A Roadmap for Reducing the Impact of Downtime Events

Every organization must know with reasonable precision how it would resume business-critical operations after any type of downtime-causing event. Natural disasters get a lot of press, and certainly every company should be prepared for these, but the more common threats come from technical failures (equipment and software failures, communication system failures, power outages, etc.) or malicious activities (hacking, ransomware, fraud, and terrorism). Regardless of the cause, any prudent organization must identify all potential threats, make efforts to reduce the possibility of these threats impacting business operations, and most of all, be fully prepared to restore business-critical operations within a predefined time window should an adverse event occur. This is what continuity planning is all about.

The emphasis is on “planning” because if you don’t have a comprehensive, well-tested, continually maintained plan that includes buy-in from your management, your chances of making a successful recovery are slim.

The Seven Core Requirements of Continuity Planning

1. **Assemble the Continuity Team**—This multi-discipline team is responsible for all aspects of continuity and recovery planning.
2. **Perform a Risk Assessment**—All continuity risks to the organization are assessed along with all current resources to mitigate these risks.
3. **Conduct a Business Impact Analysis**—The potential financial and other impacts from known risks are quantified. Recovery time and recovery point objectives are defined.
4. **Implement Mitigation Strategies**—Processes and investments are evaluated and implemented that will reduce risks and lessen impact from downtime events.
5. **Create the Continuity Plan**—A written plan is created that specifies in detail the processes for executing a recovery after various types of downtime events.
6. **Test the Continuity Plan**—This recurring process identifies weaknesses in the recovery plan by testing it in a variety of scenarios.
7. **Maintain the Continuity Plan**—An ongoing process that tracks all changes to business processes and technologies that could impact the continuity plan in order for the plan to be updated accordingly.

Naming Conventions

The terms “disaster recovery planning,” “IT continuity planning,” and “business continuity planning” are often used interchangeably, but technically disaster recovery planning (and of course IT continuity planning) refer to the ability to restore access to critical IT data and systems after a hardware failure or site disaster. Business continuity planning on the other hand, goes beyond data and IT systems to establish a detailed plan for implementing and testing redundancies for facilities, utilities (power systems, heating/cooling, bandwidth, telephone), critical equipment (beyond IT systems), critical documentation/records, and key personnel. In addition, business continuity planning includes creating protocols for communicating with clients, employees, regulators, investors, business partners, third-party providers, etc. In other words, the objective of business continuity planning is to make sure every essential business operation can recover quickly after a disaster. For the purposes of this paper, the focus will be on planning for the continuity and recovery of IT systems, and to keep things simple, we’ll refer to this process throughout as “continuity planning.”

ITIL and IT Services Continuity Management (ITSCM)

A vendor-agnostic set of standards for the management and delivery of IT services within organizations has been developed in what is known as the IT Infrastructure Library (ITIL). Shops that follow ITIL know that disaster recovery planning is a critical discipline within the ITIL Service Design framework and is referred to as IT Services Continuity Management, or ITSCM. The goal of ITSCM is to ensure designated IT resources can be recovered within tolerances for downtime and data loss as defined by the business in order to meet predefined IT service-level agreements. In addition, ITSCM endeavors to be in alignment with each organization’s greater business continuity planning process.

There are many benefits to taking the disciplined approach set forth by ITSCM, and it is certainly worth looking into. Nonetheless, the core requirements of any continuity planning process are essentially the same; the differences simply relate to the scope and rigorousness of the plan. These core requirements are:

1. Assemble the continuity planning team
2. Perform a risk assessment
3. Conduct a business impact analysis
4. Define and implement mitigation strategies
5. Create the continuity plan
6. Test the continuity plan
7. Maintain the continuity plan

The Continuity Planning Lifecycle

During the remainder of this paper, we'll examine each of the core requirements of continuity planning, but before doing so it's important to emphasize that no individual continuity-planning requirement exists on its own, but rather is part of an ever-improving process that has dependencies on many of the other core requirements. That's why effective continuity planning is never considered as a project that gets worked on for awhile and then completed; it needs to be treated as a lifecycle, not unlike other technology lifecycles. Figure 1 illustrates the continuity lifecycle, which, you will notice, includes five of the seven core requirements that were listed in the previous section (two of the requirements, "Assemble the Continuity Planning Team" and "Create the Continuity Plan" are typically treated as one-time steps). More about the continuity lifecycle is discussed in the Maintaining the Continuity Plan section of this paper.

