Storage Solutions Ensure Better Business Continuity

Fast access to mission-critical data can mean the difference between disaster and quick recovery.

You might already have a disaster-recovery plan in place. However, is it good enough to keep your business running without suffering any downtime?

"The stated objective of Business Continuity [BC] is to make sure that your business can continue to operate throughout some type of disaster," says Bob Laliberte, analyst, Enterprise Strategy Group. "This is versus just having a disaster-recovery plan where your business has the ability to recover operations some time after a disaster.

"A business-continuity solution would typically involve having a secondary site," he adds. "In the event of something happening to the primary site, that secondary site can resume operations in a matter of minutes, minimizing any downtime."

If you're going to have a successful business-continuity plan, you need immediate access to your critical business data. According to John Bennett, worldwide director for data center transformation solutions at HP, data protection and recovery are the cornerstones of business-continuity plans.

Without it you'll feel the impact of downtime including loss of revenue, decreased worker productivity and possibly damage to your reputation. For example, if e-mail goes down, will the firm loose employee work time? "You need to quantify that downtime," he adds.

With that in mind, here are five critical steps to help you build a solid business-continuity strategy.

1. The Right Protection

You need to know what you need to protect against to stay in business. It's key to understand the likely occurrences that could cause an outage.

"One of the first things you want to take into account is your geographic location," says Rick Walsworth, director of product marketing for replication products at EMC. "For example, how do I protect myself against an earthquake or a hurricane that could render my production systems useless?"

In most cases, the consideration from a strategic standpoint is knowing how to get the mission-critical data offsite to a location that is geographically dispersed from the main production center.

"There are many ways to do that," says Walsworth. "This includes putting everything on tape and trucking it off to another facility to replicating the site at another location."

You also need to look at the Recovery Time Objective (RTO) and the Recovery Point Objective (RPO). That is how fast you can recover your data, applications and servers, and how much data you are willing to lose.

2. Data Assessment

Take the time to assess your data: from e-mail to mission critical. You want to know what data is most important. To get that, you need to do a full business requirements review.

"If you look at an enterprise or mid-sized company, they have financial systems, web ordering, e-mail and other systems that have varying levels of sensitivity", says Walsworth. "Not all data is created equal."

In some cases companies are forced to protect only that data that's mission critical and nothing else. After an assessment of data, you can assign service levels.

For mission-critical applications, you can assign a very narrow recovery objective. As for second-tier applications, while they will be protected, they won't be assigned the same amount of criticality in terms of recovery.

After conducting a business-impact analysis, you can rank the most critical data and then look at the best technologies and weigh that against the cost of deploying it. "It is also important to discover any interdependencies a mission-critical application may have with other applications, and ensure those applications are protected as well," Laliberte says.

3. Storage Options

Know your storage options: from basic backup to virtualization. At the ground level, there's basic backup, where you scan a system or data from a database, and back it up to tape or disk.

"The recovery point is pretty long in this scenario," says Don Lamorena, senior manager, high availability and disaster recovery, Symantec. "If the company only backs up once a day, you'll lose an entire day's data."

Mirroring and Clustering

With the increase in the amount of data that needs to be backed up and protected, there's been a storage explosion. This adds another level of complexity to the backup and recovery processes.

"If you're dealing with multiple terabytes of data, the data-recovery process can take some time," says Bennett. "This is where clustering and mirroring come in. It lets you recover in real time or near real time the critical data that keeps your business running."

Server clustering can be used to improve performance and/or availability over a single computer. With disk mirroring two drives simultaneously duplicate and store data so if one drive fails, the system can switch to the other without loss of data or service.

"Where you have faster recovery points and times, you're going to look at mirroring and replication," Lamorena says. "Mirroring is more expensive than backup because of the real-time factor. Your RPO when you're mirroring can be as low as zero data loss."

Data Replication

Another option is replication where you move the data from one site to the next. If you want to provide site-wide failure protection, you replicate data over to the next site for failover. While you might have some data loss, it's all about the recovery time, which is how fast your users and your customers can access the service.

"That's where clustering comes into play," says Lamorena. "It monitors applications, if applications go down, it will recover that application very quickly and start it up on another server in the same data center or in a remote location."

But you also have to protect against logical corruption. "In a replicated environment you're typically replicating all of the changes," says Walsworth. "But there are some areas where companies aren't protecting themselves against logical corruption."

For example, a database that captures all the changes could, at some point, also replicate an event that corrupted the data. "Anything that came in beyond that point has become corrupted," he says. "And chances are that I've replicated that corruption from my production site to my remote site. So now I have a situation where both copies of my data are corrupted."

If you have implemented a Continuous Data Protection (CDP) model, it allows you to rewind to the point in time just before that corruption took place. The continuous paradigm allows you to replicate the data but in the case of logical corruption, you can roll back before that corruption took place and then resume operations.

