Disaster Recovery

Techniques and Topologies

By Connie Murphy

For TCM 538
Networking Concepts II

Professor Vaughn White
March 3, 2004

Table of Contents

Abstract

Every enterprise must protect its data.within the network. Network Attached Storage (NAS) implementations and Storage Area Networks (SANs) have become popular schemes for data management. They must be considered carefully in any data protection plan. Backup is the most important element in a good disaster recovery plan. Today, though system processors and tape drives have become blazing fast, they still can't keep up with the ever-growing amount of data and shrinking backup windows most companies are faced with**.** Without a solid, system-wide plan and implementation, disaster recovery is useless. This paper will outline different topologies, techniques and software available for maintaining proper backup procedures and the implementation of a solid disaster recovery plan.

Disaster Recovery - Techniques and Topologies

Introduction

Disaster Recovery is the coordinated process of restoring systems, data, and infrastructure required to support key ongoing business operations.

- Systems includes both hardware and software
- Data includes true data, log files, and audit information, as well as business procedures and rules
- Infrastructure includes phones, office space, remote access, intranets, web space, firewalls, etc.
- Business operations are whatever activities the business performs on a daily basis to generate revenue

A Disaster Recovery Plan (DRP) envelops the hardware and software required to run important business applications and all processes to restore them smoothly in the event of a natural or human-caused disaster. To prepare for a successful DRP, an assessment of mission-critical business processes and associated applications is necessary before creating the full disaster recovery plan (.Cisco White Paper, 1).

RAID Systems

No discussion of backup systems and disaster recovery would be complete without including *Redundant Array of Independent Disks* (RAID), often called *Redundant Array of Inexpensive Disks*, a reliable tool that's become more affordable in recent years and combines several physical hard disks into a unit that the host computer identifies as a single drive (RAID defined).

While RAID arrays can't prevent software or site failure, they can protect an organization from physical hardware failures at the disk level, which is a very common occurrence. RAID systems are a series of physical disks that act in unison to increase performance and/or protect the system and its data against the failure of any one disk (and in some cases, against the failure of multiple disks).

Hardware-based controllers offer the best speed and protection over software RAID systems, but they are also among the most expensive due to the number of disks needed to set it up.

Just like any other application, software-based arrays occupy host system memory, consume CPU cycles and are operating system dependent. By contending with other currently-running applications for host CPU cycles and memory, software-based arrays degrade overall server performance and are directly dependent on server CPU load (Nueffer, sec. 1)

Except for the array functionality, hardware-based RAID schemes have very little in common with software-based implementations. Since the host CPU can execute user applications while the array adapter's processor simultaneously executes the array functions, the result is true hardware multi-tasking. Hardware arrays also do not occupy any host system memory, nor are they operating system dependent.

Hardware arrays are also highly fault tolerant. Since the array logic is based in hardware, software is not required to boot. Some software arrays, however, will fail to boot if the boot drive in the array fails. For example, an array controlled by software can only be functional when the array software has been read from the disks and is memory-resident. Some software-based implementations commonly require a separate boot drive, which is not included in the array, while others such as Compaq partition a small portion of the hard drive as a boot drive during configuration.

RAID can provide several advantages; the degree to which each comes into play depends on the RAID level implemented. These advantages include:

1. **Speed**—Writing data to and reading data from multiple physical drives decreases transfer times.
2. **Reliability**—Many RAID levels provide data redundancy to varying degrees.
3. **Storage**—Combining hard disks can increase overall storage space, although various data redundancy schemes can eat into the savings.

See Appendix A for a table listing the different RAID levels with the number of disks required for performance and the pros and cons of each level.

There are a large number of vendor-specific and intermediate types of RAID hardware arrays, as well as Windows OS's that can often handle different types of RAID without hardware controllers, so RAID arrays may be created without swapping out hardware. Except for RAID 0, any form of RAID offers same-server data protection against hardware failure. All RAID systems should be supplemented with some other form of protection against larger and/or software-based failures, i.e. tape backup or replication.

However, RAID is the first step toward DRP. RAID arrays offer a lot of flexibility and a great way to protect against hardware failure.

Disk Striping

Disk striping is a technique by which data spans multiple hard drives. All hard drives involved in the stripe set are simultaneously read from and written to. For example, if a striped set of disks consists of three hard drives, then the data will be read and written about three times faster because Windows is distributing the workload among three hard drives. Creating a striped set is an inexpensive way

of dramatically increasing performance. Several levels of RAID use disk striping for fault protection (Nueffer, Sec. 2).

