

Large Scale Solutions – New Technology in Storage Systems

Storage systems have become their own unique and complex computer field and can mean different things to different people. So what is the definition of these systems? Storage systems are the hardware that store data.

For example, this may be a small business server supporting an office of ten users or less—the storage system would be the hard drives that are inside of that server where user information is located. In large business environments, the storage systems can be the large SAN cabinet that is full of hard drives and the space has been sliced-and-diced in different ways to provide redundancy and performance.

The Ever-Changing Storage System Technology

Today's storage technology encompasses all sorts of storage media. These could include WORM systems, tape library systems and virtual tape library systems. Over the past few years, SAN and NAS systems have provided excellent reliability. What is the difference between the two?

- **SAN** (Storage Area Network) units can be massive cabinets—some with 240 hard drives in them! These large 50+ Terabyte storage systems are doing more than just powering up hundreds of drives. These systems are incredibly powerful data warehouses that have versatile software utilities behind them to manage multiple arrays, various storage architecture configurations, and provide constant system monitoring.
- **NAS** (Network Attached Storage) units are self-contained units that have their own operating system, file system, and manage their attached hard drives. These units come in all sorts of different sizes to fit most needs and operate as file servers.

For some time, large-scale storage has been out reach of the small business. Serial ATA (SATA) hard disk drive-based SAN systems are becoming a cost-effective way of providing large amounts of storage space. These array units are also becoming mainstream for virtual tape backup systems—literally RAID arrays that are presented as tape machines; thereby removing the tape media element completely.

Other storage technologies such as iSCSI (learn more here:

http://www.ontrack.com/include/document.asp?file=http://www.adaptec.com/pdfs/iscsi_learn.pdf), DAS (Direct Attached Storage), Near-Line Storage (data that is attached to removable media), and CAS (Content Attached Storage) are all methods for providing data availability. Storage Architects know that just having a 'backup' is not enough. In today's high information environments, a normal nightly incremental or weekly full backup is obsolete in hours or even minutes after creation. In large data warehouse environments, backing up data that constantly changes is not even an option. The only method for those massive systems is to have storage system mirrors—literally identical servers with the exact same storage space.

How does one decide which system is best? (See *Gartner Research-Disk Storage Overview*) Careful analysis of the operation environment is required. Most would say that having no failures at all is the best environment—that is true for users and administrators alike! The harsh truth is that data disasters happen every day despite the implementation of risk mitigation policies and plans.

When reviewing your own or your client's storage needs, consider these questions:

- *What is the recovery turn-time?*
What is your client's maximum time period allowed to be back to the data? In other words, how long can you or your client survive without the data? This will help to establish performance requirements for equipment.

- *Quality of data restored.*
Is original restored data required or will older, backed up data suffice? This relates to the backup scheme that is used. If the data on your, or your client's storage system changes rapidly, then the original data is what is most valuable.
- *How much data are you or your client archiving?*
Restoring large amounts of data will take time to move through a network. On DAS (Direct Attached Storage) configurations, time of restoration will depend on equipment and I/O performance of the hardware.

Unique Data Protection Schemes

Storage System manufacturers are pursuing unique ways of processing large amounts of data while still being able to provide redundancy in case of disaster. Some large SAN units incorporate intricate device block-level organization, essentially creating a low-level file system from the RAID perspective. Other SAN units have an internal block-level transaction log in place so that the Control Processor of the SAN is tracking all of the block-level writes to the individual disks. Using this transaction log, the SAN unit can recover from unexpected power failures or shutdowns. (See *Examples of some SAN systems*)

Some computer scientists specializing in the storage system field are proposing adding more intelligence to the RAID array controller card so that it is 'file system aware.' This technology would provide more recoverability in case disaster struck, the goal being the storage array would become more self-healing. (See *On the Cutting Edge of Storage Technology-Smarter RAID Controller Cards*)

Other ideas along these lines are to have a heterogeneous storage pool where multiple computers can access information without being dependant on a specific system's file system. In organizations where there are multiple hardware and system platforms, a transparent file system will provide access to data regardless of what system wrote the data. (See *On the Cutting Edge of Storage Technology-SAN File System*)

Other computer scientists are approaching the redundancy of the storage array quite differently. The RAID concept is in use on a vast number of systems, yet computer scientists and engineers are looking for new ways to provide better data protection in case of failure. The goals that drive this type of RAID development are data protection and redundancy without sacrificing performance. (See *On the Cutting Edge of Storage Technology-RAID Implementation Improvements*)

Reviewing the University of California, Berkeley (See *How Much Information Produced in 2003?*) report about the amount of digital information that was produced 2003 is staggering. Your or your client's site may not have terabytes or petabytes of information, yet during a data disaster, every file is critically important.

