Realtime
publishers
"Leading the Conversation"

Tips and Tricks Guide™ *To*

Creating Business Continuity through Enterprise Storage Solutions

sponsored by

ca™

Chad Marshall

Introduction to Realtimepublishers

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

Note to Reader: This book presents tips and tricks for seven topics related to business continuity created through enterprise storage solutions. For ease of use, the questions and their solutions are divided into topics, and each question is numbered based on the topic, including

- Topic 1: Securing Availability and Business Continuity
- Topic 2: Maximizing Storage Resources and Solutions
- Topic 3: Information Management
- Topic 4: Cost Management
- Topic 5: Compliance
- Topic 6: Security
- Topic 7: Aligning Storage to Serve Business

Copyright Statement

Topic 1: Securing Availability and Business Continuity

Q1.1: What is the difference between business continuity and disaster recovery planning?

A: Business continuity and disaster recovery are closely related concepts that often exist as a point of contention between information technology (IT) and line of business managers. Unlike disaster recovery, which focuses almost entirely on IT infrastructure and assets post-disaster, business continuity represents the processes and procedures an organization puts in place to ensure that essential functions can continue during and after a degradation or complete loss of critical people, processes, or technology.

Business continuity planning and disaster recovery planning both address the preservation of business and involve the preparation, testing, and maintenance of plans to protect vital business processes and assets. Business continuity plans, however, are created to prevent interruptions to normal business activity and are designed to protect business *processes* from disaster. Further, business continuity planning deals with aspects of business process not generally covered within a disaster recovery plan (such as logistics). Disaster recovery planning deals almost entirely with plans to reduce the severity of an impact once the disaster has occurred.

Business Continuity Planning

Business continuity planning refers to any methodology used by an organization to create a plan for how the organization will recover from an interruption or complete disruption of normal operations. The International Organization for Standardization (ISO) and the British Standards Institute set business continuity planning best practices under “ISO/IEC 17799:2000 Code of Practice for Information Security Management” and “BS 7799 Information Security,” respectively.

The development of a business continuity plan can be divided into five major areas commonly referred to as the business continuity planning life cycle. Figure 1.1 illustrates these five areas.

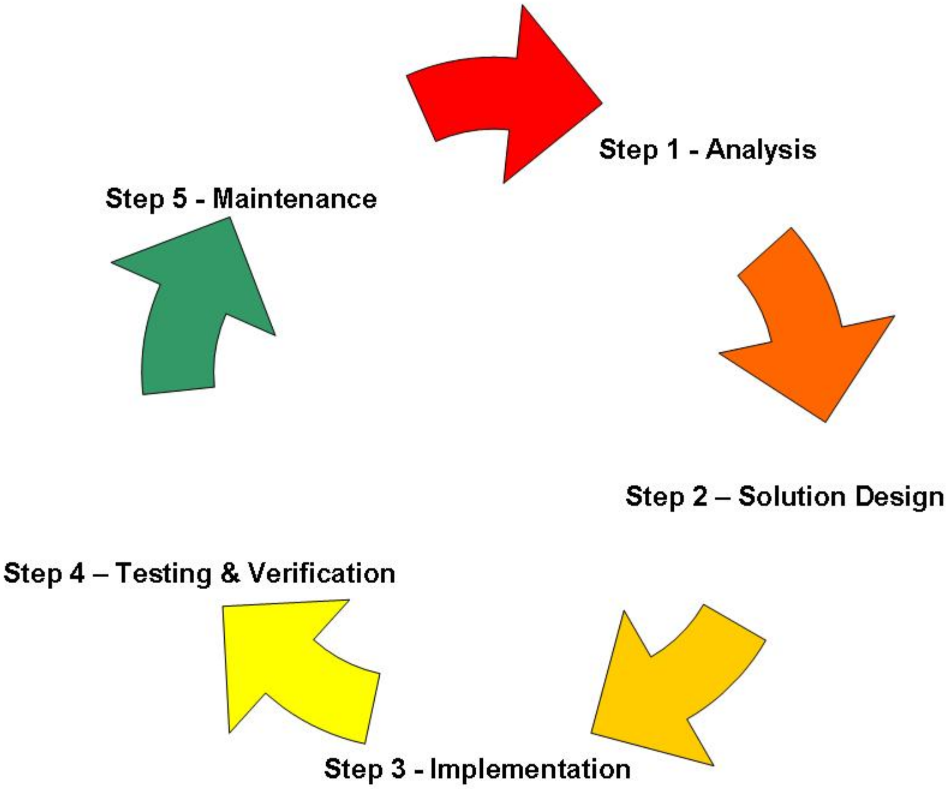

Figure 1.1: The five phases of business continuity planning.

