

Management Series

Four Steps to Disaster Recovery and Business Continuity using iSCSI

It's a fact of business life – physical, natural, and digital disasters do occur, and they interrupt operations and impact revenue. Your level of preparation and planning will determine your ability to recover from a disaster. Dramatic lessons about this topic have been learned in the past few years, yet some organizations remain unprepared. This much is certain – organizations that are prepared are able to withstand disasters, while those that are unprepared risk heavy losses.



Disaster recovery (D/R) is all about getting back to business as quickly as possible after a failure or error. But your ability to recover is directly tied to what you have done before a disaster occurs – your business continuity plans. Once trouble hits, it's too late. You need a business continuity plan that balances potential business losses against the cost of minimizing those losses. You must understand your business needs, potential disruptions, and loss tolerance in order to prepare responsibly. With that knowledge and an understanding of technology options available today, you can design your environment to avoid disasters (for example, using redundancy and hot-serviceable components) as well as develop a disaster recovery plan that's right for you.

There are four key steps to developing solid business continuity plans and ensuring rapid disaster recovery:

1. Identify your needs and define what business continuity means to you
2. Determine and document your recovery objectives
3. Evaluate technologies that support your strategy
4. Implement and test the technologies best suited to your requirements

Applying these steps to your organization will help you minimize losses when disaster strikes.

Step 1: What Business Continuity Means to You

Disasters commonly occur in two forms (Figure 1):

- Physical destruction of a location and data (or access to location and data). Examples: fire, flood, earthquake, significant power or network outage.
- Data destruction without physical destruction. Examples: hardware failure, virus/hacker attack, software malfunction, human error.



Figure 1
Location versus Data Destruction

Data destruction is far more frequent.

The ultimate goal of D/R is to get your business restarted in an acceptable timeframe. For some organizations that means within minutes, while for others it means hours or possibly days. The cost of operational downtime varies among businesses and industries. For example, financial firms often calculate that cost in millions of dollars per hour, while other industries calculate operational downtime as thousands per day. These costs include lost business transactions, employee productivity, and customers – not to mention regulatory penalties. The ability to tolerate these losses generally determines business continuity strategy.

Step 2: Determine and Document Your Downtime Tolerance

The next step adds detail to the impact of downtime on your organization. How long can you afford to be non-operational, and how much work/data can you afford to lose (Figure 2)? Does business continuity for you mean restarting business within minutes with current data? Can you tolerate restarting within hours with several hours of work/data loss? The categories below outline questions you need to answer for your organization. Documenting your answers will help you plan.

Prepare Your Strategy

To prepare your strategy, you should first identify what disasters you are vulnerable to and what losses may occur. This information will help you to build the right business continuity plan.

- Speed of restart. How quickly must you resume operations after a disaster? This is called your “recovery time objective.”
- Work/data loss. How much completed work can you lose and still function effectively, and how much productivity loss is acceptable? This is your “recovery point objective.”

In addition, you must determine what level of investment you can tolerate to protect against downtime and to recover from it. (It is important to note that higher cost does not necessarily mean higher levels of business continuity. Different technologies provide different kinds of protection, as you’ll see in Step 3.) Have you configured your environment so that more common errors (such as loss of disks or controllers) don’t become D/R scenarios? Avoiding disasters is an important part of disaster planning.

- Protection level. What is the business impact of the failure/recovery process? Have you employed sufficient redundancy and hot service capability into your infrastructure? Is the minimum level of tape-based backups enough, or is tape recovery time too long?
- Cost of being unprepared. What is the likelihood of failure or disaster? Is doing nothing an option, or are you fairly certain that you are vulnerable to things like human error, equipment failure, or viruses?
- Investment level. What are the initial and ongoing costs of being prepared?

Once you have a clear understanding of your business requirements, build a business continuity plan for both immediate disaster response and returning to regular operations. Your plan should contain:

- Current critical business processes
- Specific recovery time and recovery point objectives
- List of key personnel involved in the D/R process
- List of personnel needing information access after the disaster
- How personnel will access information – from home, from a secondary data center, from a leased facility
- What systems, applications, and data will be required, and for how long
- Chronologically, how far back you need data to conduct business as if the disaster never occurred

With this information in hand, you can review D/R technologies to identify which solutions will deliver the results you need.

Step 3: Evaluate Disaster Recovery Technologies

Remote data replication is a technology commonly used for D/R – and has become more available and affordable since the advent of iSCSI. Remote replication involves creating data copies that are sent over the network to be stored on disk at a remote location. This strategy has significant benefits:

- Data are physically protected when copies are not collocated with original data.
- Restore time is greatly improved, as data are in useable (not backup) format and are on disk (not tape).
- Work loss is diminished, since replication can be done frequently with minimal or no disruption.

COST PER HOUR, BY APPLICATION	
TYPE OF BUSINESS	LOSS OF REVENUE
Brokerage Operation	\$6.5mm
Credit Card Authorization	\$2.6mm
eBAY	\$227k
Pay-per-view	\$150k
AOL	\$125k
Home Shopping	\$113k
Catalog Sales	\$90k
Airline Reservations	\$89.5k
Teleticket Sales	\$69k
Package Shipping	\$28k
ATM Services	\$14.5k

Source: Contingency Planning Research

Figure 2
Costs of Downtime

Build a Business Continuity Plan

Once you have a clear understanding of your business requirements, build a business continuity plan for both immediate disaster response and returning to regular operations.

