Gartner

Research

Publication Date: 16 March 2005 ID Number: G00126538

Six Myths About Business Continuity Management and Disaster Recovery

Josh Krischer, Donna Scott, Roberta J. Witty

There is no "one size fits all" when it comes to developing business continuity management strategies and plans. Using someone else's requirements, which might turn out to be based on limitations or regulations that your company doesn't have, could spell disaster of another type.

The information contained herein has been obtained from sources believed to be reliable. Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner's research may discuss legal issues related to the information technology business, Gartner does not provide legal advice or services and its research should not be construed or used as such. Gartner shall have no liability for errors, omissions or inadequacies in the information contained herein or for interpretations thereof. The opinions expressed herein are subject to change without notice.

WHAT YOU NEED TO KNOW

Business continuity management and disaster recovery planning is hard work because it means addressing every aspect of your business operations in the planning, development and testing phases of a recovery plan. Start out strongly — know what is required, and what is not, by conducting a business impact analysis. Apply an integrated business and IT approach for recovery plan development, management and testing. Reduce the maintenance of your business continuity plan by using modular scenarios for disaster and recovery. Assure the integrity of your data at your secondary site through proper planning and testing, and by keeping point-in-time copies.

ANALYSIS

There are no hard and fast rules about what the "best" recovery strategy should be for a company. The recovery strategy depends on a number of issues, including:

- The maturity of the company's business continuity management (BCM) and disaster recovery (DR) programs
- The number of mission-critical business processes and applications it has
- The amount of investment the company is willing make
- If the company is in an industry that is required to recover their services, such as financial services
- The number of facilities the company has as part of its business model

When developing BCM and DR plans, there is no "checklist" that can be applied to meet the needs of all companies. However, there are well-documented approaches and best practices. Gartner responds to some of the more common misconceptions about how a company should master its BCM or DR program.

Myth: One recovery plan meets all scenario requirements. Many companies mistakenly think that there is one overall DR plan, regardless of the type of incident or crisis.

Best Practice: Think of business continuity and recovery scenarios as "modules" that fit into a broader business continuity plan. When an incident or crisis occurs, it is mapped to the appropriate business continuity scenario(s), which then dictates the appropriate recovery plan modules to be invoked. Modules can be reused for various business continuity scenarios. For example, certain types of disaster will involve making contact with external authorities, while others will not. Some types of disaster will require the involvement of a company's PR department, while others will not. See Figure 1 for an example of recovery modules that might be invoked for the business continuity scenario in the event of a gas leak. Depending on the evaluation and assessment, all modules may be invoked, or just a few.

Figure 1. Business Continuity Scenario Invoking Multiple Recovery "Modules"

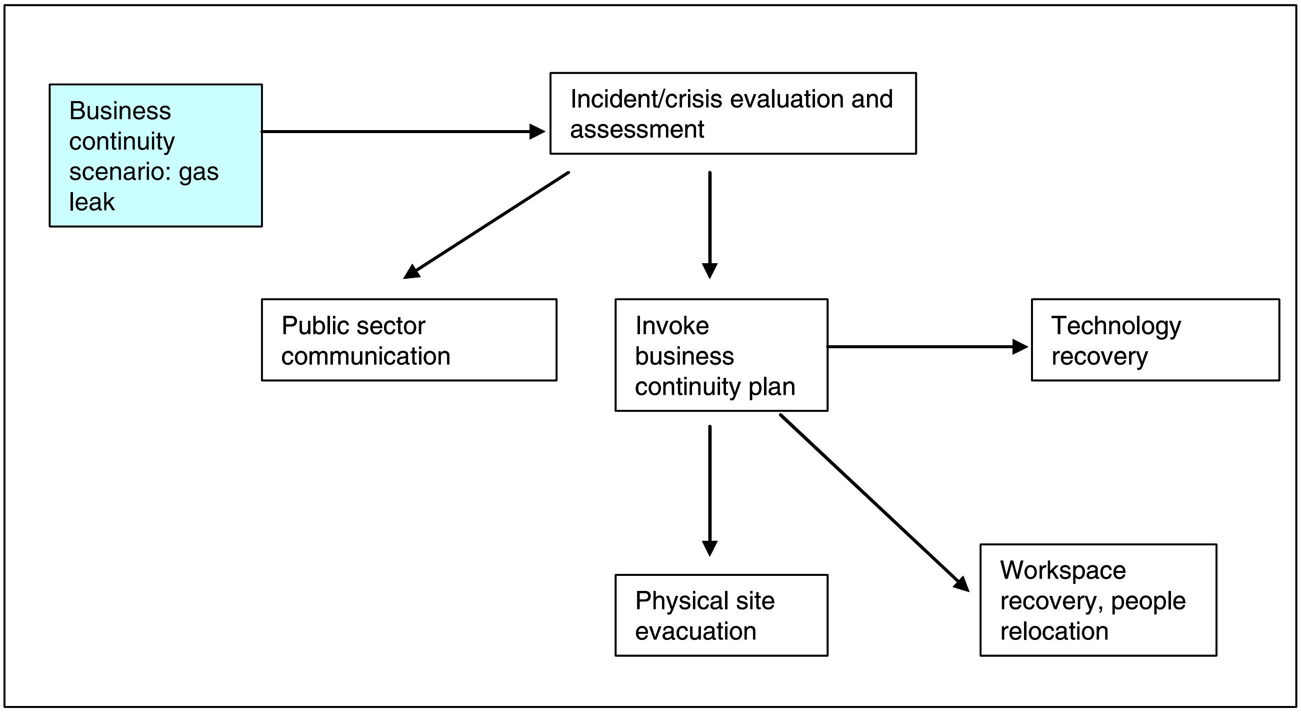
Business continuity scenario: gas leak
Incident/crisis evaluation and assessment
Public sector communication
Invoke business continuity plan
Technology recovery
Physical site evacuation
Workspace recovery, people relocation