Performance Indicators for Disaster Recovery

Performance indicators provide the means by which one can measure the success of a DRP. Performance indicators for disaster recovery are somewhat different from those used to measure network performance, because they are a blend of project status and test runs of infrastructure. Indicators of success include:

- Intervallic reports from the planning group to senior management.
- Representation of the network design team on the DRP team.
- Intervallic tests to authenticate the performance of the disaster recovery plan, including noticed gaps and risks.
- An evaluation process that includes deploying new solutions.
- Final analysis of the disaster recovery testing, and effect on the business (after a disaster occurs) (Cisco White Paper, 1).

Pre-planning Management Awareness

Management Awareness is the first and most important step in creating a successful DRP. To obtain the necessary resources and time required from each area of an organization, senior management has to understand and support the business impacts and risks. Several key tasks are required to realize management awareness:

1. **Identify Possible Disaster Scenarios**

   First, classify the top ten disasters and evaluate their impact on the company. The analysis should cover effects on communications with suppliers and customers, the impact on operations, and disruption of key business processes. Complete this pre-study prior to the evolution of the DRP process, knowing that it will require additional substantiation during the planning process.

   The following are examples of possible disasters: fire, storm, water, earthquake, chemical accidents, nuclear accidents, war, terrorist attacks and other crime, cold winter weather, extreme heat, airplane crash (loss of key staff), and avalanche. The possibility of each scenario depends on factors such as geographical location and political stability.

   Assess the impact of a disaster on the company from both a financial and physical infrastructure perspective by asking the following questions:

   - How much of the organization's resources could be lost? •
   - What are the total costs?

- What efforts are required to rebuild?
- How long will it take to recover?
- What is the impact on the overall organization?
- How will customers be affected?
- What is the impact on them?
- How much will it affect the share price and market confidence?

2. Build Management Awareness

Senior management needs to be involved in the DRP process, and should be aware of the risks and potential impact on the organization. The first study on disaster recovery should include an estimate of possible costs and time to implement a disaster recovery strategy. Once management understands the financial, physical, and business costs associated with a disaster, it is then able to build a strategy and ensure that this strategy is implemented across the organization.

3. Obtain Management Sign-Off and Funding

The senior management has to agree on the disaster recovery project, as well as provide financial and human resources for the project. The first step is the announcement of the disaster recovery project and kickoff of a planning group or steering committee, which should be led by a senior management person.

Disaster Recovery Planning Process

During the planning stage, the mission-critical, important, and less-important processes, systems, and services in a network should be identified; and plans to ensure these are protected against the effects of a disaster should be put into place. Key elements of DRP include the following:

1. **Establish a planning group to manage the development** and implementation of the disaster recovery strategy and plan. Key people from each business unit or operational area should be members of the team, responsible for all disaster recovery activities, planning, and providing regular monthly reports to senior management.

2. **Perform risk assessments and audits**. •In order to create the disaster recovery plan, the planning group needs to thoroughly understand the business and its processes, technology, networks, systems, and services. The DRP group should prepare a risk analysis and business impact analysis that includes at least the top ten potential disasters. The risk analysis should include the worst-case scenario of completely damaged facilities and destroyed resources. It should address geographic situations, current design, lead-times of services, and existing service contracts. Each analysis should also include an estimate on the financial impacts of replacing damaged equipment, securing additional resources, and establishing extra service contracts.

3. **Establish priorities for network and applications**. When the risks posed to a business's processes from each disaster scenario are considered, a priority level should be assigned to each business process. Priorities should be based on the following levels:

    - **Mission Critical**:- Any network or application outage or destruction that would cause a complete disruption to the business, cause major legal or financial ramifications, or threaten the health and safety of anyone. If the targeted system or data requires significant effort to restore, or the restoration process is disruptive to the business or other systems, it is classified mission critical.

    - **Important -** Network or application outage or destruction that would cause a ongoing disruption to the business, but not prevent it. The targeted systems or network may or may not be easily restored.

