

Realtime
publishers

The Shortcut Guide[™] To



Availability, Continuity, and Disaster Recovery

sponsored by

ARCserve.com[™]

Dan Sullivan

Chapter 2: Breaking Through Technical Barriers to Effective Recovery Management.....	18
Protecting Virtual Environments: Challenges and Solutions.....	19
Challenges Introduced by Virtual Environments.....	20
Performance Issues.....	20
Granularity of Backup and Restore Operations.....	21
Management Issues.....	21
Options for Backup and Recovery in Virtual Environments	21
Traditional File-Level Backups	22
Virtual Machine Image Backups	23
Mixed-Mode Backups.....	23
Backup by Proxy.....	24
Virtualization Vendor’s APIs and Services in Enterprise Backup Strategies	24
Benefits of Virtualization APIs.....	24
Virtualization APIs and File Restoration.....	24
Meeting Specialized Backup and Recovery Requirements of Databases and Content Management Systems	25
Email Backups.....	25
Database Backups.....	26
Content Management System Backups.....	28
Options for Addressing Specialized Backup and Recovery Requirements	29
Solving the Remote Office Backup and Recovery Challenge.....	29
Local Backup Option.....	30
Centralized Backup Option	30
Essential Features of Remote Office Backup Solution.....	30
Ensuring Continuity in the Event of a Disaster	31
The Need for Replication.....	31
Replication Failover Options	32
Summary	32

Copyright Statement

© 2010 Realtime Publishers. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers (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 or its web site sponsors. In no event shall Realtime Publishers 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.

[**Editor's Note:** This eBook was downloaded from Realtime Nexus—The Digital Library for IT Professionals. All leading technology eBooks and guides from Realtime Publishers can be found at <http://nexus.realtimepublishers.com>.]

Chapter 2: Breaking Through Technical Barriers to Effective Recovery Management

Information technologies are constantly advancing in ways that enable businesses to execute their strategies more efficiently and effectively than was possible previously. Virtualization improves server utilization, relational database provides high-performance data services, and email offers what has become a dominant form of business communication. With these advances come new levels of complexity, some of which have a direct impact on recovery management practices.

In addition to technical advances, there are organizational structures that create technical challenges to effective recovery management. The physical distribution of offices, for example, affects how we implement recovery management practices. If a company has multiple sites, it may not be practical to have skilled IT support in each office. Centralized IT support is often more economical; however, it raises the question of how to remotely provide recovery management protection. What starts as an organizational issue quickly leads to technical issues.

This chapter will examine several technical barriers commonly encountered when implementing recovery management services. These common challenges include:

- Protecting virtual environments
- Meeting the specialized backup and recovery requirements of databases and content management systems
- Solving remote office backup and recovery challenges
- Ensuring continuity in disaster recovery operations

Throughout this chapter, we will see examples of the need to adapt recovery management techniques to application-specific requirements and systems-implementation-specific requirements. These examples show that recovery management is much more than simply a matter of backing up files.

Protecting Virtual Environments: Challenges and Solutions

Server virtualization is widely adopted because it allows us to utilize computing resources more efficiently. With multi-core and multi-CPU servers, we can run compute-intensive jobs faster and on fewer servers. If we have a steady stream of these CPU-hungry applications, we can keep a server utilized. This is often not the case, though.

Typical business workflows have periods of heavy demand, such as when generating business intelligence reports for department heads and line managers, followed by periods of low demand. If we were to dedicate a single server to a single application, we would find under-utilization during off-peak periods. However, running multiple applications on the same instance of an operating system (OS) can lead to problems with incompatibilities. One application may require the latest set of link libraries while another older application is incompatible with those. Security requirements, like access to the root or administrator account, may be incompatible. The maintenance requirements of one application might require OS reboots that would unnecessarily shut down the other application. Virtualization avoids these problems.