Analysis

The analysis phase of the business continuity planning life cycle consists of four primary activities:

- Threat analysis—During threat analysis, potential threats to business operations are identified. Natural disasters, loss of utilities such as power and water, and the threat of cyber attack (hacking) or terrorism are just a few threats that may impact normal business operations.
- Impact analysis—Also known as a business impact assessment (BIA), impact analysis has a goal of identifying the business processes that will result in a significant detrimental impact to the organization if lost or degraded. How well a disruption is tolerated is often significantly influenced by the cost of establishing and maintaining appropriate technical or operational recovery solutions. Other factors influencing how disruptions are tolerated include items related to public safety and those or considered critical by law, such as regulatory requirements.

- Definition of impact scenarios—Once threat and impact analyses are complete, a series of scenarios can be developed using what you have learned. This task will further enable the understanding of business continuity requirements. For example, if your business provides financial transaction processing for other companies, imagine what would happen if the primary transactions systems failed and no transactions could be processed for 4 hours. It's easy enough to determine how many transactions are processed during an average 4-hour period; next, determine what would happen if those transactions couldn't be exercised. How many customers might go unsatisfied? What would be the potential losses? Don't forget to include potential legal and regulatory compliance concerns.
- Documentation of recovery requirements—Once the analysis phase has been completed, the business and technical requirements are documented to aide in the development of a solution.

Solution Design

During the solution design phase, you begin to take all that you have learned during the analysis phase and start to draw conclusions that lead to logical solutions. For example, if during the analysis phase you determined that a threat of a Denial of Service (DoS) attack exists, you might now take steps to design a solution to protect your storage resources against a DoS attack—for example, segregating network resources or including a rate-base intrusion prevention system (IPS) to monitor and identify abnormal rates for certain types of traffic and stop unusual or suspect activity from consuming resources.

The solution design phase needs to meet three main requirements in order to be successful:

- The minimum business process requirements are being met, including general requirements, in the form of personnel, processes, and time. This may include:
  - Identification of critical processes
  - Identification of critical personnel (or job functions)
  - Definition of the crisis management command structure
  - Identification of a contingency operations site (where necessary)
- The minimum technical requirements are being met, such as minimum software, hardware, facilities, and infrastructure requirements. This may include:
  - Identification of critical software
  - Identification of critical hardware
  - Identification of critical facilities
  - Identification of critical infrastructure
  - Definition of infrastructure relationships between primary and secondary sites
  - Definition of a data replication methodology between primary and secondary sites
- The timeframe in which the minimum business operations must be available is achievable. This may include:
  - Identification of stakeholder expectations and/or service level agreements (SLAs)
  - Time to implement process
  - Time to deploy secondary technology solutions

Implementation

Implementation is the carrying out or physical realization of something from concept to design. For example, a computer system implementation would be the installation of new hardware and system software. In the context of business continuity planning, the implementation phase merely consists of the execution of the design elements identified in the solution design phase. This might include the implementation, in the form of delivery and installation, of technology components or the official communication of personnel assignments to cover critical job functions during a crisis.

Testing and Organizational Acceptance

The purpose of testing is to gain assurance and organizational acceptance that the solution designed will satisfy all the organization's requirements. Many things can hinder an otherwise well-designed solution from achieving continuity of operations such as:

- Failure to capture critical components—This shortcoming might lead to insufficient, misunderstood, or inaccurate recovery requirements resulting in a solution that doesn't fully meet the needs of operations.
- Solution design flaws—These may include errors in capacity planning and systems design.
- Solution implementation errors—Such as misconfigured network equipment or software applications.

Testing can be broken into three major categories:

- Organizational testing—Organizational testing is testing that is explicitly designed to ensure that when an interruption to business continuity is presented, personnel respond accordingly.
- Technology testing—Technology testing should be designed to test all technological aspects linked to the business continuity plan. All hardware, software, applications, and IT management routines that will be depended upon to maintain business continuity should be tested.
- Process testing—Process testing should include tests to verify that business processes continue to function properly and encompasses portions of both organizational testing and technology testing. Specifically, these tests should ensure that processes continue and minimum business operations remain available.