    - **Minor**:- Network or application outage or destruction that would cause a minor disruption to the business. The targeted systems or network can be easily restored.
    -

4. **Develop recovery strategies**. •Just as the analysis of the business processes determine the priorities of the network, applications, and systems, the same analysis should be applied to any network design. The site priorities and location of key services contribute to a fault-tolerant design, with resilience built into the network infrastructure, and services and resources spread over a wide geography. Recovery strategy must be developed to cover the practicalities of dealing with a disaster. Such a strategy may be applicable to several scenarios; however, the plan should be assessed against each scenario to identify any actions specific to different disaster types. The plan should address the following: people, facilities, network services, communication equipment, applications, clients and servers, support and maintenance contracts, additional vendor services, and environmental situations. Any recovery strategy should include the expected down time of services, action plans, and escalation procedures. It should also determine thresholds, such as the minimum level at which can the business operate, the systems that must have full functionality (all staff must have access), and the systems that can be minimized.

5. **Prepare an up-to-date inventory and documentation of the plan**. •It is important to keep inventory up-to-date and have a complete list of all locations, devices, vendors, used services, and contact names. The inventory and documentation should be part of the design and implementation process of all solutions. Disaster recovery documentation should include:

    - Complete inventory, including a prioritization of resources.
    - Review process structure assessments, audits, and reports.
    - Gap and risk analysis based on the outcome of the assessments and audits.
    - Implementation plan to eliminate the risks and gaps.
    - Disaster recovery plan containing action and escalation procedures.
    - Training material.

6. **Develop verification criteria and procedures.** •Once a draft of the plan is created, a verification process should also be created to prove the disaster recover strategy and, if this strategy is already implemented, review and test the implementation. It's important that the plan be tested and reviewed frequently. The verification process should include an experience cycle; disaster recovery is based on experience and each disaster has different rules.

7. **Implement the plan**. Now it's time to make some key decisions: How should the plan be implemented? Who are the critical staff members, and what are their roles? Leading up to the implementation of the plan, try to practice for disaster recovery using roundtable discussions, role playing, or disaster scenario training.

Again, it is critical that senior management approves all DRP and implementation plans (Cisco White Paper)

Resiliency and Backup Services

Resiliency and backup services form a key part of disaster recovery, and these services should be reviewed to make sure they meet the criteria for the DRP. Cisco defines network resiliency as “the ability to recover from any network failure or issue whether it is related to a disaster, link, hardware, design, or network services.” (Cisco White Paper, 7).). A high-availability network design is often the foundation for disaster recovery and can be adequate to handle some minor or local disasters. Key tasks for resiliency planning and backup services include the following:

- **Assess the resiliency of your network**, identifying gaps and risks•keeping in mind the following three levels of availability:

  1. Reliable networks
  2. High-availability networks
  3. Nonstop network environments.

  This assessment helps prioritize risks, set requirements for higher levels of availability, and identify the mission-critical elements of the network.

  Be sure to evaluate the following areas of the network:

  - Network links
    - Carrier diversity •
    - Local loop diversity •
    - Facilities resiliency •
    - Building wiring resiliency •
  - Hardware resiliency
    - Power, security and disaster •
    - Redundant hardware •
    - Mean time before replacement (MTTR) •
    - Network path availability •
  - Network design
    - Layer 2 WAN design •
    - Layer 2 LAN design •
    - Layer 3 IP design •
  - Network services resiliency
    - DNS resiliency•
    - DHCP resiliency •
    - Other services resiliency

- **Review and Implement Backup Services**. The DRP should include a backup services strategy, which needs to be consistent throughout the whole organization. For example, Frame Relay services could use ISDN as a backup service. Backup scenarios are important to provide higher availability and access to main sites and/or access to existing parallel disaster recovery sites during a disaster. All system and application backup strategies depend upon network connections. Disaster handling requires communication services, and the impact of a disaster could be greatly limited by not having available communication services (Cisco White Paper, 5).

NAS/SAN Topologies /Technologies

Any discussion of DRP must start with a look at the different topologies and methods available for system-wide backups. The first area to be discussed in this paper is NAS/SAN topologies and their backup configurations. Understanding the benefits and drawbacks of the seven NAS/SAN topologies reviewed in this section will help optimize any network design in agreement with specific data protection needs.

1. **High-speed local backup (filer-to-tape).** Back up an NAS appliance (also called a filer) directly to a SCSI-attached or Fibre Channel-attached tape drive or robotic tape library. (see Figure 1)

   This topology is used if specifications require an easy-to-deploy method for rapid backup of predetermined amounts of data while keeping company LAN resources free. Because data does not traverse the LAN, data security is maximized and the network bandwidth used is negligible. Because only a single data stream is initiated for each backup job, filer-to-tape backup does not optimize hardware investment when large backups are performed. Expanding this scheme could be expensive because, as NAS appliances are added or data volumes expand, purchase of additional tape devices will become necessary. Most backup software includes some form of Network Data Management Protocol (NDMP), which is required for communications with NAS appliances, Direct Access Recovery (DAR), which substantially accelerates NDMP restores, and SAN Resource Sharing (SRS), which dynamically allocates tape resources in a SAN.

