White Paper: Physical Storage Best Practices for Cisco Unity with Microsoft Exchange

Published July 12, 2002

This document provides the best practices recommended for physical storage and disk configuration with the use of RAID (redundant array of independent disks) for Cisco Unity™ version 3.1 and later, and Cisco Unity Bridge version 2.1 and later. Use this document to plan purchases and installation of hardware.

This document does not cover installing Microsoft Windows 2000, Microsoft SQL Server, or a vendor-specific configuration of RAID.

Contents

Choosing RAID Storage Levels

Best Practices

- Investigate the different hardware-based RAID storage options available for Cisco Unity.

Corporate Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

- Base your RAID configuration choice on the system size and the type of data to be stored and accessed. To maximize performance, data integrity, and reliability for Cisco Unity, Microsoft Exchange, and SQL Server transaction logs, use RAID 1. To maximize data storage and access capacity for Cisco Unity data, use RAID 5.

Cisco Unity Storage Configurations

RAID configurations and disk volumes for Cisco Unity are detailed in the following illustrations:

- Small- to Medium-Scale Cisco Unity RAID Configurations, Figure 1 on page 3
- Medium- to Large-Scale Cisco Unity RAID Configurations, Figure 2 on page 4
- Small-Scale Cisco Unity Disk Volumes, Figure 3 on page 5
- Medium-Scale Cisco Unity Disk Volumes, Figure 4 on page 6
- Large-scale Cisco Unity Disk Volumes, Figure 5 on page 7

System-size definitions

Table 1 correlates the system-size labels used in Figures 1 through 5 with the platforms supported for use with Cisco Unity. (For the most current list of supported platforms, refer to the *Cisco Unity Supported Platforms List*, available on Cisco.com at http://www.cisco.com/warp/public/cc/pd/unco/un/prodlit/ucutp_st.htm.)

Table 1 System-Size Labels and Cisco Unity Supported Platforms

System-Size Label	Cisco-Shipped Platforms	Dell Platforms	Compaq/ Hewlett-Packard Platforms	IBM Platforms
Small-scale	ICS 7750 MCS 7825-800 MCS 7827 MCS 7835-1000 MCS 7837	OptiPlex GX150 PowerEdge 1400SC	DL320G1 DL380G1 DL380G2	None
Medium-scale	MCS 7847	PowerEdge 2500	DL380G2 ML370G2	x232 x342
Large-scale	None	PowerEdge 4600 PowerEdge 6600 PowerEdge 6650	DL580G1 ML570G1	x250 x255

Figure 1 Small- to Medium-Scale Cisco Unity RAID Configurations

Each vertical column indicates a Cisco Unity server's entire storage configuration (numbers denote disk quantity, and letters denote variations within same disk quantity).

Message storage capacity generally increases from left to right.

Refer to notes below per flagnote indicators associated throughout this drawing.

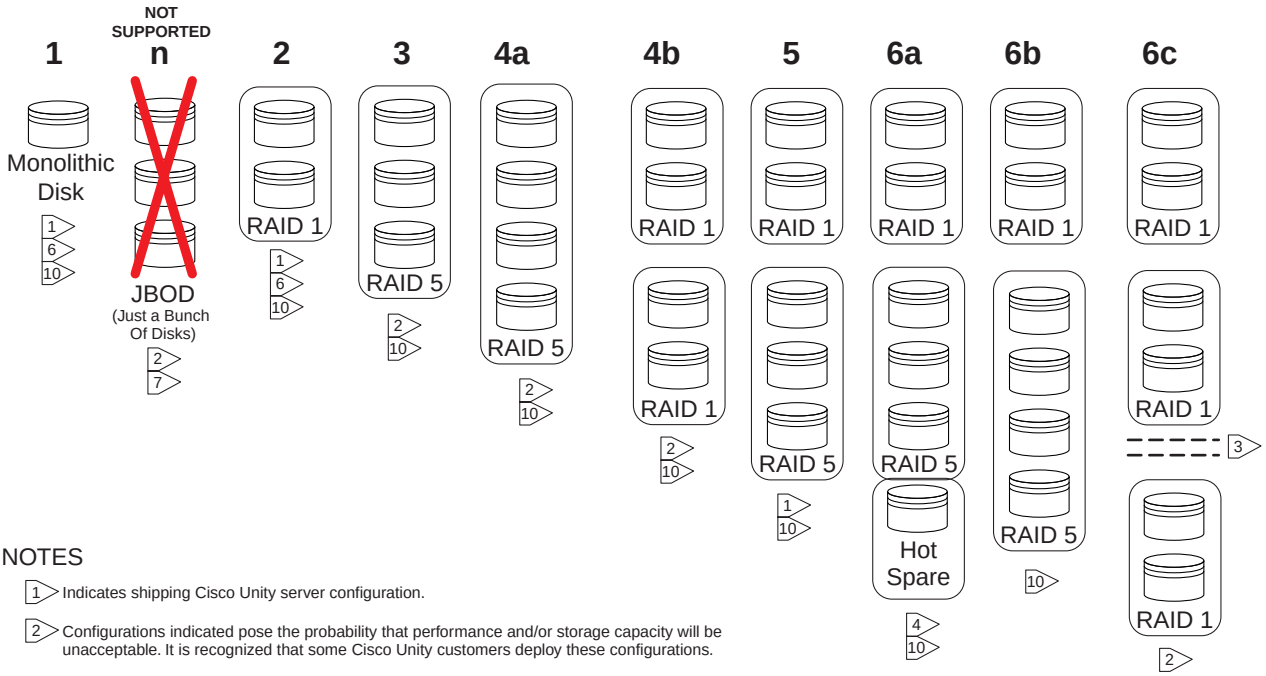

NOTES

1 Indicates shipping Cisco Unity server configuration.

2 Configurations indicated pose the probability that performance and/or storage capacity will be unacceptable. It is recognized that some Cisco Unity customers deploy these configurations.

