles McDaniels, Sameer Patel, Scott Landry

**Graphic Design:** Kristie Smith

**Production:** Sean Pohilla

## Copyright

The information contained in this document represents the current view of Microsoft Corporation on the issues discussed as of the date of publication. Because Microsoft must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Microsoft, and Microsoft cannot guarantee the accuracy of any information presented after the date of publication.

This document is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS, IMPLIED OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in or introduced into a retrieval system, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Microsoft Corporation.

Microsoft may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Microsoft, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

Unless otherwise noted, the example companies, organizations, products, domain names, e-mail addresses, logos, people, places and events depicted herein are fictitious, and no association with any real company, organization, product, domain name, email address, logo, person, place or event is intended or should be inferred.

©2003 Microsoft Corporation. All rights reserved.

Microsoft, Active Directory, Outlook, Windows, and Windows NT are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

---

# Table of Contents

Table of Contents .....	i
i	
Introduction .....	<b>1</b>
What Is Updated in This Book? .....	1
Updated Chapters .....	1
What Will You Learn from This Book? .....	1
How Is This Book Structured? .....	2
Chapter 1	
Performance Troubleshooting Tools .....	<b>3</b>
System Monitor .....	3
Performance Logs and Alerts .....	4
Microsoft Operations Manager 2000 .....	5
Event Viewer .....	7
Network Monitor .....	8
File Monitor .....	9
Notations Used in This Book .....	9
Chapter 2	
Establishing a Baseline .....	<b>11</b>
Minimal Set of Counters .....	11
Example Baseline .....	12
Questions to Answer .....	12
System Monitor Examples .....	13
Chapter 3	
Troubleshooting Performance .....	<b>15</b>
Performance Problem Origins .....	15
Understanding the Problem .....	19
Root Cause Performance Analysis: Bottleneck Identification .....	20
CPU Performance Issues .....	20
Disk Performance Issues .....	23
Memory Problems .....	25
Monitoring Non-MAPI Requests .....	28
Message Delivery Counters .....	28
Active Directory .....	28
DSAccess .....	29
Network Problems .....	29
Performance Counters .....	33
Database Counters .....	34

Epoxy Counters.....	35
Logical Disk Counters .....	36
Memory Counters.....	37
MSExchangeIS Counters .....	39
MSExchangeIS Mailbox Counters .....	43
MSExchangeIS Public Counters .....	44
Network Interface Counters .....	45
Paging File Counters .....	46
Physical Disk Counters .....	46
Process Counters .....	48
Processor Counters.....	51
Server Counters .....	52
Server Work Queues Counters .....	53
SMTP Server Counters.....	54
System Counters .....	55
TCP Counters.....	55
Thread Counters.....	56
RAID Levels.....	57
Additional Resources .....	59
Web Sites.....	59
Technical Papers.....	59
Microsoft Knowledge Base Articles.....	60

# Introduction

This book introduces the tools, concepts, and recommendations you need in order to troubleshoot Microsoft® Exchange 2000 Server performance. It also describes how to monitor the health of your servers running Exchange 2000 and establish a baseline of normal server performance to measure against when troubleshooting performance.

---

## What Is Updated in This Book?

Since the previous version of this book was released, it has been revised to include the latest information to help you troubleshoot performance bottlenecks.

---

## Updated Chapters

The following chapters are updated:

- Chapter 1, “Performance Troubleshooting Tools.” Added information about using Microsoft Operations Manager 2000, which provides comprehensive event management, proactive monitoring and alerting, reporting, and trend analysis.
  - Chapter 2, “Establishing a Baseline.” Updated information about the minimal set of recommended counters.
  - Chapter 3, “Troubleshooting Performance.” Revised entire chapter to include information to help you more easily identify and isolate the root causes of Exchange 2000 Server performance problems. The sections about CPU performance issues, disk performance issues, and memory problems were all significantly updated. To improve the flow of the book, the section about using various RAID levels was moved to the end of the Appendix.
  - Appendix. Revised each subsection of performance counters to ensure that your experience monitoring and troubleshooting performance issues is as efficient as possible.
- 

## What Will You Learn from This Book?

This book provides detailed answers to the following questions:

- What tools can I use to monitor my Exchange 2000 servers?
- How do I establish a baseline of normal server performance?
- What steps do I take to troubleshoot performance problems?
- Now that I have established a baseline, what additional areas can I monitor for performance problems?

# How Is This Book Structured?

This book is divided into three chapters and one appendix:

## **Chapter 1, “Performance Troubleshooting Tools”**

This chapter contains information about tools you can use to monitor the performance of your Exchange 2000 servers.

## **Chapter 2, “Establishing a Baseline”**

This chapter contains information about establishing a baseline of normal Exchange 2000 server performance. A baseline helps you identify system performance trends and diagnose performance issues.

## **Chapter 3, “Troubleshooting Performance”**

This chapter contains specific information about how to troubleshoot performance problems on servers running Exchange 2000 Server. This chapter provides example performance problems and captured performance data covering areas where performance problems can occur.

## **Appendix**

This section contains additional performance counters that can be monitored when establishing a baseline or monitoring the health of your Exchange 2000 servers. It also includes information about the performance impact of using specific RAID levels as part of your storage solution.

In addition to reviewing this book, you can also review the most current Exchange 2000 Server performance knowledge base articles in the Microsoft Knowledge Base at <http://support.microsoft.com>. The Microsoft Knowledge Base contains the most up-to-date and detailed information about specific performance topics. By reviewing these articles, you can often resolve known performance issues.

# Performance Troubleshooting Tools

You can use the following tools to monitor and troubleshoot Exchange 2000 Server performance:

- System Monitor
- Performance Logs and Alerts
- Microsoft Operations Manager 2000
- Event Viewer
- Network Monitor
- File Monitor


---

## System Monitor

System Monitor is part of the Performance Microsoft Management Console (MMC) snap-in administrative tool. Using System Monitor, you can measure the performance of your own computer or other computers on a network.

**Note** System Monitor may also be referred to as “Performance Monitor” or “perfmon,” which is the name of the executable.

Figure 1 shows System Monitor in action.



**Figure 1 System Monitor**

System Monitor can do the following:

- Collect and view real-time performance data on a local computer or on several remote computers.

- View current or past data collected in a counter log.
- Present data in a printable graph, histogram, or report view.
- Create HTML pages from performance views.
- Create reusable monitoring configurations that can be installed on other computers using Microsoft Management Console.

Using System Monitor, you can collect and view extensive data about the usage of hardware resources and system services activity on computers you administer. You can define the data you want the graph to collect in the following ways:

- **Type of data** System Monitor lets you select the data you want collected by specifying performance objects, performance counters, and object instances. Some objects provide data on system resources (such as memory); others provide data on application operations (for example, Exchange 2000).
- **Source of data** System Monitor can collect data from your local computer or from other computers on the network on which you have permissions. In addition, it can collect real-time or past data using counter logs.
- **Sampling parameters** System Monitor supports manual, on-demand sampling or automatic sampling based on a time interval you specify. When viewing logged data, you can also choose starting and stopping times so that you can view data spanning a specific time range.

In addition to options for defining data content, you have considerable flexibility in designing System Monitor views:

- **Type of display** System Monitor supports graph, histogram, and report views. The graph view is the default view; it offers the widest variety of optional settings.
- **Display characteristics** For any of these views, you can define the colors and fonts for the display. In graph and histogram views, you can select from many different options to view performance data, such as:
  - Provide a title for your graph or histogram and label the vertical axis.
  - Set the range of values depicted in your graph or histogram.
  - Adjust the characteristics of lines or bars plotted to indicate counter values, including color, width, style, and so on.

For more information about System Monitor, see Microsoft Windows® 2000 Server Help.

---

## Performance Logs and Alerts

Performance Logs and Alerts is part of the Performance Microsoft Management Console (MMC) snap-in administrative tool. With Performance Logs and Alerts, you can collect performance data automatically from local or remote computers. You can view logged counter data using System Monitor or export the data to a spreadsheet or database for analysis and report generation.

Using Performance Logs and Alerts, you can:

- Collect data in a comma-separated or tab-separated format for easy import to a spreadsheet. A binary log-file format is also provided for circular logging or for logging instances such as threads or processes that begin after the log starts collecting data. (Circular logging is the process of continuously logging data to a single file, overwriting previous data with new data.)
- Collect counter data that can be viewed during collection, as well as after collection stops.
- Run Performance Logs and Alerts as a service and collect data even if no one is logged on to the computer being monitored.
- Define start and stop times, file names, file sizes, and other parameters for automatic log generation.
- Manage multiple logging sessions from a single console window.

- Set an alert on a counter, thereby ensuring that a message is sent, a program is run, or a log is started when the counter's selected value exceeds or falls below a specified setting.

Similar to System Monitor, Performance Logs and Alerts supports defining performance objects, performance counters, and object instances, as well as setting sampling intervals for monitoring data about hardware resources and system services. In addition, Performance Logs and Alerts offers the following options related to recording performance data:

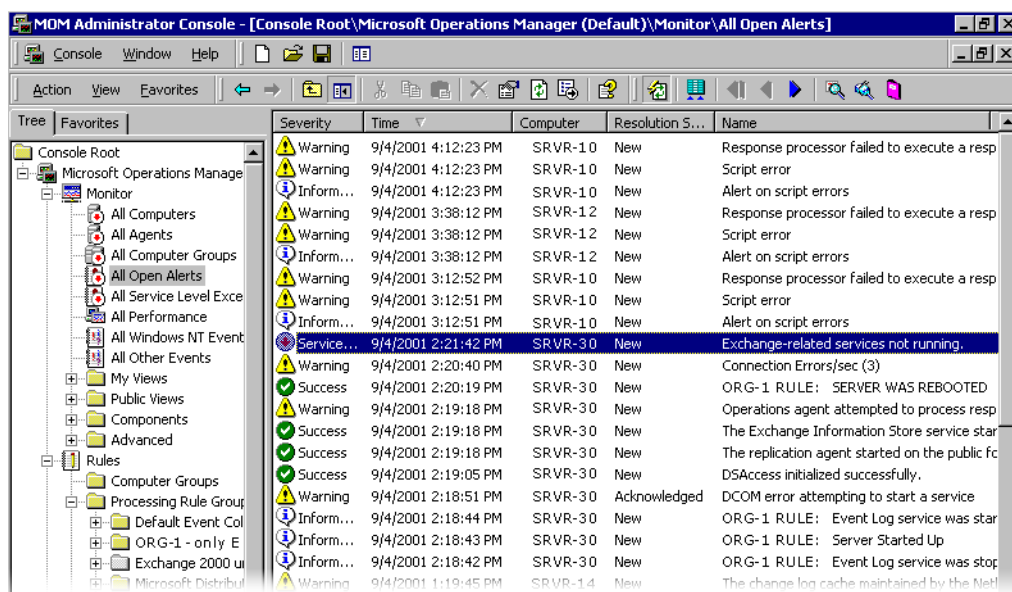
- Starts and stops logging, either manually on demand or automatically based on a user-defined schedule.
- Configures additional settings for automatic logging, such as automatic file renaming, and sets parameters for stopping and starting a log based on the elapsed time or the file size.
- Creates trace logs. Using the default system data provider or another provider, trace logs record data when certain activities such as a disk I/O operation or a page fault occur. When the event occurs, the provider sends the data to the Performance Logs and Alerts service. This differs from the operation of counter logs: when counter logs are in use, the service obtains data from the system when the update interval has elapsed, rather than waiting for a specific event. A parsing tool is required to interpret the trace log output. Developers can create such a tool using application programming interfaces (APIs) provided on the Microsoft Web site (<http://msdn.microsoft.com/>).
- Defines a program to run when a log is stopped.