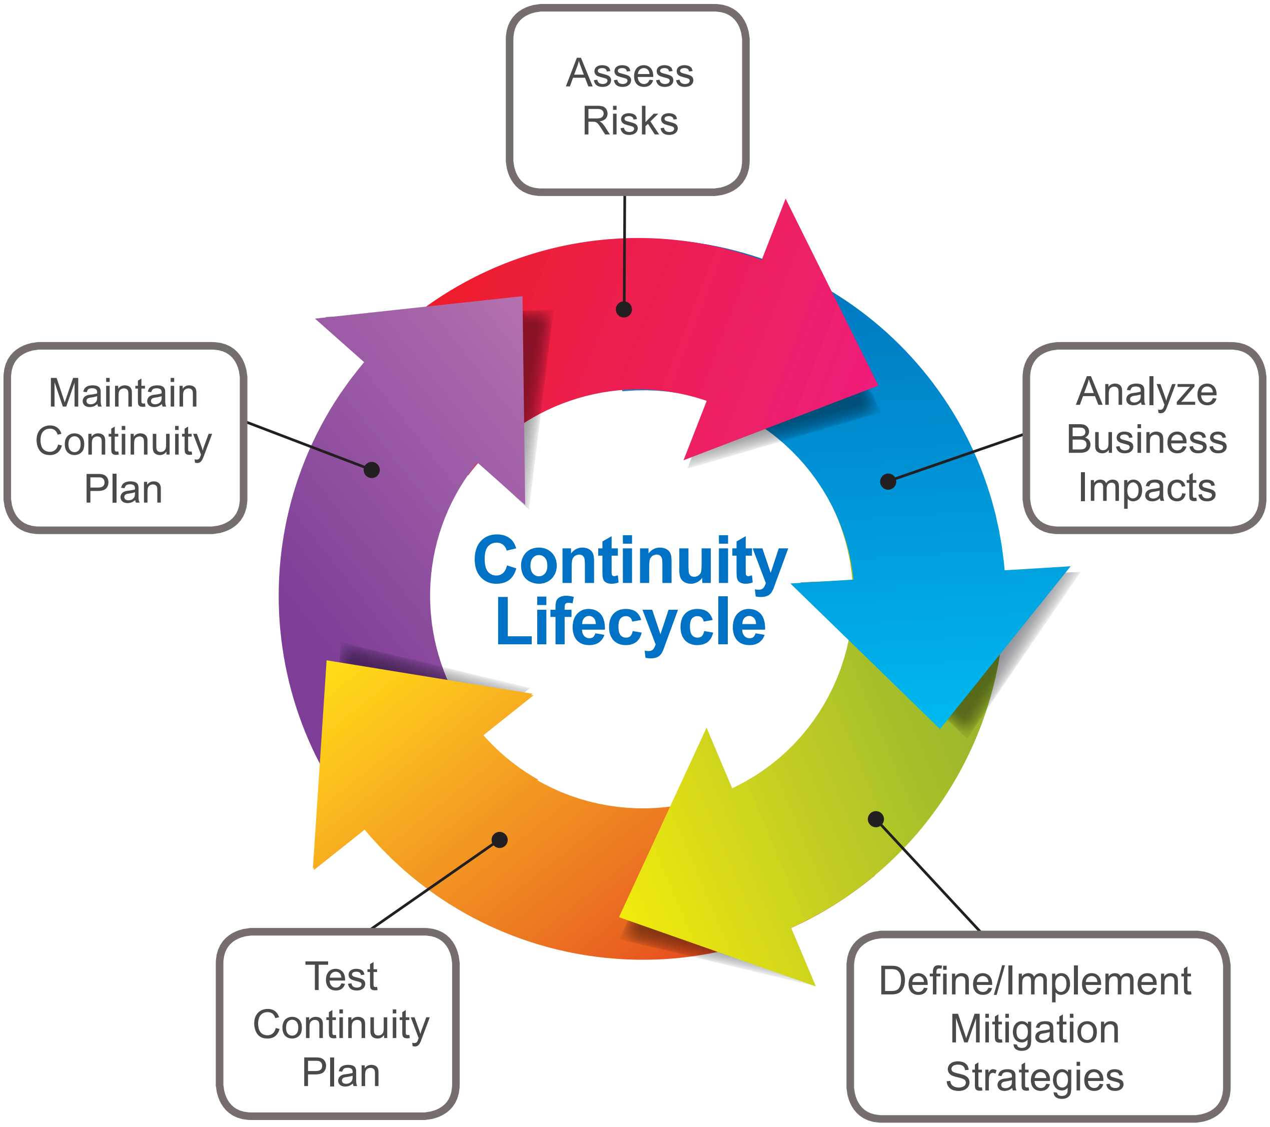

Figure 1

Assemble the Continuity-Planning Team

First and foremost, the success of any effort to create and implement a continuity plan requires sponsorship by senior management, whose responsibility it is to support the organization's continuity effort financially and otherwise, and to make sure the effort achieves its goals. Toward this end, many organizations create an executive-level steering committee that sets the agenda and provides funding and visibility for all continuity efforts as well as designates a continuity manager who in turn builds the continuity team that will ultimately make assessments, create and test continuity plans, implement needed mitigation strategies, and perform the recovery operations when needed.

Of course, a significant subset of the greater business continuity team is the IT continuity planning team, which is typically made up of one or more individuals from each key area of IT and is led by an appointed manager or coordinator who oversees all IT continuity planning and recovery processes. Clear roles and responsibilities are defined for each member of the team, which is essential given the stress and chaos that can come after a disaster.

As a rule of thumb, ideal staff members for the IT continuity team are those who exhibit the following qualities:

- Considered as an expert by peers
- Contributes actively to many IT projects
- Works well under pressure
- Self-confident but not arrogant
- Trusted by peers
- Shows a willingness to solve problems

The IT Continuity Manager