3 Double dashed line in a vertical column indicates a "split" SCSI backplane between storage volumes above and below. Ideally, connect each side of the split to its own dedicated hardware RAID channel.

4 The concept of a "hot spare" is benign and does not affect Cisco Unity performance. Customer adoption is purely voluntary to mitigate the RAID performance impact of a failed drive before it can be replaced.

6 The active partition is limited to 4 GB in systems with a monolithic disk or a single RAID 1 volume to conserve the second partition capacity for Exchange message store + logs, the SQL Server data store + logs, and all applications including Cisco Unity and its trace logs, but excluding the OS.

7 All references to "RAID" imply hardware RAID only. A software RAID is technically still a "JBOD" configuration. Neither "JBOD" nor software RAID are supported or recommended, as the risk of serious performance degradation is significant.

10 Configurations indicated are not recommended for any large-scale or heavily burdened medium-scale deployments, particularly if the voice message store is located on the Cisco Unity server ("on box" scenario).

69699

Figure 2 Medium- to Large-Scale Cisco Unity RAID Configurations

Each vertical column indicates a Cisco Unity server's entire storage configuration (numbers denote disk quantity, and letters denote variations within same disk quantity).

Message storage capacity generally increases from left to right.

Refer to notes below per flagnote indicators associated throughout this drawing.

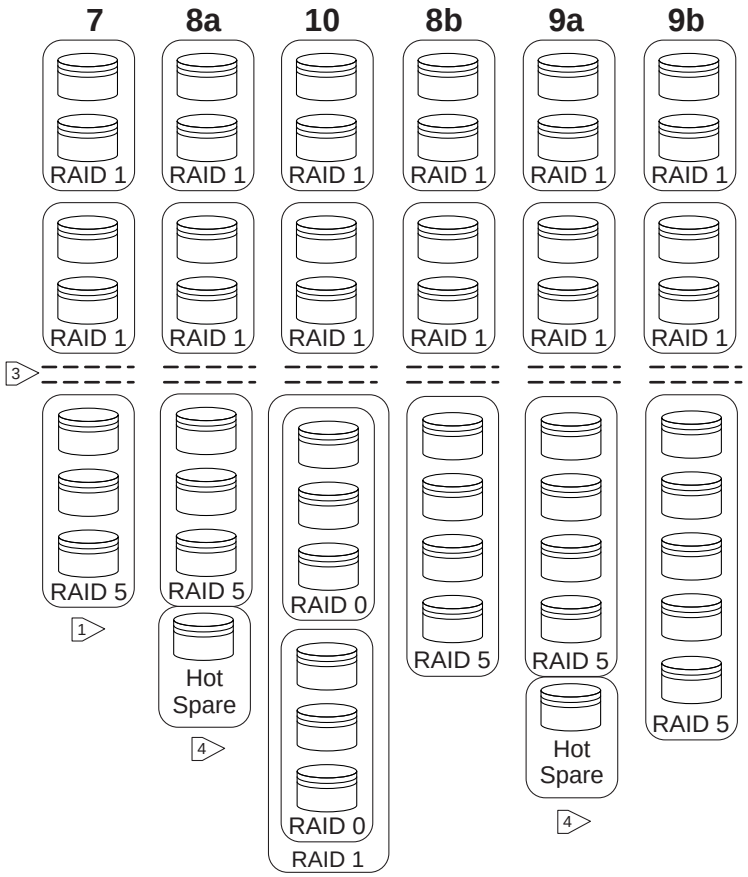

NOTES

1 Indicates shipping Cisco Unity server configuration.

3 Double dashed line in a vertical column indicates a "split" SCSI backplane between storage volumes above and below. Ideally, connect each side of the split to its own dedicated hardware RAID channel.

4 The concept of a "hot spare" is benign and does not affect Cisco Unity performance. Customer adoption is purely voluntary to mitigate the RAID performance impact of a failed drive before it can be replaced.

70000

Figure 3 Small-Scale Cisco Unity Disk Volumes

Configurations are indicated horizontally, per Figures 1 and 2.