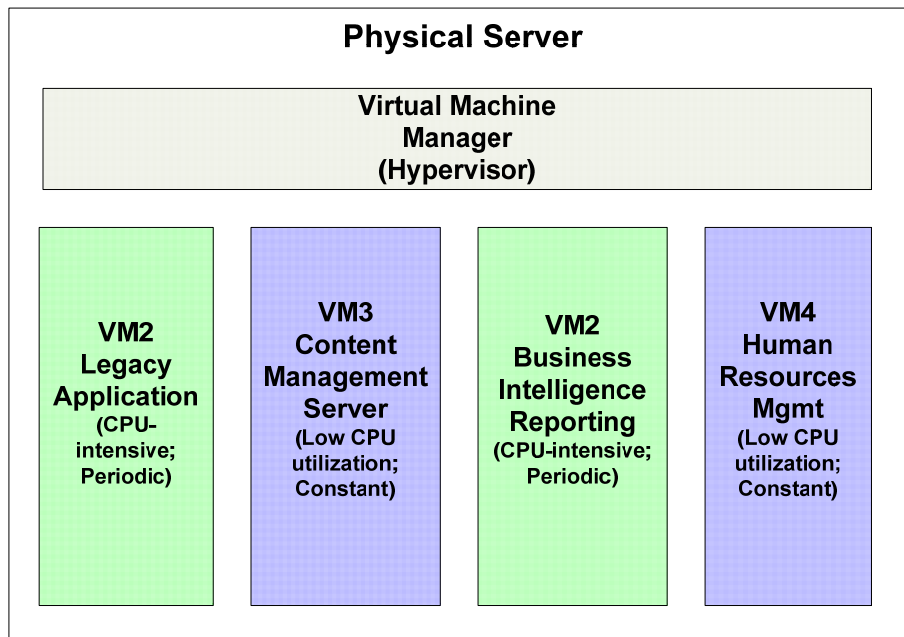


Figure 2.1: Virtualization allows for efficient mixed-workload combinations on a single physical server. Incompatible dependencies, security requirements, and maintenance requirements can make mixed workloads on a single OS instance impractical.

Of course, each virtual machine will have its own backup and recovery requirements. Meeting each set of requirements while minimizing the impact on other virtual machines is a source of new challenges we have to address.

Challenges Introduced by Virtual Environments

The types of challenges introduced by virtual environments fall into three broad categories:

- Performance issues
- Granularity of backup and restore operations
- Management issues

These challenges will strongly influence how we structure and schedule our backup operations.

Performance Issues

Backing up virtual machines can be compute-intensive, especially if deduplication and compression are done on the source system. In a virtualized environment, the guest OSs share the same physical resources, such as memory and bandwidth. If all the virtual machines were to run backup operations at the same time, there would likely be contention for these shared physical resources. Similarly, if backup operations on one virtual machine were schedule during the peak demand period of another virtual machine, such as a business intelligence reporting system, there could be contention for resources. In both scenarios, backups may not finish in the time allotted to perform them.

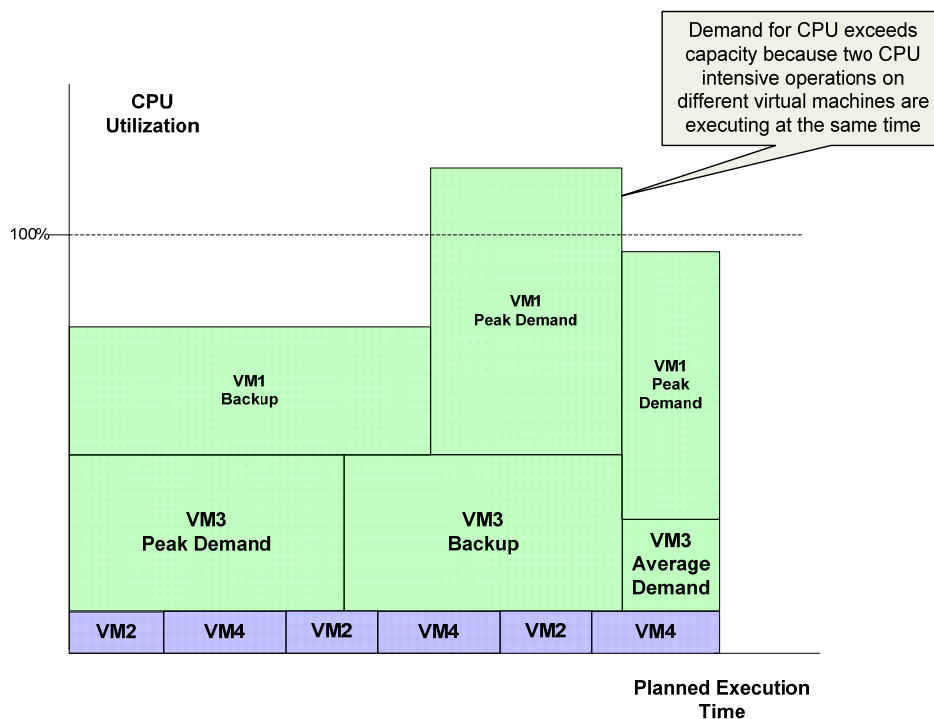


Figure 2.2: One of the challenges in managing virtual environments is understanding the distribution of workloads across virtual machines on a single server. Backup operations that execute at the same time as other CPU-demanding applications can exceed the available capacity and cause operations to run for longer times than planned.