Virtualization

Given the growing need for multiple storage devices and multiple servers to keep up with the data and access needs, many businesses are considering virtualization as a way to not only maintain continuity but save on costs.

According to Lou Shipley, group vice president and group general manager, XenServer Products Group at Citrix, virtualization is key because it decouples hardware from software.

Virtualization essentially allows you to move applications around from one physical server to another. You can provision servers much faster so you can get up and running again faster than manual recovery from tape backup.

Storage and server virtualization have elements in common. "For those doing server virtualization they're also looking to move toward network attached storage or a shared storage model for the data center," says HP's Bennett. "Direct attached storage and dedicated storage are the enemies of being able to move applications or virtual images around on virtual servers."

There's also an impact on the storage architecture of the data center associated with doing server virtualization projects. That in itself can help with data protection and recovery rather than hindering because now all your data is in a common infrastructure.

And if it's a shared infrastructure, you have the resources to do snapshots and then do electronic or tape-based backup. This is actually more advantageous in environments where you can't take the applications down for servicing.

Storage virtualization also lets you take storage arrays from different manufacturers and create a single storage pool using storage virtualization techniques. Business can not only get a network attached or shared storage utility model; they also get a lot of protection over the legacy investments they've made in storage subsystems.

4. Business-Continuity Costs

Determine the cost of deploying, maintaining and testing the system. You first need to consider the capital cost, the hardware and software — typically one-time charges. After that, understand what the ongoing operational costs will be to maintain this environment.

"You have people costs," Walsworth adds. "There are resources that need to be trained to execute the BC strategy in case of a failure. You have infrastructure costs such as looking at remote replication. There are also telecommunication costs to interconnect to remote data centers. And you have training and maintenance."

Those recurring costs can dominate the BC strategy and have prevented some businesses from deploying a plan. The dominant factor is that telecommunications costs to connect to the remote site — an ongoing cost that's driven by bandwidth — are going to be more expensive for a remote site that's far away. One solution is WAN optimization.

"One of the biggest costs of putting up a BC site is the cost of the network between the two centers," says Laliberte. WAN optimization solutions allow you to streamline and optimize and put more traffic over a smaller pipe.

Most companies do absorb these recurring costs to keep their businesses running. "It's like insurance, it's considered a good investment for guaranteeing uptime and continuity of business processes," says Citrix's Shipley. "We typically advise that you set up a strategy and put it in place and then choose your vendors."

The cost of this type of plan has to be weighed against the value of not having those applications available. "Typically all of these plans start with that business impact analysis," says Laliberte.

Another ongoing cost to factor in is testing disaster-recovery processes. If you don't test it, you won't know if you can fail over to that other environment and keep your systems running. But in many cases testing is disruptive to production systems.

"Considering the 24x7 nature of business today, taking the system down to test the BC plan isn't an option," says Walsworth. "So firms are forced to either not test it or to test it and run in degraded performance or to test only portions of the system and not others."

5. Recovery Process

Take the time to manage the recovery process from people to data. Managing both the recovery process as well as emergency response management is essential to a business-continuity plan.

"Have clear, up-to-date lists of managers and who will be called first, so you can respond and act quickly," says Bennett. "Emergency response systems are not only to react to emergencies with the data centers, but also to provide emergency response for other activities."

And while many focus on the recovery of systems and data, the employees — those who manage and run the data center and those other employees that are affected by site outages — need to be considered in the plan as well.

"Many recovery systems providers offer support not just for data center recovery but for call centers to provide services and continue operation in the event of a site outage," Bennett adds.

In addition, you also have the regular employees of the corporation, engaged in sales, marketing or manufacturing, taken care of as well. This allows the firm to move from disaster recovery from an IT perspective to true business continuity.

Another factor is focused around facilities and recovery. "One thing that people don't often anticipate is being out of a facility not just for six weeks, but for six months," says Bennett.

Sometimes disastrous events can physically destroy a facility or render it uninhabitable for health-related reasons. If you can't use your facility for a long time, how are you going to manage continuity of your operations?

A good business-continuity plan will look at all these aspects of the business. This is not just disaster recovery in the context of the data center, but also the employees of the organization, including management. And have a back-of-the-pocket plan should an extended outage take place.

Business-Continuity and Disaster-Recovery Facts

Symantec's 2008 Disaster-Recovery Research Report found that:

- Although one-third of organizations have had to execute a Disaster-Recovery (DR) plan, just under half say they are able to get fully operational within a week.
- The amount of applications that IT Managers believe are business critical has increased 20 percent over data from the previous year, and only about half of these applications are covered in DR plans.
- Virtualization is driving organizations to re-evaluate their DR plans.
- Organizations find that the biggest challenge for high availability and disaster recovery in virtual server environments is the different tools they need for their physical and virtual environments.
- Organizations report that DR testing impacts customers, sales and revenue.
- Despite increasing importance of DR, there is an alarming decrease in executive involvement.
- Thirty percent of disaster recovery plans, that are tested at least once per year, fail.