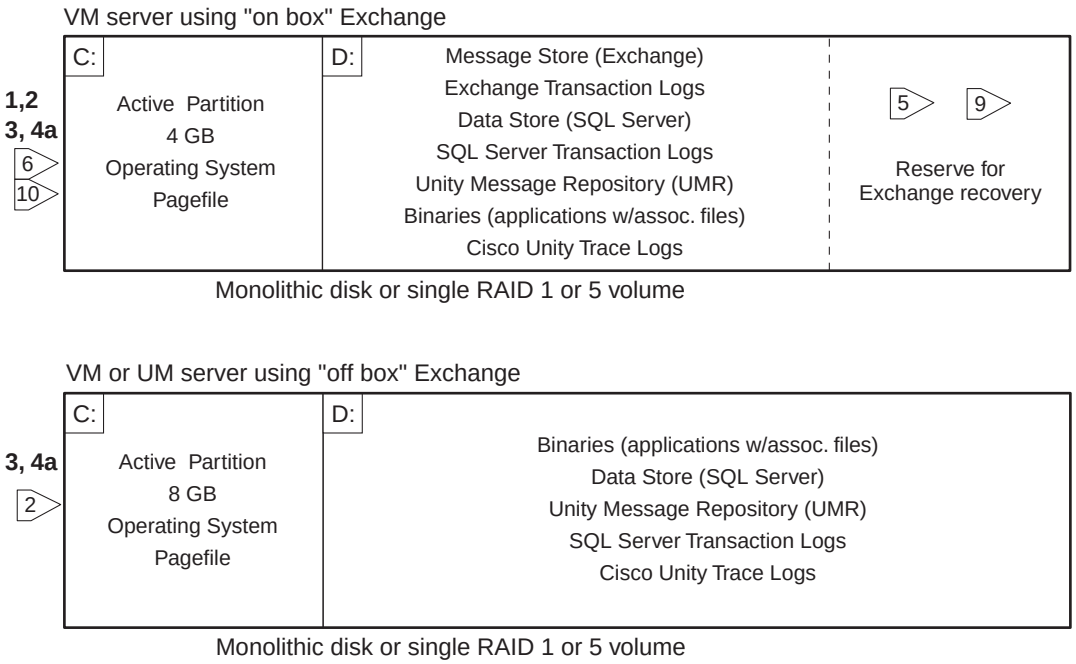

NOTES

2 Configurations indicated pose the probability that performance and/or storage capacity will be unacceptable. It is recognized that some Cisco Unity customers deploy these configurations.

5 50% of the capacity of the volume homing the Exchange message store is indicated as reserved to allow for recovery of Exchange using Eseutil.exe, given the assumption that the voice message store is located on the Cisco Unity server ("on box" scenario). If Exchange is "off box," this space need not be reserved.

6 The active partition is limited to 4 GB in systems with a monolithic disk or a single RAID 1 volume to conserve the second partition capacity for Exchange message store + logs, the SQL Server data store + logs, and all applications including Cisco Unity and its trace logs, but excluding the OS.

9 The 50% area can be partitioned (as a logical drive letter) to ensure it is reserved, but this will slow down the Exchange recovery process as the Mdb.dat file(s) will need to first be copied to the empty partition.

10 Configurations indicated are not recommended for any large-scale or heavily burdened medium-scale deployments, particularly if the voice message store is located on the Cisco Unity server ("on box" scenario).

76199

Figure 4 Medium-Scale Cisco Unity Disk Volumes

Configurations are indicated horizontally, per Figures 1 and 2.
Breaks between boxes separate physical storage volumes within configurations (monolithic, RAID 1, RAID 5).

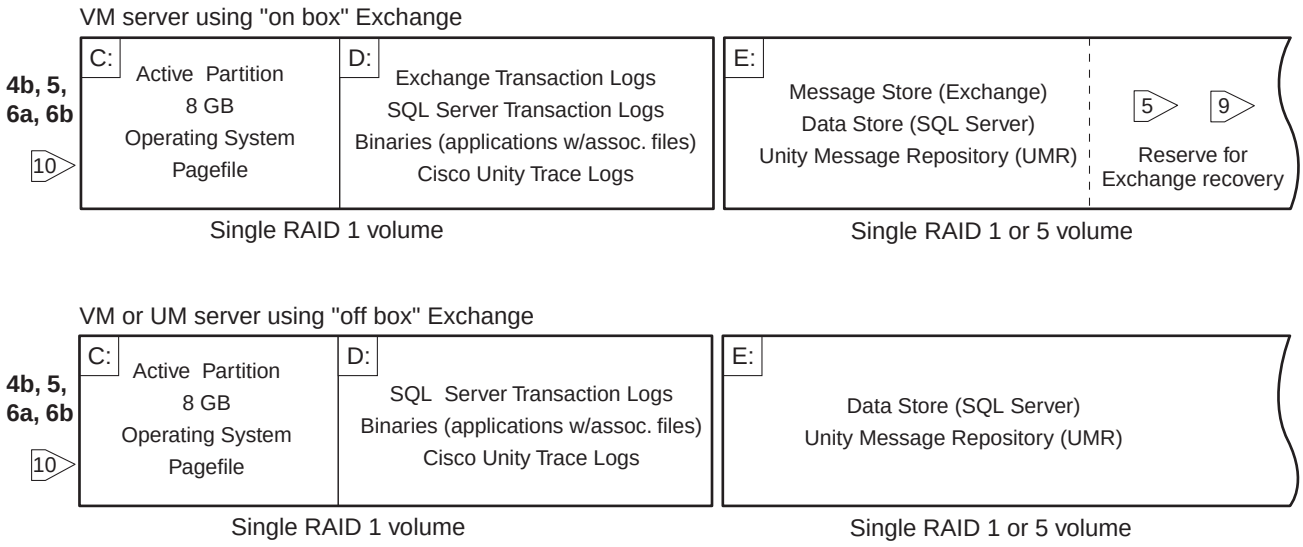

NOTES

5 50% of the capacity of the volume homing the Exchange message store is indicated as reserved to allow for recovery of Exchange using Eseutil.exe, given the assumption that the voice message store is located on the Cisco Unity server ("on box" scenario). If Exchange is "off box," this space need not be reserved.

9 The 50% area can be partitioned (as a logical drive letter) to ensure it is reserved, but this will slow down the Exchange recovery process as the Mdb.dat file(s) will need to first be copied to the empty partition.

10 Configurations indicated are not recommended for any large-scale or heavily burdened medium-scale deployments, particularly if the voice message store is located on the Cisco Unity server ("on box" scenario).

76200

Figure 5 Large-scale Cisco Unity Disk Volumes

Configurations are indicated horizontally, per preceding number from Figures 1 and 2.
Breaks between boxes separate physical storage volumes within configurations (monolithic, RAID 1, RAID 5).

