

Realtime
publishers

The Shortcut Guide[™] To



Availability, Continuity, and Disaster Recovery

sponsored by

ARCserve.com[™]

Dan Sullivan

Introduction to Realtime Publishers

by Don Jones, Series Editor

For several years now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We’ve made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book’s production expenses for the benefit of our readers.

Although we’ve always offered our publications to you for free, don’t think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you \$40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the “realtime” aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We’re an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I’m proud that we’ve produced so many quality books over the past years.

I want to extend an invitation to visit us at <http://nexus.realtimepublishers.com>, especially if you’ve received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you’re sure to find something that’s of interest to you—and it won’t cost you a thing. We hope you’ll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Introduction to Realtime Publishers.....	i
Chapter 1: The Business Case for Recovery Management	1
Keeping the Business Running: The Obvious and Not-So-Obvious Business Requirements for Recovery Management.....	2
The Obvious Business Requirements for Recovery Management.....	2
Isolated Failures	3
Natural Disaster and Catastrophic Failures.....	3
Compliance.....	4
The Not-So-Obvious Requirements	4
Data-Driven Recovery Management Requirements	6
Expectations for Continuous Availability	9
Data-Driven Requirements.....	10
Understanding Operations-Driven Requirements.....	10
Effective and Efficient Operational Management	11
Opportunities and Constraints with Virtual Environment	11
Application-Specific Backup Requirements.....	12
Understanding Business Strategy–Driven Requirements.....	12
Improving Customer Service	12
Maintaining Continuous Access to Business Services.....	13
Developing a Recovery Management Strategy	13
Assessing Risks and Threats.....	13
Threats to Data and IT Operations.....	14
Risks to Business	15
Elements of a Recovery Management Strategy.....	15
Recovery Management Policies.....	16
Applications.....	17
Summary	17

Copyright Statement

© 2009 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.realtimedpublishers.com>.]

Chapter 1: The Business Case for Recovery Management

Professionals can develop their businesses with effective strategies, stay ahead of the competition by analyzing dynamic market conditions, and build brand loyalty with exceptional customer service—and it could all turn out to be for nothing. That is, if the IT systems that store, manage, and distribute our information fail and there is no recovery management process. It would be as if we had never done our work in the first place.

Data loss is an all-too-common problem. We lose information on the small scale with damaged laptops and misplaced flash drives. We lose information on the large scale with natural disasters that destroy entire data centers. Sometimes human error is at fault and sometimes applications fail with unfortunate consequences. Regardless of the cause of the initial data loss, the ripple effects can result in redundant work to reproduce the lost data, or in the worse case, to legal liabilities, brand damage, and business disruptions.

Recovery management is a framework to mitigate the risk of lost data and lost IT systems. It includes practices such as making backup copies of essential data, maintaining stand-by systems in case primary systems fail, and establishing policies and procedures to cost-effectively protect business assets, applying appropriate procedures based on the value of the data. Of course, no business practice will eliminate all risk or guarantee we can recover from calamitous events. We can implement cost-effective measures that allow a business to continue to operate at, or near, normal operating levels in spite of adverse events.

The purpose of *The Shortcut Guide to Availability, Continuity, and Disaster Recovery* is to provide you with the information you need to understand the business drivers behind recovery management, the technical aspects of recovery management, some of the operational challenges you might face, and best practices for implementing recovery management.

The four chapters of this guide cover the following topics:

- Chapter 1 discusses the obvious and sometimes not-so-obvious business drivers behind recovery management. The chapter also describes how to develop a recovery management strategy, including assessing threats and risks and outlining policies and applications needed to implement an effective recovery management strategy.
- Chapter 2 examines the challenges posed by the increasing complexity of IT environments, including virtualization, application-specific backup requirements, and remote office protection.
- Chapter 3 delves into common issues in recovery management, including scheduling and monitoring, media options, controlling costs, the growing volumes of data, and recovering in case of disaster.
- Chapter 4 digs into the details of how different types of organizations frame their recovery management strategies. The chapter concludes with a discussion of best practices for availability, continuity, and disaster recovery.

