Realtime
publishers

# *Tips and Tricks Guide™ To*

# Windows Administration

*Don Jones and
Dan Sullivan*

## *Copyright Statement*

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.

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.

## Tip, Trick, Technique 17: Identifying Threats of Data Loss in a Windows Server Environment

No matter how robust, reliable, and secure Windows Server 2008 is, we face risks of data loss. Systems administrators know all too well how often someone accidently deletes a file and then somehow manages to accidently clear the recycling bin as well. Data loss like that falls at the easy end of the spectrum of problems we face with protecting data on Windows Server 2008 platforms. At the other end of the spectrum, we have a big challenge when it comes to data loss: disaster recovery. How will the organization continue to function if critical applications are down because servers were destroyed in a fire, flood, hurricane, or some other natural disaster that might be a familiar threat in your area. Between the simple and the complex, we have a whole range of data loss risks:

- Accidental data loss due to human error
- Accidental data loss due to application error
- Intentional data loss due to malware
- Intentional data loss due to human actions
- Data loss due to natural disaster

We can quickly see from this list, the threats underlying the risk of data loss can be roughly grouped along two dimensions: intentional versus accidental and programmatic versus human action. Natural disasters are something of a special case, and we will discuss that in different terms. One final note about the grouping, the dimensions are not mutually exclusive. We could have a situation where a problem in an application, say a bug in a patch, combines with mistakes applying the patch by the administrator to create a compounded threat. Clearly, there is no shortage of ways in which our data can be lost.

**Note About Terminology: Risks and Threats**

The terms *risk* and *threat* are sometimes used in ways that obscure their precise definitions. A risk is a hazard or potential loss, such as the risk of losing data, having data tampered with, or having login information stolen.

A threat is a means for realizing a risk. A single risk can have many threats that can bring about the unwanted outcome. For example, someone could steal your online banking credentials by overlooking your shoulder and watching you type or by installing a key logger that captures keystrokes as you type.

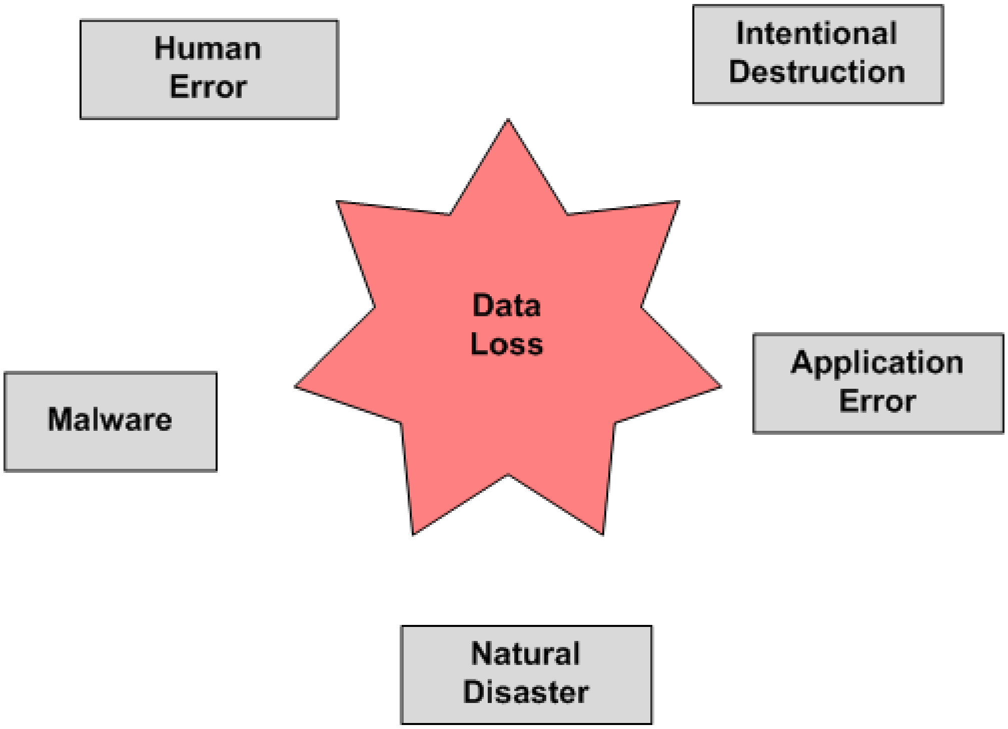

**Figure 34: The risk of data loss is due to multiple threats.**

The first step in understanding the risk of data loss is to understand how these different threats operate to undermine the integrity and availability of our data.