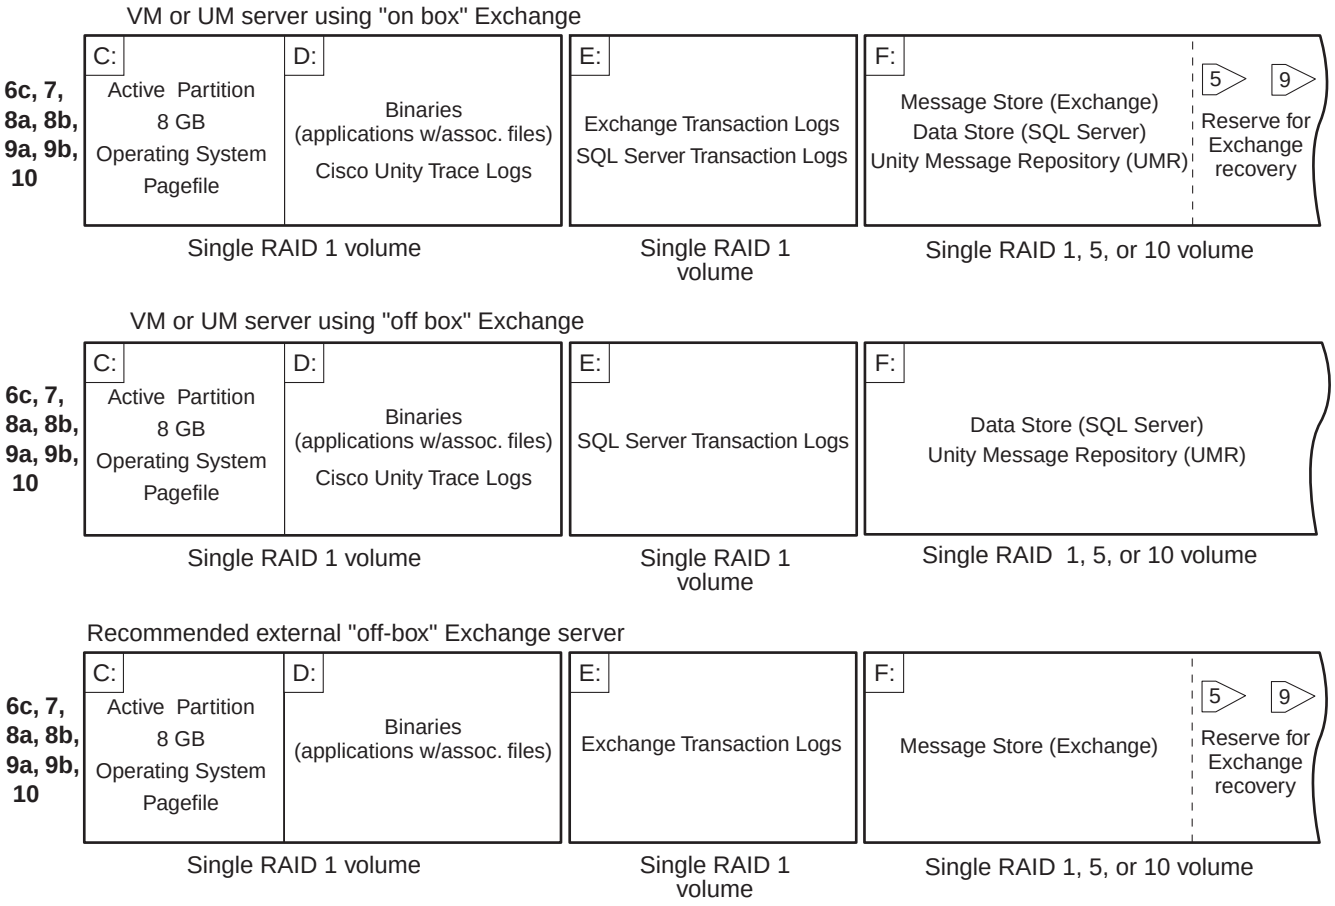

NOTES

5 50% of the capacity of the volume homing the Exchange message store is indicated as reserved to allow for recovery of Exchange using Eseutil.exe, given the assumption that the voice message store is located on the Cisco Unity server ("on box" scenario). If Exchange is "off box," this space need not be reserved.

9 The 50% area can be partitioned (as a logical drive letter) to ensure it is reserved, but this will slow down the Exchange recovery process as the Mdb.dat file(s) will need to first be copied to the empty partition.

76299

Information on RAID Storage Levels

Choosing the data storage configuration for a Cisco Unity system involves determining which RAID levels meet your organization's data storage goals.

RAID 0

RAID 0 writes data sequentially in equal-sized blocks across a set of drives, known as data striping.

RAID 0 provides better performance for random I/O applications than is possible with a single drive, because the I/O load is shared equally among the disks. However, RAID 0 does not provide data redundancy, so strict adherence to a frequent backup schedule is the only way to protect data.

Because RAID 0 is not fault tolerant and does not have a method for restoration of data from a failed drive, we recommend against implementing RAID 0 with Cisco Unity.

RAID 1

RAID 1 offers reliable data protection through the use of disk mirroring, which involves writing all of the data twice, once to each of two disks in a set. Reads can occur from either disk drive in the set.

If one disk fails, the RAID controller allows replacement of the failed disk, recovery of data using the information on the remaining disk, and remirroring while Cisco Unity, Exchange, Windows, and SQL Server remain online.

RAID 1 offers the best fault tolerance of any of the RAID levels, but it uses the most storage capacity, because two disks are required for the amount of data that would normally occupy one disk.

RAID 1 configurations are scalable and can be used successfully with any size Cisco Unity system. An even number of disks are required, with a minimum of two disks in the set.

Some literature recommends connecting the primary disk(s) and mirrored disk(s) to separate RAID controllers to provide additional fault tolerance. Such a configuration can be used for Cisco Unity at the customer's discretion.

RAID 4

RAID 4 uses data striping on an array of disks, with an additional disk dedicated solely to parity information. In the event of a disk failure, the parity information on the dedicated disk is used for data recovery.