Throughout the testing process, the goal remains to validate the solution and gain organizational acceptance. One test type that is particularly useful in this regard is a User Acceptance Test. A UAT is conducted from the point of view of the end user typically by end users or subject matter experts (SMEs) that can validate test results at the end-user level and accept changes with the authority of the line of business. Once the testing phase is complete and stakeholders have accepted the plan, the business continuity planning process can be relaxed and revisited on regular maintenance intervals to retest the solution and again validate the results.

Maintenance

Ongoing maintenance of the business continuity plan is necessary to ensure the plan remains viable and is typically conducted bi-annually or annually depending upon your organization's rate of change. The purpose of the maintenance phase is to keep the business continuity plan up to date and is generally broken down into four activities:

- Information update and testing—Includes testing accuracy of information contained within the business continuity plan.
- Testing and verification of technology—Includes tests of hardware, software, and applications for proper functionality.
- Testing and verification of recovery procedures—Includes tests of organizational recovery procedures to ensure documented recovery processes function as designed.
- Treatment of test failures—Actions taken to correct failures.

The maintenance phase links back to analysis phase. A test failure during maintenance is a sign that the requirements defined during impact analysis might no longer be valid. Once this point is reached, a new impact analysis should be conducted and an appropriate solution aligned to meet the needs of the organization.

Disaster Recovery Planning

Disaster recovery planning is specifically focused on creating a comprehensive plan of actions to be taken before, during, and after a significant loss of information systems resources. Unlike business continuity planning, disaster recovery planning assumes the worst has already occurred and major impacts, such as the loss of an entire data center, are already being felt. The disaster recovery plan will outline steps to take to recover as gracefully as possible.

During the business continuity planning process, detailed analysis was conducted of both threats and impacts. If a threat of a natural disaster, such as a hurricane or earthquake, had a potential to impact a data center, the solution aligned to mitigate that impact would align with a disaster recovery plan. There are essentially two steps in the disaster planning process they are data continuity planning and maintenance.

Data Continuity Planning

Organizations rely heavily on their ability to process data. Whether the focus of your data is simply file, print, and email services, or if your data center houses databases accessed by thousands of users, getting critical resources back up and running after a disaster is going to be a top priority. There are several options to be considered as alternatives should your data center go offline:

- Secondary (multiple) sites—A clear alternative to suffering a complete loss of data resources from the loss of a single data center is to have multiple data centers in different geographic regions. Generally, the farther apart the data centers are, the more isolated they will be from natural disasters.
- Mutual aide agreements—This option is an arrangement with another organization of similar size and computing resource needs to assist your organization in the event of a disaster. Small and midsized companies can often benefit from mutual aide agreements the most because their requirements are typically smaller and more flexible.
- Subscription services—Third-party commercial providers specializing in disaster recovery often offer subscription disaster recovery services ranging from "hot site" or full recovery site functionality to "cold site" or minimum recovery site functionality.

Disaster Recovery Plan Maintenance

One big similarity between business continuity and disaster recovery planning is how quickly the plans become obsolete. Changes in core technology, such as server platform, are major indicators that a disaster recovery plan needs to be revisited, but many minor changes can quickly add up as well that will put the disaster recovery plan out of alignment with organizational needs. Maintaining the disaster recovery plan should follow a similar bi-annual or annual schedule as that of the business continuity plan maintenance but on a generally larger scale.

Q1.2: What is storage resiliency and how can it contribute to continuous availability?

A: There are many definitions circulating about what exactly it means to be "resilient," and industry experts and non-experts alike continually tout the term "resiliency" in the most obscure (and often inappropriate) places as a synonym for "reliability." Resiliency is not reliability. It does however contribute to reliability and thus to continuous availability.

Defining Resiliency

Resiliency is a noun defined in the enterprise storage context as "an ability to recover readily from adversity," the verb form of which is to "resile." which means to "spring back, rebound or return to an original state." In business continuity, resilience is the ability of an organization, resource, or structure to sustain the impact of a business interruption, resume its normal operations, and continue to provide minimum services.

To managers of an enterprise storage infrastructure, resiliency should result from taking steps to design a reliable, scalable environment. In addition, managers should put in place plans that enable an environment with the capability to scale to meet business needs without adding complexity. In terms of storage technology and storage resource management, resiliency is derived through hardware and software features that increase reliability and scalability; features such as automated monitoring and alerting. Storage resource management (SRM) software that monitors storage events and takes a predefined action in response to a particular kind of event would contribute significantly to resiliency by automating the management process. In practice, resiliency equates to knowing the storage management boundaries of an organization, how far they can bend, and at what point the processes, people, or infrastructure will begin to unravel and, most importantly, how to shore them up before they do.