## Accidental Data Loss Due to Human Error

The threat of human error is constant; our best hope is to mitigate the risk by putting in place controls and procedures that reduce the likelihood of accidents. It is helpful to think of human error in terms of user errors, which can be bad, and administrator errors, which can be *exceptionally* bad.

End user errors lead to relatively isolated data loss: deleted files, corrupted records in a database, and overwritten files in a shared directory. These types of errors can be mitigated with access controls that limit delete and write privileges to only necessary users. In the case of application-related errors, improving usability and prompting for confirmation for destructive changes can help reduce the risk of data loss.

Administrator errors are more difficult to prevent. If you've been in systems administration long enough, you probably have tales of mistakes that still make you cringe. One way to reduce the risk of data loss from administrator errors is to document procedures and use checklists to ensure the procedures are followed. By its nature, systems administration often requires us to perform unique tasks, such as applying a particular service pack; however, after you have performed this task a few times, you can develop a pattern that can be generalized enough to create a checklist of essential steps (not the least of which is creating a backup before you start).

### Accidental Data Loss Due to Application Error

The "blue screen of death" has been a well-known phrase since the days of Windows NT. If an application made an error *and* did not properly trap for it *and* the operating system (OS) did not properly isolate the error, then it was time to abandon all hope and reboot. Today's application code and OSs like Windows Server 2008 are more resilient than their 1990s counterparts, but accidental data loss due to application error is still a problem. Not surprisingly, it is the complexity and interoperability of applications that create significant threats of data loss today. Consider some of the ways a Web application with a rich Internet application (RIA) interface may lose data:

- A bug in a browser add-on crashes Internet Explorer while executing a multistep workflow
- An error in a SQL Server stored procedure running in the application's database mistakenly corrupts data before finishing a transaction
- A misconfigured setting in Internet Information Services (IIS) causes a Web server failure in the Web server before data is posted to the backend database