   NDMP is an open standard for centralized control of enterprise-wide data management. NDMP enables backup software vendors to provide support for Network Appliance storage systems without having to port client code. An NDMPcompliant solution separates the flow of backup/restore control information from the flow of data to and from backup media (NDMP.org/).

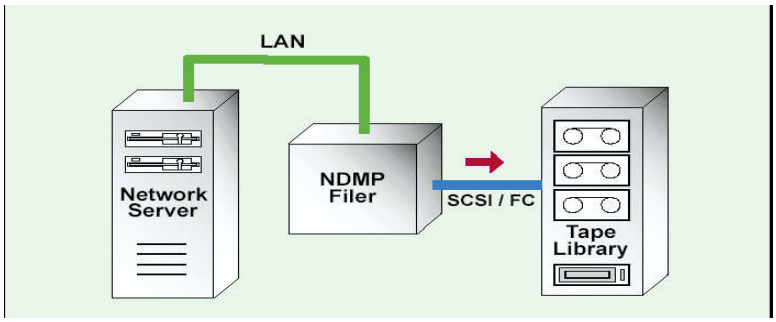

Figure 1 (SyncSort)

2. **Three-way backup (filer-to-filer).** Backup a secondary NAS appliance through the LAN to a tape device or library that is attached to a separate, primary NAS appliance. (see figure 2).

   This method is used for fast backups of data stored on more than one NAS appliance. This topology offers scalability at the file server level and leverages tape library investment. Like all NAS solutions, this is excellent for multi-platform networks. With this design, data does travel across the LAN, so bandwidth and security concerns may enter the picture.

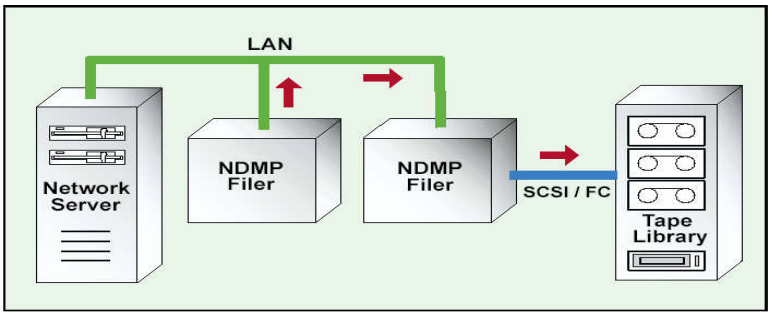

Figure 2 (SyncSort)

3. **Filer-to-server backup.** Backup a NAS appliance through the LAN to a tape device or library connected to a server. (see figure 3).

   If NAS appliances are not connected directly to a storage device, there is no need to reconfigure the network to facilitate backup and restore. Filer-to-server backup is useful for enterprises with complex networks and expectations of expansion. With this topology, filers do not need to support newer tape devices because the server drives the tape devices directly. However, one drawback may be bandwidth and server contention, as all data traverses the LAN and the device server activity will be slowed.

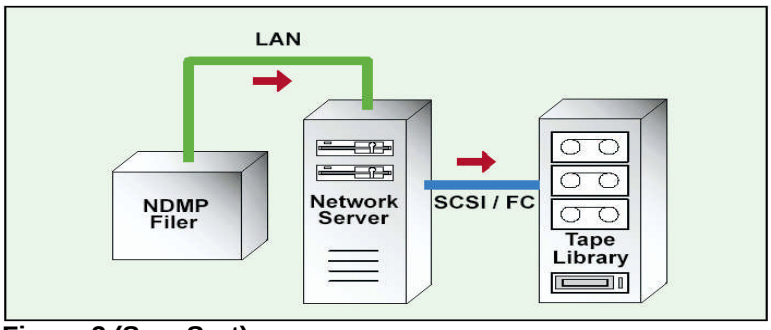

Figure 3 (SyncSort)

4. **Server-to-filer backup.** Backup data on a server through the LAN to a tape device or library that is SCSI-attached or Fibre Channel-attached to a NAS appliance (see figure 4).

   Use this topology if keeping most data on NAS appliances, but still need to back up smaller amounts of data that reside on servers. Although the data on the servers may not be backed up or restored as quickly as with other methods, the NAS appliance's critical data is available for quick return to an operable state in the event of a disaster. The data coming across the LAN from the servers are subject to security and bandwidth issues.

