

Leveraging IP Storage for Disaster Recovery

With an Introduction by
Barb Goldworm
Focus Consulting



The State of the Industry: Disaster Recovery

While disaster recovery planning has long been a topic in enterprise data centers, its rise to national headlines after the September 11, 2001 attacks brought it from a discussion within enterprise IT organizations to one amongst executive management of U.S. businesses of all sizes. Disaster recovery (DR) and data protection became top priorities. Over the past three years, with ongoing threats of terrorism along with numerous natural disasters and widespread power outages, there has continued to be strong focus on these areas.

According to a recent InfoWorld study, disaster recovery and data protection topped the list of factors driving storage spending for 2004. In addition, data availability and recovery were listed as the top two storage management challenges by nearly half of the respondents.

Storage Networks

A 2004 SNIA/Computerworld study asked IT managers to rate various reasons for implementing a storage network. Improving DR and business continuance was rated most important by more IT managers than any other reason. Next highest was improving backup and recovery. Clearly, DR and data protection are significant drivers in the move towards storage networks.

The past three years has seen tremendous growth in storage network implementation in the enterprise. Fibre Channel (FC) Storage Area Networks (SANs) and FC Arrays are now a critical part of most enterprise storage environments. With SANs, come a variety of features and benefits geared towards improving availability and manageability of storage. Hardware and software functionality including mirroring, RAID levels, snapshots, synchronous and asynchronous replication are all options for the storage manager trying to ensure data protection and business continuance.

IP Storage

For those in the Small to Medium Business (SMB) space, however, the costs and complexities of Fibre Channel have kept most from being able to take advantage of SANs. Fortunately, changes over the past several years have now brought a level of stability and confidence in SANs based on IP technology, using the iSCSI protocol. A recent study by Storage Magazine showed 45% of their readers either have implemented or are planning to implement iSCSI. According to the Enterprise Strategy Group, there are over 2000 production iSCSI implementations today.

IP SANs bring many of the advantages of Fibre Channel SANs without the cost of Fibre Channel or the need to learn a new architecture and set of protocols. Some offer many of the high-end capabilities previously available only with the high-end FC products. With this type of high-end functionality coming with IP SANs, coupled with the ease-of-use of IP storage, SMB users have new options for successfully implementing reliable data protection and disaster recovery.

Glossary

Data Protection

The ability to protect against data loss or data corruption within a single site using features such as backup, snapshot, and hardware redundancy.

Disaster Recovery

The strategy and complete set of procedures required for a business to deal with the effects of a disaster that affects an entire data center. Disaster recovery planning involves identifying essential services and information, potential types of failures/disasters, and how to resume business operations in the event of a disaster at an alternate location. A DR plan includes the precautions taken to minimize the effects of potential disasters.

High Availability

The ability of a system to perform its function continuously (without interruption) for a longer period of time than its individual components typically would suggest. Generally achieved through fault tolerance and redundancy of hardware components.

Business Continuity

The need for maintaining continuity of all critical business operations in the event of any type of failure, up to and including disasters that render complete IT sites unusable. Planning for business continuity should include consideration of both high availability and disaster recovery.

Considerations

While IP SANs bring many features that can improve data protection and DR plans, there are critical questions that must be considered as part of the DR planning process. DR planning should define the strategy and complete set of procedures required for a business to deal with the effects of a disaster. Disaster recovery planning involves identifying essential information and services, recovery objectives, potential types of failures/disasters, and how to resume business operations in the event of a disaster. A DR plan includes the precautions taken to minimize the effects of potential disasters.

While not a full guide to DR planning, this section describes key issues to consider in planning a DR strategy using IP storage.

Data Classification

The planning process starts with evaluating data into tiers, reflecting its importance to the ongoing business of the company. (Data classification is also a key element in Information Lifecycle Management (ILM), as part of the goal of storing varying tiers of information on the appropriate type of storage.) Categories can be defined based on individual business needs. Examples might include the following:

- Mission critical – Crucial to the organization's primary business processes. Critical and time sensitive recovery. (e.g., order entry, sales)
- Essential – Very important to day-to-day business. Company intellectual property. Recovery is critical, but less time sensitive. (e.g., customer info, e-mail)
- Important - Valuable to many daily operations. (e.g., employee data)
- Low-critical – Nominal to low organizational value. (e.g., market data)

Recovery Objectives

For each classification of data, key DR objectives should be defined. Objectives should include a Recovery Time Objective (RTO) and Recovery Point Objective (RPO). RTO identifies the maximum acceptable time after an outage in which recovery will be accomplished. RPO is the maximum acceptable age of the data, when recovered (i.e. the point-in-time to which systems and data must be recovered, within the DR facility). Together, RTO and RPO become the basis for compromise between the business objectives of minimum down time and currency of recovered data versus the costs associated with DR strategies.



Glossary continued...

Backup

The regular activity of copying files/databases to disk or tape to preserve them, as of a reasonably recent point in time. Used to restore data in the event of user errors, hardware/software failures or as part of disaster recovery.

Recovery Point Objective (RPO)

Used in DR planning. The maximum acceptable age of data at the time of an outage, thus determining the age of the data when it is restored.

Recovery Time Objective (RTO)

Used in DR planning. The maximum acceptable length of time to resume operations following an outage.

Types of Loss

DR planning should also include consideration of various types of loss including equipment loss (of a single disk drive, disk controller, or array), loss of an entire site/facility, and loss of an entire region. These considerations then factor into how much fault tolerance and redundancy is required for various tiers of data. It is also useful to consider the likelihood and frequency of various types of failure, and to evaluate where DR spending is most critical. For example, one strategy to consider is investing more money in the most likely failures, while creating reasonable plans for those least likely.