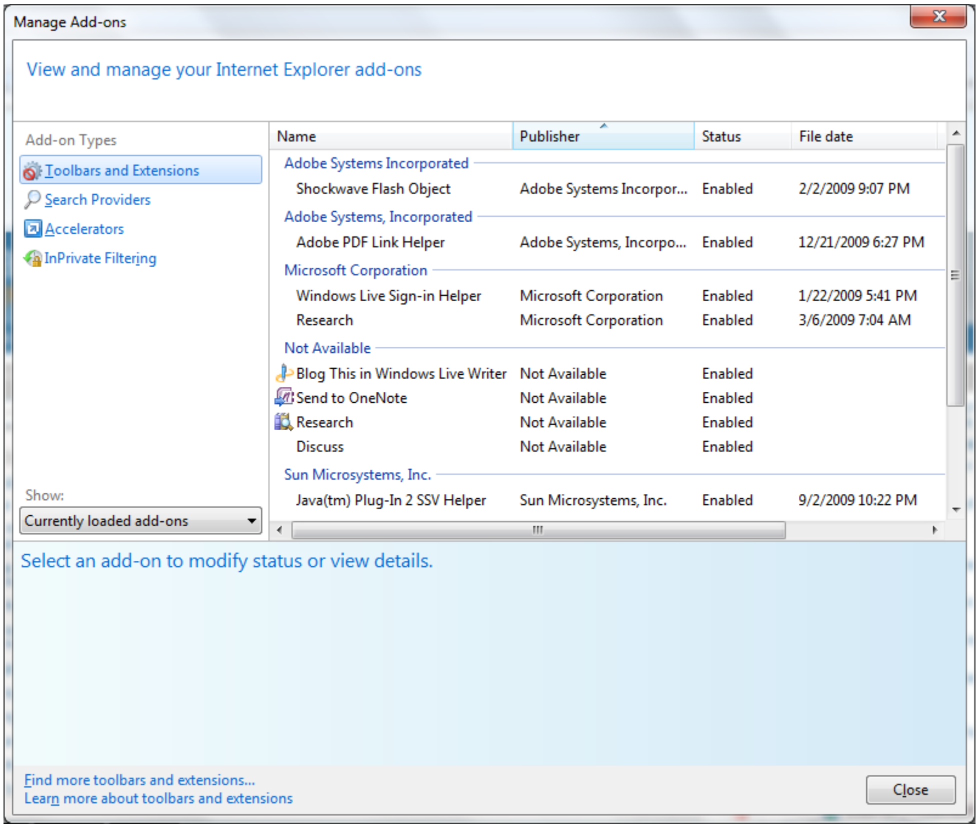

View and manage your Internet Explorer add-ons

| Name | Publisher | Status | File date |
|---|---|---|---|
| Adobe Systems Incorporated | | | |
| Shockwave Flash Object | Adobe Systems Incorpor... | Enabled | 2/2/2009 9:07 PM |
| Adobe Systems, Incorporated | | | |
| Adobe PDF Link Helper | Adobe Systems, Incorpo... | Enabled | 12/21/2009 6:27 PM |
| Microsoft Corporation | | | |
| Windows Live Sign-in Helper | Microsoft Corporation | Enabled | 1/22/2009 5:41 PM |
| Research | Microsoft Corporation | Enabled | 3/6/2009 7:04 AM |
| Not Available | | | |
| Blog This in Windows Live Writer | Not Available | Enabled | |
| Send to OneNote | Not Available | Enabled | |
| Research | Not Available | Enabled | |
| Discuss | Not Available | Enabled | |
| Sun Microsystems, Inc. | | | |
| Java(tm) Plug-In 2 SSV Helper | Sun Microsystems, Inc. | Enabled | 9/2/2009 10:22 PM |

Show: Currently loaded add-ons

Select an add-on to modify status or view details.

**Figure 35: Application stacks are becoming more complex; even browsers are collections of add-ons each of which can harbor application vulnerabilities that can lead to data loss.**

Although there have been advances in some areas of application development, especially in the area of OS robustness, the additional complexity in today's applications harbor the potential for data loss.

## Intentional Data Loss Due to Malware

Malicious software, commonly known as malware, comes in a number of forms, all of which can either directly or indirectly result in data loss if a server or client device becomes infected. Some typical types of malware include:

- Viruses, which are programs that compromise other programs in order to carry out an attack
- Worms, which are programs that can spread and operate independently of other programs (unlike viruses)
- Keyloggers are programs that capture keystrokes, including authentication details, which can be used at a later time to compromise a system
- Trojans, programs that appear to be one thing (for example, a utility) but carry out unwanted operations, such as stealing information

- Rootkits are one of the most pernicious forms of malware—they hide themselves by corrupting OS-level services; getting ride of a rootkit is difficult and victims may have to resort to wiping the hard drive or scanning and repairing from a known safe boot device
- Blended threats combine multiple types of malware in a single attack vector

Malware can cause data loss either because the malware developers designed their code to destroy data or because the malware interferes with other operations resulting in data loss. Although data loss is a problem, a bigger concern with malware is loss of confidentiality with malicious code that steals files or logs keystrokes.

### Intentional Data Loss Due to Human Actions

Disgruntled employees are nothing new. Some may not like their jobs, some may have psychological issues, and some might be looking for payback after a layoff or perceived unjust sanction. Concerns about insider abuse are probably more at the forefront of our minds during challenging economic times that inevitably lead to layoffs. The ways a disgruntled employee can cause data loss are limited only by their imagination.

One particularly difficult form is the logic bomb. This should be included with the list of malware, but we address it here because it is malicious code introduced by an insider. A logic bombs is code that is set to execute at some time after the code is introduced and will destroy, corrupt, or otherwise tamper with data or applications. The damages from a logic bomb can extend beyond the business or organization that is the initial victim. In 2008, a former systems administrator at a health services company was convicted of creating a logic bomb that would have destroyed virtually all information on the company servers, according to one report, including healthcare information (Source: Sharon Gaudin, "Medco Sys Admin Gets 30 Months for Planting Logic Bomb," *Computerworld,* January 8, 2008. http://www.computerworld.com/s/article/9056284/Medco_sys_admin_gets_30_months_for_planting_logic_bomb).

At some point, the level of data loss caused by intentional human action crosses the boundary into a more disaster-like situation. For example, if an arsonist succeeds in seriously damaging a data center, the level of data loss would approach that of a loss due to a natural disaster.