In considering technologies, a good guideline is to seek out solutions that offer you the data protection you need while minimizing complexity.

There are two major categories of remote replication: 1) Continuous Replication, in which data changes are sent continuously between locations, and 2) Periodic Replication, in which changes are sent periodically in batches. While these solutions vary in several performance and operational parameters, the key differences are in what kind of recovery they offer.

Continuous replication, common in Fibre Channel SANs, offers a single recovery point – similar to having RAID over distance. Write operations are continuously copied to the remote location, so the copy is identical to the current production data – which has pros and cons. In a physical disaster you lose very little data because the remote copy is current. However, with a virus, corruption, accidental deletion, or malfunction, the corruption is propagated to the remote location immediately. No previous restore point is available, leaving you to recover from backup tapes. This can take days or even weeks, and may result in significant amounts of data loss.

There are two kinds of continuous replication: 1) Synchronous, in which write operations are continuously copied, and the application/operating system is not released from the writing task until it completes at both locations; and 2) Asynchronous, in which writes are continuous, but with a lag so that the production application can resume operations without waiting for the remote write to finish.

With synchronous replication, application throughput and bandwidth are limited by network speed, and network latency can affect application performance. Typically this is an expensive solution that requires network extension equipment, high bandwidth, and limited distance between sites. Application recovery is needed at the remote site, so automatic restart is not possible.

Because of its built-in lag time, asynchronous replication allows applications to operate mostly unhindered by the network and remote storage. This reduces the performance impact, but at a “cost” of more lost work should a disaster occur. It is somewhat less expensive than synchronous, but still requires costly channel extenders and dedicated high bandwidth. Application recovery is even more complicated here, since data sets will be time skewed.

Because it is costly and complex, continuous replication is most often used for small portions of data that cannot be lost, such as large financial transactions.

Periodic replication is built on the premise that point-in-time copies, or snapshots, are created and sent regularly – but not continuously – to the remote location. Production performance is not impacted, and distance between sites is unlimited. This method offers multiple recovery points and a catalog from which to choose them. Because there is distance between the production and replication sites, you have protected data

on disk that is quickly accessible in case of physical disaster. In addition, should a virus, corruption, accidental deletion, or malfunction occur, you can restore “known good” data from a point in time prior to the disaster.

Application recovery is fast, because snapshots provide application data in a clean state so no additional process is required. Because only the changes between point-in-time copies are written, this method is also network and storage efficient. You decide how often to replicate depending on your business needs – typically, no more than minutes of work are lost in a disaster. This method is simple and affordable, and offers protection from all kinds of disasters.




Figure 3
Periodic Replication Enhanced by iSCSI

iSCSI adds real business benefits to periodic replication (Figure 3). First, you can connect your primary and remote

sites using standard IP-based technology. With your standard network connectivity you can define a partnership between the sites, select volumes, and schedule automatic replication. In addition, you gain the inherent benefits of iSCSI – ease of use and implementation, complete standardization and interoperability, and low cost. Also, IT personnel are familiar with the technology so there is no delay or expense to get people up to speed.

Technology Guideline

In considering technologies, a good guideline is to seek out solutions that offer you the data protection you need while minimizing complexity.

Step 4: Implement and Test your D/R Plan and Solution

Once you understand your needs, budget, and expectations and have selected the technology that suits your situation, implement it in your environment. To ensure that your business continuity plan will function as you want, you should perform a disaster recovery dress rehearsal. Walk through your plan, rebuild your application environment, and solicit business users to test their applications. Whatever the outcome, use that information as feedback to refine your process.

If you are not accomplishing your recovery time, recovery point, budget, or business objectives, you can tune the plan and the technology. Most important, as your business changes, your needs will change – and you should adjust your business continuity plan accordingly. This is not a process to put in place and forget about until a disaster occurs – review your plan and process from time to time to ensure that you are fully prepared.

Revisit the Plan

If you are not accomplishing your recovery time, recovery point, budget, or business objectives, you can tune the plan and the technology.

Summary

Disaster recovery and business continuity are important business issues that require awareness and planning. Below are some guidelines for choosing the technology or solution that's best for you.

- Look at the big picture – your business processes, systems, networks, data, and people all need to be considered when planning and implementing these processes.
- Understand your levels of tolerance for lost work, missing data, and unproductive time.
- Document and test your plans, and update them when business needs change.
- Configure your environment to minimize the likelihood of a failure escalating into a disaster.
- When evaluating technology solutions, take into account meeting your recovery objectives, kinds of disasters you're likely to face, and levels of cost, complexity, and disruption involved.
- Know the advantages and limitations of each technology, and adjust your expectations accordingly.
- Remember that backing up your data is the most reliable form of protection, without which your business is vulnerable.

For more information regarding the EqualLogic and the PS Series, please visit www.equallogic.com or contact us at 888-579-9762 ext 7792