We start our examination with the business requirements that drive the need for recovery management.

Keeping the Business Running: The Obvious and Not-So-Obvious Business Requirements for Recovery Management

Recovery management addresses the threats of different kinds of losses, from hardware failures and software bugs to stolen laptops and malicious acts. One of the surprising aspects of recovery management is the number of different situations that benefit from having a sound plan in place. Some of these are obvious, but many are not.

The Obvious Business Requirements for Recovery Management

The “obvious” drivers behind recovery management are the reasons that come to mind first when we think of file backups and stand-by servers:

- Isolated software or hardware failures
- Natural disasters and catastrophic failures
- Compliance with regulatory or internal policies

It is easy to imagine a what-if scenario in these areas, especially if you have ever had a hard drive fail or tried to recover a system after fire or water damage. A brief meeting with an auditor can dispel any lax attitudes toward maintaining the integrity and availability of essential corporate data.

Isolated Failures

Isolated failures are limited in scope affecting few people or business processes. A typical example is the accidentally-deleted file. Someone may decide they no longer need a file and delete it. Fortunately, popular end user operating systems (OSs) frequently have a staging area for deleted files (for example, the Windows Recycling Bin and the Mac OS trash can) so that removed files can often be recovered by end users. Once the staging area of deleted files has been purged, restoring a backup copy of the deleted file is the best way to recover it.

Undelete Programs Helpful But Not Enough

Utility programs are available for recovering files even after they have been purged from a staging area, like the Recycling Bin. These programs work by reclaiming the data blocks on the disk that contain the contents of the deleted file before another application overwrites those blocks. Once the contents of a block have been written over, there is no easy way of recovering it, at least by conventional standards.

Application errors present another type of isolated error. A bug in a database application may incorrectly update data. As with OSs that provide a staging area for deleted files just in case there is an error, databases often store recovery information at least for short periods of time. If an error is caught in time, the database application can help recover the correct data. After that, restoring data from a backup copy is often the preferred method for correcting the mistake.

In addition to our own applications, we need to consider the risk of malicious applications, generally known as malware. Sometimes these programs are designed to corrupt files on compromised devices. If the corruption is found in time, backups can be used to restore files to their original states. Unfortunately, not all data loss incidents are so easily remedied.

Natural Disaster and Catastrophic Failures

Many of us only think of natural disasters when we are paying our insurance premiums. Like insurance, though, we will be glad to have backups and disaster recovery plans if disaster ever occurs.

Large-scale disasters, such as Hurricane Katrina in 2005 and the Northridge California earthquake in 1994, are infrequent, but fires, flooding, and other local events are common enough to warrant disaster recovery planning. Some of the key elements one needs to consider in a recovery management strategy relate to these catastrophic failures. When formulating a recovery management strategy and defining requirements, consider questions such as:

- If the data center with production servers were unavailable, what business operations would be affected?
- If a data center was destroyed by fire and all local backups destroyed, how would we recover?

- If the data center were unavailable, where would we house a temporary data center?
- How long would it take to restore business operations?
- How should we prioritize the restoration efforts? Which business process are the most critical?
- What level of degraded performance can be tolerated in disaster recovery mode?
- What is the cost per hour when an application or data center is unavailable?

Stakeholders in a business depend on IT professionals to protect information assets of the enterprise from the worst consequences of disasters. Regulations and internal policies define collective expectations for protection. Ensuring compliance is another obvious driver behind recovery management.

Compliance

There are many dimensions of governance and one of them is ensuring that business can continue to function under a range of circumstances, including the failure of key processes and systems. Recovery management plays an important role here. In the event of technical failure, human error, or natural disaster, business has a means to recover and re-establish a normal operating mode.

Compliance often entails more than just having a “Plan B” in the event of disaster. We need to demonstrate we have that capability in place and test it periodically to ensure our recovery management policies and procedures continue to meet the changing needs of the business.

Hardware failures, natural disasters, and compliance are obvious drivers behind the adoption of recovery management practices. They are not, however, the only aspects of business operations that should drive, and benefit from, recovery management.