### Data Loss Due to Natural Disaster

There are two aspects of data loss due to natural disaster that distinguish it from other threats of data loss: the scale of data lost and the accompanying loss of infrastructure. Natural disasters do not selectively target data, the way malware might, and it is not limited to a single application or database, the way an application error might be; natural disaster can wipe it all out. When considering how to mitigate the threat of data loss due to natural disaster, we must consider how we will provide temporary servers and other infrastructure to run critical applications.

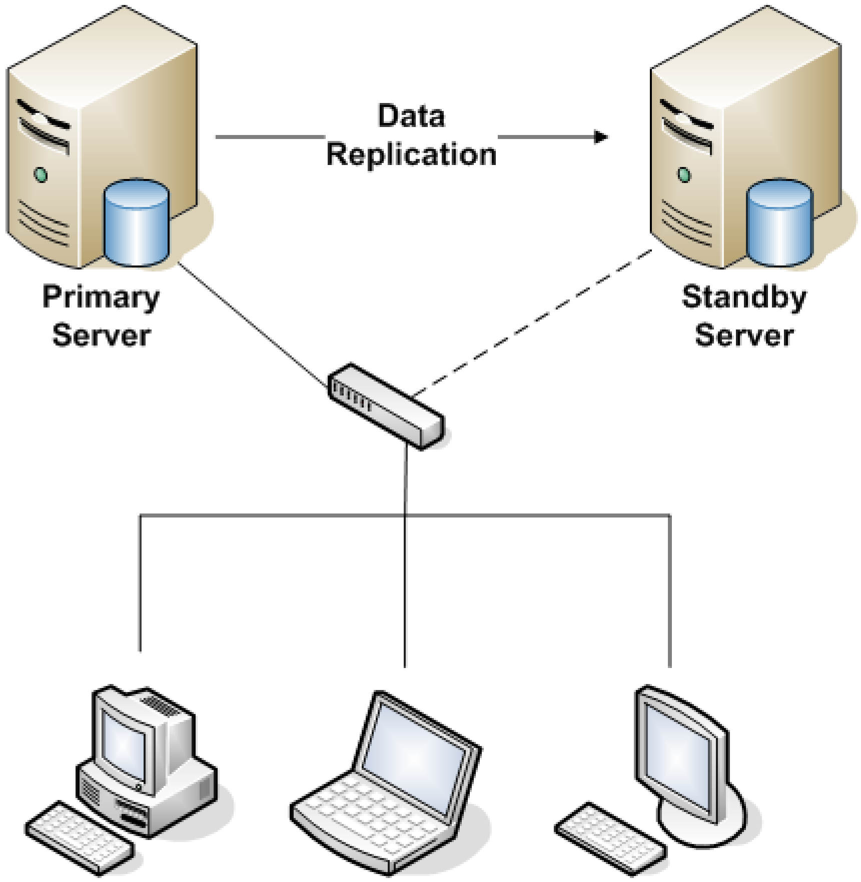

**Figure 36: Failover systems replicate data from a primary to a secondary system so that the latter can take over in the event of a failure in the primary server.**

We also need to consider how long IT services can be down before there is significant adverse effect on the business or organization. If rapid recovery is essential, then we need to consider high-availability solutions. With these systems, data is replicated from primary servers to standby servers. Primary servers may be monitored and if they fail, applications will failover to the standby server; in other cases, manual intervention is required to switch to the standby server. (We will have much more detail on high-availability and data replication services in a future volume).

Once the threats that can lead to data loss are understood, we can devise a plan to mitigate the risk. Clearly, backups will play a role in data loss protection, but as we will see, there is much more to reducing the risk of data loss than simply making backups.

## Tip, Trick, Technique 18: Understanding the Building Blocks of a Recovery Management Strategy

A recovery management strategy is a plan for reducing the chance of data loss due to any of the threats described in Tip, Trick, and Technique 17. With an overview understanding of the threats, how do we go about protecting our Windows Server 2008 servers and other infrastructure? It starts with a four step process:

1. Create a data classification system and categorize data to be protected
2. Identify critical servers and applications needed for different categories of data
3. For each category of data, determine recovery point objectives (RPOs) and recovery time objectives (RTOs)
4. For disaster recovery purposes, determine the level of performance required when operating in disaster recovery mode

At the end of the process, we have described the level of protection required to mitigate the risk of data loss balanced against the requirements and resources of the organization.