Storage and Bandwidth Requirements

Part of DR planning needs to include how much data is in each category, and how much change occurs to that data. These answers will help determine storage requirements for various types of storage, taking into account multipliers for replication and growth requirements for snapshots.

In addition, if replication across distance is required for offsite copies (e.g., to the DR site), the amount of data to be sent must be considered, and bandwidth needs must be evaluated. (This is a factor both for geocustering and synchronous replication as well as for asynchronous replication). Gigabit Ethernet upgrades may be required for reasonable performance and time windows.

Backup Window Requirements

As part of DR planning, it is important to think through backup and data protection requirements. What time period is reasonable for a backup window? Many mission critical applications today demand 24x7 availability, calling for a zero backup window. Snapshots are highly valuable here, since they can virtually eliminate the window, allowing replication and/or backups to be done from snapshots.

Features

Keeping all these considerations in mind, it is then possible to evaluate various technology features and their benefits, and how best to accomplish the DR and data protection goals that have been established. These technical features and options can be used individually or in various combinations, to provide varying levels of protection for various tiers of information. Choosing products that include strong implementations of these features provides the greatest flexibility, allowing progress towards comprehensive DR plans to occur in small steps. See Table 1.

Table 1: Data Protection and DR Features and Benefits

	High Availability	Duplicate Copy of Data	Offsite Copy of Data	File Recovery	Immediate Full Data Recovery
Snapshots	X			X	X
Asynchronous Replication	X		X	X	X
Synchronous Replication	X	X	X1		
Mirroring	X	X	X1		
Tape Backup		X	X	X	
Hot Swappable Components	X				
Hot Standby	X				

*X1- if replicated offsite

Snapshots

A snapshot is a point in time copy of a file system or storage volume, which typically contains pointers to the original blocks on a device. The advantage of using snapshots is that large datasets can be snapped with a small (seconds-long) freeze of the dataset at a point-in-time. Once a snap is taken, any changes to the data blocks will be tracked and logged along with the before-image of that data. Retrieval is done transparently to the user. Another advantage is that snapshots provide protection while using a small amount of incremental storage.

Snapshots can facilitate backup by allowing a copy to be made quickly, resuming activity to the real data, shortening the backup window. A full backup copy can then be made from the snapshot. Snapshots are also useful for recovery from data corruption and file loss. Snapshots can be copy-on-write (low capacity) or split mirror (which simplifies recovery, but is slower and requires more storage space per snapshot).

Snapshots are one of the key features enabling both smarter backups and DR. Since different products perform snapshots differently, the snapshot technology used by a product should be evaluated and understood. When used well, snapshots can eliminate the backup window, and provide a base for integrating with asynchronous replication and tape backups.



Mirroring, Replication and Tape

Part of DR planning is considering how many copies of the data are necessary, how recent they must be, and where they should be maintained. Options to consider include mirroring, replication (synchronous and asynchronous) and tape.

For mission critical data, mirroring provides a high availability option where data is duplicated on two separate disks within an array to maintain fault tolerance (Raid 1). While mirroring provides fault tolerance at the drive level, there is still a single point of failure at the controller level. Since mirroring doubles the amount of storage, cost is a key consideration.

Synchronous replication provides another option for mission critical data, also writing the data simultaneously to multiple disks. Synchronous replication refers to a process where all data is committed to both sets of storage before the client that wrote the data is informed the data has been written. With synchronous replication, if the disks have independent controllers, replication can offer a level of fault tolerance over and above mirroring, since it avoids the controller as a single point of failure. It is important to consider that both mirroring and replication require multiple times the amount of storage. This is clearly an area to weigh the importance of the data to the organization along with the likelihood of various types of loss, against the costs of multiple copies.

Asynchronous replication is a form of replication which disconnects the copy process from the process of writing of the data to the primary location. Asynchronous replication is an excellent option for providing an offsite copy of data for a DR site or for centralized backup to tape. Combined with snapshots (eliminating the backup window), async replication can leverage existing IP links to facilitate both data protection and DR. Use of existing IP links obviously will require a close evaluation of bandwidth requirements to ensure an acceptable impact on existing applications as well as a reasonable and practical implementation for replication.

While replication (both sync and async) provides good options for DR, it is not a replacement for tape. For historical data preservation, file recovery back to a certain point in time, and storing data offsite, tape backup continues to play a role in best practices. Where there is a good backup strategy in place, integration of snapshots can improve the process. In addition, async replication from distributed sites to a central site can allow centralized backup to be automated and well managed. While incorporating these options will not implicitly make a bad backup process into a good one, if used well, they can offer significant improvements.

Introduction by Barb Goldworm

High Availability Features, Failover and Failback

DR planning for protection of mission critical data and optimization of business continuance must involve evaluation of high availability options as well as failover and failback capabilities. High availability (HA) is the ability of a system to perform continuously (without interruption) for a longer time than its individual components typically would suggest. HA can be achieved through fault tolerance, redundancy of hardware components and clustering. Clustering options add both fault tolerance and scalability, allowing additional processing power to be added and included in intelligent load balancing. Failover takes effect when a failure occurs, and the system responds by switching over to a working system/component.

Evaluation of hardware and software should include the following types of questions. How will the system respond to a disk failure, a power failure, a controller failure, etc...? What are the automatic failover options? If a disk is lost, and the system continues operation, what is the performance degradation? Is there an option for a hot standby? In addition to a full hot standby, what are the options for a hot swap? What manual intervention is required, if any? What is the process for failback to resume normal operations? Are the options configurable at a volume level to allow different classes of data to run with different features?

Microsoft Integration

In general, the features described in this paper are available today, to varying degrees, in a number of IP storage products. Over time, as the foundations for these services continue to be incorporated into Microsoft Windows Server, it will be increasingly important for products to integrate with and take advantage of Microsoft Windows capabilities. Particular areas to watch in terms of this integration include Volume Shadow Copy Service (VSS) for point-in-time copy, Virtual Disk Service (VDS) for provisioning, and Multi-path I/O (MPIO) for failover and load balancing.

