

Microsoft **Exchange** 2000 **Server**

Storage Solutions for Microsoft® Exchange 2000 Server

Exchange Core Documentation



Published: August 2000

Updated: September 2003

Applies To: Exchange 2000 Server SP3

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 White Paper is for informational purposes only. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, 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.

© 2002–2003 Microsoft Corporation. All rights reserved.

Microsoft and Windows 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

- Introduction..... 1
- Planning a Storage Solution 2
 - General Storage Principles 2
 - Exchange 2000 Considerations 2
- Overview of Storage Technologies..... 3
 - RAID Levels 4
 - Storage Area Network (SAN) Solutions 6
 - Shadow Copy Backups..... 8
 - Network Attached Storage (NAS) Solutions..... 9
- Placing Exchange Data on the Storage Device..... 9
 - SMTP Queue Directory..... 9
 - .EDB and .STM Files 10
 - Transaction Log Files..... 10
- Additional Resources 11

Storage Solutions for Microsoft Exchange 2000 Server

Published: August 2000

Updated: September 2003

Introduction

As you plan your storage strategy for Microsoft® Exchange 2000 Server or any other application that stores important data, you need to balance three criteria: capacity, availability, and performance. The choices you make as you plan and implement your storage solution affect the cost associated with administration and maintenance of your Exchange 2000 environment.

- **Capacity** In Exchange 2000, your total capacity is roughly equal to the number of mailboxes multiplied by the amount of storage allocated to each mailbox. If your organization is supporting public folders, you must add the appropriate amount of disk space to accommodate public folder storage.
- **Availability** The level of e-mail availability required by your messaging system depends on your company needs. For some companies, e-mail usage is light and considered non-essential; but for many companies today, e-mail is a mission-critical service. The priority that your company places on e-mail determines the level of investment and resources allocated to a reliable e-mail solution. Overall availability is increased by redundancy. This might mean clustering applications to provide CPU redundancy or implementing a redundant array of independent disks (RAID) solution to provide data redundancy.
- **Performance** Performance requirements are also unique to each organization. This document refers to performance as it relates to throughput. With regard to storage technology, throughput is measured by how many reads and writes per second a storage device can perform when coupled with software logic.

Before you design your storage solution for Exchange 2000, determine how your company prioritizes these three criteria, especially when considering a balance between availability and performance.

This document discusses the principles of designing an Exchange 2000 storage solution. This document also compares two common storage solutions: storage area networks (SANs) and network-attached storage (NAS). However, this document does not provide procedures for configuring and deploying Exchange 2000 storage solutions, nor does it discuss storage from a clustering perspective—although the principles outlined in this article are applicable to a clustered version of Exchange. This document focuses mainly on mailbox storage, but the principles and concepts apply to public folder storage as well.

To fully understand the concepts within this document, you should have a basic knowledge of storage technology in Exchange 2000. To become familiar with basic storage technology, read the "Information Store" section in Chapter 2 of the *Microsoft Exchange 2000 Server Planning and Installation* guide.

Planning a Storage Solution

When you install Exchange 2000, all data is stored locally, by default, on the drive on which you install the application. To determine the capacity, level of availability, and performance associated with this default configuration, you must consider the following factors:

- Number and speed of CPUs
- Server type (mailbox server, public folder server, Instant Messaging server, Chat server, connector server, and so forth)
- Number of physical disks

Because of the many variables, Exchange 2000 server sizing is outside the scope of this document. In general, however, if the default configuration does not meet your requirements, you should plan a new storage solution that maximizes capacity, performance, and availability for Exchange. The remainder of this section discusses the factors you should consider.

General Storage Principles

Regardless of the application you are running, consider the following storage principles to help you maximize capacity, performance, and availability:

- You can decrease the processing required from the CPU by implementing a specialized hardware solution, such as RAID arrays or a storage area network (SAN) that incorporates RAID technology. This assumes that the hardware solution includes its own processing capabilities.
- You can also decrease CPU processing time by separating files that are accessed sequentially from files that are accessed randomly. Storing sequentially-accessed files separately keeps the disk heads in position for sequential input/output (I/O), which reduces the amount of time required to locate data.
- Multiple small disks perform better than a single large disk. For example, if you need to store 36 GB of data, consider using four 9-GB disks instead of one 36-GB disk. Depending on the type of array, this could allow information to be written as much as four times faster.

