

Realtime
publishers

The Essentials Series

Configuring Microsoft Exchange for Fast Recovery

sponsored by



by Ron Barrett

Article 1: Addressing Hardware Needs and Determining Recovery Objectives.....	1
Hardware Best Practices.....	1
Processor	2
Memory.....	3
Storage Considerations.....	3
Disk Space, Redundancy, and RAID Levels.....	3
Configuring Disk Space Correctly.....	4
Configuring Users and Storage Groups.....	5
Recovery Point Objectives.....	6
Bandwidth Needs	6
Managing Databases and Log Files	7
Summary	7
Article 2: Backup Methods Available for Exchange.....	8
Backup Methodology	8
Backing Up Domain Controllers.....	9
VSS vs. Streaming (Legacy) Backups	9
Backup Best Practices	10
Four Types of High-Availability Solutions in Exchange 2007.....	10
Local Continuous Replication	10
Cluster Continuous Replication.....	11
Standby Continuous Replication.....	11
Single Copy Cluster	12
High Availability for the Remaining Server Roles.....	12
Article 3: Exchange Recovery Recommendations	13
Database Recovery Scenarios	13
Recovering Non-Mailbox Roles	14
Restore Using Database Portability.....	14

Using RSGs.....	15
Dial-Tone Restores	16
VSS vs. Streaming (Legacy) Restores.....	16
Virtualization	17
Summary	17

Copyright Statement

© 2008 Realtime Publishers, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers, Inc. (the “Materials”) and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers, Inc or its web site sponsors. In no event shall Realtime Publishers, Inc. or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at info@realtimepublishers.com.

Article 1: Addressing Hardware Needs and Determining Recovery Objectives

The term *mission-critical application* does not capture the importance of email in the enterprise. Email is truly a *core application*; it is in many instances the lifeblood of an organization. Such being the case, the need to maintain availability is crucial.

Of course, when you plan and implement your infrastructure, it is with the idea that you will never have the need to recover from a failure. However, if you need to recover Exchange, it is best to ensure that the recovery can be done quickly. This series will focus not only on getting Exchange up and running again but also on ways to make sure it is recovered quickly. Anyone who has ever suffered through an Exchange Server failure knows 10 seconds down is 10 seconds too long. There is no way to promise that you can be back up in 10 seconds, but the information provided here will help shrink that recovery window and get you back online quickly.

Every good recovery plan starts with the hardware you choose. It is equally important, though, to make sure that hardware is properly configured if you are going to achieve fast recovery.

Hardware Best Practices

Minimum requirements for hardware would obviously not be fitting for scenarios where you are looking for fast recovery. These best practices will therefore exceed minimum and match or exceed recommended requirements.

Whether you are planning, designing, or have already implemented an Exchange environment, it is never too late to look at best practices for hardware. This is especially true because hardware costs are so incidental. Moreover, when you compare the cost of additional storage or extra RAM to the cost of an Exchange Server down for the day, you have a very compelling argument for even the stingiest CFO.

When discussing hardware for fast recovery in Exchange, many factors come into play for an organization—such as the number of server roles deployed, and whether these will be on a single server or broken down into multiple servers; how many storage groups will be created; the number of mailboxes per storage group; and mailbox size restrictions (or lack thereof for certain organizations).

A later section will talk specifically about storage considerations. Let's begin by looking at processor and memory considerations.

Processor

As Exchange 2007 works only on a 64-bit processor, this part is simple. Either the Intel EM64T or AMD 64 processor will work.

Warning

The Intel Itanium (IA64) processor will not work with Exchange Server 2007, eliminating this processor as an option.

Recommended configuration for each role is 4x processor cores except the Edge Transport server role, for which 2x processor cores are recommended. The recommended configuration for a server hosting multiple roles is 4x processor cores. Maximum configuration for Client Access, Edge Transport, Unified Messaging, and multiple roles still calls for 4x processor cores; however, for the Hub Transport and Mailbox role, 8x processor cores are recommended.

Of course, people love the idea of throwing as much hardware as possible at a mail server. Consider a few factors to help you decide whether you need the recommended or maximum configurations.