Conclusions

While many users have, up to this point, viewed DR and HA as high-end options that were out of reach for their size, staffing abilities and budgets, IP SANs are starting to change the landscape. With the ease of use of IP SANs and the high-end functions that are moving into them, DR and HA capabilities can be a part of the plan for SMB environments. Mid-range users today have successfully implemented DR plans using IP storage via snapshots, synchronous and asynchronous replication, and tape, in various combinations.

Knowing that a DR plan is possible, it is highly worthwhile to talk with other users with similar characteristics about what they have done, where they have succeeded and what challenges they have found. Ask your vendor for users who have successfully implemented DR with their products, and learn from their experience. Armed with experienced-user information, the considerations described here, good planning, and regular testing, you can join the ranks of both enterprise and (now) mid-range users who sleep better at night, knowing that they have a reasonable plan, should the worst happen.

Data Protection and Disaster Recovery with the LeftHand SAN

The LeftHand SAN architecture was designed with great flexibility to allow the user to decide how and to what extent data should be protected. Data protection and DR scenarios can range from no fault tolerance to a fully mature disaster recovery implementation, all built on the existing Ethernet infrastructure.

The LeftHand Networks IP SAN integrates a combination of synchronous replication, remote copy, snapshots and clustering to achieve varying levels of data protection within a single site or disaster recoverability across multiple facilities. Storage administrators have the flexibility to decide which scenario to use on a per volume basis. Based on data classifications and recovery objectives, volumes can be configured with varying degrees of protection. Data that is less valuable and has a lower fault tolerance priority can exist in the same SAN as data with the highest fault tolerance priorities.

The Technology Behind LeftHand's Data Protection and DR Capabilities

Overview

The LeftHand SAN is a complete networked storage solution comprised of LeftHand's SAN/iQ distributed software and standards-based storage servers. Administrators purchase any number of storage servers, called Network Storage Modules (NSMs), to achieve the capacity or spindle count required for the applications that will utilize SAN storage. The SAN/iQ software clusters the storage servers together into a single pool of storage, creates snapshots, and manages the replication levels for each volume. Volumes, or logical disks, are provisioned from the pool of storage and appear to the application server or client as a block-level storage device. See Figure 1.

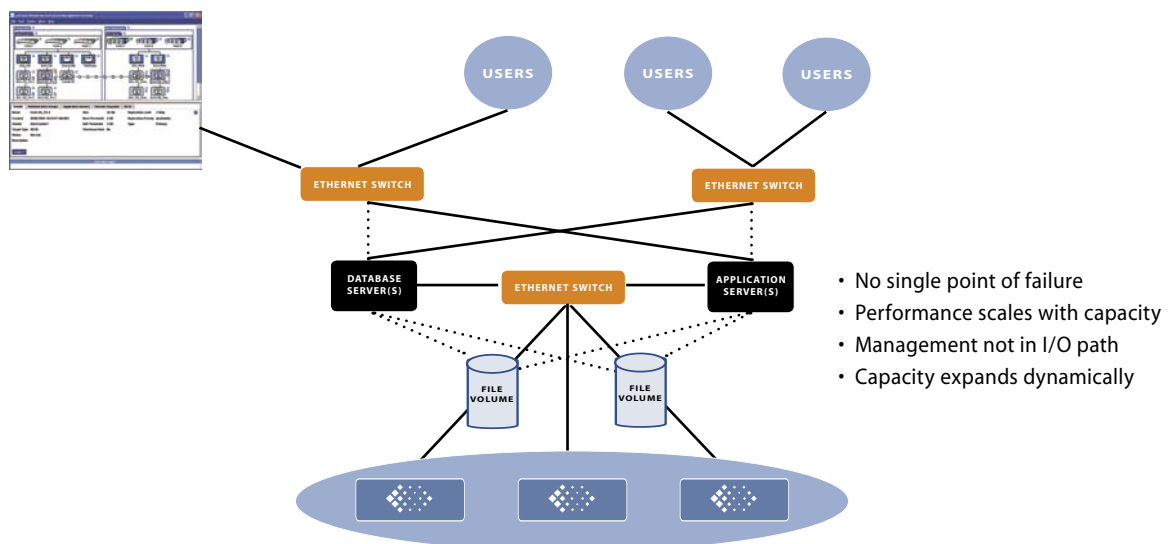


Figure 1. LeftHand Networks has implemented a second-generation distributed clustering technology to provide redundancy in an IP SAN.

Clustered Storage

The design of the NSM itself is fault-tolerant, with dual NICs, power supplies and RAID. Availability is maximized by adding multiple NSMs to form a cluster. The intelligence that enables clustering resides on each storage server; this unique distributed architecture eliminates the single point of failure common with other virtualization techniques. The cluster can survive the loss of a storage server, and additional functionality such as local replication can be enabled within the cluster. In addition, because each storage module includes its own processor, memory and network interfaces, performance scales as each new NSM is added to the storage pool.

RAID

RAID 0 and RAID 10 are supported within an NSM. If a site is implementing a single NSM, RAID 10 should be configured for data protection. RAID 10 is data striping plus mirroring so two copies of the data will reside in the single NSM. This configuration can withstand a disk failure.

If running a multiple NSM configuration, RAID 0 with replication provides the highest level of data protection and performance. RAID 0 plus replication will utilize RAID 0 (striping) for performance and LeftHand Networks' unique replication configuration for data protection. This configuration can withstand a failure of any component in the SAN and is superior in protection and performance to mirroring.