Source: Gartner Research (March 2005)

Myth: Planning and testing with IT personnel only is enough. Many companies think the endgame for business continuity is to recover the technology infrastructure, such as network, telecommunications, applications and desktops. Therefore, they do a fine job in DR, but when and if the time comes to execute the DR plan and use the recovery site for production processing, it may not be possible for business to be conducted. Small, seemingly unimportant things need to be taken into consideration by both the business and IT. For example, the phone system at the recovery site may be different from that at the production site, or the customized desktop that a user relies on in production isn't available during a recovery situation.

It's true to say that BCM had its origins in the IT department, but that doesn't mean it should stay there. All too often, the IT department decides to focus on those applications it can technically recover without first conferring with the business units about what the "right" applications may be. And worse, access to the application is tested only to the point of a successful logon; no one ensures that the correct data has been recovered — for example, that the correctly dated recovery tape was used or even that there was data on the tape. Only business personnel can test this recovery requirement successfully. However, businesses are often short sighted in saying to the IT department: "Technology is your job, you recover it." or "It's just a form of insurance, we'll never need it so I'm not going to invest much in it." Neither view is correct, and both could be seen as negligent.

Best Practice: Adopt an integrated approach to business continuity planning and testing. Both IT and the business must be involved when developing, testing and executing plans. And it starts with a business impact analysis (BIA).

Best Practice: Perform a business impact analysis when planning for business continuity. The BIA is the most critical step as it identifies what and how much the company has at risk, as well as which business processes are most critical, thereby prioritizing risk management and recovery investment. The business continuity team, which has to include the business process owners, must translate the business requirements into an overall business continuity plan. Three of the most important deliverables from a BIA are:

- **Recovery time objective (RTO)** — the length of time between when a disaster occurs and when the business process must be back in production mode.
- **Recovery point objective (RPO)** — the point in the business process to which data must be covered after a disaster occurs; for example, the start of the business day, the last backup or the last transaction that was processed.
- **Cost of downtime** — The business should calculate the potential losses incurred, both as the result of a disaster and in recreating lost data (see "Business to Consumer - Cost of Downtime Considerations").

These considerations determine the technologies and methods used to support the DR plan.

Myth: Longer distance means better disaster protection. There is no "hard and fast" rule about the minimum distance required between data centers. Rather, the distance is dictated by regulations, management's appetite for accepting risk and the location of corporate assets. For example, increasing the distance between data centers reduces the likelihood that the two centers will be affected by the same disaster. However, few disasters happen on a large scale, and putting too much distance between them increases the risk of broken links, line failures and the cost of transmitting data, and may make traveling to the recovery site more difficult for employees. These and other considerations make choosing a secondary site a complex process. And few companies can choose their secondary site freely. Instead, the choice often depends on the location of other owned or affiliated sites or service provider facilities.

Best Practice: Conduct a risk impact analysis to determine the optimal distance requirements. There is no ideal distance between primary and disaster recovery data centers. Rather, the best location is the one that minimizes the risks at an acceptable cost and meets any required industry regulations. Considerations include mitigating risks from: common outages like power, water, network and telecommunications; geophysical disasters such as earthquakes or tornadoes; geopolitical situations like riots, terrorist attacks or strikes; and a lack of people and transportation. Distance limitations in technologies that may be chosen for short recovery time and point objectives also have to be taken into account.