For more information about Performance Logs and Alerts, see Windows 2000 Server Help.

## Microsoft Operations Manager 2000

Microsoft Operations Manager 2000 provides comprehensive event management, proactive monitoring and alerting, reporting, and trend analysis. The Management Packs for Microsoft Operations Manager 2000 include an extensive product support knowledge base to help reduce day-to-day support costs associated with running applications and services in a Microsoft Windows-based Information Technology (IT) infrastructure. Microsoft Operations Manager 2000 management packs provide necessary operational knowledge about Windows 2000 Server and Exchange 2000 Server.

The Exchange 2000 Management Pack for Microsoft Operations Manager 2000 is a particularly strong tool for monitoring Exchange 2000 servers because it automates much of the analysis presented in this book. This provides for a lower cost of operating a high availability deployment and a means of simplifying performance analysis. Figure 2 illustrates typical information available from Microsoft Operations Manager 2000.



**Figure 2 Microsoft Operations Manager 2000**

Using Microsoft Operations Manager 2000, you can:

- Check system status from a Web console.
- Create sophisticated rules to respond to events.
- Generate custom reports.
- Handle basic operational tasks using one of the add-in management packs.

Microsoft Operations Manager 2000 has a full set of features that help administrators monitor and manage the events and performance of Windows 2000–based server systems.

For more information on Microsoft Operations Manager 2000, see the product Web site at <http://www.microsoft.com/mom/>. For information about monitoring Exchange with Microsoft Operations Manager 2000, go to <http://go.microsoft.com/fwlink/?LinkId=16451>.

# Event Viewer

Using the event logs in Event Viewer, you can gather information about hardware, software, and system problems, and you can monitor Windows 2000 security events.

The EventLog service starts automatically when you start Windows 2000 and records events in three types of logs as outlined in the following table.

**Table 1** Logs used by Event Viewer

Log	Description
Application log	The application log contains events logged by Exchange 2000 and other applications. Most Exchange 2000 events are logged in the application log.
System log	The system log contains events logged by the Windows 2000 system components. For example, the failure of a driver or other system component to load during startup is recorded in the system log. Windows 2000 predetermines the event types logged by system components.
Security log	The security log can record security events such as valid and invalid logon attempts, as well as events related to resource use, such as creating, opening, or deleting files. An administrator can specify what events are recorded in the security log. For example, if you enable logon auditing, attempts to log on to the system are recorded in the security log.

Event Viewer displays the types of events outlined in the following table:

**Table 2 Events displayed by Event Viewer**

Event	Description
Error	Indicates a significant problem, such as loss of data or loss of functionality. For example, if a service fails to load during startup, an error is logged.
Warning	Indicates a potentially significant problem. For example, when disk space is low, a warning is logged.
Information	Indicates the successful operation of an application, driver, or service. For example, when a network driver loads successfully, an information event is logged.
Success Audit	Indicates a successful audited security access attempt. For example, if a user's attempt to log onto the system is successful, a success audit event is logged.
Failure Audit	Indicates an audited security access attempt has failed. For example, if a user's attempt to access a network drive fails, a failure audit event is logged.

For more information about Event Viewer, see Windows 2000 Server Help.

---

## Network Monitor

Network Monitor enables you to detect and troubleshoot problems on local area networks (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.

Instructions for using Network Monitor to troubleshoot performance are in the “Troubleshooting Performance” section later in this book.

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

- 294818 – “Frequently Asked Questions About Network Monitor” (<http://support.microsoft.com/?kbid=294818>)
- 148942 – “How to Capture Network Traffic with Network Monitor” (<http://support.microsoft.com/?kbid=148942>)

## File Monitor

The System Internals File Monitor monitors and displays file system activity on a system in real time. Generally, it is a useful tool for seeing how applications use files and DLLs, or for assessing problems in system or application file configurations. It is a particularly useful tool to use to identify the files that are being written to or read from. One way to use this tool is to run it after you have first used System Monitor to identify the I/O operations that seem to be the source of problems. For more information about System Internals File Monitor, see the “Troubleshooting Performance” section later in this book and the System Internals Web site at: <http://www.sysinternals.com>.

**Note** This third-party contact information is provided to help you find the technical support you need. This contact information is subject to change without notice. Microsoft in no way guarantees the accuracy of this third-party contact information.

## Notations Used in This Book

This book covers many performance counters. Performance counters are composed of the following three parts:

- **Performance object** The part of the computer being monitored. Some of the most commonly used objects are Processor, Memory, and PhysicalDisk. When Exchange 2000 is installed, new objects such as MExchangeIS are added to the performance object list.
- **Counters** The counters available for a performance object, which are the parts of the object you can monitor. For example, on the memory object, you can monitor the available bytes, kilobytes, and megabytes of memory, as well as the page faults per second or total pages per second.
- **Instances** There may be multiple objects or counters to monitor on the computer. For example, when looking at counters under the Processor 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.

When performance counters are referenced in this book, they are listed in this format:

```
Performance Object (Instance) \Counter
```

**Note** The instance is not a requirement. For example:

```
PhysicalDisk\% Disk Time
```



# Establishing a Baseline

To help you diagnose performance problems, you should establish a baseline of the normal server usage and performance for your Exchange 2000 servers. This baseline data must be considered when your servers that run Exchange 2000 experience performance problems. More specifically, with this data you will be able to see what has changed from the time when it was performing well. For example, are there more users logged on now than before? Is the server receiving more mail now?

It is essential that this baseline data be kept current; therefore this task requires significant diligence. A baseline that is several months old is not going to be useful in helping diagnose problems.

The Exchange 2000 Management Pack for Microsoft Operations Manager 2000 automatically collects this baseline data so that it is ready for use when needed. This significantly lowers the operational cost of being ready for times when such an analysis is required.

---

## Minimal Set of Counters

The following counters are the minimal set of counters you should use to establish a baseline and monitor overall server health with accompanying descriptions.

**Note** There are many counters you can use to establish a baseline specific to your organization and to monitor the performance of your Exchange 2000 server. See the “Appendix” section later in this book for a complete list of counters, with a description and recommended value for each.

**Table 3 Minimal Set of Counters**

Counter	Description
MSExchangeIS Mailbox\ Message Opens/sec	Displays the rate at which requests to open messages are submitted to the Exchange store.
MSExchangeIS Mailbox\Folder Opens/sec	Displays the rate at which requests to open folders are submitted to the Exchange store.
MSExchangeIS Mailbox\Local Delivery Rate	Displays the rate at which messages are being delivered locally to this server
MSExchangeIS\RPC Operations /sec	Displays the rate at which remote procedure call (RPC) operations occur.
MSExchangeIS\RPC Requests	Displays the number of client requests that are currently being processed by the Exchange store.
PhysicalDisk (_Total)Disk Transfers/sec	Displays the number of completed read and write operations per second.

Counter	Description
Process (STORE.EXE)\% Processor Time	Displays the fraction of processing capacity used by the Exchange 2000 store.exe process. This counter ranges from 0 to 100 * <number of processors> %. For instance, on a four-processor system, this will range from 0 to 400%.
Processor (_Total)\% Processor Time	Displays the fraction of the total processing capacity being used by all processes running on the server. This counter has a range from 0 to 100%.
SMTP Server\ Local Queue Length	Displays the number of messages in the local Simple Mail Transfer Protocol (SMTP) queue.
SMTP Server\ Messages Delivered/sec	Displays the rate that messages are being delivered to local mailboxes.
SMTP Server\ Messages Received/sec	Displays the rate that messages are being received.
SMTP Server\ Messages Sent/sec	Displays the rate that messages are being sent.

**Note** Before troubleshooting disk problems, at the command prompt, run `diskperf -y` to activate logical disk counters. All physical disk counters are enabled by default. You must restart your computer before the logical disk counters appear.

## Example Baseline

After you begin monitoring your Exchange 2000 servers, you can use the data you capture to establish your baseline. The following sections provide questions you should answer about your normal server performance, as well as System Monitor capture examples.

## Questions to Answer

When establishing your baseline, it is important that you answer questions such as the following. Answers to these questions can help you interpret current performance data and investigate performance problems.

- What is the average number of messages that users receive per day?
- How many messages do users open, and how often do they open folders?
- What is the peak delivery rate, the peak period during the day, and the peak day of the week?
- Are there monthly or quarterly peaks?
- How many more users can your servers support?

Your goal is to compare baseline data you have gathered from typical load periods to current performance data. By comparing baseline data with your server's current performance, you can determine if the server is operating normally or if there are performance problems. Answering the preceding questions also helps you analyze current performance data and identify performance problems.