LeftHand's synchronous replication is fully automated within the SAN, and provides availability across NSMs, as opposed to the more common implementation of replication across sites. See the section "Synchronous, or Local, Replication" for a complete description of this feature.

Hot Spare

To achieve a highly fault tolerant environment, a hot spare can be configured in the storage cluster. A hot spare is a NSM that participates in the cluster only in the case of a failure. Should a module fail for any reason, the hot spare automatically takes over.

Snapshots

A snapshot is a point-in-time copy of a volume. Snapshots provide the capability to do file-level restores and to roll back to a point in time where the data was in a known "good state". LeftHand's revolutionary snapshot technology is different than any other currently implemented today. When a snapshot is taken within SAN/iQ, the volume currently being written to immediately becomes a read-only volume, and a new read-write volume is created. The new volume maintains any changed blocks and the read-only volume remains in a constant state.

Should an administrator need to go back to a point in time, for example to recover a critical file that was deleted, the read-only volume can be mounted to any server. The file is recovered by copying it back to its location before it was deleted. All of this can be done in a matter of minutes without disruption or downtime. See Figure 2.

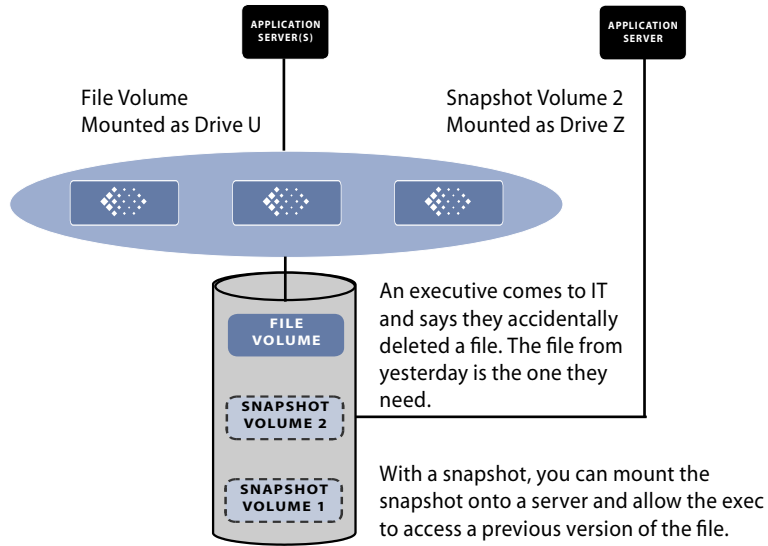


Figure 2. Attaching a snapshot for quick file recovery.

In the case of data corruption, an administrator could roll-back to a point-in-time where the data was in a known good state. An administrator would remove – or delete – volumes, that are corrupt, thus bringing the volume back to a point in time before the corruption occurred. For example, assume a snapshot occurs each morning at 3 a.m. At 8 a.m., the volume is infected with a virus, destroying all the information.

To recover from the virus, the administrator would roll the volume back to the 3 a.m. snapshot, prior to the virus, and restore the volume to a good state within minutes. During the rollback, the changed bad blocks in the read-write volume would be deleted. The data from the snapshot taken at 3 a.m. would replace the current degraded volume to become the new read-write volume, free of data corruption. See Figures 3, 4 and 5.

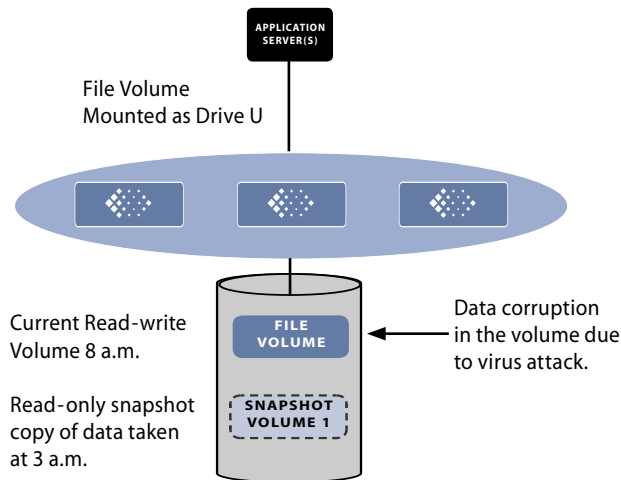


Figure 3. Depiction of a read-write volume and associated read-only snapshot.



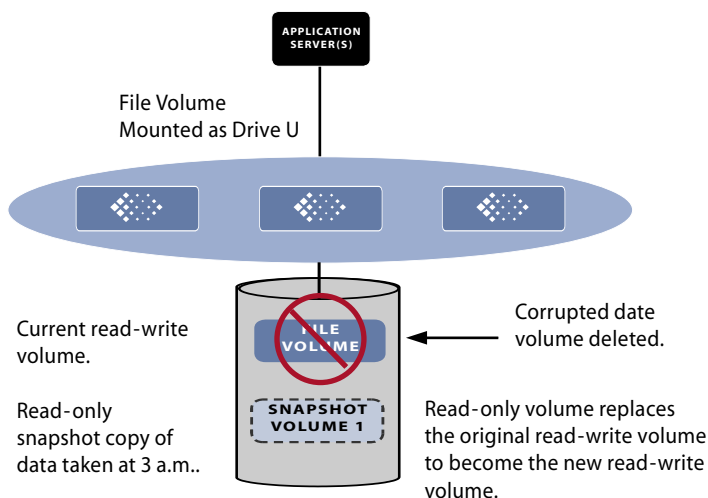


Figure 4. Perform a rollback to delete the corrupted volume.