Resiliency in Practice, Processes, and People

How rapidly can you scale your storage infrastructure an extra 2TB? What about 20TB? If you find that you cannot answer this question with a succinct process, you might be in trouble. When called upon to flex your storage muscle, processes need to be defined well enough so that those who manage the storage infrastructure have clear guidance on when, how, and under what circumstances they can manage storage to scale to meet the need and when further approvals may be required.

. When unexpected projects or business requirements push storage resources to the limit, you might need to purchase storage in a hurry. Procurement processes should be well defined for a standard as well as an expedited approval. They should be stringent enough to prevent abuse and flexible enough to ensure extra storage can be procured when necessary. One clear and excessively used de-motivator is to make expedited storage more costly to the line of business requesting the expedited service. Although cost can be effective as a de-motivational tool, it is important to not allow it to become overused.

> 🕮 One way to combat unexpected business requirements is through sound Capacity & Performance Management (C&PM). The next volume will address C&PM in-depth, answering the question: What steps can be taken to ensure a successful analysis of existing storage infrastructure and plan for growth?

Resiliency in Technology

The technologies that contribute directly or indirectly to storage resiliency are too numerous to mention by name, but they can be classified into the following areas based upon where in the storage infrastructure they reside.

Media-Level Features

Features inherit to or delivered through low-level media such as a checksum value calculated and verified at the storage level fit into this category. A checksum is a form of redundancy check, a very simple measure for protecting the integrity of data by detecting errors in data. Technologies such as Redundant Array of Inexpensive Disks (RAID) can contribute to resiliency but not all RAID levels offer the same benefits. The three most common forms of RAID are

- RAID 0: Striped Set—Splits data evenly across two or more disks with no parity information for redundancy. RAID 0 is not redundant and does not contribute to storage resiliency.
- RAID 1: Mirrored Set—Creates an exact copy of a set of data on two or more disks. This is useful for increasing performance and reliability.
- RAID 5: Striped Set with Parity—Uses block-level striping with parity data distributed across all member disks. RAID 5 has achieved popularity due to its low cost of redundancy.

Architectural Features

Storage systems can be designed to be resilient by ensuring that requirements for reliability, scalability, availability, and serviceability are being met. However, not all storage systems are created equal. Network Attached Storage (NAS) devices, for example, are often presented as a rapidly deployed solution to meet immediate need, but what they offer in decreased time to market they lack in resiliency. NAS devices, which have been historically standalone, proprietary solutions, often present management challenges in large enterprises because each unit typically needs to be managed as an independent entity. On the positive side, however, removing a dedicated server from the storage equation makes NAS more reliable than a traditional file server by simplifying the storage infrastructure. If you're operating in a small organization, scalability needs may be met quickly by installing another NAS device.

On the contrary, Storage Area Network (SAN) solutions typically result in an advanced capability of management but are usually more expensive to deploy. The difference in scalability and ultimately resiliency is that although NAS devices may need to be brought online and configured one by one, a SAN solution can be scaled through management and—with effective management software—can even be scaled to react to specific storage conditions or demands. Depending upon the size of your organization, or more directly, the size of your need, choosing the right base storage solution is critical.

Facilities Features

The resiliency of a facility is its ability to recover from adversity presented in environmental form. Power and environmental controls (heating/cooling) are factors to be analyzed as threats to business continuity. In terms of resiliency, you should also focus on ability to scale.

Management Features

It has been said that the best way to improve a process is to remove its dependency upon people. Intelligent management features, such as triggered responses to monitored events, contribute heavily to resiliency by simplifying management and automating routine management tasks within the storage infrastructure.

Q1.3: How can I create an effective business continuity plan?

A: Question 1.1 detailed all the essential elements of a successful business continuity plan. Following the business continuity plan life cycle—which covers analysis, solution design, implementation, testing and verification, and ongoing maintenance—as a framework is critical to creating an effective business continuity plan. However, the framework cannot execute itself. For your business continuity plan to be effective, the following best practices are recommended:

Best Practice #1—Gain Senior Management Buy-In

For your business continuity plan to be effective, it must gain the full support of senior management early and maintain that support. A few tips to get senior management buy-in include:

- Demonstrate the value of business continuity planning through use of industry examples.
- Develop a business case for business continuity planning that illustrates the effect an impact may have to a single line of business and present the findings to senior management.
- If your organization has suffered a recent impact, gather facts and figures from that impact, along with line of business statements, to provide a real “voice of the customer” view.
- Don’t take no for an answer.