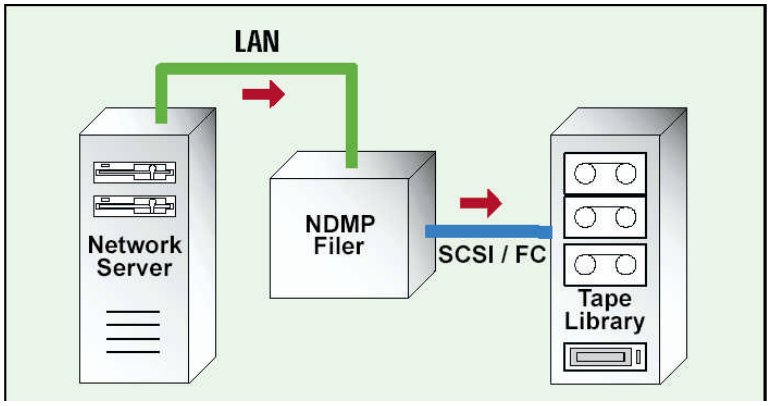

Figure 4 (SyncSort)

5. **Backup through multi-hosting.** Servers and NAS appliances in any combination are directly connected by SCSI or Fibre Channel to one or more drives in a single or multi-drive tape library. (see figure 5).

   When deployed with SCSI connections, this configuration may be somewhat difficult to expand. However, multi-hosting with Fibre Channel is rapidly gaining acceptance, which has the added benefit of enabling easy migration to a SAN. Use this topology if the duration of the backup window is short and the file servers are physically near the tape library.

   This design offers a significant decrease in network traffic and optimizes tape drive performance. All servers stream data concurrently to their

locally attached drives, increasing backup speed and data security while reducing bandwidth usage.

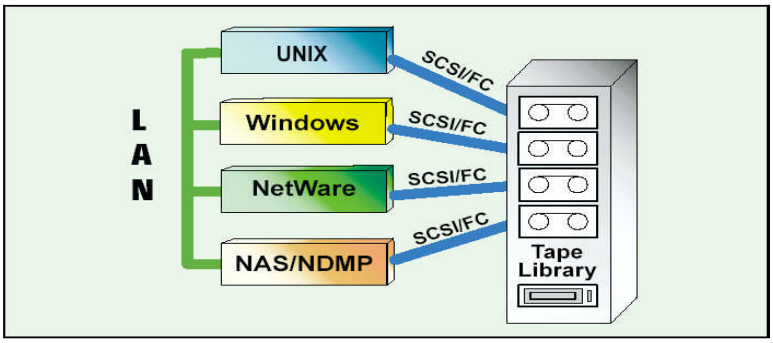

Figure 5 (SyncSort)

6. **<u>Backup through Fibre Channel SAN</u>.** Servers and NAS appliances in any combination are connected to one or more tape libraries through a Fibre Channel switch. (see figure 6).

   The addition of the switch enables any backup job to complete as quickly as possible because a sophisticated backup technology called SRS dynamically allocates tape drives to maximize hardware. If a drive is busy, SRS easily works around the use of that drive. Fibre Channel also allows for greater physical separation among the components and allows for expansion and flexibility. In addition, the SAN design keeps the LAN bandwidth free for other applications and network traffic. Deployment of a SAN can be expensive and complex, and knowledge of Fibre Channel Technology is required to maintain this topology.

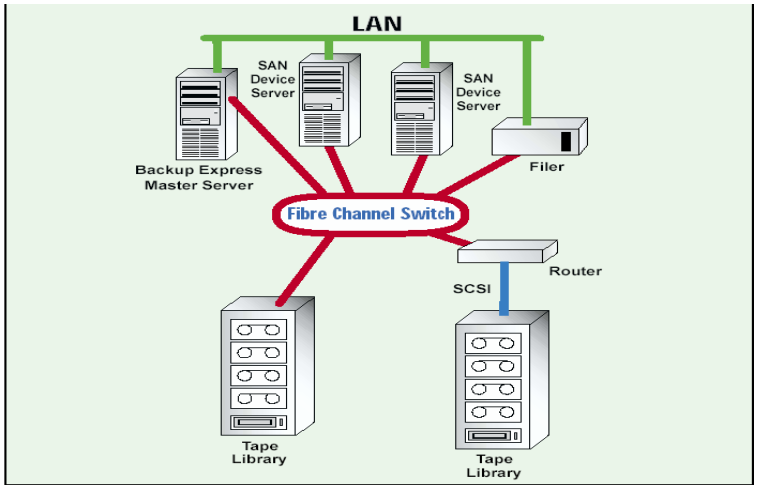

Figure 6 (SyncSort)

7. **<u>Backup through iSCSI (SCSI over IP).</u>** Using iSCSI technology, remote IP-attached storage devices are seen by a host computer as if they are local devices. (see figure 7).