The Hub and Edge Transport server roles processor utilization is based on numerous factors such as number of transport agents, average message size, message rate, configuration of antivirus, and use of third-party applications. The Client Access and Unified Messaging server roles processor utilization is based on the size of the environment. 4x processor cores will be more than sufficient in many environments for these roles.

The rule of thumb for the configuration of the Mailbox role is 1x processor core per 1000 average user mailboxes. Average user mailboxes are classified as 40 messages received and 10 messages sent per user. Even if you classify all your users as very heavy users (120 received and 30 sent), you are looking at 2000 mailboxes for a 4x processor core server.

Note

For servers hosting multiple roles, you should reduce your mailbox count by 20 percent to account for the processor utilization by the Hub Transport and Client Access roles.

Memory

Memory considerations for fast recovery require specific configuration considerations. It is recommended to have 1GB of RAM per core (with a 2GB minimum) for all the server roles except the Mailbox role. The recommended configuration for this role is 2GB plus 5MB per mailbox and 2GB for every four storage groups. As the focus is on fast recovery—and again taking into consideration the incidental cost of RAM—8GB for an average and 16GB for a heavy usage Mailbox role is best practice.

If you are hosting multiple server roles, the recommendation is 8GB plus 5MB per mailbox and 2GB per four storage groups. For fast recovery of Exchange hosting multiple server roles, best practices suggests 16GB for an average and 32GB for a heavy usage Mailbox role.

Note

Even with a new server, going beyond 32GB becomes cost prohibitive for most organizations.

Storage Considerations

Of all the hardware recommendations, storage considerations are one of the most significant. Storage considerations do not simply include how much space you need for Exchange; they also encompass how you will configure and manage that storage space for fast recovery. Therefore, you must look at disk space, disk redundancy, RAID configurations, and other best practice configurations for fast recovery in Exchange.

Disk Space, Redundancy, and RAID Levels

Let's begin with disk space considerations. When speaking about disk space for fast recovery, you are not talking about the size of the disks needed. Rather, you are talking about the percentage of available disk space in correlation to your disk size. Ideally, you want to maintain 40 percent availability for disks in a fast recovery scenario. This number is not an arbitrary value but represents a performance threshold to maintain an environment suitable for fast recovery of Exchange.

Traditionally, for mission-critical applications such as email, capacity and redundancy are considered crucial to successful implementation. Although this is true to an extent, these factors are not as significant as the need to control disk latency and optimize performance—both of which are critical for fast recovery of Exchange. Therefore, in this case, best practice is to opt for a RAID-1+0 solution. Although this setup cuts the available capacity of the disks, it provides for better I/O for read/write operations in Exchange, which is very I/O intensive. In addition, it is recommended to invest in high-performance disks and spindles. Again, this may cut down on the disk capacity (cost being a direct factor for many). Considering that disk subsystems are a greater bottleneck on system performance than processors or memory, high-performance disks are perhaps the best investment you can make in your messaging infrastructure.

Configuring Disk Space Correctly

Having the appropriate amount of available disk space and redundancy is not the only factor to consider. How you use that space is imperative. Exchange Server is I/O intensive, so you need to ensure you use the proper RAID levels and alleviate both the single point of failure as well as bottlenecks in recovery, both of which will defeat all the efforts you put forth.

The first recommendation is one you have heard for years, yet it is widely ignored or incorrectly configured. To start, the operating system (OS) should be installed on a separate volume. Separating the database and transaction logs onto separate physical drives provides both fault tolerance and fast recovery. This reduces risk of data loss if a disk failure occurs by ensuring that databases and transaction logs are not compromised by the loss of a single array—again, no single point of failure.

Another step in configuring hardware for fast recovery is the use of the `diskpart.exe` tool. Use `diskpart.exe` to create aligned disk partitions; having the tracks *sector-aligned* will improve performance by nearly 20 percent. Keep in mind that `diskpart.exe` will not work with dynamic disks and can only be used with basic disks.