Although it may seem odd that anyone would say “no” at a senior level, failure to proclaim support for business continuity planning within the organization has essentially the same effect. Thus, buy-in is so important. Make sure your leaders are out in front proclaiming the value of business continuity planning and are visible during testing and recovery exercises.

Best Practice #2—Ensure the Plan Covers the People

When facing a business-impacting incident or disaster, the most critical asset any organization has to recover from such an event is its people. Understand those threats that impact not only your business but also your associates and put plans in place to provide for their needs as well. Ensuring associates have adequate healthcare and insurance to protect their families' interests and their property is one step, although you may also consider providing for on-site healthcare or mental health programs to give associates peace of mind and avenues to pursue for stress relief. A well-insured, happy and healthy employee is the strongest business continuity asset any organization can hope to obtain.

Best Practice #3—Don't Underestimate the Adversary

Fire, flooding, hacking, viruses, and power failure are all threats to business continuity that must never be underestimated. One of the biggest mistakes that can be made during the planning process that will significantly impede the effectiveness of business continuity planning is to underestimate a threat. When considering a threat, and its impact, consult experts in the field. When considering fire as an impact, for example, local experts may be available for free or very little cost in the form of a local Fire Marshal or Fire Inspector.

Best Practice #4—Research and Collaborate

Yours is not the first organization to face the challenge of business continuity planning. Developing an effective plan includes some degree of research to study what is currently working for other organizations. Study the latest trends, technologies, and industry research on business continuity planning. Be sure to collaborate with outside vendors, consultants, and subject matter experts to ensure that your plan will result in the most effective outcome possible.

Best Practice #5—Exercise the Plan

A business continuity plan is a living entity that is continuously undergoing updates, modifications, and redesigns to suit the ever-changing state of the business, technology, and threats facing business today. To ensure that your business continuity plan is effective, development of it must never cease and exercising of the plan must be conducted on a regular basis to ensure the plan works as designed.

Q1.4: What are the top-five threats to storage infrastructure and what steps can I take to minimize their impact?

A: Threats can come in all shapes and sizes, from large natural disasters affecting entire cities to structural fires that impact a single location. In 2005, EnvoyWorldWide conducted its second annual survey "Trends in Business Continuity and Risk Management," which was conducted blindly among members of several business continuity organizations. The survey was designed to leverage a regionally diverse group of business continuity professionals to identify business continuity and disaster recovery practices and trends.

The following list highlights the top-five events that may pose a threat to business continuity and disaster recovery as they were rated in order of threat level by 140 respondents:

- Data security
- Data center hardware/software failure
- Telecom failure
- Structural fire
- Power outage

Threat #1—Data Security

Data security is a generic term designating methods used to protect data from unauthorized access. This means doing everything possible to ensure that an information system remains secure, which encompasses not only the protection of information from criminals but also from equipment malfunction and natural disasters. Data security threats also include unauthorized access to data and damage to files by malicious programs such as viruses. Part of the reason why this is number one on the list is likely due to the fact that its nature is generic and ensuring data security is a continuous end-to-end concern that makes data security an enterprise-wide concern.

Ensuring data security begins with ensuring that data is properly classified so that adequate security measures can be aligned to meet the needs of the data. Ensuring data classification is a part of information life cycle management that will be covered in great detail in the next volume of this guide; for now, understand that it is important to fully capture the business need for the data, the value of the data to internal and external resources (such as internal auditors or external regulators), and finally the classification of the data itself by data or storage architects. Data must be handled with care to ensure that its confidentiality, integrity, and availability are continually maintained as mandated by the data's classification and retention schedule. Once data is classified, the next step is to ensure that for each classification, an adequate data path exists that begins with sufficiently secured storage.

Securing data in the storage space again involves not only the confidentiality of the data but also its integrity and availability. Steps must be taken to ensure that data is not altered, disclosed, or denied which, in storage, includes steps to regularly audit data to see who is accessing it, and how, as well as ensuring that steps are taken in business continuity to ensure the data is available when needed.