Deduplication is an especially effective way to improve backup performance. Deduplication performed on the backup target relieves the source system of the CPU load associated with the process. Virtual machines tend to create a high level of duplicate data, so deduplication will result in cost savings as well.

Granularity of Backup and Restore Operations

A virtual machine can be backed up as a virtual machine image (and related configuration files) or as a set of files. Backing up a virtual machine as a single image can simplify backup procedures. All components of the virtual machine are backed up under the same criteria and a single backup image contains the entire virtual machine. This approach would be of limited value, though, if the backups could only be restored as a full image; often, restore operations are targeted to a single file or relatively small set of files.

Management Issues

There are a few management issues with regards to backing up virtual environments. Virtual sprawl, or the rapid deployment of virtual machines (sometimes outside of standard operating procedures), can cause headaches for the systems managers left to back up the growing numbers of virtual machines. Unless proper controls are in place to control provisioning and deprovisioning, the number of virtual machines active on a server can change quickly with new virtual machines adding to already complex backup schedules. Sometimes images of virtual machines that are no longer in use are left intact and included in backup operations, unnecessarily using compute, storage, and network resources.

Another management issue arises with regards to disaster recovery. Backups for disaster recovery purposes ideally could be restored to a bare-metal server in a different configuration of virtual machines. For example, a developer's server may host development and test virtual machines under normal operating conditions. Under disaster recovery conditions, the testing virtual machine may not be deployed in order to make resources available for mission-critical applications.

Virtual machines are started and shut down as dictated by business requirements. If a virtual machine is shut down during the backup window for its host, it should still be backed up. As virtual machines persist as images on disks, they should not have to be started to perform the backup. However, if the backup software used does not allow for image backups with file-level restores, systems administrators may opt to restart the virtual machine rather than forfeit the flexibility of restoring at the file level.

Options for Backup and Recovery in Virtual Environments

Systems administrators have a number of options for backing up virtual environments:

- Traditional file-level backups
- Virtual machine image backups
- Mixed-mode backups
- Backup by proxy

These options are not always mutually exclusive. For example, virtual machine image backups may or may not be performed using a backup proxy. It should be noted that this discussion is focused on how backup applications used in physical server environments can be leveraged in virtualized environment. In addition to these approaches, in a virtual environment, one may also be able to use storage area network (SAN)-specific techniques, such as creating snapshots of allocated storage units.

Traditional File-Level Backups

Virtual machines can be treated as equivalent to physical servers. Under this scenario, systems administrators would install a backup application and run whatever combination of full, incremental, and differential backup that is required for the virtual machine. The virtual machine will need to write data to a backup server, so sufficient network bandwidth will need to be available. Care must be taken to schedule backups at times that do not adversely affect other virtual machines on the same host (See Figure 2.2).

One of the advantages of the traditional file-level backups is the ability to perform fine-grained restore operations. Individual files are easily restored with this method. At the same time, however, this may not be a reliable method for full virtual machine restores.




Figure 2.3: In a traditional backup model, each virtual machine runs its own backup client.

Virtual Machine Image Backups

A second approach is to install a backup client in the host machine and back up virtual machine images.




Figure 2.4: Backup clients can also run in the service console of the virtual machine manager to provide image backups to all virtual machines on the host.

Images need to be in a consistent state throughout the entire backup operation; otherwise, we could encounter a situation where one part of the image is backed up, there is a change to the state of the virtual machine, and the rest of the image is backed up. In this case, the first and second parts of the backup may be out of sync.

One way to avoid this problem is to perform image backups only when the virtual machine is shut down. For images with frequent down time—for example, a business intelligence reporting system that generates nightly management reports and is then shut down—this model is sufficient. For virtual machines with high-availability requirements, a snapshot-based backup is a better option. With this method, a virtual machine is kept in a consistent state for a brief period of time, just long enough to make a snapshot copy. The image copy is then backed up without adversely affecting the running instance of the virtual machine.