Exchange 2000 Considerations

When planning your storage solution, consider the following information about Exchange 2000:

- All data stored on Exchange is not managed in the same way; thus, a single storage solution for all data types is not the most efficient.
- Servers that do not host mailboxes or public folders, such as connector servers, may not benefit from advanced storage solutions because they typically store data for a short time and then forward the data to another server. In some cases, you might need RAID-0 for these types of services.

- Exchange 2000 uses an Installable File System (IFS) driver. This driver requires access to physical disk characteristics that are reported by block mode storage devices. If you store Exchange 2000 databases on a device that does not appear to Microsoft Windows® as a block mode storage device, Exchange will not mount the databases. (Earlier versions of Exchange Server do not include an IFS driver and do not require block mode storage devices.)
- An Exchange 2000 server supports up to four storage groups. Each storage group has its own set of transaction logs and supports up to five databases. Your disaster recovery strategy plays an important role in determining how many storage groups and databases your storage solution should support. Generally, you should keep each storage group on its own array. However, if you want to restore individual databases, you can move each database to its own array.
- In Exchange, transaction logs are accessed sequentially, and databases are accessed randomly. In accordance with general storage principles, you should separate the transaction logs (sequential I/O) from databases (random I/O) to maximize performance and increase fault tolerance. Specifically, you should move each set of transaction logs to its own array, separate from storage groups and databases.

Overview of Storage Technologies

When planning your storage solution, it is important to familiarize yourself with the following storage-related technologies:

- **RAID Levels** Disk array implementations that offer varying levels of performance and fault tolerance.
- **Storage Area Network (SAN) Solutions** Storage that provides centralized data storage by means of a high-speed network.
- **Network Attached Storage (NAS) Solutions** Storage that connects directly to servers through existing network connections.

SAN and NAS storage solutions usually incorporate RAID technologies. You can configure the discs on the storage device to use a RAID level that is appropriate for your performance and fault tolerance needs. Use the information in the following sections to compare and contrast these storage technologies.

Important It is generally recommended that you use a Direct Access Storage (DAS) or Storage Area Network (SAN) attached disk storage solution because this configuration optimizes performance and reliability for Exchange 2000. Microsoft does not support NAS storage solutions.

For information about SAN and NAS solutions, see Microsoft Knowledge Base article [328879](http://support.microsoft.com?kbid=328879), "Using Exchange Server with Storage-Attached Network and Network-Attached Storage Devices" (<http://support.microsoft.com?kbid=328879>).

It is recommended that you contact your vendor before you deploy any storage solution for Exchange 2000 databases to obtain assurance that the end-to-end solution is designed for Exchange 2000 use. Many vendors have best practices guides for Exchange.

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 are storing data.

This article 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 an adequate 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 (see Figure 1). In this example it spans all six disks.



Figure 1 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 point-in-time restoration of the mailbox databases from the last backup.

RAID-1

RAID-1 is a mirrored disk array in which two disks are mirrored (see Figure 2).




Figure 2 RAID-1 disk array

RAID-1 is the most reliable of the three RAID 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 (see Figure 3).

