White Paper

Data Protection and Disaster Recovery Planning

Frequently Asked Questions (FAQs)

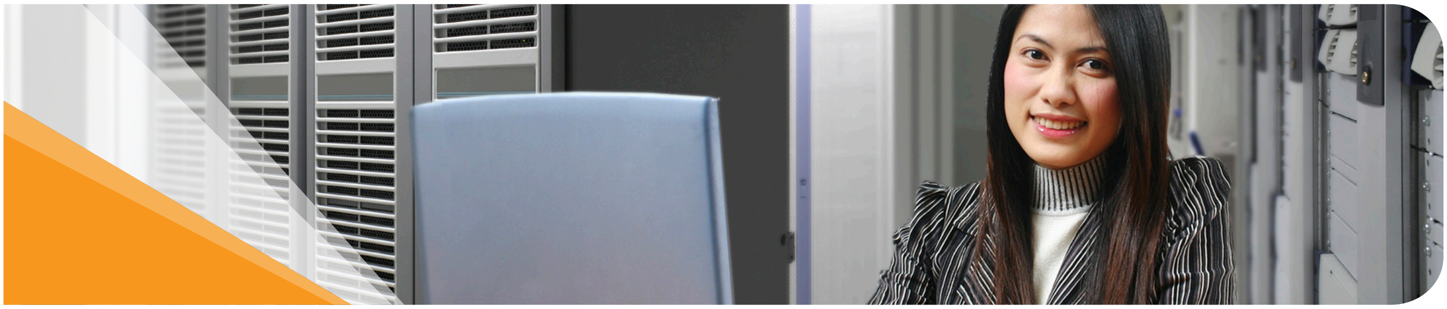

Introduction

Data Recovery is somewhat similar to life insurance – you invest your money today, to overcome a possible disaster in the future. Just like insurance, Disaster Recovery Planning is a complex and constantly changing subject. There are many options, and there are many highly-paid consultants who often have different interests than their clients. Therefore, it is no surprise that this complexity generates myths – many of which are erroneous. This white paper answers some of the most frequently asked questions in this area of technology.

1. Do we need a disaster recovery plan?

Do you have a life insurance policy? Do you have car insurance? Is your business protected with fire or theft insurance? Then why wouldn't you insure your data? Most businesses agree that their corporate data is the life-blood of the company.

We all wish to avoid disasters, but unfortunately, they do occur. In today's information age, only a few companies can afford to lose a significant amount of data or stay offline for more than a short period of time. Customers expect most companies to be able to resume service within a very short amount of time.

Research indicates that even a very short amount of downtime can have significant adverse affects on a company. Revenues are impacted. Customers are impacted. Company reputation is affected. Longer periods of time can easily shut down a business. Statistics vary by industry. However, the trend lines are all the same.

Corporate IT infrastructures, and more specifically corporate data, are critical to business operations. In most environments, almost all data is considered critical. Often, email is one of the most critical applications. Therefore, disasters are often fatal for businesses with no Disaster Recovery plan. While the size and level of detail may vary, most businesses must have some form of a Disaster Recovery Plan (DRP) in place.

2. Aren’t disaster recovery plans extremely expensive?

As mentioned above, the size, complexity and level of detail varies dramatically in DR plans. As with most plans, it is important to identify the objectives and goals and then build a plan to meet those goals. Consider what you really need, and how much you are willing to pay.

For example, one key factor to consider is your Recovery Point Objective (RPO) – which identifies the number of recovery points you can recover from. You may want to consider this on an application-by-application basis. In a simple DR plan, you may have a recovery point objective of no more than 48 hours. In other words, you can recover from any failure but not lose more than 48 hours worth of data. In other environments you may decide you cannot afford to lose more than 5 minutes worth of data. It is important to understand your RPO.

Another factor to consider is your Recovery Time Objective (RTO) – which identifies the amount of time it takes to recover. Stated differently, how much time can you afford to be offline.

Once you set your objectives, then you can begin to apply technology to meet those objectives. The sections below describe how the cost of a DRP can be dramatically reduced, to an affordable one.

3. Must the DR site be a physical mirror of the main site?

This is not typically true. It is only a very small percentage of DR implementations that require a complete physical mirror of equipment, networks, people, and other resources at the DR site. Not all applications and operations at the main site are mission critical. Therefore, not all of them have to be restored immediately. Some applications and operations can afford a longer RPO and RTO than others.

For example, if an internal R&D team develops applications for internal use, this activity is far from being critical. The DR site may therefore exclude the storage capacity, servers and working environment for this team.

Once priorities have been established for each application and operation, equipment and resources for the DR site can be identified. In most cases, the main site is well–equipped, providing adequate capacity to handle peak-loads plus some excess. However, the DR site typically does not need to be configured to support extreme loads. This allows the DR configuration to be far more cost-effective than the primary site. If a disaster should occur, the DR site will be fully functional for the short period of time it is required to support the operation. The goal is to return operation to the primary site as quickly as possible. If it is not possible to return to the primary site in a short amount of time, and peak periods are expected, upgrades to the DR site can be made.

