SearchExchange.com E-Guide

Exchange Server Disaster Recovery: Planning for the worst, hoping for the best

Email has evolved from a convenience to a business necessity. The impact of natural, random, and man-made disasters on email, and the subsequent disruption in mail-flow can cripple an organization's ability to deliver messages. This E-Guide from SearchExchange.com and AppAssure Software offers best practices and advice for protecting your existing infrastructure from disruptive email outages.

TechTarget
The IT Media ROI Experts

SearchExchange.com E-Guide

Exchange Server Disaster Recovery: Planning for the worst, hoping for the best

Table of Contents:

Exchange Server disaster recovery planning:

True disasters are unseen and insufficiently forecast. To minimize the impact of disasters on Microsoft Exchange, proper planning, testing, and investments are required. This E-Guide from SearchExchange.com and AppAssure Software provides information critical to any organization that seeks to avoid total disaster in the event of the unforeseen.

Microsoft offers a wealth of online resources that detail Exchange Server disaster recovery. But for someone new to managing Exchange, it can be an overwhelming amount of information to digest. Even for administrators with significant Exchange experience, the prospect of designing an e-mail disaster recovery plan can be intimidating.

This checklist shares a number of tried and true best practices for Microsoft Exchange disaster recovery. Most of the concepts are universal in nature, but the focus is on Exchange Server.

Defining Exchange disaster recovery

Paul Robichaux

Depending on whom you ask, the definition for disaster recovery can range from restoring data from a backup to restarting operations at an alternative business continuance site. This chapter will use the Wikipedia definition, which defines disaster recovery as "The ability of an infrastructure to restart operations after a disaster."

The first is how you know when you've successfully recovered from the disaster. There are two common metrics for recovery:

- A **recovery point objective (RPO)** specifies the point in time at which your capabilities will return when recovery is complete. Let's say you do a daily full backup at 2 a.m. If you have a failure Tuesday at 8 a.m., your RPO will probably be 2 a.m. Tuesday; in other words, your recovery will succeed if you can recover the state of your Exchange data to that particular point in time.

- **A recovery time objective (RTO)** specifies the maximum amount of time allowed for a recovery. For example, if your service level agreement (SLA) promises that you'll restore operations within six hours of a disaster, you have a six-hour RTO.

Although these metrics are clearly related, there are significant differences between them that become apparent as you start to consider how to reduce the RTO or move the RPO closer to the actual beginning of the outage. To shorten the interval between a failure and your RPO, you must make more frequent copies of your data with whatever protection mechanism you've chosen. To shorten the RTO, you need to take measures to increase the speed of your recovery. The remainder of this chapter will discuss both types of measures.

Choosing a backup type for Exchange

Paul Robichaux

Choosing a backup type may seem complicated, but it's not. The bottom line is that the amount of time required for a restore is roughly double the amount of time required to make the backup in the first place. Factor in your RTO to quickly determine how much time you can afford to do a restore, which in turn tells you how long your backup can take if you're going to hit your SLAs and RTO. You can always tweak your backup solution's hardware (for example, by adding more tape drives and striping data across them, or switching to a higher-capacity, faster solution), but the time required for the backup window will ultimately be the number one factor in determining your backup pattern.

Let's say that your RTO is eight hours and that you have a total of 120 GB of mail data evenly distributed over four servers. Thus, within that eight-hour window, you need to notice that a failure has occurred, locate any needed backup media, start the backup, wait for it to finish, and wait for any pending transaction logs to be replayed.

You must also include a fudge factor to cover you in case something unexpected happens. Suppose that you actually have only six hours worth of restore window to work with. (In fact, during most restores, IT staffers waste time trying various procedures before they decide that a restore is necessary—be sure to factor this time into your planning!) Thus, your backup time should be at or below three hours. What kind of backups should you use?

- Full backups take the longest—assuming that your backup solution can handle 10 to 20 GB per hour, you can restore one server's worth of data in one and a half to three hours—assuming that nothing goes wrong.
- Incremental backups are smaller, so they take less time to capture and restore. However, they trade time for space; in addition, if the same database page changes more than once over the time span of an incremental set, you'll end up having to play back transactions to change that page over and over, adding to your restore time.
- Differential backups give you easier management and less overhead than incrementals, at the cost of storage growth over time. It's easy to grab a full backup, plus the differential for a given RPO, but you must factor in the time it takes to restore two backups instead of just one.