The stripe size in RAID 4 is sufficient to accommodate an entire record, making it suitable for applications that use numerous simultaneous reads, but less efficient than other RAID options for database applications such as Cisco Unity that make multiple simultaneous database updates. For this reason, we do not recommend RAID 4 for use with Cisco Unity.

RAID 5

RAID 5 offers reliable data protection through the use of stripe sets and parity in an array of disks. Parity involves breaking up the data into blocks, calculating parity, then writing data in stripes to the disk drives. One stripe on each drive is reserved for the parity information, and parity is always written on a different drive than the associated data.

A read operation on the database using RAID 5 generates one I/O. A write generates four I/O operations: read the old parity information and read the old data, then write the new data and write the new parity information.

If a single disk fails in a RAID 5 array, data on the failed disk can be regenerated from the striped parity information on the other disks. However, the loss of a drive—and the resulting loss of access to the data and parity information it contains—will reduce database performance for both reads and writes until the failed drive is regenerated, because every read or write request to the failed disk initiates a verification process against all of the other drives in the parity group. If a second disk is lost before the first one can be replaced and the data recovered, the data on the second disk will be lost. The problem can be mitigated by having a hot spare available. (For more information, see the "Online Spare Drive" section on page 9.)

RAID 5 requires a minimum of three disks in an array. Although you can use RAID 5 with up to 32 disks in a set, it is not recommended because as the number of disks in a array increases, the probability of failure for any one disk also increases. For a Cisco Unity system, we recommend three to five disks per set.

RAID 5 is less expensive to implement than RAID 1. RAID 5 requires only one extra disk per array—the least amount of storage capacity required for fault tolerance of all of the RAID levels. In a read-intensive environment, especially if the percentage of reads is above 90%, RAID 5 is an ideal solution and is more cost-effective to implement than the other RAID options. Therefore, we recommend RAID 5 over RAID 1 for the Cisco Unity primary data store.

RAID 10

RAID 10 works like a combination of RAID 0 and RAID 1. Data is striped across a set of drives, then mirrored to a second set of drives. Each write operation uses two physical I/O operations to mirror the data. This method provides the mirror set data protection of RAID 1 combined with the efficient distribution of data in a RAID 0 array.

If a disk fails in a RAID 10 system, the controller switches to the remaining stripe set. There is usually no noticeable reduction in system performance while the system rebuilds the data from the failed disk onto the replacement disk.

Given the same number of physical disks, RAID 10 is faster than RAID 5. An increase in the number of disk drives provides a linear increase in I/O capacity. Because a read operation can be performed on either the primary drive or the mirror, as I/O increases, the drives become faster as a set because they can perform more multiple reads simultaneously.

Use of RAID 10 for a Cisco Unity system is supported, although it is more expensive to implement than either RAID 1 or RAID 5 due to the cost of the additional disks required.

Online Spare Drive

An online spare drive, also known as a hot spare, can further improve a system's fault tolerance. An online spare drive can be added to both RAID 1 and RAID 5 arrays. If a drive fails, the RAID controller immediately begins rebuilding the data from the failed drive to the online spare drive. The controller also sends new data directly to the online spare, minimizing the recovery impact on system I/O performance.

Configuring RAID Controllers

Best Practices

- If you have only one RAID controller, enable the write cache on the transaction logs, rather than the read cache on the database drive. Note that write caching is disabled on servers that are sold by Cisco.
- Use an uninterruptible power supply (UPS) on the Cisco Unity server(s).

Information on RAID Controllers

Most RAID controllers can be configured to support read caching, write caching, or both. A controller setting of 50:50 read:write is a typical value and is suitable for Cisco Unity.

Write caching can mitigate the additional processing overhead when using RAID 5; however, there are additional considerations. The speed of the disks must be able to support the write loads placed on them. If the number of write transactions is more than can be supported by the disks, write caching occurs, and the speed of read operations is impacted as well.

Use of a write-caching disk controller can also improve SQL Server performance, as long as the controller and disk subsystems are specifically designed for use in a data critical transactional relational database management system (RDBMS) environment. (Confirm this with your hardware vendor before using with a Cisco Unity system.) Note that write caching is disabled by default on servers purchased from Cisco.

Page-write transactions are considered atomic units of work that are either totally applied or totally rolled back. SQL Server data modification results in two logical page write transactions—one to the transaction log and another to the database itself. SQL Server writes to the transaction log almost immediately, but uses its own cache buffer system to collect and manage write transactions to the database. Because transaction-log writes occur before the database write, the transaction log is sometimes referred to as a write-ahead log. A RAID controller that performs write caching will intercept SQL Server page write transactions and buffer them to a hardware cache, and then manage the page writes for best disk performance. When tuning the read/write cache ratio, set it to 100 percent write for configurations that use a single partition.

Read caching is typically used for applications that perform a large number of sequential reads. Read caching does not provide any benefit in random read applications, such as Cisco Unity.

Optimizing the Databases

Best Practices for All Systems

- Physically separate the transaction logs and database files on different disks.
- Dedicate a high-performance disk to the transaction logs.
- Use a dedicated volume for the databases. This is especially important for RAID 5 partitions.
- It is not necessary to dedicate an entire disk solely to the pagefile. The pagefile performs well on the operating system drive, and should be left there in the case of the Cisco Unity server. If the pagefile is moved to another location, create a second pagefile to maintain operating system crash dump capability.
- Use NTFS rather than FAT for all Exchange database partitions, including partitions that contain the information store databases and transaction logs.
- Have a sufficient number of disk drives to keep latency low.
- Confirm that all disks used for transaction-log files have sufficient I/O processing power to accommodate the reads that occur in addition to the normal transaction-log write operations.
- Do not use file compression on a SQL Server database. It is not supported.
- The Temp.db database can be located with the data in most situations.
- Use 64 KB as the allocation-unit size when new drives are formatted in Disk Administrator.
- Move the transaction logs to a dedicated mirror, and back up Exchange.
- Turn on circular logging if you don't want to do full backups every time you back up the system. Otherwise, turn it off to have more flexibility with your backup schedule. For more information on circular logging, refer to the "Circular Logging in Exchange" section on page 13.