The Not-So-Obvious Requirements

It is easy to think about backups and disaster recovery in the most basic terms: Make copies of important data so that you can restore in case of an adverse event. This is certainly sound reasoning but it does not capture everything we need to consider about recovery management. The problem with this line of reasoning is that it focuses only on data and not on other business aspects that drive the creation and use of that data in the first place.




Figure 1.1: Additional requirements for recovery management become clear when we consider the business strategy and operations that drive the creation, analysis, and management of business data.

If we examine why we create, analyze, and manage the particular types of data we have, we will find that the tasks are tied to some operational process. For example, we keep customer data for order fulfillment and sales operations. Human resources data is kept to track employees' performance history, benefits, and skill sets. These operations are in turn created in order to execute a business strategy, such as increasing market share, improving customer service, and retaining top talent.

If we consider each level of this three-tier model as a source of influence on recovery management, we can ask two broad questions. First, how does each tier shape requirements for recovery management? Second, does recovery management enable new capabilities that allow us to expand or improve each tier? To answer these questions, we will start at the bottom and work our way up with:

- Data-driven recovery management requirements
- Operations-level opportunities and constraints
- Business strategy and its demands for recovery management

These levels all include a combination of business and technology issues but with varying emphasis. Data-driven requirements are dominated by technical considerations while business strategy is, not surprisingly, subject primarily business considerations.

Data-Driven Recovery Management Requirements

Rule number one of data-driven recovery management is that not all data is of equal value. Before we can define recovery management procedures, we need to understand how data falls into different groups based on:

- How long we have to recover data once an adverse event occurs before the business suffers
- How much data can be lost because it was not backed up before an adverse event
- How fast the volume of data is growing

Sometimes it is more important to recover all data than it is to get it back quickly. A company's financial database may be down for several hours without significant impact on the business, but if even a single entry in the general ledger were missing, the integrity of financial reports is lost. In other cases, the time it takes to recover data is the most important factor. For example, as long as a company's product catalog is unavailable for online purchases, online revenues stop and purchases are potentially lost to competitors.

Recovery Time Objectives and Recovery Point Objectives

The duration between a data loss event and the point at which the data should be available again is known as the recovery time objective (RTO). The point in time from which we should be able to recover lost data is known as the recovery point objective (RPO). RTO specifies how long we can tolerate being without our data; RPO specifies how much lost data (in terms of time windows) we can tolerate.

RPOs are based on how much data we are willing to lose to a data loss event. Figure 1.2 depicts a basic backup strategy employing nightly backups. At any point in time, we can recover all the data from the previous day, but any data created or modified during the day a data loss event occurred would not have been backed up. This may be sufficient for applications with a low number of transactions during a day, such as an HR database tracking changes to employees' 401(k) funds. If data is lost, it is neither difficult nor expensive to recreate it. Applications with high levels of transactions or those for which recreating data would be difficult require more robust recovery management strategies, such as continuous data protection.

Cross-Reference

We'll talk more about continuous data protection later in this chapter in the Expectations for Continuous Availability section.




Figure 1.2: In the case of a simple example, backups are performed nightly. This implements the previous day's close of business as the RPO. In this example, the business is willing to risk the need to recreate up to a full day's worth of transactions.

In addition to deciding on an RPO, we must decide how long we are willing to be without our data. Some categories of data can have relatively long RTOs. Again, an HR application may be down for a day without severe adverse consequences. Sales and customer support applications and data, however, may require near continuous availability. In the event of data loss, the business operations that depend on these systems may not tolerate the time it would take a systems administrator to find the proper backup tape, select the lost data, and restore it to the application. In this way, our RTOs and RPOs constrain our options for implementing backup and recovery.