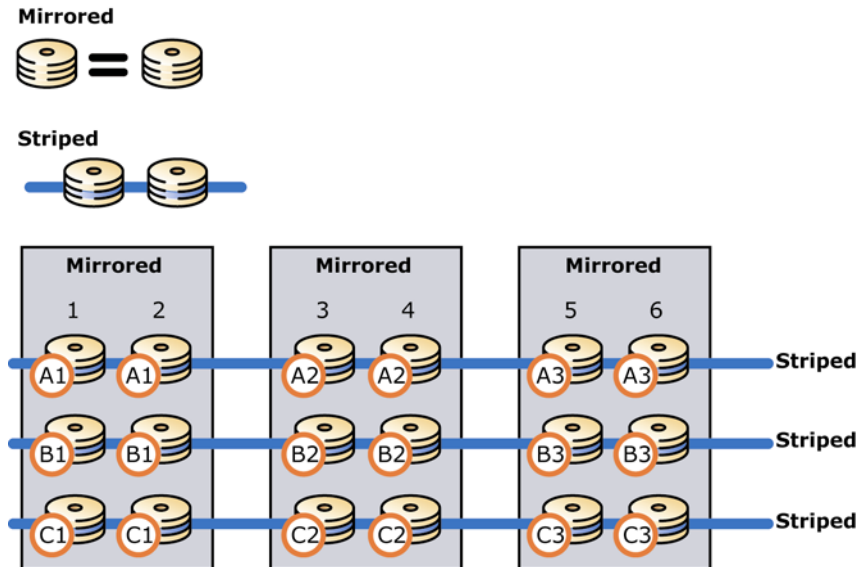


Figure 3 RAID-0+1 disk array

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 there is a mechanism that maintains the integrity of the data stored on the array, so that if one disk in the array fails, the data can be reconstructed from the remaining disks (see Figure 4). Thus, RAID-5 is a reliable storage solution.



Figure 4 RAID-5 disk array

However, to maintain parity among the disks, $1/n$ 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).

Comparing RAID Solutions

Because capacity is relatively stable, it is helpful to evaluate these RAID solutions by comparing cost, performance, and reliability against a constant capacity. Table 1 is based on the following assumptions:

- You are storing 90 GB of data.
- You are using 9-GB drives.
- Your arrays can write data to disks at the rate of 100 input/output (I/O) processes per second.

Table 1 Comparing RAID solutions

RAID solution	Number of drives (cost)	Maximum writes/second	Maximum reads/second	Reliability
RAID-0	10	1000	1000	Low
RAID-0+1	20	1000	2000	Very high
RAID-5	11	275	1100	High

Note RAID-1 is not evaluated in the table because only two disks can be implemented in a RAID-1 solution. You need two 45-GB drives to store 90 GB of data, which would result in much lower throughput.

You assess reliability by evaluating the impact that a disk failure would have on the integrity of the data. RAID-0 does not implement any kind of redundancy, so a single disk failure on a RAID-0 array requires a full restoration of data. RAID-0+1 is the most reliable solution of the three because two or more disks must fail before data is potentially lost; in other words, very specific sets of disks must fail before data is lost.

You evaluate cost by calculating the number of disks needed to support your array. The RAID-0+1 implementation is the most expensive because you must have twice as much disk space as you actually need. However, this configuration also yields much higher performance than the same-capacity RAID-5 configuration, as judged by the maximum read and write rates.

Storage Area Network (SAN) Solutions

It is recommended that you use a Storage Area Network (SAN) for the storage of your Exchange files; this configuration optimizes server performance and reliability.

A storage area network (SAN) provides storage and storage management capabilities for company data. SANs use Fibre Channel switching technology to provide fast and reliable connectivity between storage and applications.

A SAN has three major component areas:

- Fibre Channel switching technology
- Storage systems on which data is stored and protected
- Storage and SAN management software

Hardware vendors sell complete SAN packages that include the necessary hardware, software, and support. SAN software manages network and data flow redundancy by providing multiple paths to stored data (see Figure 5). Because SAN technology is relatively new and continues to evolve rapidly, you can plan and deploy a complete SAN solution to accommodate future growth and emerging SAN technologies. Ultimately, SAN technology will allow connectivity between heterogeneous systems with different operating systems to storage products from multiple vendors.




Figure 5 SAN storage solution

Currently, SAN solutions are best for large companies and for IT departments that need to store large amounts of data. A minimal deployment of a typical SAN solution may hold as much as 5 terabytes of data.

Although deployment can be expensive, a SAN solution could be preferable because the long-term total cost of ownership (TCO) may be lower than the cost of maintaining many small arrays. Consider the following advantages of a SAN solution:

- If you currently have multiple arrays managed by multiple administrators, centralized administration of all storage allows administrators to be available for other tasks.
- In terms of availability, no other single solution has the potential to offer the comprehensive and flexible reliability that a vendor-supported SAN provides. Some companies can expect enormous revenue loss when messaging services are down. If your company has the potential to lose significant revenue as a result of an unavailable messaging service, it could be cost-effective to deploy a specialized SAN solution.

Before you invest in a SAN, calculate the cost of your current storage solution in terms of hardware and administrative resources, and evaluate the company's need for dependable storage.