Leading the IT continuity planning effort as well as the effort for recovering IT systems in the event of a disaster is the IT continuity manager, who must be a respected, seasoned expert with a proven ability to manage teams and projects under stressful conditions. In addition, the IT continuity manager must be good at delegation and be able to calmly lead the team from the chaos of a disaster through each step of the recovery plan. Simply put, how well the IT continuity manager responds during a crisis can make or break a company's success in its recovery. With all of this said, it is critical that the IT continuity manager is not a senior-level manager since people at that level usually have too many other demands upon them to devote the necessary amount of time required to perform this job well. Nonetheless, the IT continuity manager needs the full support of senior management, must communicate regularly with the management team, and should have the authority to authorize the expenditures needed to replace equipment.

Perform a Risk Assessment

Conducted by selected members of the continuity planning team, the risk assessment endeavors to identify all possible points of failure as well as other vulnerabilities to IT systems and processes, including all infrastructure and external service-provider dependencies. As part of this process, an evaluation is made of all current capabilities both to reduce the assessed risks and to recover quickly should an event occur that affects a business-critical system. From this assembled information, a gap analysis is then performed that shows the current net risk exposure to IT systems.

Ultimately, the continuity planning team and senior management need to decide for each IT system whether to accept the current risk, reduce the risk to a more acceptable level, or make efforts to reduce the risk to its lowest possible level. These decisions are usually made once the business impact analysis is conducted, which is described in the next section.