Without lab testing, it's difficult to pinpoint which combination of full, incremental, and differential backups will best allow you to meet your eight-hour RTO. However, if you have sufficient hardware to support it, daily full backups offer relatively easy restoration, little additional media management overhead, and integrity checking.

Most organizations use a combination of weekly or intra-week daily backups with daily differentials, although your combination may vary. As disk space continues to drop in purchase cost, an increasing number of organizations are doing full backups to disk and intra-day differentials—doing so gives great coverage at the expense of storage space.

Online vs. offline Exchange Server backups

Paul Robichaux

Exchange online backups

An Exchange backup using the ESE API follows a predictable set of steps:

1. The backup utility asks for a list of all the Exchange Server systems.
2. The backup utility connects to the specified Exchange Server system and makes a request to back up a particular storage group or database. The ESE API allows simultaneous backup or restore of as many as four storage groups, but you can only back up or restore one database within each storage group at a time.
3. If the ESE online maintenance task is performing maintenance on any databases in the storage group, that maintenance stops.
4. ESE flushes any dirty database pages to disk. These pages are those that have been changed but haven't yet been written to the on-disk copy of the database. At this point, the checkpoint is frozen.
5. The backup utility opens the first database file to be backed up. On Exchange 2000 and Exchange Server 2003, each individual EDB and STM file is backed up separately. For a full or differential backup, the database header is updated to point to the low anchor log file.
6. The backup utility issues repeated calls to read data from the file. It can then write that data using any backup mechanism.
7. When the backup tool is finished reading, it closes the database file.
8. Steps 3 through 5 are repeated with each additional file in the selected storage group.
9. The backup utility opens the first transaction log file for the selected storage group and copies its data, closing the file when done.
10. Step 7 is repeated for each additional transaction log in the selected storage group.
11. Once all the log files have been backed up, any log files marked for truncation are removed.
12. The backup program calls the ESE API to indicate that it's done with the backup.

Sharp-eyed readers will wonder what happens to transactions created while the database is being backed up. The answer might surprise you: they're logged to the transaction logs just as they would be during normal operation. Once the checkpoint is frozen in step 3, additional logs can be generated, but their transactions will not be committed until sometime after the backup completes. This method works because the log files generated while steps 4 through 7 are taking place will themselves be backed up in steps 8 through 10.

The other kind of virtualization.

Unplanned downtime of your Microsoft Exchange® and BlackBerry® servers can cause workplace productivity to grind to a halt, and employees to lose their cool.

Virtualized disaster recovery.

With Replay Virtual*HA*, you have an always-ready mirror-image of your Exchange and BlackBerry servers, ready to stand up in just minutes.

Regardless of the type of outage, whether it was planned or unplanned, you can keep the e-mail flowing. Replay delivers a defensible and cost-effective business-continuity solution for your mission-critical applications.

Learn more at www.AppAssure.com/virtualHA

appAssure APPLICATION ASSURANCE

vmware technology alliance PARTNER

Microsoft CERTIFIED Partner

Redmond MVP Most Valuable Product

Exchange offline backups

Not every Exchange backup is performed using the ESE APIs; it's possible to copy Exchange databases under a variety of other circumstances. By convention, any backup that doesn't use the online backup APIs is called an offline backup. This categorization includes making copies of dismounted databases using xcopy and using various tricky methods to make copies of open database files without closing and dismounting them.

The Microsoft article Offline backup and restoration procedures for Exchange, describes the process that Microsoft recommends for taking offline backups that include the log files necessary to do a complete restoration.

The big downside to offline backups is that they require you to do more manually, which is a concern during disaster recovery operations; more steps mean more possible ways to make mistakes, as well as more time spent performing the steps.

For example, most savvy administrators will run eseutil with the /K switch to check the restored database's integrity; doing so can add significantly to the restore time required. Microsoft's official position is pretty much that anything other than an online backup is an offline backup; this includes point-in-time and replicated copies.