Like a SAN, iSCSI provides for a many-to-many connection among servers and storage devices. iSCSI facilitates long-distance data transfers over Intranets (or the Internet), so storage and servers can be added at suitable, remote locations. Because iSCSI networks use existing TCP/IP infrastructure, they require a low level of investment, are economical to maintain, and are easy to expand without any disturbance of network processes. And the data transfer can be extremely fast.

However, because data being backed up and restored is flowing over an Intranet, security considerations must be addressed. Further, TCP/IP protocols are not currently optimized for storage, which requires high bandwidth and relatively error-free transmission. But comprehensive standards are under development to tackle issues of infrastructure and interoperability. (Syncsort)

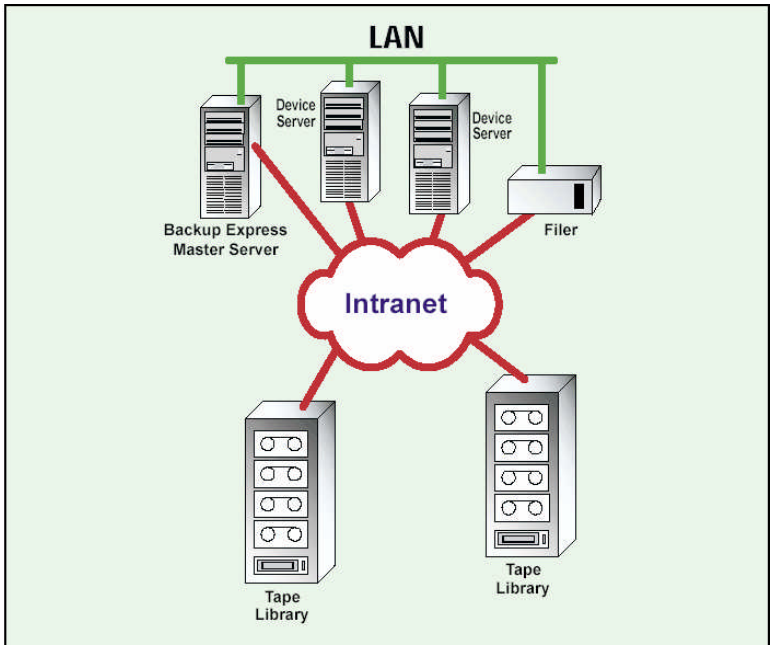

Figure 7 (SyncSort)

Backup Software Company offerings

I have listed several different packages in this section to list the different types of software that are available. The links to these company websites are available at the end of the section.

Backup Express, Syncsort's high-performance backup and restore solution, advances virtually all NAS and SAN strategies. Backup Express supports

VERITAS NetBackup 5.0 —VERITAS NetBackup provides consistent backup and rapid recovery to both disk and tape. The new version offers advanced features. For example, incremental backups can be combined with previous full

backups to create a synthetic full backup. NetBackup Advanced Client delivers multiple snapshot technologies while also providing the ability to appropriately select one that fits specific business requirements. Enhanced disk-based data protection includes integrated disk staging and the ability to make multiple copies to disk on the fly. VERITAS NetBackup 5.0 increases the speed of backup and recovery, reducing the impact to users when either operation occurs

VERITAS Data Lifecycle Manager 5.0 — This is a new solution from VERITAS. Data Lifecycle Manager manages data from creation to disposal. VERITAS claims this product is a natural extension of its data protection suite. It has the ability to remove expired data, and can thereby control data management costs, which, in turn, makes it much easier to achieve regulatory compliance while providing end users with continual access to all data. VERITAS Data Lifecycle Manager offers automatic push of templates and custom policies to local or remote sites to facilitate compliance, and features role-based administration for enhanced security. As an added plus, it also archives Microsoft Exchange mailboxes and journals.

VERITAS Backup Exec 9.1 for Windows —this updated backup and restore solution provides comprehensive, cost-effective protection for Microsoft Windows most recent versions of the Small Business Server Edition, Microsoft Exchange, and Microsoft SQL Server environments. This release also supports the new Desktop and Laptop Option as described above.