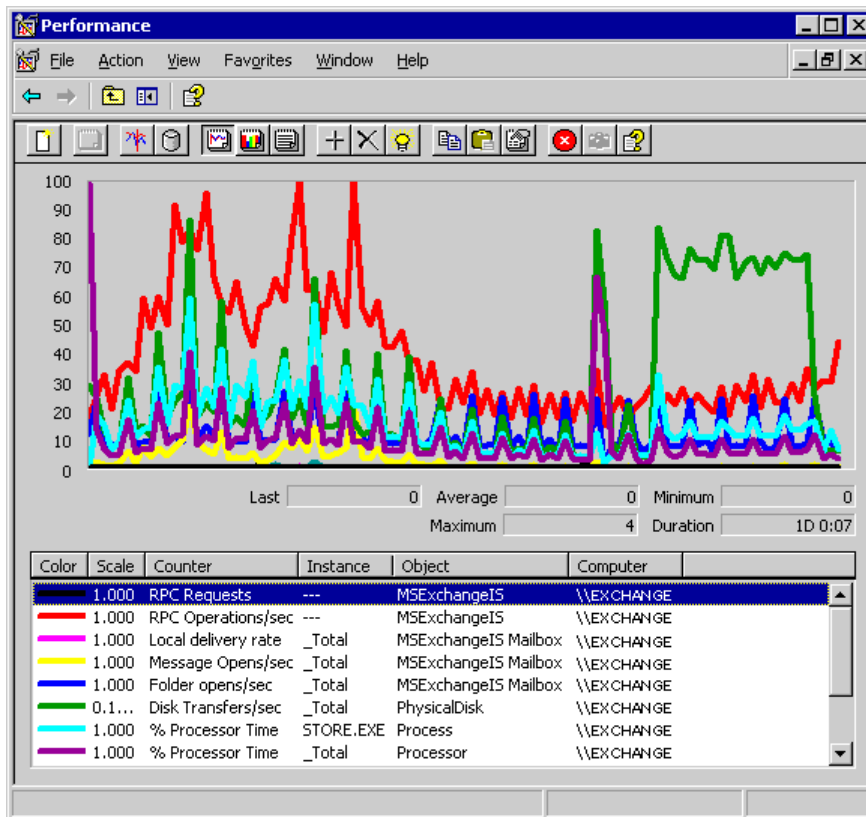
## System Monitor Examples

The following are example System Monitor performance data captures. Consider leaving System Monitor running all the time for easy access. You can do this at different collection intervals, such as:

- 900 seconds for a 24-hour view (useful for seeing daily trends)
- 60 seconds for a 1- to 2-hour view (useful for viewing recent usage and performance trends)
- 10 seconds to capture short-lived spikes (useful for viewing usage in the last few minutes)

The following System Monitor illustration was captured while monitoring a production Exchange 2000 Service Pack 3 server during business hours.

Figure 3 illustrates System Monitor capturing data with a 24-hour view.



**Figure 3** System Monitor data with a 24-hour view

By monitoring performance using each of the three collection intervals and the minimal set of counters, you can establish a baseline, as well as monitor your servers for performance problems.

**Note** You can save performance data in log files by using Performance Logs and Alerts. Performance Logs and Alerts saves the performance data in log files, so that you can compare performance data saved during typical load times to current performance data. You can then view the data in the logs files by using System Monitor.



# Troubleshooting Performance

This section outlines how to identify and isolate the root causes of Exchange 2000 performance problems. Many of these problems are the result of resource bottlenecks, in which one of the server's resources is being used to capacity. Resource bottlenecks result in longer latencies for end users.

## Performance Problem Origins

You may get early indications of performance problems from monitoring data or from users who report that their e-mail is slow. The first step in isolating an Exchange performance problem is to determine its origin. If your users report that e-mail is slow, this could indicate that a problem is occurring at one of the following periods:

- Before the request reached the Exchange server (such as a network problem)
- During Exchange processing (such as a resource bottleneck on the server)
- After the request is sent back to the client computer.

The following performance counters help you identify the cause of the user's performance problems:

```
MSEExchangeIS\RPC Requests  
MSEExchangeIS\RPC Operations/sec
```

- The `MSEExchangeIS\RPC Requests` counter indicates the number of MAPI RPC requests presently being serviced by the Exchange store. The Exchange store can service only 100 requests simultaneously.
- The `MSEExchangeIS\RPC Operations/sec` counter indicates the rate at which the Exchange store is actually servicing user requests.

The key to using these two counters is relatively simple. If the `RPC Requests` are low, and the `RPC Operations/sec` (outstanding requests) is zero, the performance problem is occurring before Exchange processing occurs. All other combinations point to a problem during Exchange processing or a problem after Exchange processing.

Figure 4 illustrates an issue with Exchange performance that was identified using the `MSEExchangeIS\RPC Requests` counter and the `MSEExchangeIS\RPC Operations/sec` counters. In this example, no operations are running for a three-minute period, but the Exchange store has outstanding requests. Because there are outstanding requests waiting to be processed (`RPC Requests`), this indicates a problem with the server running Exchange 2000. However, it is not clear from this figure, why the server running Exchange 2000 is not processing the requests (because the `RPC Operations/sec` is zero).

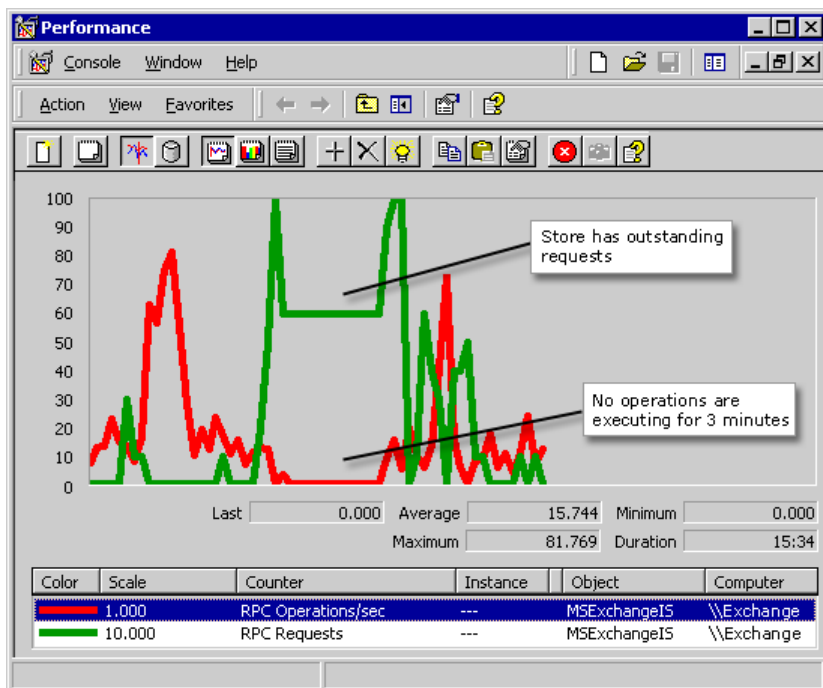
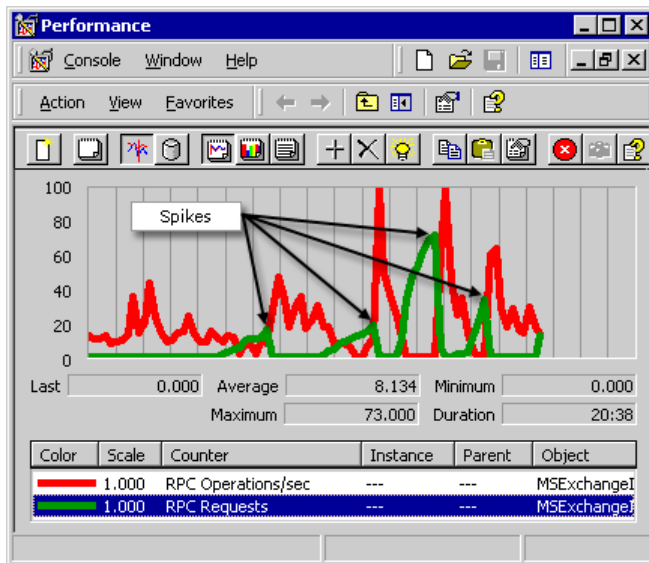


Figure 4 Example of an Exchange 2000 performance issue

Figure 5 illustrates another performance problem with Exchange that was identified using the `MSExchangeIS\RPC Requests` counter and the `MSExchangeIS\RPC Operations/sec` counter. The figure illustrates four periods of time in which outstanding requests are continuously increasing because the server cannot complete enough operations. The cause of this continuous increase may be the result of a resource bottleneck on the server. For more information see “Understanding the Problem” immediately following Figure 7.



**Figure 5** Example of an Exchange 2000 performance issue

Figure 6 illustrates a client problem identified using the MExchangeIS\RPC Requests counter and the MExchangeIS\RPC Operations/sec counters. The RPC Operations/sec and the RPC Request rate are growing simultaneously. A client may be running a utility or script that is making many requests of the Exchange store and the Exchange store is struggling to keep up. In this situation, you could use the Network Monitor tool to find the computer from which the requests are coming.

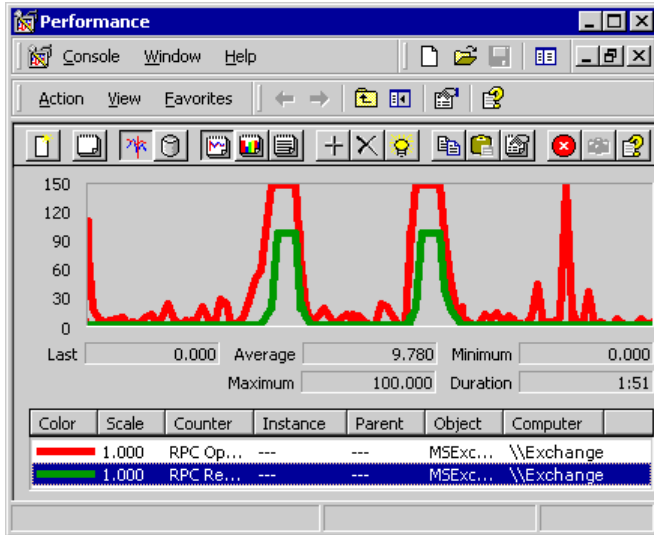


Figure 6 Example of a client performance issue

Figure 7 illustrates a network problem identified using the MSEExchangeIS\RPC Requests counter and the MSEExchangeIS\RPC Operations/sec counters. In two cases, RPC Operations/sec and RPC Requests are both zero. In this situation, something is preventing the requests from arriving at the Exchange store. You can use the Network Monitor tool to determine whether requests are arriving at the server.