- Do not store backups of the Cisco Unity server (includes Active Directory, Exchange Directory, Exchange Information Store, Exchange mailboxes associated with the local Cisco Unity server, SQL Server, and System State) on the same disks as your data files or log files. Use a network location for backups if available. For more information on backing up the Cisco Unity server, refer to *White Paper: Maintaining a Cisco Unity System (with Microsoft Exchange as the Message Store)* available on Cisco.com at http://www.cisco.com/univercd/cc/td/doc/product/voice/c_unity/whitpapr/index.htm.

Additional Best Practice for a Medium -Scale System

- Put the MTA database and tracking logs on the system disk, not on the database volume.

Additional Best Practices for a Large-Scale System

- Place all log files on their own dedicated disk, so that they are not impacted in the event of a disk failure that corrupts the database. Exchange and SQL Server logs can share the same disk. For information on moving Exchange and SQL Server logs, refer to the "Moving Logs and Files" section on page 15.
- Place the Cisco Unity application and diagnostic trace logs on their own dedicated disk. Use the Cisco Unity Diagnostic Tool in Tools Depot to move the Cisco Unity diagnostic logs.
- Mirror set 1: system disk, including binaries, swap file, and MTA database.
- Mirror set 2: SQL Server and Exchange (if present) transaction logs only.
- RAID 5 volume: Exchange information store and directory databases only.

Information on Cisco Unity, SQL Server, and Exchange Databases

Database Overview

The Exchange information store is two separate databases. The private information store, created as Priv.edb on a new installation, contains data in user mailboxes. The public information store, created as Pub.edb on a new installation, contains data in public folders. The access patterns for these database files are approximately 60% random read and 40% random write.

The basic unit of a database file is a database page. Most SQL Server I/O processing is done at the database page level. Database pages are 8 KB in size. An extent consists of eight contiguous pages and is the allocation unit for SQL Server objects, such as tables and indexes. An NTFS allocation unit consists of 256 pages and is the largest unit of storage.

Reads are done by application processes, and multiple reads can be done simultaneously directly to the database. Writes are not done directly to the database by the application. Instead, writes are sent to a cache, which in turn is executed by an independent background database process that inserts new data and deletes unwanted data. Write activities usually occur in a steady stream for transaction-log pages and in periodic asynchronous batches for other types of data.

Transaction-log files are a sequence of files that provide a safety net for data in memory until it is written to the disk, as well as after it has been written. When a change is made to the database, a record of the transaction is written to the log so that the transaction can be done again in the event of a system failure. The size of a transaction-log file is limited, and when the file reaches the maximum size, Exchange rolls over to a new log file and assigns a generation number to it for tracking purposes.

Database Systems

Online transaction processing (OLTP) systems contain a large amount of data. New records are constantly being added, and data is accessed by a large number of users concurrently. The I/O activity typically is a heavier volume of reads than writes, with the exception of the transaction logs. Database locks are of short duration, with moderate lock contention. Examples of OLTP data access on a Cisco Unity system include looking up subscriber information by using the Cisco Unity Administrator, call processing, message notification, use of distribution lists, as well as use of the Cisco Unity Assistant and Cisco Unity Visual Messaging Interface (VMI).

Decision support systems (DSS) contain data gathered from OLTP systems and handle multiple simultaneous queries. The I/O activity in DSS systems is predominantly random reads of the database and may include complex or lengthy queries that result in long duration database locks. DSS data access on a Cisco Unity system occurs when reports are run.

Sequential and Random I/O

The time required to perform an I/O task depends on whether it is random or sequential. Random I/O takes more time to complete than sequential I/O because most of the time spent doing random I/O is seek time—the time it takes to find the data on different parts of the disk. Using Windows Performance Monitor can help determine how many I/O operations are generated per second, but it does not count the extra I/O generated with fault tolerance (writes multiplied by two when using RAID 1 or RAID 10; writes multiplied by four when using RAID 5).

Sequential I/O is good for performance. If data is distributed around the hard-disk platter in nonsequential locations, it takes a significant amount of time for the hard disk to move the disk arm (seek time) and to spin the read/write heads (rotational latency) to locate the data. If data is co-located on one contiguous physical section of the hard disk platter, the disk arm and read/write heads move a minimal amount to perform the necessary disk I/O. Most hard disks deliver performance that is twice as fast when processing sequential operations as when processing nonsequential operations.

With sequential I/O, it is possible to run a 9-GB disk drive at 200 I/O per second without regard to latency effects. However, tests have shown that a limit of 75 to 80 I/O per second should be used for random I/O operations.

Because it takes almost as much time to read or write 64 KB of data as it does to read or write 8 KB, it is beneficial to try to perform larger KB transfers. SQL Server data in index pages are 8 KB in size. The log manager performs sequential writes in 32-KB chunks whenever possible. When it is necessary to retrieve a large amount of data, SQL Server uses read-ahead at 64 KB I/O.

RAID controllers divide data reads and writes into 16- to 128-KB slices that are spread across all disks participating in the RAID array. Splitting the data across physical drives distributes the workload evenly and increases I/O performance because disks on the whole are kept equally busy, as opposed to some disks becoming a bottleneck due to uneven distribution of I/O.