Exchange disaster recovery planning

Paul Robichaux

This guide focuses on Exchange disaster recovery and availability, so it's not the appropriate place to provide a complete guide to disaster recovery planning. However, disaster recovery planning is so important that it's worth mentioning, even if only briefly. There are three components to a successful recovery plan:

- Having a plan—your plan must account for every possible contingency that might necessitate a recovery. At a minimum, this plan will include hardware failures, corruption or loss of your Exchange data, failure of the infrastructure components (such as Active Directory—AD—and electrical power) that Exchange requires, and interruption of physical access to your servers.
- For each of these contingencies, you need to have a response. This response might be simple (for example, if non-critical hardware breaks, you wait for the vendor's service technician) or complicated (if your Los Angeles data center is damaged by an earthquake, you fail over its operations in your Denver data center). The point is to accurately describe the potential problems you might run into, and to have solutions identified for them.
- Being able to follow the plan—just having a plan is fairly useless if you don't also have the ability to put your plan into action. This action will probably require a combination of money, persuasion, education, management support, and acquisition. For every solution you identify in your disaster recovery plan, you must have the necessary mix of equipment, skills, and preparation to make it actually happen.

Every cliche' you've ever heard about the value of prior planning applies here, in spades. The best way to make

sure that your disaster recovery plan includes both of the necessary components is to write down the plan and then practice it. Writing down the plan is important because it sets out everything that you think should be included—and that makes it easier to identify what's not included but should be. Practicing the plan is important because prior testing will make it much easier for you to identify shortcomings in the plan, in your equipment or infrastructure, or in the people who have to implement it.

The third component of a successful disaster recovery plan is perhaps the most often overlooked—keeping the plan up to date as your IT operations, staffing, and business requirements change. For example, a disaster recovery plan originally written for Exchange 5.5 doesn't take advantage of some of the best new features in Exchange Server 2003, such as recovery storage groups (RSGs). A plan that assumes restore windows of 12 hours might not work well when the actual current SLA only allows for 6 hours of downtime. Performing regular and frequent tests of your disaster recovery plan will act as an antidote to this problem by highlighting areas of the plan that need to be brought up to date.

The golden trine of Exchange Server disaster recovery

Brien M. Posey

If you are like most administrators, you probably dutifully back up your Exchange databases each night and store the backup tapes in a secure location. That's great—but those steps alone may not be enough to get you through a serious disaster. In this article, I explain three of the most important disaster recovery best practices for Microsoft Exchange.

Documentation

Regardless of your organization size or server complexity, it is critical to document as much information as possible about your server configuration—and keep it up to date!

Let me give you a case in point: A few years ago, I had an Exchange server catch fire (yes, this actually happened). I smelled the smoke and was able to put the fire out quickly. The rest of my network was OK, but that server was trashed. I was then at the mercy of my insurance company and backup tape.

At the time, I didn't have a full server backup available; all I had were the Exchange databases. This wouldn't be such a big deal—except, if you're restoring Exchange databases (not a full system backup) to a new server, the server has to pretty much mimic the old one. The server's name has to be an identical match of the old server's name. It also helps if the volumes are structured similarly to the way they were on the old server.

This is why documentation is so important. At a minimum, you should document:

- Your server's name
- IP configuration
- Disk configuration

- What each volume is used for
- Hardware specifications

That way, if you ever have to replace the server, it will be relatively easy to configure the new hardware in a similar manner to the old system.

Back up critical systems

Another thing you should do to prepare for a disaster is perform full system backups (including the System State)—at least once a month—of your most critical servers.

There are a couple of reasons for this. First, Exchange Server is dependent on Active Directory. I was fortunate in that a fire only destroyed one server. But imagine if all my servers had been destroyed. A backup of the Exchange databases wouldn't have done me much good without a functional Active Directory.

Ideally, you want to create full system backups of all your servers. But if operational requirements prevent this, make sure you at least get a monthly backup of:

Your DNS server. At least one domain controller in each domain (preferably the domain controller that holds the operations master roles for the domain)

A global catalog server. If you are a small organization, and your AD configuration hardly ever changes, you might be wondering why the once-a-month requirement is so important. Windows considers AD-related backups out of date after 60 days. There is a way to recover AD information older than 60 days, but it involves manipulating various tombstone settings, and can be difficult to accomplish. It's easier to just make sure you have a current backup.

Test your backups

Perhaps the most commonly overlooked disaster recovery planning step is to test your backups periodically.

When I worked for the military, we ran Exchange Server 5.0. The information store on one of the mail servers contained corrupted data that nobody knew about. Eventually, the corruption spread and the server crashed.