Figure 7 Example of a network performance issue

## Understanding the Problem

After determining if the problem is occurring during Exchange 2000 processing, before Exchange 2000 processing, or after Exchange 2000 processing, you must try to identify the next step of troubleshooting the root problem. Before beginning troubleshooting, you should have the answers to the following questions about clients and hardware on the server on which the problem is occurring:

- Are clients acting sluggish or have they stopped responding?
- Is the problem occurring with a particular client operation?
- Do all clients experience the problem at the same time?
- What is the frequency of the problem?
- What hardware is being used on the Exchange 2000 Server?
- Will the bandwidth support what is being attempted? (For example, are you trying to use the Site Connector over a 56-Kbps line?)
- Is the network the problem? For example have you confirmed all IP information, including Windows Internet Name Service (WINS), Domain Name System (DNS), and global catalog or domain controller communication?

It is also essential to know the configuration of the server in detail, such as:

- How many processors are there on the server?
- How much memory is there on the server?
- For each physical disk volume, how many disks exist and how are they configured (such as RAID-0, RAID-1, RAID-5)?

- What versions of Exchange, Windows, and their respective service packs are installed? Are those versions the most current and supported versions?

---

# Root Cause Performance Analysis: Bottleneck Identification

The process for identifying the root cause of performance problems involves first identifying the most likely sources of performance problems, and then considering each of the potential bottlenecks that can inhibit server performance. The primary sources of such performance bottlenecks are CPU, disk, and memory.

---

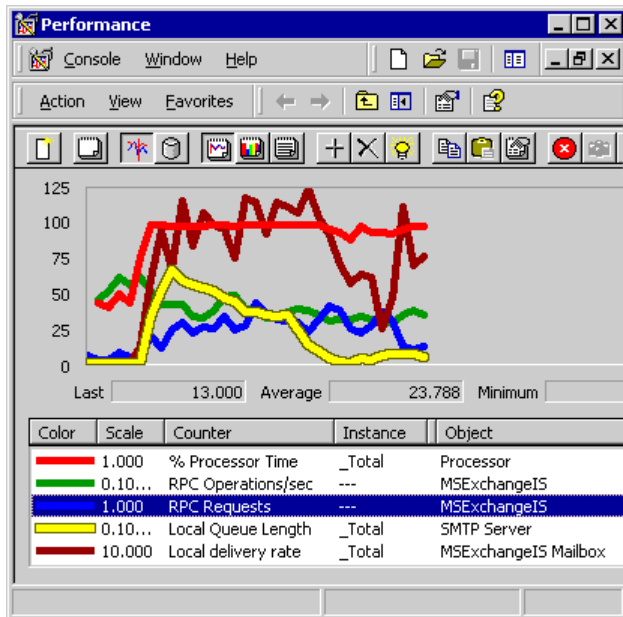
## CPU Performance Issues

CPU bottlenecks are the easiest bottlenecks to detect. If the `Processor(_Total)\% Processor Time` counter is approaching 100 percent, then that indicates a CPU bottleneck.

**Important** If you are running Indexing Service on the server, it will use all available CPU when indexing, so disable it when trying to verify a potential CPU bottleneck. Indexing Service appropriates all idle CPU processing power and uses it. If another process requests additional CPU power from the system while Indexing Service is running, the Indexing Service engine relinquishes the CPU.

If the `Processor(_Total)\% Processor Time` counter is high, check to see if the `MSExchangeIS\RPC Requests` counter is increasing. If the `MSExchangeIS\RPC Requests` counter reaches the maximum of 100, it causes client time-outs. The Exchange store can handle only 100 simultaneous RPC requests.

Figure 8 illustrates a CPU performance issue. It shows a sudden increase in the local delivery rate. As a result, CPU usage has risen to 100 percent. In this situation, the CPU is working at capacity to deliver local messages.



**Figure 8** An example CPU performance issue

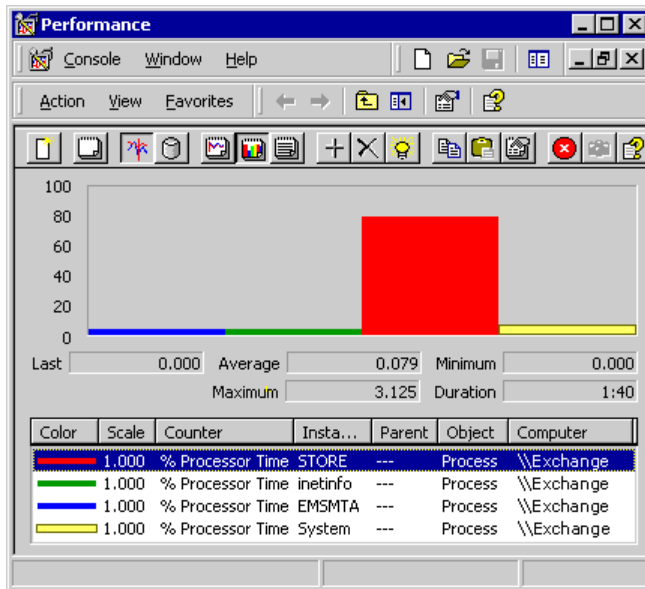
## CPU Consumption

After you have determined the problem is with the CPU, you should determine what is consuming the CPU. The counters below are the most likely suspects for this problem, from most likely first to least likely fourth. These four counters normally add up to 90 percent of the CPU being used.

```
Process(STORE.EXE)\% Processor Time
Process(inetinfo)\% Processor Time
Process(emsmta)\% Processor Time
Process(system)\% Processor Time
```

**Note** Process counters count 100 percent for each CPU on the server. On an eight-processor computer, the value of each of the processor counters above would be between 0 percent and 800 percent.

Figure 9 illustrates a histogram view of the processes that are most likely to consume the CPU. The Exchange store process is using up most of the CPU. If you suspect that other processes besides the four most likely may be hanging up the CPU, you should include them in this histogram view.



**Figure 9 A histogram view of processes most likely to consume the CPU**

**Note** Viewing multiple counters in histogram view in System Monitor is a quick way to isolate the counter indicating a problem.

The following are other common processes that can consume the CPU:

- Backup utilities
- Monitoring utilities
- Remote access tools

## Isolating Threads

An advanced step to help further determine what process is consuming the CPU is to monitor the individual threads using the CPU. This can help isolate the thread or threads in a specific process that are consuming the CPU.

Use the same histogram view technique in System Monitor to isolate the thread consuming the CPU, as you did to isolate the process. Add all `Thread(process/threadnumber)\% Processor Time` counters for the target process to a histogram view of System Monitor. You can identify the thread using the `Thread\process(threadnumber)ThreadID` counter.

## Disk Performance Issues

Unlike CPU performance issues, disk performance issues cannot be diagnosed with a single counter that indicates that you have a disk bottleneck.

**Note** A disk bottleneck can also be a symptom of a memory issue. In cases where a memory issue is the actual root cause of a performance issue, adding more disk throughput capacity will resolve solve these issues. For information about troubleshooting memory issues, see “Memory Problems” section later in this book.

Ensure that when you size your Exchange 2000 disk configurations, size for I/O capacity and not for disk space only.

### Method 1: I/O Capacity

The first approach to determining if you are encountering a disk bottleneck is to monitor the following counters for each of your physical drives:

```
PhysicalDisk(drive:)\Disk Writes/sec
PhysicalDisk(drive:)\Disk Reads/sec
```

Look at each drive and compare to the total instance to isolate where the I/O is going. You can use the recommendations below to assist with the comparison and determine if you have a bottleneck:

- Raid-0: Reads/sec + Writes/sec < # Spindles x 100
- Raid-1: Reads/sec + 2 \* Writes/sec < # Spindles x 100 (Each write has to go to each mirror on the array.)
- Raid-5: Reads/sec + 4 \* Writes/sec < # Spindles x 100 (Each write requires two reads and two writes.)

**Note** This scenario assumes disk throughput is equal to 100 random I/O per spindle.

For more information about RAID, see the following “RAID Levels” section in the appendix.

### Method 2: Disk Queues

The second approach to determining if you are encountering a disk bottleneck requires looking at the I/O requests waiting to be completed, using the following disk queue counters:

```
PhysicalDisk(drive:)\Avg. Disk Queue
PhysicalDisk(drive:)\Current Disk Queue
```

The `PhysicalDisk(drive:)\Avg. Disk Queue` counter shows the average queue length over the sampling interval. The `PhysicalDisk(drive:)\Current Disk Queue` counter reports the queue length value at the instant of sampling.

You are encountering a disk bottleneck if the average disk queue length is greater than the number of spindles on the array and the current disk queue length never equals zero. Short spikes in the queue length can drive up the queue length average artificially, so you must monitor the current disk queue length. If the queue length drops to zero periodically, the queue is being cleared, and you probably do not have a disk bottleneck.

**Note** When using this approach, correlate the queue length spikes with the `MSExchangeIS\RPC Requests` counter to confirm the effect on clients.

### Method 3: Disk Latencies

For the third approach, to determine if you are encountering a disk bottleneck, look at I/O latency, which can give you an indication of the health of your disks:

```
PhysicalDisk(drive:)\Avg. Disk sec/Read
PhysicalDisk(drive:)\Avg. Disk sec/Write
```

A typical range is .005 to .020 seconds for random I/O. If write-back caching is enabled in the array controller, the PhysicalDisk(drive:)\Avg. Disk sec/Write counter should be less than .002 seconds.

If these counters are between .020 and .050 seconds, there is possibly a disk bottleneck. If the counters are above .050 seconds, there is definitely a disk bottleneck.

## Which Process is Causing the I/O?

In Microsoft® Windows® 2000, you can use these counters to help determine which process is causing the disk I/O:

```
Process(process name)\IO Read Operations/sec
Process(process name)\IO Write Operations/sec
```

**Note** These counters include more than file I/O performance. They can also help you determine the process that is causing the I/O.

## To Which File is the I/O Going?

In Exchange deployments that isolate certain types of files on specific drives, it is simpler to determine the file that is the source of the disk bottleneck. However, if there are multiple files on a given volume to which I/O operations could be going, you can use the System Internals File Monitor to determine which file or files are showing I/O activity. Choose the logical disks that need investigation and show all disk reads and writes. This procedure is particularly useful for multi-use disks, such as drive C, which may have several major files on it that are used by the system or applications.

Figure 10 illustrates the System Internals File Monitor.