Isolating Sequentially Accessed Data

Three database operations are primarily sequential: transaction logging, bulk updates to the database, and batch reports. On a Cisco Unity system, the transaction logs are the predominant sequential I/O activity. For example, the Microsoft SQL Server transaction log, whose I/O is almost exclusively sequential write, is a candidate for isolation on its own set of disk drives. The disks that contain the transaction-log files can then efficiently perform the transaction-logging operations without interruption by other nonsequential I/O requests.

Although the transaction logs may grow significantly, they may not require exclusive use of an entire physical disk, and can share the space with other files that are not heavily used.

Bulk updates to the Cisco Unity subscriber database usually occur when the system is first set up, and subsequent updates can be planned to occur at less busy times. Batch reports can also be scheduled to run at off-peak hours. Therefore, bulk updates and batch reports are not significant considerations for sequential I/O segregation for a Cisco Unity system.

Circular Logging in Exchange

Exchange 5.5 and Exchange 2000 store data in Jet databases. The most current transactions are stored in the Edb.log file. When the Edb.lg file reaches 5 MB in size, Exchange checks to see if the oldest transactions in the file have been committed to the database. If they have, then the old transactions are overwritten with new transactions, preventing the Edb.lg file from growing beyond 5 MB in size. Exchange will not, however, overwrite any transactions that have not been committed to the database. This process is called circular logging. When circular logging is enabled, full backups of Exchange are required.

When circular logging is disabled, and the Edb.lg file reaches 5 MB in size, Exchange creates a new log file of the same name, and renames the inactive file. The number of log files increases as more transactions are logged, and the system continues to create new log files until a backup is completed. If you disable circular logging, Exchange backups must be done, or the transaction logs can fill up all of the available disk space. We recommend that you set a backup schedule to do periodic full backups and incremental backups all other times.

Do not disable circular logging if you are not going to back up Exchange. If you decide to disable circular logging, confirm that backups continue to run and that they do not fail or stop working for some reason. If backups are not done, the transaction logs can cause the hard disk to fill up, which in turn causes Exchange to stop working and subsequently can cause Cisco Unity to stop working.

Choosing a RAID Type

Best Practice

Use hardware RAID only. Software RAID is not supported for use with a Cisco Unity system.

Information on Hardware RAID vs. Software RAID

It is helpful to have an understanding of the advantages of hardware-based RAID over software-based RAID implementations, in order to understand why hardware RAID is the method of choice for Cisco Unity installations.

Software-based RAID requires operating-system and application-specific components that add complexity, processing overhead, and maintenance overhead. A software RAID engine must handle all I/O requests and share CPU resources with application components. The extra processing and transfer of data to and from physical memory can consume a large amount of CPU resources and degrade system performance.

Hardware-based RAID offers complete independence from the operating system and the application. The RAID firmware runs on its own dedicated processor and does not share the application CPU. This enables the application CPU to perform application activity while the RAID array adapter’s processor

simultaneously manages disk I/O and fault tolerance. In the event of a malfunction in the RAID hardware or firmware, the application CPU can usually continue to operate and inform the system administrator of the malfunction. If the Cisco Unity server crashes due to some unexpected event, hardware-based RAID protects against loss of subscriber and transaction data by using mirroring and/or parity. RAID battery backup modules maintain the coherency of the caches and complete outstanding operations.

For these reasons, software RAID is not supported for use with a Cisco Unity system.

Planning System Capacity and Performance

Best Practices

- Size conservatively. For information on system sizing, refer to the *Cisco Unity Supported Platforms List*, available on Cisco.com at http://www.cisco.com/warp/public/cc/pd/unco/un/prodlit/ucutp_st.htm.
- Continually monitor the available disk space and rate of database growth.
- Monitor the Cisco Unity CPU context switches per second and processor interrupts per second as indicators of potential system performance bottlenecks.
- Attempt to keep the disk queue value less than two.
- Set SQL Server transaction logs to auto-grow, but size them so that they should not need to grow.
- Specify and manage to a 20-to-30-minute per-user storage quota. Cisco Unity is codec-dependent. For more information on storage implications associated with codecs, refer to *White Paper: Audio Codecs and Cisco Unity*, available on Cisco.com at http://www.cisco.com/univercd/cc/td/doc/product/voice/c_unity/whitpapr/codecs.htm.
- Do not install Exchange into the active partition. If the active partition fills up, Windows 2000 and Cisco Unity will not run.

Information on Database Planning

Disk Fragmentation

To reduce the impact of memory paging to the hard disk, regularly check for fragmentation on the Cisco Unity volume. If the total fragmentation appears to be more than 10 percent or if Cisco Unity files have become overly fragmented (more than 500 fragments), it is advisable to defragment the volume. For more information on disk defragmentation and other resource performance monitoring, see *White Paper: Monitoring Cisco Unity Performance*, available on Cisco.com at http://www.cisco.com/univercd/cc/td/doc/product/voice/c_unity/whitpapr/perfmon.htm.

Transaction Log Sizing

The optimal size is based on your recovery model, the level of logged activity in the database, and the time interval between backups. If the transaction log expands too frequently, performance can be affected.

Filling a disk beyond 80% capacity will result in significant performance degradation. If the disk fills up completely, the database cannot grow, and Cisco Unity will stop.

Reliability Planning

For a Cisco Unity system, some items add significant cost without significantly increasing reliability. For example, redundant NICs can be used with Cisco Unity in Adapter Fault Tolerance mode, and redundant power supplies for the RAID controllers can be considered; however, redundant RAID controllers are typically too expensive to be justified for the small gain in reliability they offer.

Moving Logs and Files

As part of database planning, you may want to move the Windows paging file when you install a new RAID 1 array that is faster than the disk currently storing the paging file.

To move the Windows paging file

Step 1 In Control Panel, double-click **System**.

Step 2 Click the **Advanced** tab.