Configuring Users and Storage Groups

The configuration of users and storage groups is an area where you will see variance depending on the user base, mailbox quotas, and other configurations. Microsoft recommends one storage group to one database. However, the Standard edition of Exchange allows five storage groups with as many as five databases per storage group. The Enterprise edition allows 50 storage groups with 50 databases per storage group.

In keeping with the objective of fast recovery, it is most beneficial to plan the storage group, databases, and, by extension, user mailboxes in such a way to allow for better performance. For example, a 4x processor core with 12GB of RAM would meet the requirements for hosting 2000 mailboxes in a single storage group on a server hosting only the Mailbox role.

Considering a maximum mailbox size of 250MB per user, you are looking at roughly 504GB of storage. However, that is only part of the equation; add in space for the database dumpster (104MB) and white space (7.3MB), and you now have a mailbox size of 306MB. 2000 mailboxes now require 844GB for databases and 198GB for log files.

In this example, the recommendation is to have eight storage groups and databases with 250 mailboxes per storage group. This would bring you to 130GB per database with log files. Now you can very quickly see the logic in dividing up users into storage groups. Even if you need to recover multiple databases of 105GBs, it will go much quicker than a single 844GB database recovery. Of course, if you only have the standard version of Exchange, eight storage groups is not an option for you, so creating five storage groups of 400 users is the best practice.

You can drive yourself crazy and do the math alone or you can download the Exchange 2007 Mailbox Server Role Storage Requirements Calculator from the Exchange team blog site at

<http://msexchangeteam.com/files/12/attachments/entry438481.aspx>.

Just to Drive Home the Point

Other storage considerations that have not been included to the individual mailboxes are

- Content Indexing—5 percent database size
- Maintenance—10 percent database size
- Log files—12 to 36MB per mailbox
- Log file overhead—20 percent of total log file space allocated

Although breaking up storage groups does not eliminate the disk space needed to support 200 users, it certainly gives you a much better way of managing them to meet your recovery point objectives (RPOs).

Recovery Point Objectives

Howard Aiken said, “Time has a wonderful way of weeding out the trivial.” Not a bad philosophy when you talk about RPOs. Although the goal is to lose nothing at all, the objective of fast recovery is to get up and working quickly with minimal data loss. No data loss would be phenomenal, and achieving that goal comes with a price tag. Albeit the price has dramatically been reduced thanks to many of the technologies introduced in Exchange Server 2007.

Simply put, RPO is your company’s policy for the acceptable amount of data loss in a distressed situation. RPO is measured in time, so it is determined in this case by both the amount of acceptable data loss (emails) and recovery time.

Bandwidth Needs

Although RPO is considered more of a hardware issue, it is fitting that to look at the requirements for bandwidth along with RPO. Available bandwidth on both the LAN and WAN side will affect fast recovery. From a LAN perspective, best practice is to have teamed Gigabit network cards for load balancing and fault tolerance. In addition, optimize link speed and duplex settings. In most instances, when you talk about bandwidth needs for fast recovery, you are talking about restoring a mail database(s) that has become corrupt or unreachable locally. Occasionally, you need to think in terms of a complete unrecoverable disaster. The amount of time it takes to back up or replicate data to an alternative off-site location becomes a bandwidth issue. Allocating sufficient bandwidth depends on many factors, including which method or backup is used and the organization’s RPO, RTO, and availability of high-bandwidth communications. Large cities have the availability of affordable fiber optic 100Mbit Internet speed. Others are limited by the traditional 1.54Mbit connectivity, so this will be a factor.

Managing Databases and Log Files

One factor to keep in mind for database and log file management in Exchange is the need to back up the environment. Ultimately, the size of the database, backup windows, Service Level Agreements (SLAs), and available resources are going to dictate the size of the database and frequency of maintenance. Another effective method of managing databases and log files is utilizing mailbox quotas. For some organizations, this is easier said than done.

Managing log files is best accomplished by backing up the database and dumping committed logs. Monitoring backup time and performance issues may make moving storage group paths another option. Moving the storage group from a disk that is more than 40 percent full will enhance performance overall and provide a better scenario for fast recovery.