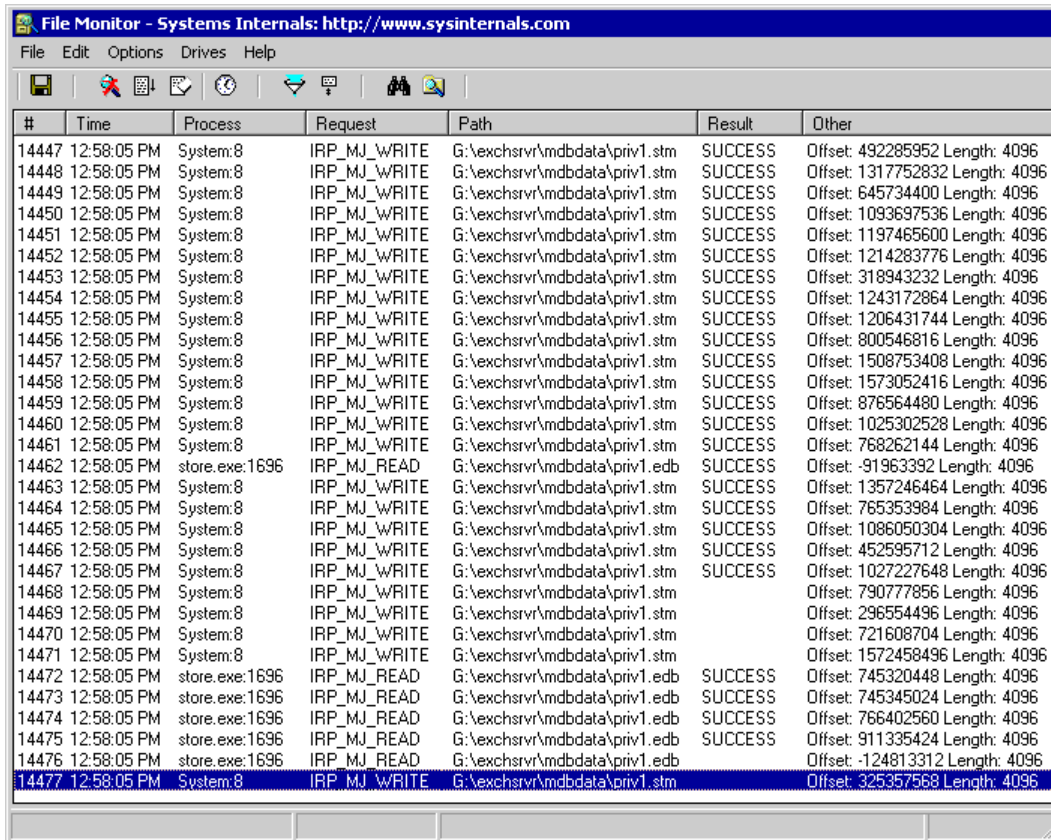


Figure 10 System Internals File Monitor output showing the I/O going to priv1.stm and priv1.edb

**Note** This third-party contact information is provided to help you find the technical support you need. This contact information is subject to change without notice. Microsoft in no way guarantees the accuracy of this third-party contact information.

---

## Memory Problems

When investigating memory problems, the first counter to use to monitor physical memory usage is `Memory\Available MBytes`. If this counter goes below 4 MBs, Windows aggressively starts cutting the working sets of running processes. The server is generally healthy if the `Memory\Available Mbytes` counter is greater than 4 MBs.

---

### Primary Memory Counters

The following counters are the primary counters to use when investigating memory problems. They help you determine if there are paging problems. These counters provide information about hard pages, pages that are causing information to go to and from the disk.

```
Memory\pages/sec  
Memory\page reads/sec  
Memory\page writes/sec
```

Memory\pages/sec reports the total number of pages going to disk, while Memory\page reads/sec and Memory\page writes/sec provide the rate of paging read and writes.

**Note** Paging I/O is normal because Exchange 2000 uses the Windows system cache to back the .stm file.

When the paging to and from disk gets high enough, eventually a disk bottleneck will occur, and consequently performance will suffer. The disk bottleneck can be identified as discussed in the previous section. If the Memory\pages/sec indicates that paging I/O is responsible for most of the disk I/O, then the real problem is memory, and the disk bottleneck is just a symptom.

---

## Additional Memory Counters

There are additional counters you can use to further investigate memory problems:

```
Memory\Page Faults/sec
Memory\Cache Faults/sec
Memory\Transitions Faults/sec
Process(process)\Page Faults/sec
```

The Memory\Page Faults/sec counter is not necessarily an indication of a memory problem because it also includes the Memory\Cache Faults/sec counter, and cache faults are a normal part of Exchange 2000 operation due to the .stm file. Also, both the Memory\Page Faults/sec counter and the Memory\Cache Faults/sec counter include transition faults indicated by the Memory\Transition Faults/sec counter. Transition faults are faults that do not go to the disk because the memory manager has the pages on the standby list.

The Process(process)\Page Faults/sec counter can be useful to identify processes with high page faults. Using System Monitor, add processes in a histogram view to quickly identify the process with high page faults.

**Note** This counter should be used as a guide. Page faults do not necessarily indicate a memory problem. However, a process with high page faults is probably also generating many page read and write operations.

---

## Where Did The Memory Go?

To determine where memory is being used, monitor the following counters, which are the most likely suspects for memory consumption:

```
Process(STORE.EXE)\Working Set
Process(inetinfo)\Working Set
Process(emsmta)\Working Set
Memory\Cache Bytes
```

The Exchange store process indicated by the Process(STORE.EXE)\Working Set counter tends to consume most of the committed bytes. This is due to the Exchange store, which maintains a large cache. You can use the Database\Cache Bytes counter to confirm the size of this cache.

---

## Virtual Memory

One of the most problematic areas of Exchange scaling is the fragmentation of virtual memory – otherwise known as address space – in the STORE.EXE process. As you scale a server to accommodate more users and more usage, the server may run low on virtual memory. This problem is signified by the presence of MExchangeIS 9582 events in the application log, which can come in warning and error severities depending on how fragmented the virtual memory has become.

The Information Store service logs the following events if the virtual memory for your Exchange 2000 server becomes excessively fragmented:

```
EventID=9582
Severity=Warning
Facility=Perfmon
Language=English

The virtual memory necessary to run your Exchange server is fragmented in such a way that performance may be affected. It is highly recommended that you restart all Exchange services to correct this issue.
```

**Note** This warning is logged if the largest free block is smaller than 32 MBs.

```
EventID=9582
Severity=Error
Facility=Perfmon
Language=English

The virtual memory necessary to run your Exchange server is fragmented in such a way that normal operation may begin to fail. It is highly recommended that you restart all Exchange services to correct this issue.
```

**Note** This error is logged if the largest free block is smaller than 16 MBs.

Adding more physical memory does not solve errors that indicate virtual memory is very fragmented. Monitoring virtual memory fragmentation is most crucial on active/active clusters because if the virtual memory becomes sufficiently fragmented on one node, a failover to that node may not be successful if there is not enough contiguous virtual memory.

Virtual memory problems can be substantially, though not entirely, mitigated by enabling 3 GB of virtual memory on Windows 2000 Advanced Server. If your server is running Windows 2000 Advanced Server and more than 1 GB of physical RAM is installed, this is done by adding the /3GB switch in the boot.ini file and rebooting the server.

More information about virtual memory issues is in Microsoft Knowledge Base articles [317411](#) “XADM: How to Gather Data to Troubleshoot Exchange Virtual Memory Issues” and [302254](#) “XADM: Computer That Is Running Exchange 2000 and Windows 2000 Server May Run Out of Virtual Memory with Event ID 12800.”

### To troubleshoot virtual memory problems

1. Check the application log for 9582 warnings (less than 32-MB virtual memory blocks available) or 9582 errors (less than 16-MB virtual memory blocks available). On some large systems, it is usual to drop below the 32-MB threshold during peak activity; however, the available virtual memory should rise significantly during non-peak activity.
2. Check the application log for other errors that indicate that you are out of memory, such as 12800 Multipurpose Internet Mail Extensions (MIME) processing errors, in addition to 9582 warnings. If the warnings are accompanied by other errors indicating that you are out of memory, users may be unable to access mail. If no other processing errors occur and users are able to access their mail, it indicates that the 9582 warnings may be relatively harmless. However, you should investigate 9582 warnings for possible action.
3. Monitor the `MSExchangeIS\VM Largest Block Size` counter. Using this counter is the best way to investigate virtual memory issues. You can monitor this counter in real time or monitor one-minute intervals. Collect 18 to 24 hours of data to determine if a trend indicates that memory is being released. Monitor the minimum value to see what the drop is. It can be normal on large servers if this minimum value is around 55 MB.
4. Be aware that other store-related processes, such as virus scanning, can tip the threshold. However, as long as user performance is not affected and the virtual memory block grows again during non-peak activity, corrective action may not be necessary. However, if you expect the user load to increase, you may want to reduce overall virtual memory consumption so that the server can accommodate a greater load.
5. To reduce virtual memory consumption, consider the following steps:

- a. Ensure that the server is running Exchange 2000 Server Service Pack 3 (SP3). Exchange SP3 has specific virtual memory optimizations.
- b. If 9582 warnings are still being logged, then you must perform a registry change. This registry change is acceptable as long as an adequate amount of RAM is available on the server. Monitor the `Memory\Available Bytes` counter. Make sure the counter indicates more than 200 MB. Change `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\HeapDeCommitFreeBlockThreshold` to equal 262144.
- c. If you are still experiencing virtual memory issues, it is possible that you are experiencing a memory leak. This can be investigated by monitoring the `Process(STORE.EXE)\Private Bytes` counter to determine if it is growing over time.

**Note** If doing the preceding steps does not reduce virtual memory consumption, you must reduce the load on the server by moving users to another server.

## Monitoring Non-MAPI Requests

In the same way that you used the RCP counters to examine the use of the Exchange store by MAPI clients, such as Outlook, you can use another set of queue counters to examine the use of the Exchange store by Post Office Protocol (POP3), Internet Message Access Protocol (IMAP4), Simple Mail Transfer Protocol (SMTP), Distributed Authoring and Versioning (DAV), and Network News Transfer Protocol (NNTP) clients. These counters are contained in the Epoxy performance object. These are queues in which information is passed out of Internet Information Services (IIS) to the Exchange store and then returned from the Exchange store to IIS. These queue counters include:

```
Epoxy(protocol)\Client Out Que Len
Epoxy(protocol)\Store Out Que Len
```

The `Epoxy(protocol)\Client Out Que Len` counter indicates the number of requests waiting to be processed by the Exchange store, and the `Epoxy(protocol)\Store Out Que Len` counter indicates the number of requests waiting to be processed by the IIS protocol handlers. You can use these counters to investigate whether information is being successfully passed between IIS and the Exchange store.