How a SAN Benefits Exchange

The following are advantages to implementing a SAN solution in your Exchange 2000 organization:

- Exchange 2000 requires high I/O bandwidth that is supported only by a channel-attached disk storage system, such as a SAN. In contrast, network storage solutions that rely on access to Exchange 2000 database files through the network stack can increase the risk of data corruption and performance loss.
- Exchange 2000 also requires mailbox and public folders stores to exist on a drive that is local to the Exchange server. This requirement is met by SAN solutions, which connect to Exchange servers through a local Fibre Channel connection. Other storage solutions that rely on a network redirector to process disk resources do not meet this requirement.
- SANs are highly scalable, which is an important consideration for Exchange. As mail data grows and mailbox limits are continually challenged, you must increase storage capacity and I/O rates. As your organization expands, a SAN allows you to easily add discs and spindles. Select a SAN that incorporates storage virtualization, which allows you to easily add storage and quickly reallocate it to your Exchange servers. With storage virtualization, you can purchase storage discs in accordance with your budget; even if the discs are of various capacities, a SAN that features storage virtualization is capable of immediately using all available disc space.
- The scalable nature of SANs also allows you to expand your Exchange organization by adding servers. SANs allow you to connect multiple Exchange servers to the same storage device, and then divide the storage among them.
- Through the use of volume mirroring and shadow copy backups, backup, recovery, and availability are all enhanced with a SAN (shadow copy backups are discussed in detail in the following section). Because SANs allow multiple connections, you can connect high-performance backup devices. SANs also allow you to designate different RAID levels to separate storage partitions.

Shadow Copy Backups

The Exchange 2000 online backup application programming interface (API) automatically synchronizes and gathers the Exchange 2000 database and transaction log file data that is required for successful restoration. An online backup of Exchange 2000 databases occurs through the same channel as normal database access. If this access is across the network, backup and restore operations might greatly increase peak bandwidth requirements.

To provide rapid backup and restore functionality, several SAN solutions bypass the Exchange 2000 online backup API. These backups are known as shadow copy backups. When considering a storage solution vendor, ensure that their custom shadow copy solution backs up and synchronizes all of the appropriate Exchange 2000 data files, and that it captures these data files in the correct state. If the vendor's solution does not meet these requirements, the shadow copy backup processes may cause issues with database reliability and consistency.

Important If you implement a shadow copy backup solution for Exchange 2000, the vendor of your shadow copy solution is your primary support provider for backup and recovery issues.

For more information about shadow copy backups and the limited support that Microsoft Product Support Services (PSS) may provide, see Microsoft Knowledge Base article [311898](http://support.microsoft.com?kbid=311898), "XADM: Hot Split Snapshot Backups of Exchange Server" (<http://support.microsoft.com?kbid=311898>).

Network Attached Storage (NAS) Solutions

Network attached storage (NAS) refers to products that use a server-attached approach to data storage. In this approach, the storage hardware connects directly to the Ethernet network through small computer system interface (SCSI) or Fibre Channel connections. A NAS product is a specialized server that contains a file system and scalable storage. In this model, data storage is decentralized; the NAS appliance connects locally to department servers, and therefore, the data is accessible only by local servers.

Important Exchange 2000 has local data access and I/O bandwidth requirements that NAS products do not generally meet. Incorrect use of Exchange 2000 software with a network-attached storage product might result in data loss, including total database loss. Therefore, Microsoft does not support using NAS with Exchange 2000.

For more information about why NAS solutions are not currently supported, see Microsoft Knowledge Base article [317173](http://support.microsoft.com?kbid=317173), "XADM: Exchange 2000 Server and Network-Attached Storage" (<http://support.microsoft.com?kbid=317173>).

Placing Exchange Data on the Storage Device

Exchange stores data in three main locations:

- Simple Mail Transfer Protocol (SMTP) queue directory
- .edb and .stm files
- Transaction log files

SMTP Queue Directory

The SMTP queue stores SMTP messages until they are written to a database (private or public, depending on the type of message), or sent to another server or connector.

Typically, messages stored in the SMTP queue are there for a short time. Therefore, your storage solution for the SMTP queue should optimize performance before capacity and reliability. However, in some situations, when downstream processes fail, the SMTP queue could be required to store a large amount of data. For that reason, do not assume that a RAID-0 array is the best solution for SMTP queues. Generally, RAID-0 is acceptable only if mail loss is acceptable. RAID-1 is a good solution because it gives some measure of reliability, while providing adequate throughput.