Mixed-Mode Backups

Sometimes a combination of file-level and image-level backup is the optimal combination to meet recovery management requirements. For example, a weekly image backup followed by daily incremental file backups, has advantages over just file-level or just image-level backups. The weekly image backup provides the ability to rapidly restore a fully functional virtual machine. It does, however, often require more storage space to keep the entire image when compared with incremental file backups. If a relatively small percentage of files in a virtual machine change each day, incremental backups will capture those changes without unnecessarily duplicating data that is unchanged since the weekly image backup.

Backup by Proxy

Proxy backup servers relieve servers by taking on the load imposed by backup operations. Snapshot images of virtual machines can be copied to a proxy server where they are backed up. An advantage of the proxy model is that it allows for centralized management of backups. This method can also alleviate some performance issues by performing snapshots during non-business hours and then creating backups from the snapshots during business hours. Also, jobs can be scheduled to accommodate the particular requirements of each virtual machine and systems administrators can use a single management console to monitor and administer backup operations.

Note

It should be noted that although virtual machine vendors provide backup applications, advanced performance and management features may only be available with third-party backup applications designed to support virtual machines.

Virtualization Vendor's APIs and Services in Enterprise Backup Strategies

Advances in virtualization backup progress on two fronts: with virtualization vendors and with backup application vendors. (It appears that the age-old economic principle on the specialization of labor applies to new technologies such as virtualization and backups as well.)

Benefits of Virtualization APIs

Virtualization vendors are improving the performance and functionality of their hypervisor platforms. In addition to improving speed and reliability, vendors are providing application programming interfaces (APIs) that allow third parties to programmatically access key functionality. Enterprise application vendors, including recovery management vendors, can take advantage of these APIs to extend the functionality of their products to include support for virtual machines. For example, recovery management vendors will provide more support for advanced management features than one would expect from a virtualization vendor more focused on implementation details of their hypervisor.

Virtualization vendor's APIs are a critical linchpin that enable third parties to provide the same kinds of recovery management features provided for physical servers as well as to offer specialized functionality needed only in virtualized environments. Consider, for example, the challenge of recovering individual files from a virtual machine image backup.

Virtualization APIs and File Restoration

Restoring a single file from a virtual machine image typically requires a number of steps:

1. Restoring the backup virtual machine image to a physical server
2. Restoring the file from the restored virtual machine to temporary storage
3. Copying the file to target location
4. Shutting down the instance of the virtual machine from which the file was restored

Ideally, we should be able to restore individual files from an image backup. Although the functionality is not typically part of virtualization vendor offerings, this type of advanced functionality can be incorporated by third-party providers if the appropriate APIs are made available.

Protecting virtual environments pose plenty of challenges with regards to performance, granularity of backup and restore operations, and manageability. There is no single best way to back up virtual environments. By combining different modes of backups (for example, file vs. image), taking advantage of virtualization-specific techniques, such as backup by proxy, and exploiting extended functionality enabled by virtualization APIs, businesses can choose from a variety of options to find the best backup model for their needs.

Meeting Specialized Backup and Recovery Requirements of Databases and Content Management Systems

In its most basic form, backup operations are about making copies of files. Files many of us work with on a day-to-day basis, such as word processing documents, spreadsheets, and presentations, are easy to back up. These kinds of files are self contained and, unless they are open for update operations, are not going to change during the backup process. Not surprisingly, this simple model of file usage begins to break down as we consider more complex applications:

- Email
- Databases
- Content management systems

Challenges begin to arise in applications such as these because they utilize multiple files with different functions and different rates of change and varying levels of dependencies between them. It is not just the underlying design or technology that makes backups of these systems more difficult; sometimes, it is the way we use these systems and the organizational requirements we impose on them that stymie simplistic backup models.

Email Backups

Although we still use the term email for applications such as Microsoft Exchange and Lotus Notes, the name does not capture the extent of collaboration and communication functions provided by these applications. Even a basic email system today is likely to include:

- Email services
- Calendar
- Contact management
- Task lists
- Notes

Backup applications must be able to capture these different types of data and restore them to a consistent state. Although a disaster recovery situation may require a complete restore of an email system, a more common task may be restoring a small set of deleted messages, users' folders, or other data structures. In such cases, the ability to rapidly identify and restore selected data is essential to meeting recovery time objectives (RTOs). Backup applications can take advantage of metadata about the structure of user messages and other data to provide email-specific functions, such as single user restore. The requirements for email backup are now much more than basic copy and restore operations.