## Message Delivery Counters

The Exchange store responds preferentially to user requests as opposed to delivering mail. If your servers begin to build delivery queues, you have an overbooked server. User requests are arriving at such a high rate that the server cannot efficiently process the e-mail. Use the following counters to monitor message delivery:

```
SMTP Server\Local Queue Length
SMTP Server\Messages Delivered/sec
```

The `SMTP Server\Local Queue Length` counter should not grow continuously. This counter grows during peak lead periods, and anywhere from 0 to 1000 is a reasonable length. The `SMTP Server\Messages Delivered/sec` counter should be continuous. However, gaps of zero delivery followed by spikes of delivery indicate a bottleneck.

## Active Directory

Exchange 2000 Server is dependant on Microsoft® Active Directory® directory service. You can investigate CPU, disk, and memory bottlenecks on your Active Directory servers. Most techniques used to identify and

investigate problems with Exchange 2000 servers are equally applicable to Windows® 2000 Active Directory servers.

---

## DSAccess

DSAccess is the cache on the server running Exchange that caches frequent Active Directory queries from the same server. By caching Active Directory information, the server running Exchange does not have to contact an Active Directory server each time a query is needed. The following counters are useful for investigating problems with DSAccess:

```
MSExchangeDSAccess Caches\Cache Hits/Sec
MSExchangeDSAccess Caches\LDAP Searches/Sec
```

You should compare the current data from these counters with baseline data from other servers that are operating normally.

---

## Network Problems

Network problems can result in information not getting to the server running Exchange. The following counters are useful for investigating network problems:

```
Network Interface(netcard)Bytes Received/sec
Network Interface(netcard)Bytes Sent/sec
```

In data center environments or in environments in which there are high bandwidth connections, network problems are rare. However, you could possibly create a network problem by, for example, scheduling backup operations during the day when you should have scheduled them at night.

---

## Using Network Monitor

If client traffic is not getting to your server running Exchange, you can use the Network Monitor tool to examine the traffic. Network Monitor is a network diagnostic tool that monitors local area networks and provides a graphical display of network statistics. While collecting information from the network's data stream, Network Monitor displays the following types of information:

- The source address of the computer that sent a frame to the network. (This address is a unique hexadecimal (or base 16) number that identifies that computer on the network.)
- The destination address of the computer that received the frame.
- The protocols used to send the frame.
- The data, or a portion of the message being sent.

The process by which Network Monitor collects this information is called “capturing.” By default, Network Monitor gathers statistics on all the frames it detects on the network into a capture buffer, which is a reserved storage area in memory. To capture statistics on only a specific subset of frames, you can single out these frames by designing a capture filter. When you have finished capturing information, you can design a display filter to specify how much of the captured information is displayed in Network Monitor's Frame Viewer window.

To use Network Monitor, your computer must have a network card that supports promiscuous mode. If you are using Network Monitor on a remote computer, the local workstation does not need a network adapter card that supports promiscuous mode, but the remote computer does.

Once data has been captured either locally or remotely, you can save it to a text or capture file that can be opened and examined later.

**Note** To fully troubleshoot possible network issues using Network Monitor, consider configuring Network Monitor to capture not only what the client sends and receives, but also what the server is sending and receiving. Performing both a client and server-side trace of network traffic further helps you troubleshoot network issues.

---

### **Creating an Address List**

To use address pairs in a capture filter, you should first build an address database. After this database is built, you can use the addresses listed in the database to specify address pairs in a capture filter.

### To create an address list

1. From the Capture menu, click Start. Optionally, open a .cap file in the Frame Viewer window.
2. When you finish capturing information, click **Stop and View** from the **Capture** menu to display the **Frame Viewer** window.
3. From the **Display** menu, click **Find All Names**. Network Monitor processes the frames and then adds them to the address database.
4. Close the **Frame Viewer** window, and display the **Capture** window.
5. From the **Capture** menu, click **Filter** to display the **Capture Filter** dialog box.
6. In the **Capture Filter** dialog box, double-click **Address Pairs**. Or, click **Address** in the **Add** dialog box.
7. Network Monitor displays the address database you created. You can use the names in this database to specify address pairs in the capture filter.

### To monitor traffic between two computers

1. From the Capture menu, click Filter to display the Capture Filter dialog box.
2. Double-click ANY<->ANY to display the **Address Expression** dialog box.
3. In the left window of the **Address Expression** dialog box, select the address of a computer.
4. In the right window of the **Address Expression** dialog box, select the address of a computer.
5. In **Direction**, select one of the symbols:
6. Select the <-> symbol to monitor the traffic that passes in either direction between the addresses that you selected.
7. Select the --> symbol to monitor only the traffic that passes from the address selected in the left window to the address selected in the right window.
8. Choose the <-- symbol to monitor only the traffic that passes from the address selected in the right window to the address selected in the left window.
9. Click **OK**.
10. In the **Capture Filter** dialog box, click **OK**.
11. From the **Capture** menu, click **Start**.

---

### Tracing in a WAN Environment

When troubleshooting network problems, you may need to create a capture of network traffic between two specific computers that are separated by one or more routers. In this case, you may want to analyze all network traffic between the first computer and its nearest router and all network traffic between the second computer and its nearest router. Most of the time, this analysis is done to check whether network packets are being lost or corrupted somewhere between the routers. To make these traces consistent and to be able to read these traces simultaneously, the system clocks must be synchronized between the two computers before making the trace.

### To synchronize time between two computers

1. From the computer against which you want to synchronize the time, at the command prompt, type `net time \\ComputerName /set /yes`, where `ComputerName` is the name of the computer to which you want to synchronize.
12. Verify the computers have the same time by typing **TIME** at the command prompt for each computer.
13. Proceed with the trace.



# Appendix

---

## Performance Counters

The following are additional performance counters that you can use to monitor the health of your Exchange 2000 servers or to establish a baseline. They are grouped by their performance object area. When investigating a performance problem, you can use these counters to gather more information or add them to the minimum list of counters to use when establishing a baseline.

**Note** Some of the counters do not have recommended values, as the values are specific to your organization or provide additional information only.

## Database Counters

The following are Database (Exchange store) performance object counters. These counters are monitored using the Information Store instance.

**Table 4 Database (Exchange store) Counters**

Counter	Description	Recommended Value
Database Cache Size	Displays the amount of system memory the database cache manager uses to hold commonly used information from the database files in order to prevent file operations. If the database cache size seems to be too small for optimal performance and little memory is available on the system (see the Memory/Available Bytes counter), adding more memory to the system may increase performance. If a lot of memory is available on the system and the database cache size is not growing beyond a certain point, the database cache size may be capped at an artificially low limit. Increasing this limit may increase performance.	This counter may grow to 900 MBs by default.
Log Record Stalls\sec	Displays the number of log records that cannot be added to the log buffers per second because they are full. If this counter is not zero most of the time, the log buffer size may be a bottleneck.	Generally, this counter should remain at zero.

**Table 4 Database (Exchange store) Counters (continued)**

Counter	Description	Recommended Value
Log Writes/sec	Displays the number of times the log buffers are written to the log files per second. If this number approaches the maximum write rate for the media holding the log files, the log may be a bottleneck.	This counter is useful for showing how busy ESE is. The value of this counter is specific to your organization.
Table Opens/sec	Displays the number of database tables opened per second.	This counter is useful for showing how busy ESE is. The value of this counter is specific to your organization.

## Epoxy Counters

The following are Epoxy performance object counters.

**Table 5 Epoxy Counters**

Counter	Description	Recommended Value
Client out Que Len	Displays the number of requests waiting to be processed by the Exchange store.	Generally, this counter should be zero.
Store out Que Len	Displays the number of requests waiting to be picked up by the IIS protocol handlers.	Generally, this counter should be zero.

## Logical Disk Counters

The following are Logical Disk performance object counters.

**Table 6 Logical Disk Counters**

Counter	Description	Recommended Value
% Free Space	Displays the ratio of the free space available on the logical disk unit to the total usable space provided by the selected logical disk drive.	A recommended threshold for % Free Space is 15 percent.
Free Megabytes	Displays the unallocated space on the disk drive in MBs.	Alerts must be configured on disks that contain Exchange databases or log files that notify you as soon as they approach capacity. Exchange stops if its log files or databases have no space to grow.

## Memory Counters

The following are Memory performance object counters.

**Table 7 Memory Counters**

Counter	Description	Recommended Value
Available Bytes	Displays the amount of physical memory, in bytes, available to processes running on the computer.	You should keep this counter above 4 MB.
Committed Bytes	Displays the size of virtual memory (in bytes) that has been committed (as opposed to simply reserved). Committed memory must have backing (disk) storage available or must be assured never to need disk storage (because main memory is large enough to hold it). This is an instantaneous count, not an average over the time interval. Acceptable average range is less than the amount of physical RAM on the server. However, before making such an assumption, check Memory\Pages/sec and Memory\Page Faults/sec. If Memory\Pages/sec is large enough to cause a disk bottleneck, and Memory\Page Faults/sec is greater than Memory\Cache Faults/sec, then there is too much paging.	This counter should remain below the amount of physical RAM on the server.

**Table 7 Memory Counters (continued)**

Counter	Description	Recommended Value
Page faults/sec	Displays the overall rate at which the processor handles faulted pages. A page fault occurs when a process requires code or data that is not in its working set (its space in physical memory). This counter includes both hard faults (those that require disk access) and soft faults (in which the faulted page is found elsewhere in physical memory). Most processors can handle large numbers of soft faults without consequence. However, hard faults can cause significant delays.	This counter should never show a consistently high single figure amount.
Pages/sec	Displays the number of pages read from or written to disk to resolve hard page faults. (Hard page faults occur when a process requires code or data that is not in its working set or elsewhere in physical memory. The code or data must then be retrieved from disk). This counter was designed as a primary indicator of the kinds of faults that cause system-wide delays. It is the sum of the numbers in the Memory\Page Reads/sec and Memory\Page Writes/sec counters. It includes pages retrieved to satisfy faults in the file system cache (usually requested by applications).	This must be controlled to a level such that there is no disk bottleneck to and from the disk.
Pool Nonpaged Bytes	Displays the number of bytes in the nonpaged pool, an area of system memory (physical memory used by the operating system) for objects that cannot be written to disk, but must remain in physical memory as long as they are allocated.	This counter should remain level. If this counter is steadily increasing, it can indicate a memory leak.