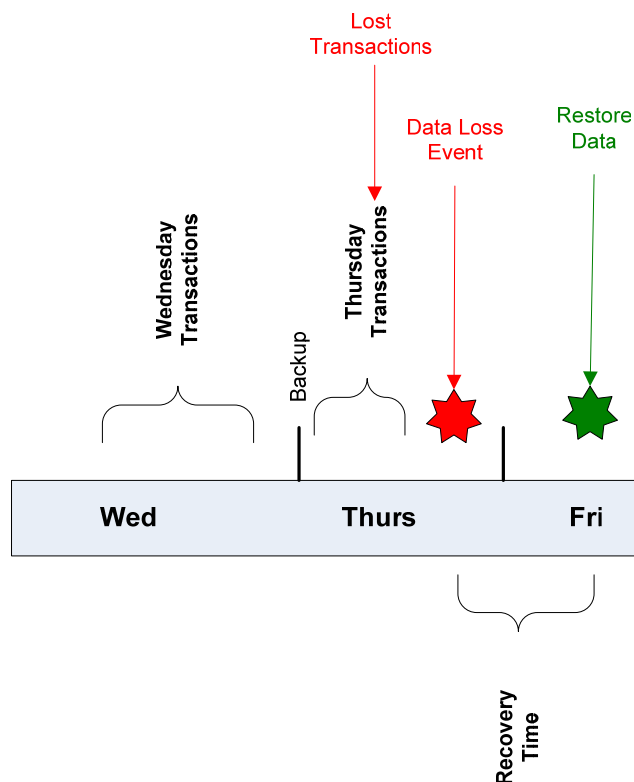


Figure 1.3: RTOs are defined by the amount of time that can pass between a data loss event and the restoration of data before there are adverse consequences for business operations.

Data Growth and Its Impact on Recovery Management

Another constraint on how we implement backups and disaster recovery procedures is the rate at which data volumes grow. There are many sources for increasing volumes of data:

- More automated contacts with customers, such as through self-service systems and online account management
- Increased use of business intelligence techniques to analyze sales, marketing, and operations data
- Detailed account and auditing of transactions in support of compliance and security management efforts
- Innovations in products and services offered to customers that generate new data
- Increased use of email and other collaboration systems

The rapid growth in data volumes is driving the adoption of better data management techniques, such as more efficient management of network storage devices and the use of deduplication in backup systems.

Data is easily duplicated. Database records, email messages, and multiple versions of a document can all be data structures that result in redundant data. An obvious question is: Why backup up and store redundant copies? Why not backup up one copy and track references to where the data is re-used? This is exactly what data deduplication does.

Deduplication processes operate either at the source system being backed up or at the target system receiving the copy of the backup. As each block of data is processed, the deduplication process determines whether a block of data with the same content has already been backed up. If it has, the system stores a reference to the copy that was made earlier instead of making another copy of the block.

Expectations for Continuous Availability

As our expectations for continuous availability grows, acceptable RTOs shrink. It is difficult to find maintenance windows to update applications, patch OSs, and perform other routine maintenance because customers are coming to expect 24 hour a day, 7 day a week access to applications. Again, the answer depends on the type of data and its level of criticality for business operations, but it is safe to say, for many customer-facing applications, the tolerance for downtime is close to zero. Businesses look to continuous data protection to ensure continuous availability. If data is so critical that we cannot tolerate virtually any downtime, data replication is probably the appropriate strategy.

With replication, as data is written to a primary system, it is copied to a stand-by system that maintains a close to real time copy of data from the primary system. If the primary system fails, operations switch to the stand-by system and continue as normal. When the primary system is restored, data that had been updated on the stand-by server is copied to the primary server and then operations can be shifted back to the primary system.

Key considerations include:

- Time required to update the stand-by system once a change has been made to the primary system
- Time to switch from a failed primary server to the stand-by server; should it be done automatically or can it be done manually and still meet RTO objectives?
- Tolerance for degraded performance with the stand-by server. For example, could the stand-by server be a virtual server sharing a physical server with several other stand-by virtual machines?
- Can the stand-by server be used for read-only operations, such as loading a data warehouse or generating reports, to reduce load on the primary server?
- Cost of additional hardware, and possibly software licenses, to perform replication




Figure 1.4: Replication duplicates all transactions on a stand-by server. In the event of a server or storage failure on the primary devices, the stand-by devices can be rapidly deployed.

Replication supports disaster recovery as well as continuous availability. Stand-by servers may be located in different offices or data centers from the primary site. This helps to mitigate the risk of site-specific threats, such as fire and flooding, to the primary site.