LEGATO NetWorker 7.1. Version 7.1 is designed to meet the growing pressures of data recovery and backup, and budgetary constraints of large and small environments by enabling:

- Instant backup Recover and protection of massive, increasing volumes of data faster
- Increased service levels and reduced downtime costs
- Increased IT staff productivity and lower TCO of storage resources
- Greater reliability, performance, and scalability of data protection operations across the enterprise
- Increase ROI of DAS, NAS and SAN investments
- Centralizes management of information assets and ensures fast recovery in the event of disaster.
- Maximizes utilization of key resources by automating data migration and protection.
- Improves operational efficiency and ensures business continuity through comprehensive protection and preservation of data.
- Includes backup and recovery, HSM/archiving and media and device management.

Advanced indexing and media management, cluster support, high speed parallelism, cross-platform tape interoperability, comprehensive NDMP support, backup to disk, tape cloning, archive, serverless backup, and dynamic drive

sharing are among key components that enable IT organizations to protect storage assets and minimize downtime with this product.

Vendor Support Services

Having support services from your major vendors in place adds a strong value to disaster recovery planning. For example, specific managed hot standby sites or on-site services with rapid response times can significantly ease disaster recovery. Key questions regarding vendor support include:

- Are support contracts in place? •
- Has the disaster recovery plan been reviewed by the vendors, and are the vendors included in the escalation processes?
- Does the vendor have sufficient resources to support the disaster recovery?

Most vendors have experience handling disaster situations and can offer additional support, assisting with limiting downtime in the case of an unexpected outage (Cisco White Paper, 7).

Completing the DRP

When creating disaster recovery plans, an administrator often deals with strict "black and white" factors. In other words, they look at data in terms of whether it needs to be protected and develop policies based specifically on that decision.

However, particularly in the case of end-user file servers and network-attached storage (NAS) devices, many files and data sets fall into the "gray area"--data that isn't important enough to include in the DR plan, which corporate policies still allow to exist.

The most frequent examples are temporary files created by software, temporary and/or testing databases or code for the development staff, and personal files that corporate policies allow users to store on corporate data systems. While most of these files aren't vital for corporate operations, they require space on corporate data systems.

When developing a DR plan, the first priority is to create a system that can protect the vital data and make it available if a disaster destroys the original systems. In this process, important decisions must be made about these gray-area files to ensure nothing vital is left out.

The first decision is whether to replicate or back up gray-area files. Backup is the easiest way to make sure you protect all the vital files, but it also means you're going to either transmit and/or back up a great deal of non-vital data in many cases.

This requires extra tape resources, extra disk space, and potentially greater bandwidth costs. It could also mean higher software and hardware costs as you move toward "enterprise" or "advanced" editions of backup and/or replication tools.

Perform a comprehensive study to determine what is and what isn't a gray-area file within the organization. Once this distinction is made, determine the best way to exclude those files. This isn't as easy as it sounds.

With the use of software backup tools, it is possible to selectively exclude many files by their extensions or other file patterns. However, many files can be indistinguishable from vital files. In these cases, segregate these files--presumably by file directories--to ensure the system protects only vital files. Replication can have issues of its own.

Software-based replication tools often allow the same level of file recognition as tape systems, so it is possible to use the same techniques to protect only vital data. Hardware-based replication systems generally can't distinguish one file from another, so here is where segregation of files into directories or even data volumes will be helpful.

Since the IT staff doesn't usually make the decision of what users can store on data systems, especially file servers and NAS devices, there will most probably be large numbers of these files. Corporate standards, executive privileges, and other factors will always create gray-area files. How these files are handled can make the difference between a successful and an unsuccessful DR plan when it's time to fail over.

Conclusion

Today's business environment places demands on the IT organization that are more challenging than ever. Enterprises are global, and IT operations must be always on. Additionally, companies are faced with regulatory requirements that have them scrambling to archive unstructured information for multiple years, with the ability to selectively retrieve it on request. These combined — and sometimes conflicting — requirements make it critically important for IT departments to backup data with virtually no backup window. Developing a storage utility means that IT resources can be delivered as a service based on business demand. Certainly, many IT managers make DRP a long-term strategic goal.

It is important to understand the purpose of the DRP, define specific requirements, determine success criteria, and develop a plan to ensure that success. Updating and refining this plan is critical, as many changes occur in every business evolution. DRP can be a challenging and expensive undertaking, but one that will be well worth the effort as it could determine the future of the company in the event of a disaster.