Summary

Addressing hardware needs and considering your RPOs is the first step in successfully implementing a fast recovery solution in Exchange Server. The next article will consider the different backup methods available for Exchange.

Article 2: Backup Methods Available for Exchange

To be sure you have the ability to perform fast recoveries in Exchange Server, you need to be sure you have a good method of backup. The backup must be fast, complete, and, of course, recoverable. This article will look at different backup methodologies and consider some of the new features in Exchange 2007 that make fast recovery possible. We will then look at considerations for backing up the various server roles and domain controllers. Finally, we will finish up with backup best practices.

Backup Methodology

Estimates show that 80 percent of all corporate knowledge is stored in email. Gartner estimates that businesses lose \$25 per hour, per employee when email is down. Suppose your organization supports one database of 250 mailboxes, which is down for 1 hour; you are looking at a cost of \$6250 per hour.

Considering these numbers, you can see the importance of making sure that your Exchange backup plan includes not only database and logs but also the operating system (OS) and all volumes on the server. The first article in this series addresses recovery point objective (RPO); to reiterate, it is a determined acceptable data loss measured in time. To meet fast recovery objectives, you must also consider Recovery Time Objectives. RTOs are defined as the duration of time that the systems need to be restored.

RTO is tied to service level agreements (SLA), which are part of the business continuity plan. Thus, the RTO encompasses not only how fast the system needs to be back online but also the acceptable level of operation. Thinking along these lines, you need to consider a few factors: What media will you back up to? Will it be tape, SAN, NAS? How often will you back up? Will nightly backups be performed or do you need more frequent intra-day backups? Each of these decisions will affect the RPO and RTO and ultimately your plan for fast recovery.

In addition, each of these factors is crucial and comes with a pain threshold (how much you want to spend versus how much you're willing to lose). When you look at the potential costs associated with a critical data loss, fast recovery just makes sense.

Backing Up Domain Controllers

It is important to think about backing up domain controllers in an Exchange environment. Since Exchange Server 2003, the importance of Active Directory (AD) to Exchange makes it necessary to ensure you have at least one domain controller backed up for fast recovery. Performing a system state backup on a domain controller will back up the necessary AD files. Remember that domain controllers have circular logging enabled for AD; therefore, any data written to AD after a backup will be lost. Depending on the frequency of changes to your domain controllers, you should back up at least one domain controller nightly. Doing so will ensure that you can meet RPO and RTO levels and ensures that you will not run into trouble by having an “old” AD backup, which would be un-restorable.

VSS vs. Streaming (Legacy) Backups

Exchange Server offers two options for backing up data. Both support the four backup methods (full, incremental, differential, and copy). The options are the traditional streaming (legacy) backup, which utilizes the ESE API and has been the available option for backing up Exchange Server using NTBackup, Windows Server Backup, and many third-party backup solutions. Highlights of streaming (legacy) backup:

- Backup is taken from the active copy of the database
- Can perform backup at the database level
- Only one backup running against a single storage group
- Separate storage groups can be backed up concurrently

Warning

Windows Server 2008 does not support streaming backups and is not Exchange-aware. Therefore, a third-party solution that utilizes VSS is required for Exchange backup in Windows Server 2008.

The second method utilized is the Volume Shadow Copy Service (VSS), which provides a point in time “snapshot” of your data. In subsequent backups, it looks at the last snapshot and then backs up *only* the changes. VSS was introduced to Exchange 2003; although it provides the ability to take shadow copies, these copies were made at the file level and were not Exchange-aware. Snapshot backups in Exchange are fast and consistent due to the use of checksums to the database pages. Highlights of VSS:

- Backup can be taken from the active and passive copy of the database
- Can perform backup at the storage group level
- Separate storage groups can be backed up in parallel

Backup Best Practices

Shrinking the backup window in Exchange can be achieved with the use of multiple storage groups. Keeping databases to a manageable size will also shrink the time it takes to back up those databases. The high-availability solutions considered in the next section can help to shrink backup windows by requiring less need for full backups. Replication can be the primary fast recovery option with backup as the secondary option.

Microsoft has a great document titled "[What Needs to Be Protected](#)," which is a good gauge for how to setup your backup strategy.