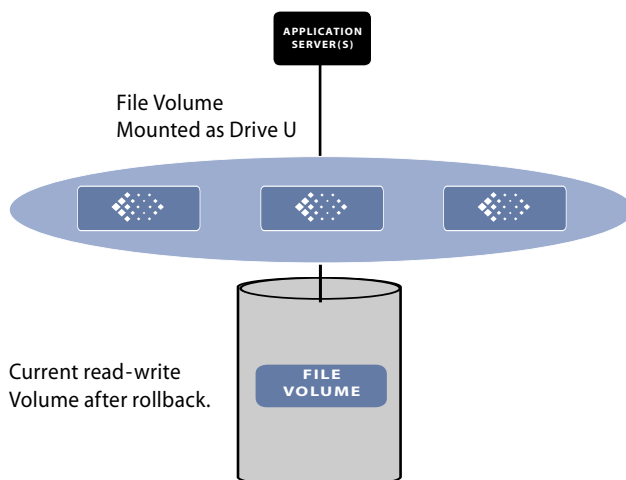


Figure 5. The new read-write volume is ready to access.

Synchronous, or Local, Replication

The core technology behind SAN/iQ's data availability is a replication-based redundancy scheme called chained declustering. Chained declustering offers superior protection and performance over RAID 1 mirroring because it evenly spreads data blocks across NSMs in a storage cluster, ensuring that the loss of an NSM will not result in loss of data. In contrast, RAID 1 mirroring can only be configured within one storage module. If the module experiences a failure, access to data is lost. Through chained declustering, LeftHand Networks has eliminated any single point of failure, creating an extremely fault tolerant data storage system.

SAN/iQ manages the replication-based redundancy at the volume level, thus for *each volume* created an administrator selects the replication level. Replication is synchronous and can be set to one of three ways – no replication, two copies of data, or three copies of data. Figure 3 below demonstrates the block level write scheme for each replication level in a cluster of four storage servers. The ability to choose replication by volume optimizes storage utilization, using additional storage only for that data that warrants the extra protection. See Figure 6.




Figure 6 - Chained declustering volume-level redundancy scheme.

- **B1/B2/B3/B4:** Identifies data storage blocks being written to storage pool
- **No replication:** Because each block is written once, only one copy of the data exists in the storage pool, providing no data redundancy.
- **Two-way replication:** Because each block is written twice, two copies of the data exist in the storage pool, providing fault tolerance against the loss of a single NSM, or two non-consecutive NSMs.
- **Three-way replication:** Because each block is written three times, three copies of the data exist in the storage pool, providing fault tolerance against the loss of any two NSMs.

In the example presented above, if an administrator has four NSMs in the storage cluster and chooses two-way replication, the SAN can sustain the loss of one NSM. As illustrated, identical blocks are written synchronously across multiple NSMs in the storage cluster. Since more than one NSM contains a copy of a data block, loss of an NSM does not equate to loss of data access. Three-way replication offers double the protection of two-way replication.

Replacing a failed NSM in the cluster after a loss is a simple process that requires no downtime or disruption to users. As storage modules are added to the storage pool, the data is “re-striped” across the cluster, according to the replication level specified, with no administrator involvement. This ensures an even distribution, eliminates “hot-spots”, and increases performance. This is a radically different approach than a DAS or NAS environment where all data movement associated with adding additional captive storage or a new NAS box requires the administrator to manually move and balance data among the servers.

The option to set the replication level on storage volumes is part of the basic SAN/iQ install and does not include additional costs or licenses.

Remote IP Copy

The SAN/iQ Remote IP Copy technology is built off of the snapshot implementation. When a snapshot is taken, a read-only volume is created, as described in the “Snapshots” section above, containing only the changed blocks. This read-only snapshot volume is then copied, block for block, to a designated cluster, which can reside anywhere IP communication is possible. The remote copy is an asynchronous operation which can use existing IP infrastructure. By capturing and copying only changed blocks, SAN/iQ limits the amount of bandwidth required.

Administrators create and schedule Remote IP Copy jobs by selecting a specific start time, occurrence interval, and retention policy. Understanding bandwidth limitations and the amount of data to be copied is imperative to creating a successful Remote IP Copy implementation and to meeting organizational data protection and recovery mandates. As with all other features, SAN/iQ manages Remote IP Copy at the volume level. SAN/iQ allows administrators to select the individual volumes to copy, frequency to perform the operation, and retention policies

Remote copies can serve many purposes depending on their intended use. Typical uses include disaster recovery at a failover site, off-site or centralized backups, and a “split-mirror” configuration for data migration and content distribution.

In the next example – see Figure 7 – a business has two sites of operation. Remote IP Copy is being used between the sites as both a means of disaster recovery and for off-site backups. In order to establish a DR site, the primary snapshot in Site 1 is being remote copied to a remote snapshot in Site 2, and the primary snapshot in Site 2 is being remote copied to a remote snapshot in Site 1. This creates an identical read-only environment in a geographically separate location. Should a disaster take down Site 1, the data that was remote copied to Site 2 is ready to be promoted to primary, read-write volumes and resume activity. Backups are being performed by doing a remote copy of the primary snapshot at Site 1 to a remote snapshot at site 2. The remote snapshot, made up of Site 1 data, and the Site 2 primary snapshot, made up of Site 2 data, are backed up to tape at Site 2. With this scenario, not only are there redundant copies of data and means to a quick and simple recovery, but there is also no need to configure another backup environment.