### Creating a Data Classification Scheme

Think about all the different types of data in a typical midsize business. (The principles we develop apply equally well to non-business organizations, but for simplicity, we'll use a business example here). There is transaction data about sales, customer details and account summaries, HR data about employees, data warehouses and executive reporting data, emails, documents, and other unstructured data. Now we need to ask, Is all this data equally valuable? Another way to think about it is, How would the business be affected if the data were lost?

- Would there be a sudden and significant negative impact on the business? Losing a sales order database would probably fall into this category.
- If the entire data warehouse were lost, how bad would that affect the business? As data warehouses are traditionally used for management reporting but not core operations, the impact would be limited. Furthermore, some parts of the data warehouse could be reconstructed from data in transactional systems, although some historical data may be lost.
- If an HR content management system containing employee work plans were lost, there would be some need to recreate this data. This task could be done over a longish period of time without having a significant adverse impact on the businesses.

From these examples, we can see three categories for data classification: critical, important, and optional. Critical data deserves the greatest level of protection (we'll define what that means in operational terms shortly), important data should be protected but not at the expense of critical data, and finally, optional data should be protected if possible but is of lower priority than the other types of data.

A key benefit of having a data classification scheme is that is allows us to prioritize how we commit resources to protecting data, and that priority is based on business, not technical, requirements. For example, a business may have two Windows Server 2008 systems running SQL Server; one is hosting an orders database and the other is used for a data warehouse. It is the type of data in the database, not the fact that the server is used for SQL Server, that determines its data protection priorities.

| Data Classification Category | Description | Priority | Example |
|---|---|---|---|
| **Critical** | Data that is essential to the continued operation of the organization. If the data were lost, it would severely and adversely impact the organization. | Highest | Financials, Customer database |
| **Important** | Data that is needed for normal operations. If the data were lost, it could be recreated with some effort. Its loss would not have an immediate adverse impact although long-term loss would. | Medium | Data warehouse, Marketing data |
| **Optional** | Low value data that would not adversely impact the organization if it were lost for an extended period. Is easily recreated at low cost. | Low | Copies of publically available data (for example, census data used in marketing) |

**Table 1: Data classification schemes partition data by value and impact on business operations.**

### Identifying Critical Servers and Applications

Just as some data is more important than other data, some servers are more important to a business or organization. To identify which servers are most important, we need to understand what functions the server carries out, in terms of business processes, and what data the server stores. Often, but not always, there will be overlap; critical business data resides on critical servers. This is not always the case. For example, a development server is critical to a software development group but it does not (or at least should not) have any critical organizational data stored for anything but development purposes.

#### Critical Servers Host Critical Applications

The first step in identifying critical servers is to create a high-level map of where different types of data reside. For example, servers could be labeled as storing critical, important, or optional data as well as a combination of multiple types. In the case of multiple types, the server should be considered as having the higher priority category of data.