Appendix A

Level	Description	# of disks	Pros	Cons
0	Uses striping but not redundancy of data; often not considered "true" RAID	2	Provides the best performance because no parity calculation overhead is involved; simple and easy to implement	No fault tolerance; failure of one drive will result in all data in an array being lost
1	Duplicates but does not stripe data; also known as *disk mirroring*	2	Faster read performance, since both disks can be read at the same time; provides best fault tolerance, because data is 100 percent redundant	Inefficient high disk overhead compared to other levels of RAID
2	Disk striping with error checking and correcting information stored on one or more disks	Many	Very reliable; faults can be corrected on the fly from stored correcting information	High cost; entire disks must be devoted to correction information storage; not considered commercially viable.
3	Striping with one drive to store drive parity information; embedded error checking (ECC) is used to detect errors	3	High data transfer rates; disk failure has a negligible impact on throughput	Complex controller design best implemented as hardware RAID instead of software RAID
4	Large stripes (data blocks) with one drive to store drive parity information	3	Takes advantage of overlapped I/O for fast read operations; low ratio of parity disks to data disks	No I/O overlapping is possible in write operations, since all such operations have to update the parity drive; complex controller design
5	Stores parity information across all disks in the array; requires at least three and usually five disks for the array	3	Better read performance than mirrored volumes; read and write operations can be overlapped; low ratio of parity disks to data disks	Most complex controller design; more difficult to rebuild in case of disk failure; best for systems in which performance is not critical or that do few write operations
6	Similar to RAID 5 but with a second parity scheme distributed across the drives	3	Extremely high fault tolerance and drive-failure tolerance	Few commercial examples at present
7	Uses a real-time embedded operating system controller, high-speed caching, and a	3	Excellent write performance; scalable host interfaces for connectivity or increased transfer bandwidth	Very high cost; only one vendor (Storage Computer Corporation) offers this system at present

	dedicated parity drive			
10	An array of stripes in which each stripe is a RAID 1 array of drives	4	Higher performance than RAID 1	Much higher cost than RAID 1
53	An array of stripes in which each stripe is a RAID 3 array of disks	5	Better performance than RAID 3	Much higher cost than RAID 3
0+1	A mirrored array of RAID 0 arrays; provides the fault tolerance of RAID 5 with overhead for fault tolerance of RAID 1 (mirroring)	4	Multiple stripe segments enable high information-transfer rates	A single drive failure will cause the whole array to revert to a RAID 0 array; is expensive to implement and imposes a high overhead on the system

References

Seven Effective Strategies for Data Protection: A Practical Guide to NAS and SAN Deployment. White Paper from Syncsort 2002. Retrieved February 15, 2004, from the World Wide Web. http://www.syncsort.com.

Disaster Recovery: Best Practices White Paper. Cisco Systems updated Jan 27, 2004. Retrieved February 18, 2004 from the World Wide Web. http://www.cisco.com/warp/public/63/disrec.html#topic16

"Raid defined." ND. Retrieved February 16, 2004 from the World Wide Web. http://searchstorage.techtarget.com/sDefinition/0%2C%2Csid5_gci214332%2C00.html

"Enterprise Data Protection Bill of Rights." Quantum Corporation, 2003. Retrieved March 1, 2004, from the World Wide Web. http://quantum.treehousei.com/Surveys/publishing/survey_148/graphics3/prod_pgs/whitepaper.pdf

Nueffer. "What is Raid?" Retrieved March 1, 2004, from the World Wide Web. http://www.uni-mainz.de/~neuffer/scsi/what_is_raid.html

http://www.netapp.com/tech_library/ftp/3066.pdf

Wold, Goeffrey, Robert Shriver. (2001) *Disaster Proof Your Business*. New York: McGraw-Hill Primis Custom Publishing.

Sources of data protection and backup software

Legato http://www.hastorage.com/BackupSoftware_Legato.htm

Networker http://www.hastorage.com/BackupSoftware_Legato_Networker.htm

Backup Express http://www.hastorage.com/BackupSoftware_SyncSort.htm

Dantz http://www.hastorage.com/BackupSoftware_Dantz.htm

Retrospect http://www.hastorage.com/BackupSoftware_Dantz.htm

Bocada http://www.hastorage.com/BackupSoftware_Bocada.htm

BackUp Report http://www.hastorage.com/BackupSoftware_Bocada.htm

Veritas http://www.veritas.com/Products/www