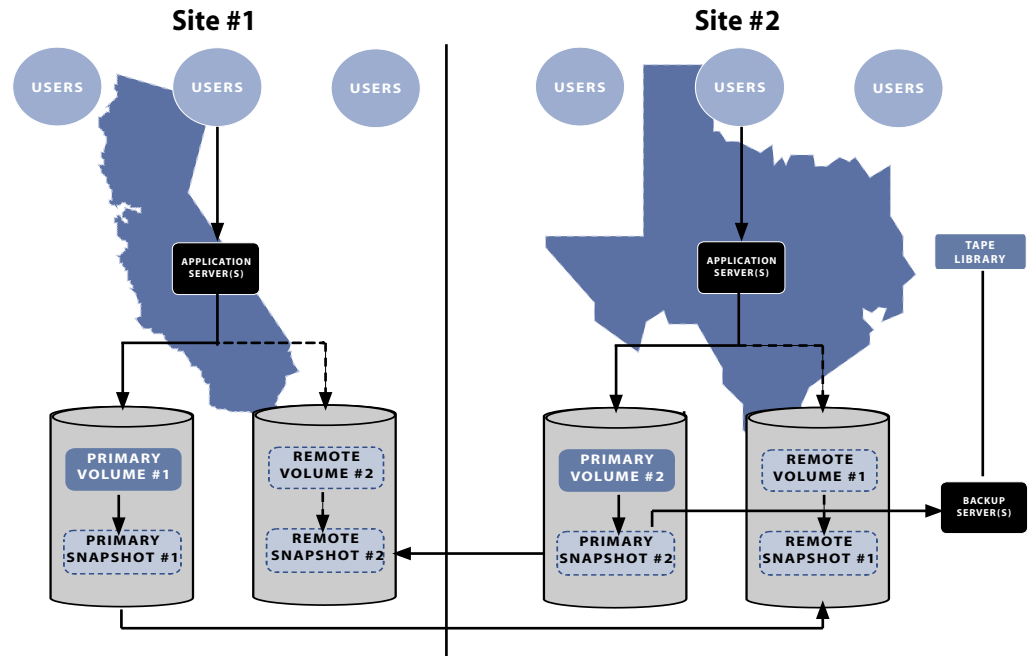


Figure 7. Graphical representation of Remote IP Copy.

Customer Data Protection and Disaster Recovery Scenarios

End users can combine the features described above – clustered storage, snapshots, synchronous replication, and remote copy – to build a data protection/DR scenario that’s appropriate for their environment. Companies with a single site can build a highly redundant and available environment utilizing sophisticated yet easy-to-use technology. Companies with multiple sites can add remote copy functionality to move backup offsite, centralize backups, or build a secondary site to protect against complete site failure. Below are some real-world customer configurations leveraging these technologies, ranging from simple to complex, and the process to recovery. Scenarios begin on the next page.

Scenario 1- Data Protection within a Single Site

A small, budget-conscious, privately owned business needed to consolidate their direct-attached storage (DAS) environment. Having “islands” of SCSI disk was making backups difficult to complete, growing their storage was a weekend-long event, and recovery from tape took up valuable network resources during periods of heavy user load. The solution had to be affordable, scalable, easy to manage, and provide data protection. The following solution addressed their requirements:

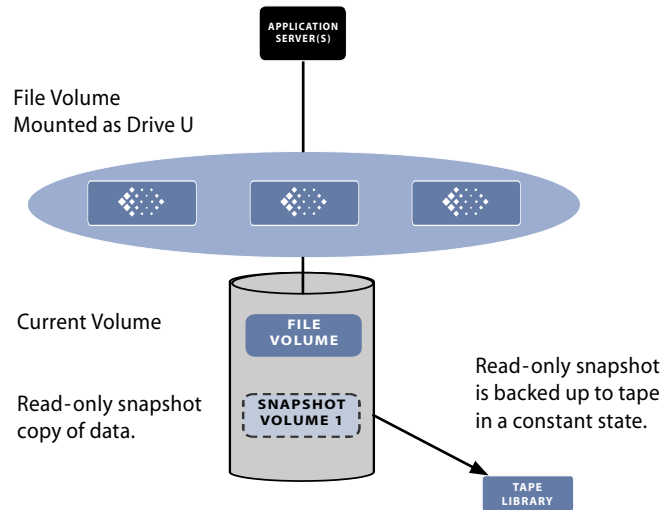


Figure 8. LeftHand SAN designed for small business.

The above design offered the following benefits:

- The customer was able to move completely off of direct-attach storage, simplifying their storage management by allowing them to manage all their storage through a single, intuitive GUI.
- They implemented snapshots to create a read-only, impenetrable copy of their data. The customer can perform a file-level restore in the case of accidental file deletion or a complete rollback in the case of data corruption.
- The customer now has restore options without having to go to tape backups.
- Using snapshots alleviates their backup problem. Because snapshots create a read-only copy of data, that copy can be mounted on the backup server, removing the load from the application server and the network.
- Two-way replication within the cluster allows for the loss of a storage server with no data loss or availability.

Scenario 2 – Centralized Backup between Multiple Facilities

Businesses with multiple facilities face a unique set of challenges surrounding backup and recovery. Frequently, remote sites don't have resident IT staff, impacting the types of backup and recovery procedures that can be put in place and reducing the success rate of backup. By centralizing backup, a university was able to utilize the IT staff at the main data center to manage backup/recovery and consolidate backup hardware. The same remote copy capabilities also allowed the school to affordably maintain multiple copies of data in multiple remote sites. In the case of a disaster, user and application traffic can be re-routed to an available copy of the data, allowing business to continue to function.

Benefits of this scenario include:

- Using Remote IP Copy, data from the satellite campuses is copied to the main data center for backups.
- The copies of data are mounted to the backup server at backup time and the data is archived to tape. This alleviates not only the need for personnel to monitor backups at remote sites, but also speeds up the backup process by taking it off of the network and offloads the backup process from the application servers.
- The university consistently maintains reliable backups, completed using the higher-end tape library available at the main data center.
- The university's main goal was centralizing backup, but they now also have a secondary copy of data at the main data center which can be utilized in the case of disaster.