Conduct a Business Impact Analysis

Following the risk assessment (or in conjunction with it), a business impact analysis (BIA) is conducted that identifies and quantifies financial and other impacts (e.g.: reputation, regulatory breach, etc.) that could result from disruption to business-critical IT systems. In addition, the BIA estimates the maximum downtime and data loss that can be tolerated by the organization for each critical system. When this information is combined with the data gathered from the risk assessment process, the continuity team and senior management can begin to evaluate and implement strategies and technologies that will lessen these assessed risks and business impacts.

Quantify Financial and Other Impacts to Each Critical System

When it comes to making sound business decisions for investing in resilience technologies and strategies, knowing the true cost of having each high-priority IT function unavailable for extended periods is critical information. Of course, as part of this estimate, all tangible costs should be calculated, such as equipment replacement, lost sales, penalties, lost discounts, etc.; however, it is equally important to estimate intangible costs, such as potential loss of customers, regulatory violations, negative publicity costs, loss of stock value, etc.

For a good overview of technologies that provide varying levels of RTO and RPO for IBM Power systems, see the Vision Solutions white paper, *An Overview of Disaster Recovery and Availability Technologies for IBM Power Systems.*

Quantify RPO and RTO Expectations for Each Critical IT Function

An important step of the BIA process is to determine for each system the recovery time objective (RTO), which is the amount of downtime that can be tolerated, and the recovery point objective (RPO), which is the amount of data loss that can be tolerated. The IT continuity team simply cannot make a solid recovery plan, nor make good resilience investments, if it doesn't know the RTO and RPO that the company needs to achieve. For instance, if potentially losing up to 24 hours' worth of data is not acceptable to your company, your IT continuity plan isn't going to rely on daily tape backups as a recovery strategy.

Identify and Prioritize for Recovery All IT Functions, Services, and Processes

In the event of a disaster, your business is thrust into triage mode, where high-pressure decisions must be made as to what functions need to be recovered first. Without a plan that specifies recovery priorities, conflicting decisions are likely to be made, and a great deal of time and resources could be wasted recovering entire processes and systems in which only a portion is needed to get the business back on its feet. That's why as part of the process of determining your RTO, it is critical make decisions about the recovery priority of each IT system. Of course, this information will become a key piece of your written continuity plan.

When prioritizing systems for recovery, it is helpful to put these into four categories:

Critical These are the systems, applications, and other IT processes that will cause the business to completely stop functioning should they become unavailable. When considering what resources are critical, it is beneficial to look at things in a granular way; for example, in the interest of recovery time, you may choose to first restore critical IBM i libraries from a particular system and then restore non-critical libraries at another time. In another example, you may decide to recover both critical and non-critical applications on a backup system that has less power than the original system, but allocate fewer computing resources to the non-critical applications.

Essential These resources aren't required for staying open for business in the near term, but could cause long-term disruption to the business if not restored within a reasonable amount of time. For instance, these might be resources that you choose to recover only once production systems are restored rather than recovering these under your hotsite contract or on a redundant server at one of your satellite locations.

Necessary These resources contribute at some point to the business running smoothly and are useful to employees, but they're neither critical nor essential. Examples in this category might include reporting and software development.

Optional These include resources that can wait until your production environment is fully back to normal. Examples in this category might include testing environments and company intranets.

Define and Implement Mitigation Strategies

Once the BIA and risk assessment have been conducted, your organization now has the actionable information it needs to not only begin creating a formalized continuity plan (see next section), but to also begin evaluating, budgeting for, and implementing processes and technologies that reduce risk and lessen the business impacts of downtime. Using this information, your organization can run various cost-benefit analyses to determine where the most difference can be made in ensuring the continuity of IT systems.

Top of the list should be reducing any security vulnerabilities and implementing redundant resources that can take over should a failure occur. High availability configurations in which real-time replication of critical data is made to an offsite system that can take over in the event of hardware failure or site disaster can have a huge impact on improving RTO and RPO metrics. When properly configured and managed, high availability systems can reduce the downtime of critical systems from one or more days to an hour or less. If during your cost-benefit analysis it is determined you don't have the budget to implement high availability, then at minimum, having a contract to use backup resources at a hotsite provider will lessen the risk of being down for several days or longer if a primary computing facility is damaged or destroyed.