Step 3 Click **Performance Options**.

Step 4 In the Virtual Memory section, click **Change**.

Step 5 Select the new drive that you want to use to store the paging file. Note that if you move the paging file from the system partition, Windows cannot write debug information to the disk in the event of a Stop Error message (blue screen error).

Step 6 Set the **Initial** size (MB) and the **Maximum** size (MB). For the Initial size, use the Microsoft-recommended value shown in the Total Paging File Size for All Drives window.

Step 7 Click **Set**, then click **OK** twice.

Step 8 Click **Restart Now** for the change to take effect.

You may also want to move the Exchange 2000 transaction or system logs to another drive.

To move the Exchange 2000 transaction or system logs

Step 1 Start Exchange System Manager.

Step 2 Right-click the Storage Group for the log files you want to move, then click **Properties**.

Step 3 Click the **General** tab.

Step 4 Click **Browse**.

Step 5 Locate the type of logs you want to move. Note that both system logs and transaction logs can be moved in the same operation.

Step 6 Click the new location for the log files, or create a new folder.

Step 7 Click **Apply**.

Step 8 Click **Yes** twice.

Step 9 Do a full online backup to complete the operation.

You can also move the SQL and MSDE transaction logs. If you move the transaction logs, you will move both the Unity database and the Reports database.

To move the SQL and MSDE transaction logs

Step 1 Go to a command line prompt.

Step 2 Detach the databases by entering:
EXEC sp_detach_db 'UnityDb'

EXEC sp_detach_db 'ReportDb'

Step 3 Copy the **UnityDb.mdf** and **UnityDb_log.ldf** files to the new destination drive. For example, to move the files from the C drive to the D drive, enter
Copy c:\Program Files\Microsoft SQL Server\MSSQL\Data\Unity.mdf d:\<path>\data copy c:\<path>\data\UnityDb_log.ldf d:\<path>\data

Step 4 Copy the **ReportDb.mdf** and **ReportDb_log.ldf** files to the new destination drive. For example, to move the files from the c: drive to the d: drive, enter
Copy c:\Program Files\Microsoft SQL Server\MSSQL\Data\ReportDb.mdf d:\<path>\data copy c:\<path>\data\ReportDb_log.ldf d:\<path>\data

Step 5 Reattach the data to SQL Server by entering
EXEC sp_attach_db 'UnityDb', 'd:\<path>\data\UnityDb.mdf', 'd:\<path>\data\UnityDb_log.ldf'

EXEC sp_attach_db 'ReportDb', 'd:\<path>\data\ReportDb.mdf', 'd:\<path>\data\ReportDb_log.ldf'

The Cisco Unity trace logs and Unity Messaging Repository (UMR) can also be relocated. If you move the UMR, you will move both the UnityMta (storage) and the Failed (undeliverable messages) directories.

To move the Cisco Unity trace logs

Step 1 On the Cisco Unity server desktop, double-click the **Tools Depot** icon.

Step 2 In the left pane of the Tools Depot window, under Diagnostic Tools, double-click **Cisco Unity Diagnostic Tool**.

Step 3 In the Cisco Unity Diagnostic Tasks window, click the **Configure Log Settings** icon.

Step 4 In the Logging Properties window, enter the new location.

Step 5 Click **OK**.

Step 6 Restart the Cisco Unity server.

To move the Cisco Unity UMR

Step 1 Start Regedit.

Caution Changing the wrong registry key or entering an incorrect value can cause the server to malfunction. Before you edit the registry, confirm that you know how to restore it if a problem occurs. (Refer to the “Restoring” topics in Registry Editor Help.) Note that a typical backup of the Cisco Unity server

does not back up the registry. Also note that for Cisco Unity failover, registry changes on one Cisco Unity server must be made manually on the other Cisco Unity server, because registry changes are not replicated. If you have any questions about changing registry key settings, contact Cisco TAC.

Step 2 If you do not have a current backup of the registry, click **Registry > Export Registry File**, and save the registry settings to a file.

Step 3 Expand the key HKEY_LOCAL_MACHINE\Software\ActiveVoice\UnityUMR\1.00.

Step 4 Double-click **MtaFailedPath**, and set the Value to the new location.

Step 5 Click **OK**.

Step 6 Expand the key HKEY_LOCAL_MACHINE\Software\ActiveVoice\UnityUMR\1.00

Step 7 Double-click **MtaStoragePath**, and set the Value to the new location.

Step 8 Click **OK**.

Step 9 Restart the Cisco Unity server.

For information on moving SQL Server databases, refer to Microsoft Knowledge Base article Q224071. For information on moving Exchange 2000 databases, refer to Microsoft Knowledge Base articles Q257184 and Q155761. Contact Cisco TAC if you need more information prior to moving either of these databases.

Using Storage Area Networks

The storage area network (SAN) is a type of fault-tolerant storage for both networked applications and data. A SAN solution is typically used in very large-scale networks, where storage needs exceed 5 terabytes of data. SANs typically use fiber channels and fiber switches to connect to application servers and provide centralized data storage in a separate network that is not vulnerable to currently known types of attacks. SANs also simplify administration of backup and restore, and keep backup traffic off of the data network. If the cost can be justified, a SAN also provides increased scalability over multiple RAID arrays.

A Cisco Unity system does not require such a level of storage capacity. While SANs may be used in some installations with an active-active Exchange cluster, SAN performance and sizing recommendations with Cisco Unity are outside of the scope of this white paper.

Many good references also exist on Cisco.com about the use of Cisco storage routers to complement SAN implementation.

Additional Reference

To define generally accepted RAID levels and educate users, the RAID Advisory Board was formed in 1992. The RAID Advisory Board maintains a website that serves as a resource for both vendors and users.