Another best practice is to create reliability for backups. This can be done several ways. One way I like to emphasize is the use of disk backups for Exchange. After you have a disk backup of the Exchange Server, you can use tape backup to create redundancy.

Four Types of High-Availability Solutions in Exchange 2007

Exchange 2007 SP1 provides four high-availability solutions to ensure a percentage-based uptime to the Mailbox role in Exchange. First, let's look at a quick description of each replication type and then explore the pros and cons of each.

Local Continuous Replication

As the name suggests, Local Continuous Replication (LCR) creates a continuous local copy of the active storage group to a second set of disks using log shipping and replay.

Pros:

- Cheapest solution requiring only a disk or disks and secondary controller
- Can implement on a single server
- Can perform VSS backups on the passive side

Cons:

- Failover needs to be implemented manually
- Hardware failures could make data unavailable
- Only one database can be added to a storage group running LCR

Cluster Continuous Replication

Cluster Continuous Replication (CCR) creates a continuous copy of the data over to a second server, and uses log shipping to maintain a copy of a storage group onto non-shared storage. Clustering adds automation to replication.

Pros:

- Provides automatic failover
- Provides both disk and server redundancy
- Utilizes the transport dumpster to maintain data that has not been synchronized between clusters
- Can perform VSS backups on the passive side

Cons:

- Requires knowledge of clustering services
- Requires an additional server and software to implement (that is, additional cost)

Standby Continuous Replication

Introduced in SP1, Standby Continuous Replication (SCR) uses the same log shipping and replay technology as LCR employs. Replication is server to server, but it does not require clustering services to provide redundancy.

Pros:

- Provides both disk and server redundancy
- Can be used in conjunction with the other three replication methods
- Replay delay protects from database corruption to the SCR copy

Cons:

- Failover needs to be implemented manually as with LCR
- Management is done through the Exchange Management Shell (EMS)

Single Copy Cluster

Similar to Exchange Server 2003 clustering technology, Single Copy Cluster (SCC) uses a single copy of the storage group that is located on shared storage in a cluster.

Pros:

- Provides automatic failover
- Supports up to eight cluster nodes

Cons:

- Requires knowledge of clustering services
- Requires an additional server and software to implement

High Availability for the Remaining Server Roles

As replication is used for the Mailbox server role only, how do you ensure that the remaining roles are backed up properly? Let's talk about how to achieve high availability in the remaining server roles. Providing high availability for the Client Access, Hub Transport, Edge Transport, and Unified Messaging servers is for the most part similar:

- Client Access Server—Deploy multiple client access servers and use Network Load Balancing (NLB) to provide high availability.
- Hub Transport Server—Resilient by default due to the fact that all Hub Transport servers are registered within AD. You can achieve high availability by deploying multiple hub transport servers.
- Edge Transport Server—Deploy multiple edge transport servers and use DNS MX records to achieve high availability.
- Unified Messaging Server—Deploy multiple unified messaging servers and place them into the same dial plan so that the VoIP gateways can retrieve a list of servers within the dial plan. Configure VoIP gateways to round-robin calls to ensure high availability if a unified messaging server is down.

Once you have a reliable backup plan that meets your RPO, RTO, and SLA requirements, and you have the ability to test that backup to ensure reliability, you need to move on to the last part of the fast recovery plan—finding out both how to recover Exchange quickly and ensure you have what you need in your environment once Exchange is recovered. The next article will explore this topic and consider recommendations for recovery in Exchange Server.

Article 3: Exchange Recovery Recommendations

Good hardware and configuration provide the foundation for fast recovery, and backup methods provide the how-to for fast recovery. However, all the configuration and backup means nothing if you cannot restore the Exchange Server system and meet your recovery time objectives (RTOs) and Service Level Agreement (SLA) requirements. Thus, the need to look at recovery scenarios and different methods of recovery, including restoring non-Mailbox roles and using database portability, and Recovery Storage Groups (RSGs). In addition, this article will look at VSS recovery compared with streaming (legacy) backups. Finally, we will close with a word on how virtualization can be used in fast recovery.