.

Engaging Application Service Providers to Reduce Risk and Business Impact

Your company may have good software and other technologies in place that reduce risk and business impact, such as high availability and security solutions; however, these technologies are only as good as the experts who configure, monitor, and manage them. To further reduce their exposure, some companies implement a managed service contract with their application provider in which the vendor oversees their own technologies to help them run optimally. By regularly monitoring systems, fine-tuning configurations, and performing any needed troubleshooting, the likelihood these solutions will prevent downtime increases dramatically. In the case of high availability solutions, the vendor's experts can also perform regular switchover and failover tests so that you are assured the solution will recover systems quickly in the case of an adverse event.

Vision Solutions provides high availability and security managed services to many of its customers.

Create the Continuity Plan

Proper continuity planning means creating detailed, well-organized documentation that describes the procedures and processes for performing a successful recovery should it become necessary to do so. Your plan should cover various scenarios (e.g.: hardware issues, security issues, site disasters, etc.), defining for each scenario the procedures that are to be followed to recover each type of IT system along with a description of the final recovery goal.

Creating your plan requires a significant effort by the continuity team to gather and organize all needed information as well as to come to a consensus on when/how a disaster is declared (activation of the plan) and the follow-on recovery procedures. In addition, the plan should be reviewed and approved by the enterprise-wide business continuity planning team (and included in their business continuity plan) as well as by senior management.

Assign sections of the plan to be written by different members of the team who have strength in their assigned areas. Provide a guideline for how each person/team should write its section to keep things clear and to ensure it integrates well with other sections. Define who will execute each recovery component. As the recovery procedures are detailed, it should always be assumed that the hardware and other IT infrastructure would need to be rebuilt from scratch.

While the continuity plan needs to include as much detail as possible, it is equally important that the content is easy to understand and well-organized, particularly the sections that give detailed recovery instructions as these will likely be read under very stressful conditions. In fact, it is important to emphasize in the written plan that those who use the plan must carefully read through each procedure before it is executed. Under the stress of a disaster, it is very easy to overlook important details.

A hard copy of the plan should be given to all team members with the instructions for each member to store their copy somewhere away from their office building so it can be accessed in the event of disaster that affects access to their office.

Key Elements of the Continuity Plan

- Criteria for who is notified when a disaster is declared, who activates the recovery plan and how, and what the criteria is for assessing damage
- Specific duties for each member of the continuity team
- Recovery priorities for servers and applications. Include the information from the prioritization exercise of the business impact analysis.
- Contact lists for all continuity team members and company managers
- Communication protocols that define who should be contacted and when, and what level of detail should be communicated
- Hardware and software dependencies, including all necessary license keys. Not having this information can significantly delay the recovery effort.
- Records of the location of critical documents, alternate sites, equipment suppliers, data-storage locations, portable power generators, supplies, etc.
- Listing of key vendors and customers that are to be notified in the event of a disaster

Phases of a Recovery

When an event occurs that makes a recovery of one or more systems necessary, several predictable phases take the recovery process from detection of an incident all the way through to its final resolution. As the continuity plan is created, it should describe for each disaster scenario the people and resources required for each of the following phases:

Response phase When an incident is detected, a recovery plan coordinator or their delegate must be present at the incident site to do an assessment of the damage and its impact to business operations. If it appears that significant downtime of critical IT systems will result from the incident, then the recovery portion of the continuity plan can be invoked on the spot; otherwise, the assessment is communicated to the IT continuity manager, who decides whether to activate the recovery process. If the recovery process is activated, the recovery team is mobilized and communications are initiated with relevant staff, management, vendors, etc.