Steps must also be taken to ensure the data is protected in transit, which may include the use of firewalls, intrusion detection and intrusion prevention systems, and encryption technologies. Data, when in transit, is subject to interception and alteration through various forms of information, or cyber, attacks.

Once you're assured that data is secure both in storage and in transit, the final state is to ensure data is secured when outside of the system or when the system is being manipulated by authorized users. Education is absolutely critical to data security. Users of data must be educated on data classification, data use, their responsibilities for data protection and retention, and how to react when data is mishandled. Further, all employees should be subject to mandatory information protection training that covers data protection in depth, including social engineering prevention techniques.

Threat #2 Data Center Hardware/Software Failure

Hardware and software failure is not a matter of "if" but simply a matter of "when." Steps must be taken to ensure data center hardware and software are resilient enough to handle the challenges placed upon it by operations and by internal and external threats to stability.

As businesses continually adapt to realize the potential of infrastructure consolidation, more and more server resources are being consumed by differing processes. The result is that a single rack of servers within a data center may contain literally dozens, if not hundreds, of applications. To reduce the threat of hardware and software failures, your organization should focus on redundancy and ensuring that the appropriate level of monitoring and evaluation are in place to ensure a timely response to hardware or software failure.

Redundancy

For enterprise infrastructure components to be as resilient and successful as possible in the face of adversity, redundancy must be deployed to protect your vital assets. Dual processors, dual memory modules, redundant storage, redundant network connections, and redundant power supplies are a good start, but they're not the end. Care must be given to ensure that up-level and down-level relationships are redundant throughout the enterprise infrastructure so that no one, single, point of failure can cause a complete failure. Allowing for a single point of failure within an infrastructure is virtually the same as having none at all.

Monitor and Evaluate

When hardware or software failure occurs, time is of the essence. Effective IT management requires a monitoring process to ensure that the appropriate IT team is promptly informed of systems outages and can rapidly respond to incidents. Start your monitoring by defining relevant performance indicators, then establish a systematic report process.

Threat #3—Telecom Failure

The best way to remediate the threat of telecom failure is to ensure redundancy through primary and secondary means. Through primary means is to have a redundant way to conduct operations through the primary circuit type. For example, if your facility requires a single T1 network line, you may consider going with two so that one is always on standby, but be certain not to procure the secondary line from the same provider. Redundancy is concerned with eliminating all points of failure, so signing up for two circuits from the same provider is going to do little good if the provider, itself, experiences problems that impact your circuits.

Take steps to procure redundant circuits from separate providers and be sure to research the routes. Oftentimes telecom providers will sell and resell each other's products and service offerings or rely upon the same third-party (or in some cases fourth, fifth, sixth, and seventh party) vendors to provide up-level or down-level services. The ideal state is for the redundant circuits to be completely redundant from start to finish with no chance of a single point of failure.

Redundancy through secondary means is to have a second, ancillary form of communication. Satellite communication, although expensive for day-to-day use, can serve quite effectively as a backup to regular communication during long periods of time when normal connectivity is not able to be restored due to lack of power or some other mitigating factor. It also has the benefit of being wireless, which means that so long as the site has power they can achieve connectivity. Other options include cable service providers (cable modems) and DSL lines. Although neither are likely to provide the same amount of bandwidth a site is accustomed, most will agree that some connectivity is better than no connectivity at all.

Threat #4—Structural Fire

Structural fires can occur in nearly any environment at any time and can cause a tremendous amount of damage. Throughout history, fire has delivered tremendous blows to data, from courthouse fires that wipe out vital birth, marriage, and death records to warehouse fires that consume countless financial documents beings stored for compliance purposes. Fire is likely to be the largest adversary your organization faces as a threat. Why? Mostly because many people fail to realize how many fires actually occur each year, how devastating they can become, and how long it takes for help to arrive.