Database Recovery Scenarios

So what are some of the recovery scenarios you might face in an Exchange environment? They fall into three categories: hardware, software, and site. To begin with, you must consider what can be called database items, such as email that is permanently deleted or perhaps the loss of a mailbox in Exchange. Next, you have lost databases or storage groups on a server. Moving up from here, you can face a lost server where the database and log files are intact or a lost server where database and log files are also lost. If clustering is used, you might have the loss of a storage group, database, or servers in a cluster, loss of external services such as DNS or AD, or even the loss of an entire site. No matter how small or how big the problem, a loss of email from a single email up to an entire site is costly to any organization.

Recovering Non-Mailbox Roles

We will start with recovery of non-Mailbox roles in Exchange. These roles are in many ways easier to recover for several reasons. First, the configuration for these roles exists in Active Directory (AD). Simply using Setup /m: RecoverServer will care for most of the non-Mailbox servers. The Edge Transport server requires a little bit of a different recovery because it does not exist within the AD structure. Certain configurations will need to be restored manually, such as the voice recordings for the Unified Messaging role or metadata within IIS for the CAS role. When restoring the Hub Transport server, you should add the /DoNotStartTransport switch and check the configuration before putting it into production. To restore this server role, you need to export the configuration information with a script using EMS. Install Exchange, and use the same name for the Edge Transport server. Run a script to create an answer file from the information you exported. Adjust the answer file, and run the script again, applying the changes to the Edge server. Lastly, run the EdgeSync process to synchronize the Edge Transport and Hub Transport servers. Using the methods discussed in the second article of this series, you can provide resiliency and uninterrupted operation until the failed server is restored or re-created.

Restore Using Database Portability

If you are going to have a solution for fast recovery, you also need to have flexibility to achieve your goals. Exchange Server offers you that kind of flexibility in a few ways. One notable way is the ability to use database portability. In the past, recovering an Exchange Server meant that you needed to do one of the following:

- Recover to a server with the same name as the original server from which the database came
- Recover another server within the same Administrative group (Exchange Server 2003 and 2000 also had conditions that needed to be met in order to perform the move)
- Recover to a Recovery Storage Group

Database portability now makes it possible to recover to any server within the same organization. Database portability gives you the ability to relocate a database in the case of a server that is having maintenance performed or is being upgraded.

Note

Upgrading from Server 2003 to Server 2008 will cause the database to rebuild its secondary indexes and will delay mounting the database. The time is dependent on the size of the database.

Another use for database portability is to restore a database to a new server. Using this method, you would need to use either NTBackup or third-party backup/restore software to restore the database to the new server. This is achieved by creating a new database with the same name as the original on the target server (do not mount that database) on the next restore of the database and allow it to overwrite the existing database. Logs will replay after the restoration, and you can choose to mount the database after the restore.

Using RSGs

Not all recovery solutions call for full database restores. In situations in which you need to extract data rather than completely restore, you can use RSGs. RSGs allow you to mount a database and extract data from them while having the server remain in production. Common uses for RSGs:

- Same server dial-tone recovery
- Alternative server dial-tone recovery
- Recover a mailbox, database, or group of databases within a storage group
- Recover deleted items that are past the deleted items retention policy

Note

Do not look for the RSG inside the Exchange Management Console. Even after the database is mounted, RSGs are created and managed using the Exchange Management Shell (EMS) or the Exchange Server disaster recovery analyzer (ExDRA). Only one RSG can exist per server.

Using RSG, you can restore all visible and hidden folders. RSG gives the ability to restore all mail items within Exchange. On a mailbox restore, you can restore the extracted data into an existing mailbox (merge) or into a subfolder of the mailbox (copy). RSG cannot be used to mount Public Folders; neither can you mount mailbox databases from previous Exchange versions. You would not use RSG for a full-server restore or for restoring multiple storage groups. Once an RSG is created, you can then perform a dial-tone restore. You have the option of either performing the restore to the same server or to a different server.