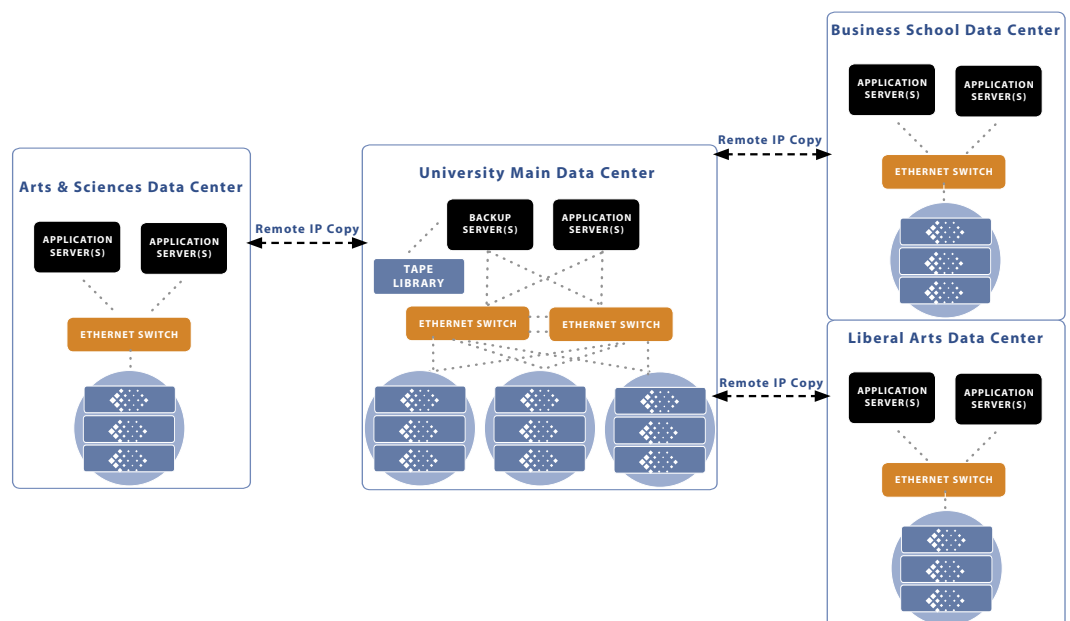


Figure 9. Multi-site centralized backup infrastructure.

Scenario 3 – Protection Against Full Site Failure

A city government needed to move away from DAS and find a disaster recovery solution that would meet the needs of their entire infrastructure. Their applications included a mail server and several databases that store important information about the city government and emergency services. In addition, because of their direct attached storage environment, backups weren't completing in the window and were bogging down the network for users around the city. The solution had to offer highly reliable data availability, protection and recovery. The following SAN was designed and implemented – see Figure 10.

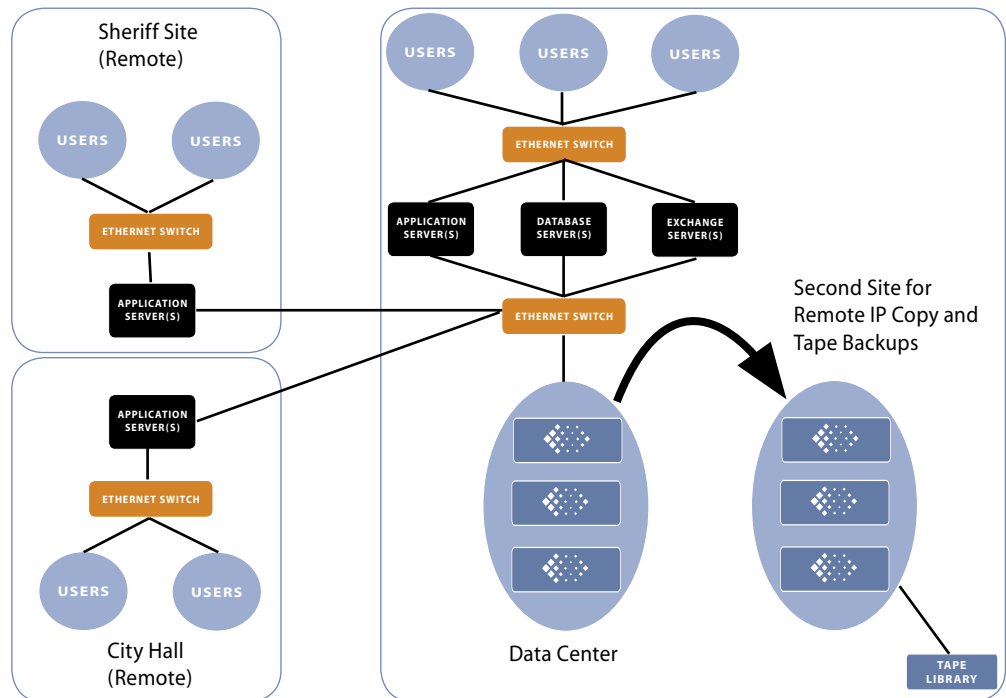


Figure 10. LeftHand SAN designed for a local city government.

The design depicted above offered several DR options to the customer:

- Synchronous replication is used within the main data center, ensuring multiple copies of data are always available to protect against loss of any component of the SAN.
- Snapshots are scheduled to occur twice a day, allowing an administrator to perform a file-level restore in the case of accidental deletion or a full rollback to a point in time if data should become corrupted.
- Remote IP Copy copies the snapshots to a geographically separate site. Should anything happen to the data center, the data is available at the secondary site.
- Tape backups are also being handled at the secondary site, alleviating the load from the primary network at the data center.

Scenario 4 – Full Site Recovery with Clustered Servers and Clustered Storage

A regional entertainment management group's disaster recovery plan was reevaluated post 9/11. Evaluations uncovered a need for maximum uptime and access to data for several of their applications. In order to achieve this, redundancy, high availability and recovery had to be considered for both the servers and the SAN. Following is the solution architected for this customer.