Recovery phase The recovery phase involves executing the procedures as outlined in the recovery process to get systems back online. This includes ordering necessary hardware, executing failover if high availability is in place, getting backup tapes from offsite repositories, activating the hotsite with the contracted provider, etc. Once the needed recovery assets are in place, then the recovery team configures hardware, restores data, re-routes networks, brings systems back online, etc. Of course, the goal of the recovery is to have it be completed within the time periods as defined in the continuity plan, which in turn were set from the RTO and RPO defined in the business impact analysis.

Restoration phase The recovery might necessitate running essential IT operations for a period of time in temporary facilities, such as at satellite offices and/or with a hotsite provider. When that’s the case, the restoration phase focuses on the work needed to repair or replace infrastructure and equipment at the primary facility in order to fully resume normal operations there; in fact, many of the same tasks in the recovery phase are repeated here, but this time at the primary location.

Testing the Continuity Plan

Once the continuity plan has been created, the work doesn't stop there; in fact, the real discipline of continuity planning happens only after the plan is created, for it is during this time that the plan matures and the skills of the recovery team are honed. As mentioned at the beginning of this paper, successful continuity planning is a lifecycle that requires a continual cycle of assessment, documentation, deployment, testing, and improvement.

The Continuity Plan Review

Once the first version of the continuity plan is completed, it is important to bring the members of the continuity team together to perform a thorough walkthrough of the plan to ensure everyone fully understands their role as well as all recovery and communication procedures and protocols in the event of a disaster. As part of this process, each team member should review in detail their components of the plan in order to ask questions, clarify procedures, and identify any overt issues for correction. As a result of this initial walkthrough, it's not uncommon to find enough issues so as to require the creation of a new draft of the continuity plan (and perhaps a second walkthrough) before actual testing of the plan commences.

Beginning the Testing Process

Once the plan walkthrough has been completed, the next step is to begin testing the plan by undergoing a variety of processes that validate its effectiveness to restore IT resources. During testing, it's not unusual for many weaknesses and even failures to be discovered. This is to be expected; in fact, it's precisely the point of doing testing as the continuity team won't know what needs to be changed until it runs a test recovery from the plan. In addition to finding holes, another objective of testing is to have each member of the team become familiar with the recovery processes and their role in it, which will make for a more efficient and effective recovery when it's done under the difficult conditions of a true disaster.

Testing of the continuity plan can be performed in a variety of ways, but these methods typically fall into two categories: passive and active.

- Passive tests are sometimes referred to as a "table-top exercise" and usually occur in a conference room with all members of the team bringing their own copy of the plan to the meeting and then working deliberately through a specific type of disaster scenario. Passive tests are more often than not a means to make team members familiar with the plan and to get a sense of how a recovery might unfold.
- Active tests are those in which some portion of IT backup systems are activated so as to see the effectiveness (or lack thereof) of these resources in various scenarios. A key part of measuring effectiveness is determining not only whether the activation is successful, but also whether it occurs within predefined RTO and RPO parameters. Common components of active tests include activating a hotsite with the contracted provider and then restoring systems from the most recent tape backups, performing a high availability failover, and activating LAN and WAN failover mechanisms. In the case of high availability failover testing, many companies will engage the services of the high availability vendor to assist during the test and to make recommendations for optimizing the solution and associated environments so any future failover attempt goes smoothly.

Because any circumstance that involves the recovery of business operations is performed under very stressful conditions, a carefully crafted, well-tested plan is crucial. On top of having management breathing down your neck during a recovery, you'll likely have to deal with one or more of the following circumstances: restoring operations at another location or facility, having access to resources that are stored only offsite, having to recover within a predefined period of time, and having some key personnel unavailable. It's a tall order, and the only chance for success is if your continuity team puts in the necessary upfront work to regularly and thoroughly test the plan.