For more information about moving the SMTP queue directory from its default location, see Microsoft Knowledge Base article [318230](http://support.microsoft.com/?kbid=318230), "XCON: How to Change the Exchange 2000 SMTP Mailroot Directory Location" (<http://support.microsoft.com/?kbid=318230>).

.EDB and .STM Files

An Exchange database consists of a rich-text .edb file and a native multimedia content .stm file. The .edb file stores all of the MAPI messages, tables used by the store process to locate all messages, and checksums of both the .edb and .stm files. The .stm file contains messages that are transmitted with their native Internet content. Because access to these files is generally random, they can be placed on the same disk volume.

As you plan your storage solution for these files, you should assume a certain amount of reliability; in other words, RAID-0 is not a recommended option. After reliability, your storage solution is based on a choice between optimizing performance (RAID-1) and optimizing capacity (RAID-5). If possible, use RAID-1 (or 0+1) for these files.

For public folders, you could store these files on a RAID-5 array, because data on public folders is usually written once and read many times. RAID-5 provides better read performance than write performance.

Transaction Log Files

Each storage group generates its own set of transaction log files. Transaction log files maintain the state and integrity of .edb and .stm files. As new transactions occur, the transactions are simultaneously written in the log file and in memory. Log file transactions are not recognizable as Exchange messages, but they contain transaction data and specify where in the .edb file the data should be written. Before the transactions are committed to the .edb file, users access the transactions from memory. Then, when the load on the server has decreased, transactions are committed to the .edb file for permanent storage. The process of caching transactions in memory and deferring the update of the physical disk is referred to as a "lazy write."

If a disaster occurs, and you must rebuild a server, you use the latest transaction log files to rebuild your databases. If you have access to the transaction log files and the latest backup, you can recover all of your data. However, if you lose the transaction log files, the data is permanently lost.

You can significantly improve the performance and fault tolerance of Exchange servers by placing each set of transaction log files on a separate drive. Because each storage group has its own set of transaction logs, the number of dedicated transaction log drives for your server should equal the number of planned storage groups. With a SAN solution, select a product that allows you to easily partition the virtualized space into separate virtual drives for storage groups and transaction log files. In addition, because transaction log files are critical to the operation of a server, you should protect the drives against failure, ideally by hardware mirroring using RAID. A RAID level of 0+1 (in which data is mirrored and then striped) is recommended.

Tip Distribute the database drives across many SCSI channels or controllers, but configure them as a single logical drive to minimize SCSI bus saturation.

An example disk configuration is as follows:

C:\ System and boot (mirror set)

D:\ Pagefile

E:\ Transaction logs for storage group 1 (mirror set)

F:\ Transaction logs for storage group 2 (mirror set)

G:\ Database files for both storage groups (multiple drives configured as hardware stripe set with parity)

Note The file system for transaction log drives should always be formatted for NTFS.

For more information about transaction log files, see the technical paper *Disaster Recovery for Microsoft Exchange 2000 Server* at <http://go.microsoft.com/fwlink/?linkid=1714>.

Additional Resources

- Disaster Recovery for Microsoft Exchange 2000 Server
<http://go.microsoft.com/fwlink/?linkid=1714>
- 296787 XADM: Offline Backup and Restore Procedures for Exchange Server 4.0, 5.0, and 5.5
<http://support.microsoft.com/?kbid=296787>
- 296788 Offline Backup and Restoration Procedures for Exchange
<http://support.microsoft.com/?kbid=296788>
- 311898 XADM: Hot Split Snapshot Backups of Exchange Server
<http://support.microsoft.com/?kbid=311898>
- 317172 XADM: Exchange Server 5.5 and Network-Attached Storage
<http://support.microsoft.com/?kbid=317172>
- 317173 XADM: Exchange 2000 Server and Network-Attached Storage
<http://support.microsoft.com/?kbid=317173>
- 318230 XCON: How to Change the Exchange 2000 SMTP Mailroot Directory Location
<http://support.microsoft.com/?kbid=318230>

- 328879 Using Exchange Server with Storage-Attached Network and Network-Attached Storage Devices
<http://support.microsoft.com/?kbid=328879>

Did this document 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.](mailto:exchdocs@microsoft.com)