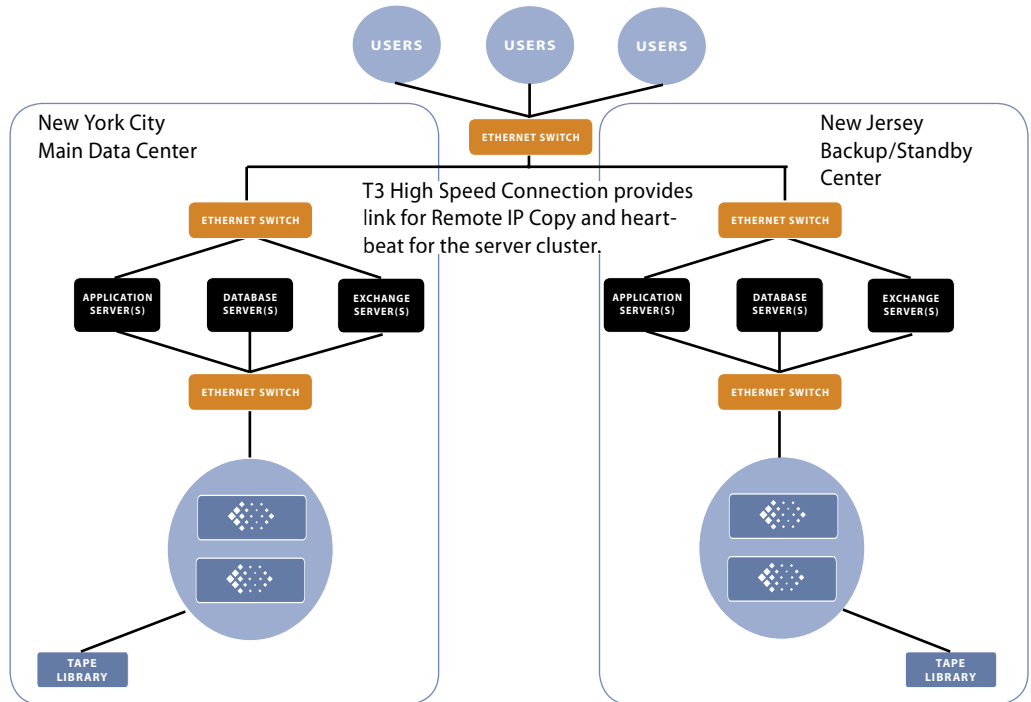


Figure 11. LeftHand SAN designed for server and data availability.

In the example above, environments were exactly replicated to achieve the following:

- The main data center in New York is accessed by the client machines most of the time. Barring a failure or disaster, this is the principal data center.
- The data from the New York facility is copied over a T3 line, using Remote IP Copy, each night to the New Jersey backup/standby center 30 miles away. This ensures that there is an exact replica of the data at both sites.

- In addition to having two identical copies of their data at different sites, it was crucial for this customer to also plan for server availability. The primary servers at the New York site are clustered with their respective standby servers at the New Jersey site. A heartbeat connection via the T3 line checks the availability of each of the servers every few seconds. If a problem is detected, the clustering software will fail a primary server over to the standby server or vice versa. Following are some failure/recovery example scenarios for this configuration.
 1. Primary storage is inaccessible - Servers in the primary site can be pointed to the storage at the New Jersey site in a matter of seconds.
 2. Primary servers are inaccessible – Standby servers automatically take over and access the storage at the primary site
 3. Primary site is completely lost – Servers at the standby site access the storage at the standby site.
 4. Both primary and standby sites lost – restore from tape.

Summary

Data is the most valuable commodity companies possess in this day and age. Take away access to data and you take away profits, productivity, business opportunities, government compliance and customers. While there's general concurrence that data needs to be protected, what's not as readily agreed upon is the best way to protect data. The reality is that there is no singular answer to this dilemma; each company has their own requirements and budgets to consider.

The LeftHand SAN has the ability to comply with virtually any company's data protection and DR needs. Based on the availability requirements for a particular set of data, the LeftHand SAN can provide data protection ranging from technologies that apply to a single site – from RAID 10 in a single NSM, to file-level restores, to synchronous replication – to a complete replication of data stored in an alternate location to protect against site failure due to natural disasters, a malcontent employee, or an act of terrorism. This flexibility lets customers choose a data protection and disaster recovery strategy appropriate for their needs and budget, all easily implemented using LeftHand's SAN/iQ software and the familiar Ethernet infrastructure.

About Barb Goldworm

Barb Goldworm is an independent analyst and consultant with over 25 years experience in the computer industry in systems and storage management, in various technical, marketing, industry analyst and senior management positions with IBM, StorageTek, Novell, Enterprise Management Associates and other successful startup ventures. She has been a frequent speaker at industry conferences worldwide for over 10 years, and was the creator and track chair for the Networld+Interop track on Networked Storage. She has been a regular columnist and contributor for various trade publications including NetworkWorld, ComputerWorld and Storage Networking World Online, as well as being frequently quoted in the press. She also provides consulting to the institutional investment community, advising clients on business issues and trends occurring within the high tech industries. Barb can be reached at

barbgoldworm@focusonstorage.com.

WHITE PAPER



LeftHand Networks, Inc.

1688 Conestoga Street

Boulder, CO 80301

Phone: 1.866.4.IPSANs

www.lefthandnetworks.com

© Copyright 2004 LeftHand Networks Inc. All rights reserved.

LeftHand Networks, the LeftHand Networks logo, and AEBS are registered trademarks of LeftHand Networks, Inc..
SAN/iQ, Distributed Storage Matrix, and Remote IP Copy are trademarks of LeftHand Networks, Inc..



Storage as it should be