Archiving and e-discovery have grown in importance as email has become more central to the communications of business. E-discovery is the process in civil litigation in which electronic information is reviewed for details of relevance to the case at hand. We only need to look to the now-famous case of [Qualcomm v. Broadcom Corp.](#) for the importance of e-discovery. In that case, Qualcomm was severely sanctioned (at a cost of more than \$9.25 million) for failure to produce emails relevant to the case. The cost of backing up and archiving email can pale in comparison to the cost of insufficient e-discovery.

Key requirements for email backup with regards to e-discovery include the ability to:

- Create and maintain a comprehensive archive of all email messages
- Search messages for particular terms
- Search based on message metadata, such as date of messages, sender names, recipients, and so on

Email archiving programs may support these features natively in ways sufficient for small and midsize businesses. Large organizations may require specialized e-discovery software that imports email messages and other content into a content repository for specialized analysis and document classification.

Email and e-discovery demonstrate how organizational requirements can spur the development of advanced backup functionality. Databases are good examples of applications that provide plenty of technical drivers to backup innovation.

Database Backups

Databases are pervasive in today's software environment. Relational databases, in particular, have become the persistence mechanism of choice for many kinds of developers. A combination of features drives the adoption of relational databases:

- Ease of development with relational database
- A standard query language, SQL
- Broadly supported programming interfaces, such as ODBC and JDBC
- Reliability
- Scalability
- Flexibility

To realize these features, especially the last three, relational database systems depend on a complex storage management system that ultimately depends on low-level file system support. (Actually, in some cases, file systems can be bypassed in favor of “raw” disk access, but that is a much less common implementation option).

Databases Require Multiple Types of Files

The details of the storage management system will vary from one database management system to another, but the following list offers typical elements that affect how we perform backups:

- Files for storing data and indexes; they are sometimes called tablespaces
- Files for keeping a temporary copy of records as they are updated so that processes consistently read the same data, even if another processes is updating it; these files are sometimes called redo/transaction logs
- Files for auditing operations on data; these files may include username, times, client software information, and the IP address of the process making the changes
- Parameter files with configuration information for the database
- Error logs
- Temporary file space for sort operations

As with email systems, databases require specialized files to implement the full range of features and non-functional characteristics, such as scalability, that we have come to expect. Figure 2.5 shows a simplified version of a single update operation can result in changes to several types of files underlying the database system.

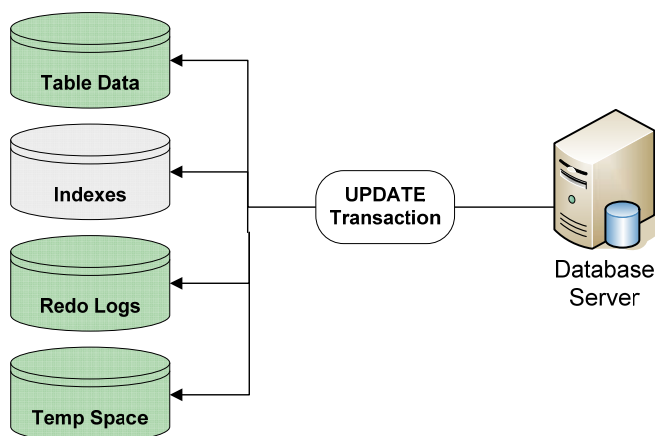


Figure 2.5: A single database transaction can update multiple files (green) as well as read from multiple files (blue). To capture a consistent state of the database, backups must be performed (a) when no changes are made to the database or (b) by a backup application that can track dependencies between the different database components and ensure a consistent database state is captured in the backup image.

The complexity of the underlying storage system creates difficulties for some types of backup and restores operations.

Restoring a Single Logical Record

Let's consider how we would restore a customer record that was accidentally deleted from a database. The customer record might include identifying information, such as name and address, purchase history, credit rating information, and account balance details. If all this information were stored in a single file, restoring it would be trivially easy. Relational databases efficiently manage large volumes of data, in part, by distributing the data in a single logical record across a number of tables. Names and addresses, for example, may be in one table, while purchase history is kept in several tables, including tables with order summary, order items, product codes, and other data. When a logical record is deleted, information may be deleted from a number of tables.