Data-Driven Requirements

The type and volume of data we have drives some recovery management requirements. Factors such as how long we can function without certain types of data and how long we can wait before data is restored have long been fundamental issues. The increasing volumes of data are also driving the need for more cost-effective storage strategies such as data deduplication. Expectations for continuous availability and the needs of disaster recovery are well met by replication technologies. In addition to these data-driven requirements, the day-to-day operations required to maintain an IT infrastructure are also the source of recovery management requirements.

Understanding Operations-Driven Requirements

Operations-driven requirements focus on the implementation aspects of recovery management. These are the issues that systems administrators and IT managers have to consider when formulating the best way to implement the data-driven and business strategy-driven requirements. Three commonly-encountered types of requirements are:

- Effective and efficient operational management
- Opportunities and constraints with virtual environments
- Application-specific backup requirements

We will delve into the technical details of these and other operation issues in Chapters 2 and 3, so we will just introduce some of the most salient elements of these issues here.

Effective and Efficient Operational Management

In an ideal world, recovery management procedures require minimal manual intervention, especially with backups, replication, and other ongoing tasks. Even in our less-than-ideal world, backup procedures and disaster recovery preparations should be as automated as possible for two reasons. The more automated the procedure, the less opportunity for human error. Backups can be scheduled for automatic execution. Logs can be generated for future reference. Errors can be flagged and generate alerts to notify systems administrators.

Restoring data and services can be automated as well. When replication is used, automatic failover to a stand-by server can be enabled, at least with some replication systems. A drawback of this type of rapid failover is a potentially more complex configuration. For example, an additional service may be required to detect the failure of the primary server and automatically redirect service to the stand-by server. Alternatively, if automatic failover is not used, a systems administrator could manually update a local domain name server to map a domain name to the stand-by server instead of the primary server.

Opportunities and Constraints with Virtual Environment

Virtualization can significantly increase server utilization and reduce costs but it also introduces new variables into the recovery management equation. The most basic question is how should we back up our virtual machines? There are several options:

- Treating virtual servers the same way we treat physical servers
- Shutting down the virtual machine and copying the machine's files to backup media
- Making a snapshot copy of the virtual machine while it is running and backing up that copy

Treating virtual servers as physical servers can simplify backup procedures, but it requires installing a backup client on each virtual machine. The second option eliminates the need for installing a client in exchange for shutting down the virtual machine. This may be acceptable depending on the function of the server. Snapshot copies require a staging area to store the snapshot and can briefly degrade performance of the virtual machine while the snapshot is made. The best option will depend on the specific business requirements of the virtual machine.

Application-Specific Backup Requirements

Backup and restore operations become more complex when we are working with files that are used with applications such as email servers and database servers. Consider some of the characteristics of database servers, for example. Databases typically use a small number of large files to store data about a large number of transactions. This has a number of implications for backup and recovery operations:

- If a single record in the database must be restored, will the entire contents of the storage file have to be restored?
- A single transaction may cause updates to multiple storage files. What is the optimal method for organizing those files to reduce the number of files that are updated by a single transaction?
- Can backups be performed while the database is active or must it be offline to ensure consistency across all transactions?

As we can see, the way applications use file storage to implement services, such as data management and email, can have an impact on the way systems administrators implement recovery management operations. Recovery management requirements are shaped in part by operational considerations, such as the efficiency of day-to-day procedures, the increasing use of server virtualization and the implications for backup operations, and application-specific constraints on the way we perform backups for databases and email systems.

Understanding Business Strategy–Driven Requirements

Unlike data- and operations-driven requirements, business strategies are as varied as businesses themselves, so there is no universal set of requirements we can all adopt as our own. Instead, in this section, we will consider two broad strategies that can provide examples to help elucidate the types of business strategy–driven requirements in your own business.

Improving Customer Service

A business may decide that providing online access to detailed, historical account data is crucial to improving customer service. Implementing this strategy will require increasing amounts of storage to support the customer service application, but it will also increase the demands on recovery management services. These increased demands include additional backup storage and increased throughput to continue to meet RTO and RPO with larger volumes of data. Meeting these demands can be done with a combination of additional hardware and network services as well as improved backup techniques, such as deduplication.