We soon realized that we had been backing up corrupt data for weeks, and none of our backups were any good. Had we occasionally tested the backups, we would have found out we had a problem and could have taken steps to repair the information store before it crashed.

Exchange Server 2003 is much more resistant to data corruption than Exchange 5.0, but it is still extremely important to test your backups from time to time. You never know when you might have a bad tape or some other unforeseen problem.

It's better to discover that you have a backup problem while your servers are still functional than while trying to recover from a disaster.

appAssure
APPLICATION ASSURANCE

Disaster can be costly.

According to research conducted by Gartner, two out of five companies that experience a disaster will go out of business within five years as a result of the event.

Preventing disaster is not.

Replay for Exchange® and BlackBerry® continuously validates your application, resulting in fewer incidents of unplanned downtime.
Guaranteed fast recoveries in the event of unplanned downtime get your team back online within minutes, and with virtually zero data loss.

Your e-mail is mission-critical, so treat it that way. Don't let data corruption and unplanned downtime affect the bottom line.

Learn more at www.AppAssure.com/prevent-disaster

Microsoft CERTIFIED Partner

vmware technology alliance PARTNER

Redmond MVP Most Valuable Product

Conclusion

Disaster recovery planning goes way beyond backing up your Exchange databases each night, and there's no such thing as over planning for it. Follow the three best practices I've outlined here, and you'll be well on your way to disaster preparedness.

Best Practice: Understanding Exchange databases

The Extensible Storage Engine (ESE) and its associated files are probably the least understood components of Exchange for many administrators. This can unfortunately lead to a number of misconceptions and bad backup and restore practices.

Here are a few basic database concepts all administrators should understand before creating a disaster recovery plan:

Database: In Exchange 2000/2003, the database is comprised of two files—an .edb (rich text data) and .stm (native content) files. Databases can be either mailbox stores or public folder stores.

Transaction logs: Before data is written to the database, it is first written to a sequential set of transaction logs. These logs are not only essential to the performance of Exchange, but they are also a key ingredient in an Exchange disaster recovery plan.

Consistent database: When a database is dismounted, or the information store service (store.exe) is shut down, all logs are committed to the database. After all logs have been committed, the header information in the database is flagged as "consistent."

Inconsistent database: When a database is stopped prior to all transactions being committed (because of a power outage, for example), the header of the database is not flagged and is therefore regarded as "inconsistent." To make the database consistent, the transactions that were not committed must be replayed into the database. A database must be consistent in order for it to be mounted.

Best Practice: Building your plan around the technology at hand

The sheer variety of disaster recovery technology to choose from can make you dizzy. Unless you have a virtually unlimited budget, you will need to develop a disaster recovery plan that is primarily based on the technology at hand.

If this is true for you, then your backup hardware will dictate the type of backups you will need to perform. The technology you have deployed will also control the schedule necessary to complete those backups.

There are some key factors in determining if the current technology should be replaced as part of a disaster recovery plan. Consider the following:

- Allowed backup times
- Allowed recovery times (a.k.a. downtime)
- Acceptable data loss (a.k.a. data loss tolerance)

If your backup and restore procedures cannot be configured to meet these requirements using your existing hardware and software, an upgrade may be required.

Many organizations will utilize incremental and differential Exchange backups to reduce the amount of time a backup takes. While this works for one side of the equation (backup times), it may not satisfy the other side (recovery times), since incremental and differential backups can actually take longer to restore then a full backup.

When it comes to data loss tolerance, simply running multiple backup types can provide you with an appropriate solution. For example you might perform the Microsoft recommended Normal (Full Online Backup) of Exchange Monday through Friday.

You could also perform an offline backup on Saturday. The offline backup is performed while the databases are dismounted or the information store service is stopped. This copy of the database, which is made while the database is in a consistent state, can be used in the event that you are unable to restore from an online backup. If seven-day data loss exceeds your organization's tolerance threshold, you may want to consider performing additional offline backups throughout the week.

Best Practice: Keeping e-mail in perspective

An e-mail disaster recovery plan should be an integrated part of an overall business continuum model/plan that in enacted for your entire organization. If you lose a key dependency of your e-mail systems, the first symptom your end users will see is that e-mail is down. However, Exchange Server might not be the problem. It could be a DNS, Active Directory, or network issue.