From a backup perspective, the critical question in logical record backup is, What rows from which tables need to be restored? The answer is highly application specific. Often, the answer can only be reliably determined by understanding the code that manipulates the database. Alternatively, rather than trying to restore a single logical record, one could restore to a particular point in time. The assumption here is that the database is always in a consistent state, so backing up to a point prior the time the logical record was deleted will allow you to restore to a consistent state in which the logical record is in place. A drawback is that any changes made since the restore point will be lost.

Online vs. Offline Backups

When databases are online and actively updated, there is the potential to make a backup copy of database files that are in inconsistent states. One method to avoid this is to take the database offline before performing the backup. This may not be practical when the database needs to be constantly available. Another approach is to export data using a database-specific utility that can capture a consistent state of the database using a combination of table data and redo log information. The export files can then be backed up using standard file-based methods. Content management systems, like databases, have application-specific backup and recovery issues we must take into account.

Content Management System Backups

Content management systems, such as Microsoft SharePoint and a variety of open source wikis, are increasingly used to manage unstructured data in analogous ways to which databases are used to manage structured content. (In fact, much content and configuration data is now stored in databases.) One of the advantages of content management systems is that they allow content owners to define groups of users with varying privileges to content. For example, a content creator may grant read and update privileges on a document to members of her department but only read access to other employees. Backing up content in these systems requires that we capture this type of access control information as well as the content itself.




Figure 2.6: Content management systems consist of both unstructured data in content repositories and structured metadata, including access controls. Both types of data must be captured by backup programs and must be consistent with each other.

Here again, we have an example where a logical unit of business information is stored across multiple files within a complex application.

Options for Addressing Specialized Backup and Recovery Requirements

The options for dealing with application-specific backup and recovery requirements fall into three broad categories:

- Shutting down applications and backing up all files used by the application. This ensures a consistent copy of data but at the cost of system availability.
- Using a specialized application to backup application data or export data to a file or files that are then backed up using a general backup application. This approach requires additional storage space for the exported files(s) until they are copied to a backup device.
- Replicating data to a standby system. The standby system could be taken offline to allow for backup without interrupting availability of the primary system. This approach also provides for rapid recovery in the event the primary system fails. This feature requires additional hardware and possibly software licenses.

The best option will depend on a combination of factors, including cost and RTOs.

Solving the Remote Office Backup and Recovery Challenge

One of the challenges with managing IT infrastructure in remote offices is the lack of onsite technical support. Often, there is no justifiable case for having full-time IT staff in each remote office. At the same time, we cannot expect the non-IT staff in those offices to suddenly become an on-call reserve for taking care of IT operations. Remote offices should have the appropriate level of recovery management services as required by their business operations; this is the same standard that should be applied to central offices. The question is, How do we deliver those services without busting the IT budget?

Local Backup Option

One option is to maintain remote office resources for backing up and restoring data. This could entail having a backup server in each remote office as well as sufficient disk space for all backups. For disaster recovery purposes, replicated data will have to be maintained on some type of off-site storage. In cases where there are few remote sites or sufficient network bandwidth is not available, this may be a reasonable option. As the number of remote sites grows, the economies of scale inherent in backup infrastructure come into play and a centralized approach would be more economical.

Centralized Backup Option

With a centralized backup system, data is copied from remote offices to a central location where it is backed up and stored. Centralized backups have a number of advantages over local backup options. First, the marginal cost of adding sites is low compared with a local backup option. Additional licenses for backup agents are required and sufficient bandwidth must be in place between the remote office and the central location. Additional storage and backup hardware is not required onsite. Second, by consolidating backup services, the central site can share resources among multiple offices. The chance of having unused storage capacity is reduced. Servers are more likely to be utilized because they can manage backups for multiple sites. Finally, systems administrators are on site and can quickly respond to hardware failures, networking problems, or other issues that require a knowledgeable person on site to correct.