According to the US Fire Association in a single year there were more than 52,000 confined structural fires in the United States and local fire response time (from the time the call is actually received until a first responder is on-site) is less than 5 minutes 50 percent of the time. In the life of a fire, 5 minutes is enough time to cause a great deal of damage and although you can take steps to minimize the impact a fire can have, so long as you have a mixture of air, fuel, heat, and people, a potential for fire must always be accounted for. The following list highlights structural fire preparation best practices:

- Ensure employee awareness and participation—Fires can start quickly and spread rapidly. From the time of the first detection, there can be no confusion—everyone must know exactly what to do, and how to do it. Ensure that stairs and exits are properly marked. Identify a meeting place for employees and ensure that everyone knows where it is. Hold regular training and fire drills, and practice fire awareness regularly.
- Partner with your local fire/police representatives—Fire departments typically offer consulting at no (or very little) charge; after all, the training they provide your organization may prevent them from one day needing to risk their lives combating a structural fire. Consider their involvement mandatory in all fire planning, policy development, and training exercises and celebrate their involvement.
- Keep things clean—Dust and rubbish are major contributors to fires both in the residential and commercial space. Taking steps to ensure that work areas are kept clean is the first step in fire prevention.
- Ensure adequate fire protection and alerting systems—Data centers should be protected with an up-to-date and regularly tested fire protection and alerting system. Consult with local fire experts and security contractors to assess the needs of your data center and ensure they're properly met.

Threat #5—Power Outage

Combating a power outage usually involves ensuring that adequate onsite power generation abilities exist and that uninterruptible power supply (UPS) systems are in place to handle the power load during the switch from "street" power to internal generators. Beyond this point, ensuring power remains available is largely a matter of logistics. If the power outage lasts for days, weeks, or months, supplies of fuel will need to be regularly delivered and, if the power outage affects employee homes as well, accommodations for employees must be made a top priority. If your organization conducts business in an area with a high flooding threat, ensuring that generators and major electrical panels are located in a dry space (not in the basement) is important.

Q1.5: What steps can I take to ensure continuous availability within a growing storage infrastructure?

A: Ensuring continuous availability in an environment of growth requires developing and deploying an IT infrastructure tuned to provide high availability and well prepared for a business continuity or disaster recovery event.

High Availability

Over the past few decades, great leaps and bounds have been made in hardware, software, and storage technology, but while all have improved, none is currently available off-the-shelf that can meet the promise of continuous 24/7 availability. Servers still suffer hardware failures, software still requires regular maintenance, and storage can still become corrupted. So what can be done? There are a few best practices that can be followed to help ensure high availability.

Best Practice 1—Operate on the Most Stable, Hardened Platform Possible

Beginning with a solid foundation is the first step in ensuring high availability. The concept of system “hardening” is one that is characterized by identifying all unnecessary or high-risk features within a platform (either as an operating system—OS—platform, hardware platform, or software application platform), and eliminating those that are not required. Some OSs lend themselves to such manipulation easier than others, but the end result should be focused on achieving the most stable production platform possible.

Best Practice 2—Control Change

Once a stable production platform is in place, you need to keep it that way. All changes must be rigidly controlled to ensure that no potentially damaging change is deployed to a production system. In an environment of rapid growth, no change should be allowed to be overlooked. Changes should be reviewed, tested, and authoritatively approved for production prior to being deployed.

Best Practice 3—Deploy High-Availability Technology

Certain technologies are simply more resilient and capable of handling adversity. For example, deploying a Redundant Array of Independent Disk (RAID) 5 solution has been a big step forward from a deployment using a single disk solution, but even RAID in and of itself cannot protect the data on the array. A RAID array has one file system. This creates a single point of failure and the array's file system is vulnerable to a wide variety of hazards other than physical disk failure, such as a virus and user error. Research technologies that are high in redundancy and resiliency and learn what benefits they can bring to your infrastructure.

Best Practice 4—Control User Access

A large, yet often overlooked, threat to high-availability is user error. Users who are inadequately trained and/or possess overprivileged accounts can cause a great deal of damage. Your organization should regularly audit user privileges to ensure not only that users are not being assigned more access than they need but also that as users change position within the organization, their access is re-evaluated accordingly. As an outcome of auditing, specific attention should be given to examine a user's need to perform manual tasks. Automating tasks will remove the user from the equation and thus prevent user error from ever becoming a factor.

Best Practice #5—Monitor and Evaluate

Ensuring high availability requires a good deal of proactive monitoring and evaluation to identify and work to eliminate potential problems before they develop into production-altering events. Ensure that appropriate monitors are in place and learn from past events. For example, if through a previous incident, the root cause was discovered to be a run-away process on a server that brought down production, ensure that steps are being taken to monitor that process (in addition to getting the vendor on the line to fix it so it doesn't happen again).