Dial-Tone Restores

In an Exchange environment where being down even 30 minutes can be disastrous, there is a solution that can allow you to quickly recover email operations while you work on resolving the issues with the database. Simply put, dial-tone restores allow you to quickly recover email functionality without the need to wait for the mail database to restore. Mailboxes exist in AD, so you can use a dial-tone restore to quickly get users functioning.

Dial-tone restore provides users the ability to send and receive email and gives users limited access to email using the offline (.ost) file. Users that have Outlook set up to use Exchange in cached mode will have access to some historical data. For users with Outlook 2007, auto discovery allows the new mailbox to be found instantly when a dial-tone restore is done on an alternative server.

Warning

When you perform a dial-tone restore, messages are not the only thing missing. The new mailboxes will not have views, rules, forms, or any other metadata from the former mailbox. These are restored once the mailboxes are merged.

When a database restore has been completed, you can then 'merge' the data from both mailboxes to achieve congruency.

VSS vs. Streaming (Legacy) Restores

Just as Exchange supports both streaming (legacy) and VSS methodologies for backup, it uses both methods for restoration. Restoring databases using streaming (legacy) backups consists of creating a storage group and database with the same name as the original storage group and database. Streaming (legacy) incremental backups cannot be taken after a full VSS backup. With Windows Server 2008, there is no support for streaming backups. Therefore, you need to use VSS.

VSS offers additional options for restoring Exchange Server. With VSS, you can restore backups to the active copy of the database. You can perform restorations to the same storage group or an alternative storage group. You can restore to the same server or an alternative server. You even have the option to restore the database to a non-Exchange location using file-level restores.

Virtualization

When we talk about virtualization, we often look at it in terms of server consolidation. We often talk about reducing costs for servers and energy by implementing new servers in virtualized environments. However, when we talk about virtualization for fast recovery, we are taking a different tact.

Virtualization offers two advantages in the Exchange environment. First, in the area of restoring Exchange databases, virtualization allows you to stretch your available resources in an emergency. For example, if you lose the physical server that hosts the Mailbox role, you can use virtualization to provision an underutilized server in your network and perform a dial-tone restore. This would allow you to return email to users quickly while you wait for the server to be repaired or replaced.

Another good use for virtualization is to perform restores to the virtual server away from a production environment. This affords you the opportunity to test your backup and restore solutions to ensure they are working.

You can also use virtualization to test various restore solutions, such as dial-tone restores, database portability, and RSGs. Using the high-availability solutions such as LCR or SCR, you can work out configurations that meet your RPO and RTO plans.

Summary

So what is the ideal formula for fast recovery? There are probably quite a few ideal and even third-party solutions to help you configure the right environment for your organization. A best practice is to start by making sure you have a good hardware base. Then plan your deployments, or tweak an existing deployment if you already have Exchange 2007 in production, to maximize your uptime and recovery potential.

You can use the high-availability options of replication and clustering (if it fits within the budget) as great ways to create a fast recovery infrastructure. The near-time replica of your active data offers countless possibilities for recovering quickly from minor to somewhat serious issues. Coupling high-availability options with a backup plan that meets your RPO, RTO, and SLA requirements gives you the possibility to offload that replicated data quickly using either disk or tape backup along with VSS (remember the rules though; you need some additional help to do this right and get an Exchange-aware backup).

Next, you have several ways you can restore your Exchange environment. Again, you can get your users up quickly using these fast recovery methods. And, if necessary, you can return email functionality while you wait for the database to return to a usable state. Virtualization gives you some extra options for scenarios in which a physical replacement may not be available for a lost server or site.

Finally, test that backup and restore process (here's where that virtualization piece fits in nicely again). With virtualization, you have the ability to not only test restores but do so in a controlled environment and test every sort of imaginable event to determine in 'real time' the effects and responses. That is where the real fast recovery comes in, by being able to create, test, and manage infrastructure. Then using the things you have learned and the tools available, you have the ability to tweak the things that can make your infrastructure better and improve the metrics that matter to your organization.

When you plan your Exchange environment, you do so with the hope that you will never need to worry about recovering mailboxes, databases, or worse. By planning for such events and following the prescribed steps in this essentials series, you can create a solution for fast recovery in Exchange Server.