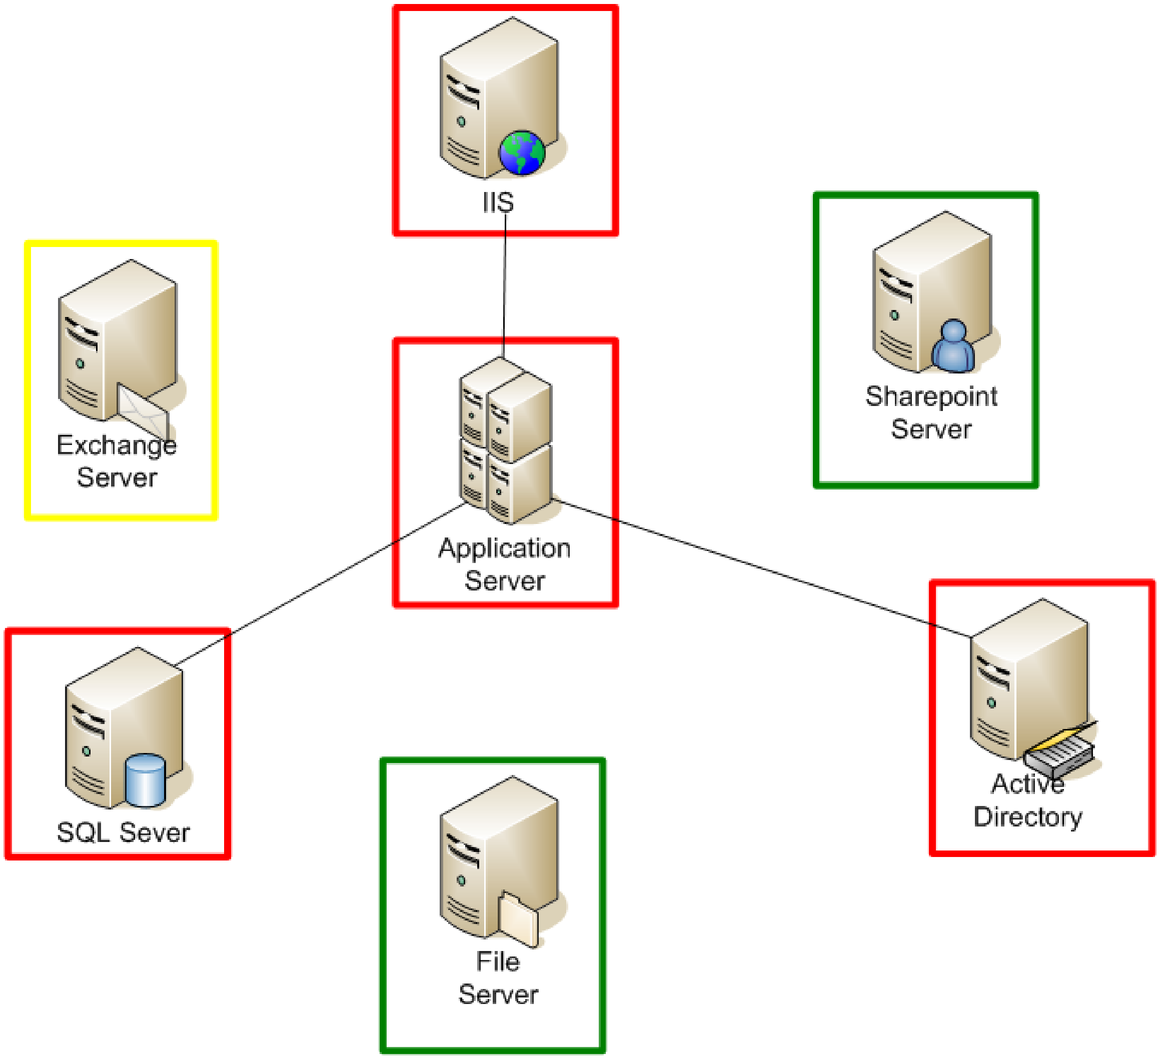

**Figure 37: Servers are considered critical if their core function is a critical business function or if another critical server is dependent on them.**

Consider a simple scenario of a small business or a department within a larger organization. There are several applications running on an application server and all of those systems are considered critical. That adds the application server to the critical list. But we can't stop there. The applications running on that server depend on a SQL Server database that is hosted on another server. Authentication to the applications depend on an Active Directory (AD) server on yet another server. Finally, the applications are of little use without the Web interface that allows users to interact with the application. What started as one critical server quickly became four because of application dependencies.

When creating these prioritized lists of servers, it helps to have someone with knowledge of the application architecture to help run down all the dependencies. The last thing any of us want is to get into a disaster recovery situation only to learn we missed a critical dependency.

Of course, when we speak of data on the server, we really mean data that is logically managed by that server. The data may actually reside on a storage array that is shared by multiple servers. From the perspective of protecting against data loss, that does not matter. If the server is down, the data it manages is not readily accessible even if the storage array is functioning.

### Critical Servers Support Critical Business Processes

If we conducted a survey and asked IT and business professionals to define the critical applications used in their business, we would probably get many answers about sales order systems, customer relationship management systems, financials, and other back-office applications. These certainly fit into the critical category, but they do not cover the full spectrum of essential systems. Take messaging, for example. Many of the applications listed make minimal use of email services yet email systems are essential in many organizations. We just have to ask how long we could continue to operate without a functioning email system. Chances are it would be longer than if our online sales system was down, but we would not want to go long without email. This demonstrates the point that critical systems come in many forms.