One could argue that the storage applications responsible for mirroring the data to the DR site, operate only with the same storage systems, and therefore, lower cost solutions would not be available. This was indeed true several years ago when storage vendors held more of a closed system mentality. Today, however, several independent software vendors offer mirroring between heterogeneous storage devices, providing the option to choose the most suitable storage system for the DR site – regardless of the vendor. Additionally, the efficiency of the mirroring solutions support mirroring from high-performance devices to lower-performance solutions.

4. Is the DRP only for physical disasters?

No. If this was the case, you would not be able to recover from a massive virus attack or a total data corruption due to a rogue application.

Research has shown that more than 93% of errors are “logical” failures rather than “physical” failures. While some of these errors are not considered “disasters”, others (e.g., viruses, data corruption, or accidental file deletion) can be just as devastating as a physical catastrophe. The net result is the same – the business is taken offline.

A good data protection plan needs to enable rapid recovery from both physical and logical failures at the main site, while guaranteeing that logical errors are not propagated to the DR site.

The most common and effective protection against logical failures is the use of snapshots. Several vendors now offer low-capacity snapshots that consume storage space only for changes made after the most recent snapshot. These snapshots can be used to recover within seconds, eliminating the need to spend hours recovering from traditional tape (or disk) backup solutions.

5. Do synchronous mirrors provide better protection than asynchronous mirrors?

Both synchronous mirrors and asynchronous mirrors have advantages and disadvantages. The best solution depends on what you are trying to accomplish. With a synchronous mirror, the data at the DR site is identical to the data at the main site. However, there is no guarantee that the application using this data will be successfully restored. For example, databases cannot always resume operations from any point-in-time snapshot of the data. They typically require data integrity (data consistency) in order to resume operation. Some database applications have time-consuming internal procedures to restart with non-integrity data, and statistically are not entirely reliable.

There are some vendors that offer asynchronous mirroring solutions that are based on low-capacity snapshots. At pre-defined time intervals a snapshot is created that collects the changes. Those changes are transferred to the DR site. The low-capacity snapshots are standard snapshots that can be used by servers on either the main or DR sites. They also provide protection against logical failures (see above).

It is suggested that you pick a vendor that offers both synchronous and asynchronous mirroring to give you the flexibility to apply the technology that best meets your needs.

6. How far should the DR site be located?

Recent events such as September 11th, the Tsunami and Hurricane Katrina show that a DR site, close to the main site may not be a sufficient approach. In these cases, entire regions were hit by disasters and the only way to resume operation was from a DR site far away from the main site.

The use of asynchronous mirroring can extend a mirror to the opposite side of the earth. Asynchronous mirroring solutions are very efficient and can work over low-bandwidth IP lines. Some vendors are even offering this option without any further hardware. On the other hand, synchronous mirroring usually requires Fibre Channel (FC) lines that are indeed limited to roughly 10 kilometers. Nevertheless several vendors are offering today FC extensions that can reach more than 1000 miles.

Again, it is recommended that you pick a vendor that offers both synchronous and asynchronous mirroring to give you the flexibility to apply the technology that best meets your needs.

7. Can the DR site be used while the mirror is taking place?

Yes. If this were not true, the DR site would be a tremendous waste of resources the vast majority of the time. In business, the goal is to get the maximum return on your resources. Having servers, storage, networking and human resources all sitting idle waiting for a failure would not be a good use of resources. On the other hand, if you could put these resources to work, you can begin to pay for the DR plan over time.

Example #1 — Remote Backup: Backup tapes are usually shipped to a remote location, in many cases they are shipped to the DR site. Moreover, backups consume system, storage, network, SAN and LAN resources. Since the DR site has an updated version of the data, backups to tape can be accomplished at the remote location, freeing those resources at the primary location.

Example #2 — Development and Test Teams: Many organizations have R&D teams that are developing and testing applications and utilities for internal and external use. These teams often require an updated version of the data to perform tests. Since the DR site already has the necessary equipment and the required data, it can be put to use rather than buying more equipment for the main site.

These are just two of many examples of how the DR site can be used if the DR tools support the reading and writing of data. Make sure your solution supports this feature.

Summary

Implementing a DR Plan is much like having insurance – the business can operate without it, but the risk is just too high. Data protection is a very real requirement in the market and many new solutions are now available that provide advanced functionality not previously available. Many of the myths from years ago prevail in discussions, but are no longer true. Problems that not long ago were considered unsolvable (or too expensive to solve) now have affordable solutions. By looking beyond the myths and choosing the right solutions, every business can create a cost-effective DRP that best matches their requirements.

This document is provided by SGI,
an authorized reseller of LSI Corporation

For more information and sales office locations, please visit the LSI web sites at: lsi.com lsi.com/contacts

North American Headquarters
Milpitas, CA
T: +1.866.574.5741 (within U.S.)
T: +1.408.954.3108 (outside U.S.)

LSI Europe Ltd.
European Headquarters
United Kingdom
T: [+44] 1344.413200

LSI KK Headquarters
Tokyo, Japan
Tel: [+81] 3.5463.7165

LSI Corporation reserves the right to make changes to any products and services herein at any time without notice. LSI does not assume any responsibility or liability arising out of the application or use of any product or service described herein, except as expressly agreed to in writing by LSI; nor does the purchase, lease, or use of a product or service from LSI convey a license under any patent rights, copyrights, trademark rights, or any other of the intellectual property rights of LSI or of third parties.