**Table 7 Memory Counters (continued)**

Counter	Description	Recommended Value
Pool Paged Bytes	Displays the number of bytes in the paged pool, an area of system memory (physical memory used by the operating system) for objects that can be written to disk when they are not being used.	This counter usually stops increasing at 196 MB on a server that has the /3GB switch set (270 MB without it). When this counter reaches its maximum, the server can become unresponsive. A continuously growing value can be indicative of handle leaks (check progress handles counters) or a growing SMTP queue.

## MSExchangeIS Counters

The following are MSExchangeIS performance object counters.

**Table 8 MSExchangeIS Counters**

Counter	Description	Recommended Value
Active Connection Count	Displays the number of connections to the Exchange store that have shown activity in the last 10 minutes.	The value of this counter is specific to your organization.
Active User Count	Displays the number of user connections that have shown activity in the last 10 minutes.	The value of this counter is specific to your organization.
Connection Count	Displays the number of client processes connected to the Exchange store.	The value of this counter is specific to your organization.

**Table 8 MExchangeIS Counters (continued)**

Counter	Description	Recommended Value
RPC Averaged Latency/sec	Displays RPC latency in milliseconds averaged for the past 1024 packets.	The counter is typically less than approximately 20 milliseconds in normal operations.
RPC Operations/sec	Displays the rate that RPC operations are occurring.	The value of this counter is specific to your organization.
RPC Requests	Displays the number of client requests currently being processed by the Exchange store.	This counter should typically be less than 10. If it is larger than 25, this is a likely indicator of a resource bottleneck. Only 100 requests can be handled at a time. If the RPC Requests reach 100, the client will experience refused connections.
User Count	Displays the actual count of users (not connections) currently using the Exchange store. Performance measurement must always be correlated with current user numbers when interpreting this counter.	The value of this counter is specific to your organization.
Virus Scan Queue Length	Displays the current number of outstanding requests that are queued for virus scanning.	The value of this counter is specific to your organization.

**Table 8 MExchangeIS Counters (continued)**

Counter	Description	Recommended Value
VM Largest Block Size	Displays the size in bytes of the largest free block of virtual memory. This counter is a line that slopes down as virtual memory is consumed. When this counter drops below 32 MB, Exchange 2000 logs a warning in the event log (Event ID=9582) and logs an error if this number drops below 16 MB.	This counter should remain above 32 MB.
VM Total 16-MB Free Blocks	Displays the total number of free virtual memory blocks that are greater than or equal to 16 MB. This line forms a pyramid as you monitor it. It starts with one block of virtual memory greater than 16 MB and progresses to smaller blocks greater than 16 MB. By monitoring the trend on this counter, you can predict when the number of 16-MB blocks is likely to drop below 3, at which point restarting all the services on the node is recommended.	This counter should remain above three 16-MB blocks.

**Table 8 MExchangeIS Counters (continued)**

Counter	Description	Recommended Value
VM Total Free Blocks	Displays the total number of free virtual memory blocks regardless of size. This line forms a pyramid as you monitor it. This counter can be used to measure the degree to which available virtual memory is being fragmented. The average block size is the Process\Virtual Bytes\STORE.EXE instance divided by MExchangeIS\VM Total Free Blocks.	The value of this counter is specific to your organization.
VM Total Large Free Block Bytes	Displays the sum in bytes of all the free virtual memory blocks that are greater than or equal to 16 MB. This line slopes down as memory is consumed. This counter monitors store memory fragmentation.	This counter should stay above 50 MB.

## MSExchangeIS Mailbox Counters

The following are MSExchangeIS Mailbox performance object counters.

**Table 9 MSExchangeIS Mailbox Counters**

Counter	Description	Recommended Value
Active Client Logons	Displays the number of clients that performed any action within the last 10-minute time interval.	The value of this counter is specific to your organization.
Message Opens/sec	Displays the rate that requests to open messages are submitted to the Exchange store.	The value of this counter is specific to your organization.
Receive Queue Size	Displays the number of messages in the mailbox store's receive queue.	This counter should remain generally at zero during normal operations.
Send Queue Size	Displays the number of messages in the mailbox store's send queue.	This counter should remain generally at zero during normal operations.
Local Delivery Rate	Displays the rate at which messages are being delivered locally.	The value of this counter is specific to your organization.

## MSExchangeIS Public Counters

The following are MSExchangeIS Public performance object counters.

**Table 10 MSExchangeIS Public Counters**

Counter	Description	Recommended Value
Folders Open/sec	Displays the rate that requests to open folders are submitted to the Exchange store.	The value of this counter is specific to your organization.
Message Open/sec	Displays the rate that requests to open messages are submitted to the Exchange store.	The value of this counter is specific to your organization.
Receive Queue Size	Displays the number of messages in the public store's receive queue.	Generally, this counter should remain at zero during normal operations.
Send Queue Size	Displays the number of messages in the public store's send queue.	Generally, this counter should remain at zero during normal operations.

## Network Interface Counters

The following are Network Interface performance object counters. These counters are monitored using all instances.

**Table 11 Network Interface Counters**

Counter	Description	Recommended Value
Bytes Received/sec	Displays the rate at which bytes are received on the interface, including framing characters.	The value of this counter is specific to your organization.
Bytes Sent/sec	Displays the rate at which bytes are sent on the interface, including framing characters.	The value of this counter is specific to your organization.
Bytes Total/sec	Displays the rate at which bytes are sent and received on the interface, including framing characters.	The value of this counter is specific to your organization.
Output Queue Length	Displays the length of the output packet queue. A queue length of 1 or 2 is often satisfactory. Longer queues indicate that the adapter is waiting for the network and therefore cannot keep pace with the server.	This counter should remain below 1 or 2.

## Paging File Counters

The following are Paging File performance object counters.

**Table 12 Paging File Counters**

Counter	Description	Recommended Value
% Usage	Displays the amount of the paging file that is in use during the sample interval, as a percentage. A high value indicates that you may need to increase the size of your Pagefile.sys file or add more RAM.	Microsoft recommends keeping this value below 75 percent.

## Physical Disk Counters

The following are Physical Disk performance object counters.

**Table 13 Physical Disk Counters**

Counter	Description	Recommended Value
Avg. Disk sec/Transfer	Displays how fast data is being moved, in seconds. A high value might indicate that the system is retrying requests due to lengthy queuing or, less commonly, a disk failure.	Watch this counter for significant variances from baseline data.
Avg. Disk sec/Write	Displays the average time in seconds of a write of data to the disk.	This counter should remain below the manufacturer's specifications. A general threshold is well below 20 milliseconds. If a disk system has a write cache, then typical values are about 1 millisecond per write.
Avg. Disk sec/Read	Displays the average time in seconds of a read of data to the disk.	This counter should remain below the manufacturer's specifications. A general threshold is well below 20 milliseconds.
Current Disk Queue Length	Displays the instantaneous value of the disk queue for a particular physical disk.	If this is not hitting zero periodically there is likely to be a disk bottleneck.

---

<b>Counter</b>	<b>Description</b>	<b>Recommended Value</b>
Average Disk Queue Length	Displays the average value of the disk queue for a particular physical disk.	This should typically be less than the number of spindles in the RAID array.

## Process Counters

The following are Process performance object counters. Select the different Exchange processes that you want to monitor as the instance of these counters.

**Table 14 Process Disk Counters**

Counter	Description	Recommended Value
% Processor Time	Displays the percentage of time the processor is running non-idle threads for a given process. You can use this counter to monitor the percent each Exchange service is using the processor.	The value of this counter is specific to your organization and the process in question.
Elapsed Time	Displays the number of seconds a process has been running. It gives you a quick way to see whether a server or service has recently been restarted without having to look through the event log. A zero value indicates a non-active process.	The value of this counter is specific to your organization.
Handle Count	Displays the total number of handles currently open by this process. This number is the sum of the handles currently open by each thread in this process.	The handles opened by System Attendant, message transfer agent (MTA), and Exchange store should remain fairly constant. Inetinfo handles can grow radically during queue buildup.

**Table 14 Process Disk Counters (continued)**

Counter	Description	Recommended Value
Page faults/sec	Displays the rate Page Faults occur in the threads running in this process. A page fault occurs when a thread refers to a virtual memory page that is not in its working set in main memory.	Use this counter to monitor for processes lacking virtual memory.
Page File Bytes	Displays the current number of bytes this process has used in the paging files. Paging files are used to store pages of memory used by the process that are not contained in other files. Paging files are shared by all processes, and lack of space in paging files can prevent other processes from allocating memory.	The value of this counter is specific to your organization.
Pool Nonpaged Bytes	Displays the number of bytes in the nonpaged pool, an area of system memory (physical memory used by the operating system) for objects that cannot be written to disk, but must remain in physical memory as long as they are allocated.	The value of this counter is specific to your organization.
Private Bytes	Displays the current number of bytes this process has allocated that cannot be shared with other processes.	System Attendant, MTA, and Exchange store private bytes should remain constant except when background tasks run. Inetinfo private bytes can grow radically during queue buildup.

Table 14 Process Disk Counters (continued)

Counter	Description	Recommended Value
Virtual Bytes	Displays the current size in bytes of the virtual address space the process is using.	Virtual bytes is most important for the Exchange store process, where it only has 2 GB or 3 GB of virtual address space to work with whether running with the /3GB switch or not. On a large server with the /3GB switch, this counter should stay below 2.8 GB.
Working Set	Displays the current number of bytes in the working set of this process. The working set is the set of memory pages used recently by the threads in the process. If free memory in the computer is above a threshold, pages are left in the working set of a process even if they are not in use. When free memory falls below a threshold, pages are trimmed from working sets. If they are needed, they are then soft-faulted back into the working set before they leave main memory.	System Attendant, MTA, and Exchange store working sets should remain constant except when background tasks run. Inetinfo working set can grow radically during queue buildup.

---

## Processor Counters

The following are Processor performance object counters.

**Table 15 Processor Counters**

Counter	Description	Recommended Value
% Processor Time	Displays the percentage of time the processor is being used by processes running on the server.	The value of this counter is specific to your organization.

## Server Counters

The following are Server performance object counters.

**Table 16 Server Counters**

Counter	Description	Recommended Value
Pool Nonpaged Bytes	Displays the number of bytes of non-pageable computer memory the server is using.	The value of this counter is specific to your organization.
Pool Nonpaged Failures	Displays the number of times allocations from nonpaged pool have failed. If this number is high, either the amount of RAM is too small or the paging file is too small, or both. If this number is consistently increasing, increase the physical RAM and the size of the paging file.	The value of this counter is specific to your organization.
Work Item Shortages	Displays the number of times the STATUS_DATA_NOT_ACCEPTED message was returned at receive indication time. This occurs when no work item is available or can be allocated to service the incoming request. This counter shows whether the InitWorkItems or MaxWorkItems parameters might need to be adjusted.	If the value reaches the recommended threshold of 3, consider tuning the InitWorkItems or MaxWorkItems entries in the registry (in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\lanmanserver\Parameters).