Make sure you document all Exchange dependencies and have a documented strategy of how those services can be recovered. The following are just some examples of the dependencies you should be prepared for:

- Storage Area Network (SAN)
- Windows Server 2003 (OS)
- IIS metabase
- Certification authorities (CA)

- Active Directory domain controllers
- Firewalls

Best Practice: Configuring Exchange for disaster recovery

Exchange Server 2000/2003 has built-in capabilities that can be taken advantage of to reduce single points of failure and facilitate the disaster recovery process. However, the Enterprise Edition of Exchange Server 2000/2003 includes extra features not available in Standard Edition.

- Enabling each of the following Exchange features could be considered a best practice:
- Adjust the Deleted Items retention to match your organizations data retention policy. By default, it is set to seven days in Exchange Server 2003.
- Adjust the Deleted Mailbox retention to match your organizations data retention policy. By default, it is set to 30 days in Exchange Server 2003.
- Put your transaction log files on a dedicated mirrored (RAID 1) volume.
- Put the information stores for each storage group on a dedicated RAID 5 or RAID 0+1 volume. (Enterprise Edition only).
- Divide your users' mailboxes up and spread them across multiple storage groups and mailbox stores (Enterprise Edition only).
- Configure Status and Notifications to monitor server resources like memory, disk, and processor so that you can be notified when resource failure is looming.

Best Practice: Simulating a disaster

Disaster recovery plans often fail because there was no testing of the recovery process. The best way to test the recovery process is to stage a simulated disaster.

There are many lessons that can be learned from this drill. You can document how much time it takes to respond to a given problem. You can determine if your procedures are accurate or if they need to be modified. Most importantly, you can get comfortable with restoring your data from backups and learn to trust the backup solution.

Prior to Exchange Server 2003, the biggest stumbling block for many small organizations in simulating a disaster was the need to dedicate recovery servers to the recovery process. Recovery servers can still be used to create disaster recovery labs, but are no longer necessary. This is good news for administrators on a tight budget.

appAssure
APPLICATION ASSURANCE

Downtime got you crazy?

Microsoft research reports that 42-percent of Exchange® recoveries from tape fail. Even worse, 20-percent of unplanned downtime is caused by Exchange corruption.

Don't sweat recoveries.

With Replay, you can rollback your Exchange or BlackBerry® environment to any point in time and recover from any type of failure within minutes, guaranteed. No more worrying about failed backups or incremental restores.

Reclaim your lost weekends and let Replay do the work for you.

Just point, click and restore.

Learn more at www.AppAssure.com/point-click-restore

Microsoft CERTIFIED Partner

vmware technology alliance PARTNER

Redmond MVP Most Valuable Product

Now the recovery storage groups in Exchange Server 2003 can be used to test restore procedures. A recovery storage group can be created on each Exchange 2003 server you are running. You must have enough disk space available to restore at least one of the information stores on each server.

Placeholder databases are created in the recovery storage group; they represent the database(s) you want to restore. After that, you go through the restoration process as if you were restoring to the live production server. The recovery storage group will intercept the restore destined for the production database and overwrite the placeholder database with the restored data.

Using Exchange System Manager, you can explore the recovered mailboxes. Exmerge can then export mail from the recovery storage group to .PST files and use Microsoft Outlook to view the .PST file and verify that data is intact.

This process should be incorporated into your daily operations as a sanity check of your backups. It should be performed by junior- and senior-level administrators so the procedures can be learned and reinforced regularly.

Resources from appAssure

Free download: Full-power trial version of Replay to help prevent disaster

Watch on-demand web casts on virtualizing disaster recovery

Access additional eGuides on Microsoft Exchange Backup & Recovery

About appAssure

AppAssure Software's Replay 2007 product suite combines best of breed protection for your Exchange environment delivering fast recoveries, e-discovery, and high availability, in a single solution that leverages your existing backup and storage investments. Using an industry-first virtualization-based approach, Replay's unique technology easily scales to support the consolidation trend of very dense Exchange servers and continually captures and self-validates the Exchange application at the volume-block level enabling fast and predictable recoveries. Replay's powerful features are available with unparalleled ease-of-use via a 'point-click-restore' interface that takes only minutes to install, configure and begin protecting.

www.appassure.com