Best Practice: Invest in infrastructure to ensure availability of resources that are usually beyond your control. Increasing the distance between the primary and secondary sites will mean higher telecommunications costs and the deployment of appropriate techniques. It may also reduce performance and increase the chances of disruption. In most cases, regardless of the distance between the sites, each data center should have a separate main power supply (different providers or at least different transformers) and separate telecommunications paths. It would be even better if each data center had redundant power generators and an uninterruptible power supply (UPS). If both sites are to be connected by fiber-optic cables, redundancy should be provided by using two separate routes.

Myth: The most important issue in remote copy design is to keep the data losses to a minimum (small RPO). Keeping data losses to a minimum is critical for some applications. But a more important issue is assuring data consistency and integrity at the recovery site. If the data is not consistent at the recovery site, a time-consuming backup is usually required, which may take days. Also, hunting down conflicting data and reconciling the status of key information can mean a much longer recovery time. Many companies mistakenly believe replication technology suppliers that say there will always be data consistency in a disaster.

Best Practice: Ensure data consistency and integrity at the site to influence a speedy recovery. If the data is not consistent at the recovery site, a time-consuming diagnosis and reconciliation process may ensue or recovery from tape may be required. This could take days and put the business at risk. We recommend that companies fully understand how their chosen replication technology works, what its limitations are, and how it will react in various disaster scenarios, such as loss of network, physical site disaster, component failure and application failure. Only then can they put in place a strategy to assure data recovery with integrity, and still meet their RPOs.

Companies should also ensure they have a plan to recover within an acceptable time frame, as data corruption is always a risk (even with proper planning) because of potential application logic errors or rolling hardware failures. Typical strategies include maintaining point-in-time copies of data on disk and replicating databases at the secondary site a few minutes or even hours behind those at the primary site. This will enable recovery prior to the point of corruption. Companies should test various types of scenario at least once a year to ensure that business applications have access to consistent data.

Best Practice: Review storage vendors' data consistency techniques. Ensure that these techniques are deployed and tested as part of your data recovery plan.

Myth: One copy of mirrored data at the recovery site is sufficient. If synchronous remote copy is suspended (as a result of link failure, for example) and then activated, the updates to the recovery site will be sent sequentially and not in the order in which they arrive to the primary system. The act of starting a resynchronization activity between the primary system and the remote system will temporarily compromise the consistency of the remote data until resynchronization is complete. If a disaster strikes in the meantime, it will result in lack of consistency and data integrity at the recovery site.

Best Practice: Maintain storage-controller replication by keeping two copies at the recovery site — a main copy (target of the replication) and a point-in-time copy. The reasons are:

- If the remote copy operation is suspended, a split between the target (secondary disk) and the point-in-time copy should be performed. If a disaster strikes during the resynchronization process, data on the secondary disk may not be consistent, but the point-in-time copy will contain the last consistent image. The local copy is reestablished after resynchronization is complete.
- In a disaster it can happen that, during the recovery process, some data will become unusable. If the recovery is done from the local point-in-time copy, it will not damage the source data and a new local copy can be made at any time.
- A DR infrastructure without testing is useless. Testing done with the local copy on the secondary site is less risky than with the "main" copy.

Gartner recommends keeping a consistent, restartable image of the data volumes on the recovery side for both synchronous and asynchronous remote copy. For more information, see "Consider Data Consistency When Planning Disaster Recovery."

Myth: The planned telecommunications bandwidth should exceed the peak data transfer requirements.

Best Practice: Ensure that only the bandwidth in synchronous remote copy exceeds peak data transfer requirements. For asynchronous remote copy, the bandwidth for average activity is sufficient.

This research is part of a set of related research pieces. See "Use Good Business Continuity Management to Prepare for a Disaster" for an overview.

REGIONAL HEADQUARTERS

Corporate Headquarters
56 Top Gallant Road
Stamford, CT 06902-7700
U.S.A.
+1 203 964 0096

European Headquarters
Tamesis
The Glanty
Egham
Surrey, TW20 9AW
UNITED KINGDOM
+44 1784 431611

Asia/Pacific Headquarters
Gartner Australasia Pty. Ltd.
Level 9, 141 Walker Street
North Sydney
New South Wales 2060
AUSTRALIA
+61 2 9459 4600

Japan Headquarters
Gartner Japan Ltd.
Aobadai Hills, 6F
7-7, Aobadai, 4-chome
Meguro-ku, Tokyo 153-0042
JAPAN
+81 3 3481 3670

Latin America Headquarters
Gartner do Brazil
Av. das Nações Unidas, 12551
9° andar—World Trade Center
04578-903—São Paulo SP
BRAZIL
+55 11 3443 1509