### Important But Non-Critical Servers

Depending on your organization and its dependence on email, an Exchange Server might be considered critical or important. Important servers, like the category of data classification, indicate a lower priority than critical servers. Some examples of important servers include:

- Email servers
- Database servers hosting non-critical applications data
- Collaboration servers, such as SharePoint servers with non-critical data
- File servers hosting shared directories

At the end of this exercise, we have a breakdown of the types of data and servers by criticality. This allows us to organize the servers based on their importance to essential operations. There is just one more step before we can create a summarized, consolidated report of our recovery management needs that will allow us to define an informed set of backup and disaster recovery procedures.

## Determining RPOs and RTOs

RPO and RTO are a couple of terms that are frequently used when discussing recovery management, backup, and recovery. Let's start with a couple of definitions. RPO is the maximum amount of data that can be lost expressed in the time from a data loss event to the time of the last backup. For example, an RPO of one day means we can accept the loss of one day's worth of data. RTO is the maximum amount of time that data or systems may be unavailable while restoration or disaster recovery procedures occur. If we must have a database backed up and fully restored within one hour of a failure, then we have a one hour RTO.

The last step in putting together the pieces required to formulate a recovery management strategy is to define the RPOs and RTOs for each server or application. Table 2 shows an example of a summarized report with RTOs and RPOs assigned.

| Server Description | Applications | Data Categories | Server Category | RTO | RPO |
|---|---|---|---|---|---|
| Web Server | IIS | N/A | Critical | 1 hour | 1 day |
| Application Server | Financials | Critical | Critical | 1 hour | 1 day |
| | Executive Reporting | Important | | 8 hours | 5 days |
| | Human Resource Mgmt | Important | | 8 hours | 5 days |
| Database 1 | Financials | Critical | Critical | 1 hour | 1 day |
| | Marketing | Important | | 8 hours | 5 days |
| Database 2 | Human Resource Mgmt | Important | Important | 8 hours | 5 days |
| Directory Server | Active Directory | Critical | Critical | 1 hour | 1 day |
| File Server 1 | Windows Server | Important | Important | 8 hours | 5 days |
| File Server 2 | Windows Server | Important | Important | 8 hours | 5 days |
| Collaboration Server | SharePoint Server | Important | Important | 8 hours | 5 days |
| Email Server | Exchange Server | Critical | Critical | 1 hour | 1 day |
| Development Server | Windows Server | N/A | Important | 8 hours | 5 days |

**Table 2: Example summary assessment of data and server classifications and associated RPOs and RTOs.**

In this example scenario, we have limited ourselves to two levels of RPOs and RTOs. In practice, we should craft these objectives as the business requirements demand, but we need to balance those with management considerations. The more variation we have, the more backup policies we will need to define and manage. Remember, complexity is often the enemy of reliability. Keeping backup schedules as simple as possible, but not simpler, can help reduce management overhead and the potential for errors.

Let's use the information we've compiled here to help formulate a disaster recovery policy.

### Creating a Disaster Recovery Policy

RPOs and RTOs define how long we can be without critical and important systems and their data. These objectives guide our decision making when it comes to deciding a range of issues regarding data loss prevention:

- How often should we back up servers?
- What types of backups should we use (full, incremental, or differential)?
- How long should we keep each type of backup?
- How long do we have to restore data from a backup?
- Is a tape-based solution fast enough to meet recovery objectives?
- Is disk storage needed for short-period RPO systems?

This information also guides us in disaster recovery, especially with regard to how long we have to restore services. That, in turn, influences our choice of architectures for implementing disaster recovery. Key questions that arise in disaster recovery include:

- Can systems and data be restored from backups in time to meet RTOs?
- Are backups performed frequently enough to meet RPOs?
- What servers will the applications run on?
- How will the backup media, such as tapes, be transported to the disaster recovery site?
- Will the applications need similarly configured servers in disaster recovery mode or is some level of degraded performance acceptable?

Depending on the answers to these questions, we can formulate the disaster recovery plan. At the highest levels, a disaster recovery plan will document the following:

- The location of a disaster recovery site. This site may be a remote office, a dedicated facility, or an on-demand infrastructure service, such as a cloud provider.
- If backup media will be used to restore systems, the plan should include a procedure for ensuring backups are kept up to date at the disaster recovery site.
- When, if at all, data is replicated from primary servers to standby servers because restoring from backup media will either take too long or will not meet RPOs.
- A schedule showing how applications and data in the primary environment will map to servers in the disaster recovery environment. This schedule is especially important if virtual servers will be used to host multiple applications in disaster recovery mode.
- A plan for synchronizing data from the disaster recovery center back to the primary servers once they are restored.
- A list of decision makers responsible for determining when to switch to the disaster recovery center and switch back to the primary servers.