Essential Features of Remote Office Backup Solution

When evaluating remote office backup options, look for features that increase manageability and limit demands on network resources:

- Centralized management console that allows administrators to schedule backup jobs, examine logs, receive alerts, and generate reports on backup operations
- Ability to back up remote servers efficiently through the use of deduplication and other methods to minimize demand on bandwidth
- Allow for centralized control of installation and updates to backup agents running on remote servers
- Ability to back up to either a central facility or to local storage at the remote site

Also, be sure to consider how the options under consideration affect your ability to meet RTOs and recovery point objectives (RPOs). Depending on requirements, you may have very short RTOs that demand off-site backup services.

Ensuring Continuity in the Event of a Disaster

An essential part of recovery management is preparing for disaster—that is, the loss of compute, storage, and network services which prevents the ability to deliver essential IT services. Disaster can be isolated to a single business, such as a fire that destroys a data center; regional, such as hurricane or earthquake damage; or widespread, like the Northeast Blackout of 2003 that caused widespread power failure throughout the US northeast and parts of Canada. When disaster disrupts IT services, by what means can they be restored?

Backups can be used to restore data and applications assuming servers and network devices are in place. Reconfiguring servers, installing OSs, and restoring backups can be a time-consuming operation subject to human, and technical, error. An alternative to waiting until there is a need for disaster recovery infrastructure is to maintain standby servers and keep them up to date with a continuous data protection process.

The Need for Replication

Replication services are used to keep standby servers up to date with primary servers. Consider an example of how these systems work: An order fulfillment database is continually updated during the day as new orders arrive. Orders are processed, customer credit is verified, inventory is checked, and shipments are readied from this system. The system is capturing business-critical transactions, so each time the database is updated, a copy of the update is sent to a standby server located in a separate data center. The standby server is running the order fulfillment system as well and is ready to take over for the primary server should it fail.

Replication systems can reduce their impact on production systems by minimizing additional computation or I/O on the production server. For example, rather than implementing a custom procedure to copy every transaction as it is executed in the application, low-level I/O operations can be duplicated instead. There is no need to execute full extent of the programming logic required to calculate the final results; duplicating the results is sufficient.

Another consideration with regards to performance is the demand on LAN and WAN resources. Replication technologies that minimize the amount of data replicated between primary and standby servers reduce overhead on the network.

Even with the additional overhead on production systems, the benefits of replication and continuous data protection can outweigh the costs. Replication reduces the time to recover by eliminating the time required to restore data from backup storage. By configuring the standby server prior to a disaster, there is more time to correct errors in configuration and diagnose other problems that may arise when setting up the standby server. The last thing any systems administrator wants when recovering from a disaster is debugging an unanticipated configuration error.

Replication also supports stringent RPOs. Replication processes can be configured to commit changes to the standby server on a frequent basis. This reduces the amount of data lost when the primary system fails. Only the data generated since the last update to the standby server would be lost, and that is typically far less data than the amount generated since the previous night's backup.

Replication Failover Options

With data replicated to standby servers, systems administrators have a number of options for failing over to the standby server. High-availability options for failover include:

- Server monitoring with a high-availability application—If a failure is detected, the system automatically sends traffic to the standby server.
- Manually redirecting traffic to backup servers—This can be done by systems administrators, for example, by updating local domain name services entries to map server domain names to the secondary server.
- Another option is to replicate data only and not have a standby server in place. This option reduces the delay in recovering data but does not tie up a standby server for a specialized purpose. This option could be used, for example, if a virtual machine image needs to be started to act as a standby server.

A general rule of thumb with failover options is that the faster and more automated the failover, the more complex and costly the solution. Such guidelines have to be considered with respect to business requirements. When rapid recovery is needed, high-availability options with continuous data protection are a sound option.

Summary

If recovery management were just a matter of backing up and restoring files, our professional lives would be much simpler. Complex IT systems have complex recovery requirements. Virtual machines are a boon to improving server utilization, but they introduce several challenges for backup and recovery operations, especially with regards to capturing a consistent state of the virtual machine without adversely affecting availability. Applications, as diverse as email, databases, and content management systems, make efficient use of file systems but in ways that challenge simple backup strategies. Remote offices need recovery management protection but cannot afford onsite, dedicated IT staff. Disaster recovery and the need for constant availability are driving the adoption of high-availability solutions. In all of these cases, we need to adapt recovery management practices and backup applications, including replication services, to these challenges. As the examples in this chapter show, a combination of recovery management practices and the most advanced backup and replication applications can be combined to meet these demanding challenges.

Download Additional eBooks from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit <http://nexus.realtimepublishers.com>.