Business Continuity and Disaster Recovery Planning

Both business continuity planning and disaster recovery planning are covered in detail in Question 1.1. In specific relation to an organization experiencing growth, it is important to re-evaluate business continuity and disaster recovery plans on a regular basis. As the storage infrastructure grows, so will the need of business continuity and disaster recovery plans to meet those needs and ensure critical services and infrastructure are available when needed.

Q1.6: What is recovery management and how does it differ from disaster recovery?

A: Disaster recovery planning is specifically focused on creating a comprehensive plan of actions to be taken before, during, and after a significant loss of information systems resources. Recovery management is specifically focused on the carrying out of those actions to ensure an expeditious and successful recovery. Recovery management begins when a disaster is declared and is usually handled by a Computer Emergency Response Team. A CERT should be established within every organization with computer or data assets that may need to respond to computer and data-related emergencies and engage business continuity or disaster recovery plans. Recovery management begins with a CERT or other authorizing body within an organization declaring a need for disaster recovery operations. Once declared, the process of recovery management begins. Recovery management responsibilities may include:

- Ensuring all line of business, technology, and other concerned parties are made aware that a recovery operation is in process
- Coordinating and communicating between primary and secondary (recovery) sites
- Intervening during disaster recovery plan failures
- Coordinating with government agencies (as required)
- Reporting to senior management on the status of operations and the recovery effort

Recovery management teams should be comprised of leaders and subject matter experts that are empowered to make judgments on the best course of action should a disaster recovery plan fall short of actual real-world demands.

Q1.7: What needs to be in place to ensure continuous availability?

A: Ensuring continuous availability encompasses all the areas covered in detail in Question 1.5, such as ensuring that high availability and adequate business continuity measures are in place. The following list highlights additional requirements for ensuring continuous availability.

The Best People

Ensuring continuous availability doesn't happen automatically, and despite having all the same ingredients and all the same equipment and utensils, there is still a great deal to be said about the cook. Ensuring continuous availability requires the very best personnel that are installing, upgrading, monitoring, evaluating, and managing the environment.

The Best Platform

The best platform does not necessarily mean the best hardware or the very best software. Ask any high-end user, and they will tell you that matching the "very best" memory, motherboard, CPU, and video card may not always result in the "very best" platform. A platform takes into account all the major pieces such as CPU, motherboard, and supporting technologies. To this end, hardware developers are paying particular attention these days to how their hardware works within a platform.

Intel, for example, has introduced a Centrino platform for mobile laptop computers that marries a specific motherboard, processor, and wireless network adapter that work together for the best result. Intel makes this same parallel in its server lines. Within your own organization, you, as well, may add to these platforms with your own standards of architecture being certain to take into account high-availability technologies and business continuity planning as you develop your own standards.

The Best Team (of Processes)

To provide the information that an organization needs to achieve its objectives, IT resources need to be managed by a set of naturally grouped processes. People need to speak the same language and understand how to measure IT processes in a standard way. This is accomplished through strong IT governance. Many organizations already have some sort of management controls in place, either through the Capability Maturity Model (CMM), Six Sigma, IT Infrastructure Library (ITIL), or Project Management Professionals (PMP); although these are all good individual players, they don't provide a one-size-fits-all solution for any organization.

Six Sigma, which was derived from a manufacturing process improvement effort, is a great way to improve processes. ITIL provides a good way to manage the delivery and support of infrastructure services. The Project Management Institute (PMI) has a specific certification for PMP that is without question the most comprehensive work on the subject in the industry. All these are tools can work together and complement each other within an organization, much like players on a team. Through a Six Sigma project, for example, a project may develop that needs to be "passed" over to a PMP who can drive it to fruition. Each player has its part. A player that is becoming more common to the field is COBIT.

Control Objectives for Information and related Technology (COBIT) is a framework of best practices for information management. This framework was created by the Information Systems Audit and Control Association (ISACA) and the IT Governance Institute (ITGI). Currently in version 4, the COBIT framework is generally accepted as one of the most comprehensive works for IT governance, organization, and process and risk management. COBIT provides good practices for the management of IT processes in a manageable and logical structure. COBIT strives to meet the multiple needs of enterprise management by bridging the gaps between business risks, technical issues, control needs, and performance measurement requirements.

Leveraging the best team of processes from IT governance accomplished through COBIT to process improvement performed through Six Sigma's approach will provide a common framework of understanding for IT managers to follow and help to ensure a common taxonomy is adopted enterprise wide. This has the effect of reducing management complexity, which directly contributes to ensuring continuous availability.

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.