Maintaining Continuous Access to Business Services

Availability is a fundamental attribute of online services. It would be hard to imagine running a factory without a steady supply of power; it is equally difficult to imagine running a modern business without continuous access to the applications and data that provide business services. To mitigate the risk of lost services, businesses can implement redundancy at multiple levels:

- Redundant disk arrays
- Replicated data
- Stand-by servers
- Multiple points of access to the Internet
- Redundant power sources in data centers
- Physically distributed servers
- Well-defined backup and restore procedures

We must remember that there are many ways a business service can fail, so there will be multiple techniques required to mitigate that risk. Both traditional backup operations and replication services should be considered when there is a need to maintain continuous access to business services.

By considering recovery management from the perspective of data, operational, and business strategy requirements, we can identify essential aspects of business processes that need protection. The next logical step is to develop a plan for addressing those needs.

Developing a Recovery Management Strategy

The first stage in developing a recovery management strategy is assessing threats and risks to services. This is followed by assigning RPOs and RTOs as well as defining the policies and applications needed to address those threats and ultimately implement the recovery management strategy.

Assessing Risks and Threats

Risks are adverse outcomes that we typically want to protect against, such as data loss, system failure, or security breaches. Threats are ways in which a risk can be realized, for example, a data loss (the risk) could occur if a poorly-developed application inadvertently deleted files from a server (the threat). Although there are many types of threats, we will consider several with obvious impact on recovery management.

Threats to Data and IT Operations

Threats that disrupt the functioning of IT services fall into several categories, all of which must be addressed in a recovery management strategy:

- Hardware and software failures
- Malware and other security threats
- Natural disaster
- Human error
- Power failure

Hardware and Software Failures

Hardware failures are better understood than software failures. Consider the fact that when we buy hard drives, we can get estimated mean time between failures. This metric does not tell us when a hard drive will fail, but it at least gives us some idea of how long we can expect the device to function, at least on average. Software, including OSs, is more complex and diverse as well as developed under widely varying levels of quality control. There are no well-established metrics comparable to mean time between failures for measuring the reliability of software. From a recovery management perspective, it is safe to assume that both will fail and could corrupt relatively isolated sets of data or damage entire disks of data; given that assumption, we backup appropriately.

Malware and Other Security Threats

Security threats can pose significant threats to information systems. Threats such as viruses, worms, Trojan horses, and blended threats (multiple attack vectors in a single package) can all be used to destroy or tamper with data. Data breaches that result in large numbers of disclosed records are well documented in the popular press. Security threats to the integrity and availability of data are less frequent topics of discussion but still dangerous to businesses. Reliable and timely backups can make a significant difference in the cost and time it takes to recover from a security breach. Of course, if files on backups are corrupted by malware or other security threat, this is not an option,. Often the best strategy is to have a security management strategy in place to mitigate the risk from malware and other security threats. One way to mitigate malware risks is to use backup software that contains antivirus software, which can scan files during both backup and restore operations.

Natural Disasters

Natural disasters need little explanation. The key questions we need to answer about disaster recovery include where to store backup copies of data and stand-by servers, how long are we willing to tolerate service disruption, and what procedures need to be in place to ensure services can be started at a disaster recovery site. In addition, what are the steps to resuming normal operations once the primary site is up and functional?

Human Error

Human error will always be with us, so we must design systems in ways to minimize the potential impact of error. Programmatic techniques, such as validating input and prompting for verification of destructive operations, are one way. Organizational techniques, such as separation of duties and requiring authorizations from multiple individuals are another way to mitigate the risk of human error.

Power Failures

Disruption caused by power failures can be mitigated with multiple power supplies. Large data centers may employ a redundant source of primary power, including on-site generators, which may not be practical for smaller facilities. Facilities of any size should consider uninterruptible power supplies (UPS) for temporary power. A UPS can provide power during brief outages and allow time for a controlled shutdown of systems in the event of long outages. These types of risks are just some of the ways risks to business can be realized.