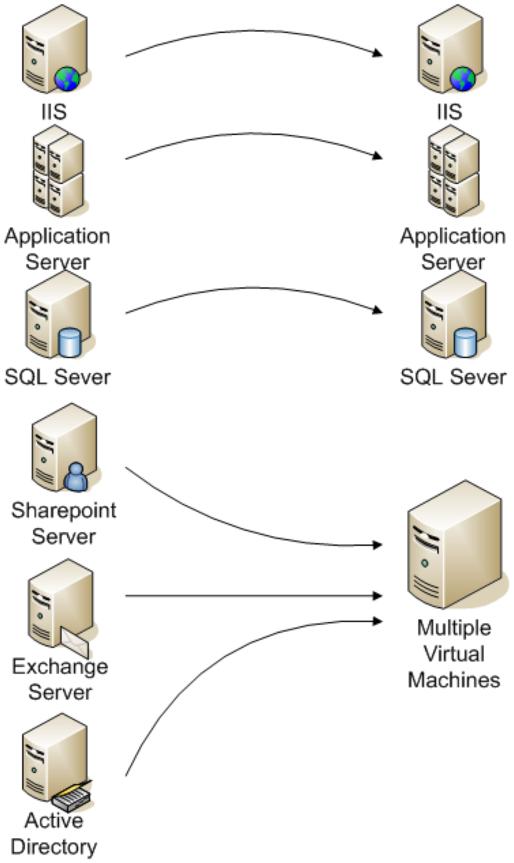

**Figure 38: In disaster recovery mode, multiple virtual servers can be hosted on a single physical host, reducing the cost of maintaining disaster recovery infrastructure.**

This tip, trick, and technique has outlined the basic building blocks of a recovery management strategy; unfortunately, analyzing the parts does not always give a comprehensive picture of the whole. One aspect of recovery management that was not addressed here is security, so we will turn to that next.

## Tip, Trick, Technique 19: Understanding Security Issues with Backups, Archives, and Disaster Recovery

We expend a lot of effort to keep our data secure. We set up access controls, implement authentication mechanisms, and limit privileges to reduce the risk of someone tampering with data or accessing data they should not see. Many of the mechanisms we use are not able to protect data once it moves from the servers that normally house it to backup media. For example, file-based operating system (OS) access controls do not limit access to files in a backup set. The need to protect against data loss has implications that conflict with our need to protect the confidentiality of data.

### Protecting Confidentiality of Backup Data

There are a few key drivers for the need for confidentiality of data. Depending on the type of business or organization, there may be regulations that proscribe levels of privacy protection that should be ensured for personal information. Healthcare and financial services industries are obvious examples where such is the case. Even in industries without well-defined regulations, there are still incentives to preserve the privacy of customer or client data. A well-publicized data breach can damage a business' image, lead to the loss of customers, and eventually impact the bottom line. Heartland Payment Systems and the TJX Companies, Inc. received quite a bit of press about their record-breaking data breaches in 2009 and 2007, respectively. Intellectual property is a particularly important target in some industries in which high research and development costs provide incentive to steal rather than develop intellectual property.

**Building a Business Case for Security**

The Open Security Foundation maintains the Data Loss Database at http://datalossdb.org/. The site has a wealth of information about data loss incidents that may be useful if you need statistics to justify a business case for the need for information security. For threats to intellectual property, see Kim Zetter's "Report Details Hacks Targeting Google, Others," *Wired,* February 3, 2010 (http://www.wired.com/threatlevel/tag/apt/), and the *Christian Science Monitor's* "US Oil Industry Hit by Cyberattacks: Was China Involved," January 25, 2010 (http://www.csmonitor.com/USA/2010/0125/US-oil-industry-hit-by-cyberattacks-Was-China-involved).

Security breaches come in many forms, including the loss or theft of backup media. Confidential and private data should be encrypted when it is backed up and the backup media leaves the control of the organization. With encryption, even if the media is lost or stolen, there is little chance the data will be compromised.

One other point to keep in mind about encryption is that the definition of strong encryption changes over time. Use strong encryption algorithms and long encryption keys to maximize the protection provided by encryption.

## Download Additional eBooks from Realtime Nexus!