With the above in mind, the following are some best practices for testing your plan:

- Use potentially real disaster scenarios for your testing, and introduce potential complications such as not having access to the building or having telephone services being down. This gets team members thinking outside the box, while discovering undocumented vulnerabilities and how to resolve them.
- Have team members change roles during different test scenarios so they can get a sense of what might be expected should some person in a key role turn up missing during a disaster.
- Try a test with a significantly reduced recovery team. Of course, all team members who were not selected for the reduced team test should be in attendance during the test to observe and learn from the exercise.
- Rather than trying to test everything at once, consider doing a variety of different tests during the year, with each test working to recover a different component.
- Have some of your continuity plan tests be conducted unannounced as this brings in the element of surprise that will uncover issues you may not have considered. Testing shouldn't always need to accommodate everyone's schedule.
- Detailed notes should be taken and distributed, along with specific time-based action items for both resolving all discovered vulnerabilities and updating the plan.
- In advance, level set with management that tests will not always go well. That's the point; it's all about finding and resolving points of failure.

Maintaining the Continuity Plan

Once your continuity plan has been written and tested, it doesn't mean the work of the continuity planning team is done; in fact, effective continuity planning is not something that is finished. Your plan needs continual updating based not only on the results of regular testing, but also on changes to IT systems as well as changing processes and priorities within the business. And then, of course, there's documenting the experience that's gained should actual downtime events occur that trigger a recovery situation. Without doing regular maintenance of the plan, it can become quickly outdated, which significantly increases the chances of an unsuccessful recovery after an incident.

The best practice is to consider continuity plan maintenance like any other IT change-management process. This means processes are put into place across IT departments, as well as the entire organization, to ensure any changes that might affect the business and its IT environments are communicated in advance to the continuity planning manager or another designated member of the continuity team so that steps can be taken to maintain system resiliency and update the plan in light of the changes. For instance, a server upgrade means that changes will likely need to be made to the company's contract with its hotsite provider to accommodate the need for expanded resources. And when high availability systems are in place, it is particularly critical that any hardware and software updates on production servers are accommodated on the backup servers; otherwise, a failover will run into complications.

In addition to regularly testing the continuity plan and having the continuity planning process be part of IT change management, there should also be a process in place in which, at regular intervals, the IT continuity planning team reviews and updates information in the continuity plan. This includes making sure things like contact lists and equipment inventories are kept updated with current information.

Coming Full Circle

In addition to maintaining the plan, it is important for your continuity planning team to periodically conduct a new risk assessment and business impact analysis. From that information you will likely find it necessary to augment your strategies (or implement new ones) to mitigate risk and business impact, and of course, all of this will trigger more changes to your continuity plan, which means a new round of testing of your plan. With this, the continuity planning process again comes full circle, which is why effective continuity planning needs to be treated a lifecycle, as described in the beginning of this paper.

Don't Leave It to Chance

Successful IT continuity planning, not to mention enterprise-wide business continuity planning, requires a well-trained team, a thorough assessment of risks and business impacts, and a carefully crafted plan that's regularly tested and continually maintained. Any enterprise that does less than this takes a big chance on its ability to effectively recover after a disaster. Given the many surveys publicized over the years showing that a significant percentage of companies will go out of business within a year or two after a major disaster, no prudent organization should leave its fate to chance.

Easy. Affordable. Innovative.
Vision Solutions.

Vision Solutions is a leading provider of business resilience solutions – high availability, disaster recovery, migration, data replication and security – for IBM Power Systems. For more than 25 years, customers and partners have trusted Vision to protect and modernize their environments, whether on-premises or in the cloud.

Visit visionsolutions.com and follow us on social media, including Twitter, Facebook and LinkedIn.

Find us on:

Facebook:	http://www.facebook.com/VisionSolutionsInc
Twitter:	http://twitter.com/VSI_Power
YouTube:	http://www.youtube.com/VisionSolutionsInc
Vision Solutions Blog:	http://www.visionsolutions.com/blog

VISION® SOLUTIONS

15300 Barranca Parkway
Irvine, CA 92618
1.949.253.6500
1.800.683.4667

visionsolutions.com