Risks to Business

When systems are down and information is unavailable, businesses are adversely affected. Some of the most immediate concerns we have about loss of business continuity are:

- Lost productivity and backlog of work
- Loss of revenue because sales cannot be completed, pre-sales information cannot be provided, orders cannot be processed
- Loss of customer confidence and brand damage that can arise from the inability to access systems and account information or execute transactions; there is also the potential for lost confidence in the business to provide reliable, robust services
- Cost to restore operations to a normal state; without proper planning and disaster recovery management, the task of recreating data and reinstalling systems under tight deadlines can be costly
- Cost of fines for compliance or e-discovery violations resulting from being unable to produce data as required.

Businesses face a host of risks to their operations and many risks can be realized by multiple types of threats. It is prudent, and cost effective, to plan ahead and develop a recovery management strategy before an adverse event occurs.

Elements of a Recovery Management Strategy

A sound recovery management strategy is a combination of (1) policies that address the various data, operations, and business requirements with respect to the risks and threats a business faces and (2) applications and technologies that enable the implementation of those policies.

Recovery Management Policies

The purpose of recovery management policies is to document and put into practice methods for mitigating the risks facing businesses. Five types of policies are essential:

- Backup policies
- Continuity and failover policies
- Disaster recovery policies
- Testing policies
- Security policies

Policies should define the scope of what should be done to mitigate risks; technical implementation details are defined after policies are formulated. They are codified as procedures that are executed by systems administrators and other IT professionals responsible for day-to-day operations.

Backup Policies

Backup policies specify what types of data and applications should be backed up, the RPO for each type of data, and the RTO for each as well. For example, an HR database may have an RPO of the previous business day and an RTO of 4 hours. Procedures for this policy may call for a combination of weekly full backups plus incremental nightly backups.

Continuity and Failover

Continuity and failover policies focus on critical data and applications. The purpose of these policies is to ensure that systems that should be available at all times are protected with high-availability techniques. For example, a sales database may have an RPO of the last 10 minutes and an RTO of 10 minutes as well. These demanding constraints warrant a replication-based solution.

Disaster Recovery

Disaster recovery policies specify what disaster recovery procedures should accomplish and who should be involved. These policies specify criteria for establishing disaster recovery sites or services, such as location in separate buildings or different localities depending on the criticality of the data and services protected. They should also specify the RPOs and RTOs of different categories of services. The policy should also include some description of when disaster recovery procedures are implemented, typically when service infrastructure is so compromised that normal services cannot be maintained.

Testing

Disaster recovery policies should also indicate the need to test disaster recovery procedures and systems at regular intervals. Modifications to procedures should be tested when they are implemented and then tested again during regularly scheduled test operations.

Security

Security policies must take into account much more than recovery management but should include directives on the appropriate use of the Internet and restrictions on installing non-authorized software on company devices.

In addition to policies defining what is required of disaster recovery procedures, we need applications to meet those needs.

Applications

Disaster recovery depends on two types of systems: backup and restore applications and high-availability solutions. Backup and restore applications give us the means to recover from a wide array of adverse events, from hardware failures that lose data and software bugs that corrupt the integrity of data to natural disasters that destroy entire data centers. It is important to consider backup and recovery operations when deploying new systems and implementing new business services. We must be able to back up and restore critical data with the time ranges allotted to us by the business. Growing volumes of data make this more difficult; however, techniques like deduplication can help us keep pace with the growth in data volumes.

High-availability solutions allow us to replicate services and data on stand-by servers and keep them up to date. These solutions are essential when we must maintain 24 × 7 systems and allow for extremely short RTOs.

Summary

A recovery management strategy should take into account a variety of requirements. Some of these requirements are a function of the criticality of the data we have, some are dictated by operational and efficiency considerations, and others are derived from business strategy. Regardless of the source of the requirements, a sound recovery management strategy starts with codifying those requirements in policies that can be used to develop operational procedures to protect business services and data. Applications such as backup and restore systems and high-availability solutions play a critical role in implementing those policies and procedures. We will turn our attention to those implementation issues in the next chapters.

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>.