## Server Work Queues Counters

The following are Server Work Queues performance object counters.

**Table 17 Server Work Queues Counters**

Counter	Description	Recommended Value
Active Threads	Displays the number of threads currently working on a request from the server client for this CPU. The system keeps this number as low as possible to minimize unnecessary context switching. This is an instantaneous count for the CPU, not an average over time.	The value of this counter is specific to your organization.
Queue Length	Displays the current length of the server work queue for this CPU. A sustained queue length greater than four might indicate processor congestion. This is an instantaneous counter; observe its value over several intervals.	This counter should remain below 4.
Read Bytes/sec	Displays the rate the server is reading data from files for the clients on this CPU. This value is a measure of how busy the server is.	The value of this counter is specific to your organization.
Write Bytes/sec	Displays the rate the server is writing data to files for the clients on this CPU. This value is a measure of how busy the server is.	The value of this counter is specific to your organization.
Write Operations/sec	Displays the rate the server is performing file write operations for the clients on this CPU. This value is a measure of how busy the server is.	This value should always be 0 in the Blocking Queue counter instance.

## SMTP Server Counters

The following are SMTP server performance object counters.

**Table 18 SMTP Server Counters**

Counter	Description	Recommended Value
Categorizer Queue Length	Indicates how well SMTP is processing LDAP lookups against global catalog servers. This should be at or around zero unless you are expanding distribution lists. This is an excellent counter that tells you how healthy your global catalogs are. If access to your global catalogs is slow, this counter can increase.	This counter should remain at or around zero.
Local Queue Length	Displays the number of messages in the local SMTP queue.	The value of this counter is specific to your organization.
Messages Delivered/sec	Displays the rate that messages are being delivered to local mailboxes.	The value of this counter is specific to your organization.
Messages Received/sec	Displays the rate that messages are being received.	The value of this counter is specific to your organization.
Messages Sent/sec	Displays the rate that messages are being sent.	The value of this counter is specific to your organization.

## System Counters

The following are System performance object counters.

**Table 19 System Counters**

Counter	Description	Recommended Value
Processor Queue Length	Displays the number of threads in the processor queue. There is a single queue for processor time, even on computers with multiple processors. This counter shows ready threads only, not threads that are currently running. This value should be 2 or less.	This counter should remain at or below 2.
System Up Time	Displays the elapsed time (in seconds) that the computer has been running since it was last started.	The value of this counter is specific to your organization.

## TCP Counters

The following are TCP performance object counters.

**Table 20 TCP Counters**

Counter	Description	Recommended Value
Segments Received/Sec	Displays the rate at which segments are received, including those received in error. This count includes segments received on currently established connections. A low value means that you have too much broadcast traffic.	A low value means that you have too much broadcast traffic.
Segments Retransmitted/Sec	Displays the rate at which segments containing one or more previously transmitted bytes are retransmitted. A high value can indicate either a saturated network or a hardware problem.	A high value might indicate either a saturated network or a hardware problem.

## Thread Counters

The following are Thread performance object counters.

**Table 21 Thread Counters**

Counter	Description	Recommended Value
% Processor Time	Displays the percentage of elapsed time that a thread used the processor to run instructions.	Watch for threads that consume a high amount of processor time.
ID Thread	Displays the unique identifier of this thread. ID Thread numbers are reused, so these numbers only identify a thread for the lifetime of that thread.	None.
Thread State	<p>Displays the current state of the thread. States include:</p> <ul style="list-style-type: none"> <li>0 for Initialized</li> <li>1 for Ready</li> <li>2 for Running</li> <li>3 for Standby</li> <li>4 for Terminated</li> <li>5 for Wait</li> <li>6 for Transition</li> <li>7 for Unknown</li> </ul> <p>A Running thread is using a processor; a Standby thread is about to use one. A Ready thread wants to use a processor, but is waiting for a processor because none are free. A thread in Transition is waiting for a resource in order to run, such as waiting for its execution stack to be paged in from disk. A Waiting thread does not use the processor because it is waiting for a peripheral operation to complete or a resource to become free.</p>	None.

**Table 21 Thread Counters (continued)**

Counter	Description	Recommended Value
Thread Wait Reason	<p>Thread Wait Reason is only applicable when the thread is in the Wait state (see Thread State). States include:</p> <p>0 or 7 when the thread is waiting for the Executive,</p> <p>1 or 8 for a Free Page</p> <p>2 or 9 for a Page In</p> <p>3 or 10 for a Pool Allocation</p> <p>4 or 11 for an Execution Delay</p> <p>5 or 12 for a Suspended condition</p> <p>6 or 13 for a User Request</p> <p>14 for an Event Pair High</p> <p>15 for an Event Pair Low</p> <p>16 for an LPC Receive</p> <p>17 for an LPC Reply</p> <p>18 for Virtual Memory</p> <p>19 for a Page Out</p> <p>20 and higher are not assigned at the time of this writing.</p> <p>Event Pairs are used to communicate with protected subsystems.</p>	None.

## RAID Levels

Although there are many different implementations of RAID technologies, they all share two similar aspects. They all use multiple physical disks to distribute data, and they all store data according to a logic that is independent of the application for which they store data.

This section discusses four primary implementations of RAID: RAID-0, RAID-1, RAID 0+1, and RAID-5. Although there are many other RAID implementations, these four types serve as a representation of the overall scope of RAID solutions.

### RAID-0

RAID-0 is a striped disk array; each disk is logically partitioned in such a way that a “stripe” runs across all the disks in the array to create a single logical partition. For example, if a file is saved to a RAID-0 array and the application that is saving the file saves it to drive D, the RAID-0 array distributes the file across logical drive D, as in the following figure. In this example, it spans all six disks.




**Figure 11 RAID-0 disk array**

From a performance perspective, RAID-0 is the most efficient RAID technology because it can write to all six disks at once. When all disks store the application data, the most efficient use of the disks occurs.

The drawback to RAID-0 is its lack of reliability. If the Exchange mailbox databases are stored across a RAID-0 array and a single disk fails, you must restore the mailbox databases to a functional disk array and restore the transaction log files. In addition, if you store the transaction log files on this array and you lose a disk, you can perform only a restoration of the mailbox databases from the last backup.

## RAID-1

RAID-1 is a mirrored disk array in which two disks are mirrored as in the following figure.




**Figure 12 RAID-1 disk array**

RAID-1 is the most reliable of the three RAID disk arrays because all data is mirrored after it is written. You can use only half of the storage space on the disks. Although this may seem inefficient, RAID-1 is the preferred choice for data that requires the highest possible reliability.


## RAID-0+1

A RAID-0+1 disk array allows for the highest performance while ensuring redundancy by combining elements of RAID-0 and RAID-1 as in the following figure.

### Mirrored



### Striped



**Figure 13 RAID-0+1 disk array**

In a RAID-0+1 disk array, data is mirrored to both sets of disks (RAID-1), and then striped across the drives (RAID-0). Each physical disk is duplicated in the array. If you have a six-disk RAID-0+1 disk array, three disks are available for data storage.

## RAID-5

RAID-5 is a striped disk array, similar to RAID-0 in that data is distributed across the array; however, RAID-5 also includes parity. This means that a mechanism maintains the integrity of the data stored in the array, so that if one disk in the array fails, the data can be reconstructed from the remaining disks as in the following figure.

Therefore, RAID-5 is a reliable storage solution.



**Figure 14 RAID-5 disk array**

However, to maintain parity among the disks,  $1/n$  gigabyte (GB) of disk space is sacrificed (where  $n$  equals the number of drives in the array). For example, if you have six 9-GB disks, you have 45 GB of usable storage space. To maintain parity, one write of data is translated into two writes and two reads in the RAID-5 array; thus, overall performance is degraded.

The advantage of a RAID-5 solution is that it is reliable and uses disk space more efficiently than RAID-1 (and 1+0).

For more information on comparing RAID solutions and RAID levels, as well as Storage Area Network (SAN) and Network Attached Storage (NAS) solutions, see the *Storage Solutions for Microsoft® Exchange 2000 Server* white paper at <http://go.microsoft.com/fwlink/?LinkId=1715>.

---

## Additional Resources

The following technical papers and Microsoft Knowledge Base articles provide valuable information about troubleshooting Exchange 2000 performance.

---

### Web Sites

- Microsoft Operations Manager  
<http://www.microsoft.com/mom/>
- Exchange 2000 Management Pack For Microsoft Operations Manager  
<http://go.microsoft.com/fwlink/?LinkId=16451>

---

### Technical Papers

The following technical papers are available on the Web at <http://www.microsoft.com/exchange>

- Microsoft Exchange 2000 Internals: Quick Tuning Guide  
<http://go.microsoft.com/fwlink/?LinkId=9942>
- Storage Solutions for Microsoft® Exchange 2000 Server  
<http://go.microsoft.com/fwlink/?LinkId=1715>

## Microsoft Knowledge Base Articles

The following Microsoft Knowledge Base articles are available on the Web at <http://support.microsoft.com/>.

- 294818 – “Frequently Asked Questions About Network Monitor”  
(<http://support.microsoft.com/?kbid=294818>)
- 148942 – “How to Capture Network Traffic with Network Monitor”  
(<http://support.microsoft.com/?kbid=148942>)
- 317411 – “XADM: How To Gather Data to Troubleshoot Exchange Virtual Memory Issues”  
(<http://support.microsoft.com/?kbid=317411>)
- 296073 – “XADM: Monitoring for Exchange 2000 Memory Fragmentation”  
(<http://support.microsoft.com/?kbid=296073>)
- 266096 – “XGEN: Exchange 2000 Requires /3GB Switch with More Than 1 GB of Physical RAM”  
(<http://support.microsoft.com/?kbid=266096>)
- 253251 – “Using Diskperf in Windows 2000”  
(<http://support.microsoft.com/?kbid=253251>)

**For more information:** <http://www.microsoft.com/exchange>.

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

[mailto:exchdocs@microsoft.com?subject=Troubleshooting Microsoft Exchange 2000 Server Performance Problems](mailto:exchdocs@microsoft.com?subject=Troubleshooting%20Microsoft%20Exchange%202000%20Server%20Performance%20Problems)