Avoiding Storage System Failures

There are many ways to reduce or eliminate the impact of storage system failures. You may not be able to prevent a disaster from happening, but you may be able to minimize the disruption of service to your clients.

There are many ways to add redundancy to primary storage systems. Some of the options can be quite costly and only large business organizations can afford the investment. These options include duplicate storage systems or identical servers, known as 'mirror sites'. Additionally, elaborate backup processes or file-system 'snapshots' that always have a checkpoint to restore to, provide another level of data protection.

Experience has shown there are usually multiple or rolling failures that happen when an organization has a data disaster. Therefore, to rely on just one restoration protocol is shortsighted. (See *Backup Basics*) A successful storage organization will have multiple layers of restoration pathways.

Ontrack has heard thousands of IT horror stories of initial storage failures turning into complete data calamities. In an effort to bring back a system, some choices can permanently corrupt the data. Here are several risk mitigation policies that storage administrators can adopt that will help minimize data loss when a disaster happens:

- **Offline storage system** — Avoid forcing an array or drive back on-line. There is usually a valid reason for a controller card to disable a drive or array, forcing an array back on-line may expose the volume to file system corruption.
- **Rebuilding a failed drive** — When rebuilding a single failed drive, it is import to allow the controller card to finish the process. If a second drive should fail or go off-line during this process, stop and get professional data recovery services involved. During a rebuild, replacing a second failed drive will change the data on the other drives.
- **Storage system architecture** — Plan the storage system's configuration carefully. We have seen many cases with multiple configurations used on a single storage array. For example, three RAID 5 arrays (each holding six drives) are striped in a RAID 0 configuration and then spanned. Keep a simple storage configuration and document each aspect of it.
- **During an outage** — If the problem escalates up to the OEM technical support, always ask "Is the data integrity at risk?" or, "Will this damage my data in any way?" If the technician says that there may be a risk to the data, stop and get professional data recovery services involved.

The Leader in Storage System Recoveries

Ontrack Data Recovery has been successfully recovering data from large storage systems for many years. Ontrack's unique approach is what sets us apart from other data recovery companies.

A recovery of a data volume implementing a RAID configuration starts out with a Senior Engineer evaluating each hard disk involved and analyzing the data structures to determine the proper recovery path. There is no standard configuration for these systems and each OEM implements RAID configurations differently, making every job unique and challenging. The final step is verifying the file system is correctly pointing to the data, validating the file system information and data.

These types of recoveries are the pinnacle of engineering challenges. It is amazing to see one of these systems come together after hours of hard work – going from a data disaster to a complete and successful recovery. Often times, these recoveries result in the original files being recovered and archived without any hardware or software manipulation required on the part of the customer.

We applaud the storage industry for continuing to find better ways to preserve data and maintain business continuity. Some failures are beyond the soft recovery methods that the hardware can handle. This is where Ontrack fits into your or your client's Data Availability plans. Ontrack has services available to accommodate your or your client's time requirements for original data restoration.

Ontrack Data Recovery is the leader in storage system recoveries because of our experience, development resources, and engineering staff. Ontrack is the recovery company of choice for users, partners, and IT professionals who have high requirements for data recovery.

References

Examples of some SAN systems

EMC Clarion SAN

http://www.emc.com/products/systems/clariion_cx700.jsp

EMC Symmetrix SAN

http://www.emc.com/products/systems/DMX_series.jsp

Hitachi Lighting SAN

http://www.hds.com/products_services/storage_systems/enterprise_storage/

Hitachi Thunder SAN

http://www.hds.com/products_services/storage_systems/modular_storage/

On the Cutting Edge of Storage Technology

Smarter RAID Controller Cards

<http://www.usenix.org/publications/library/proceedings/fast04/tech/sivathanu.html>

San File System

http://adic.com/us/collateral/TechArticle_infostor_SAN_file_system.pdf

RAID Implementation Improvements

<http://www.usenix.org/publications/library/proceedings/fast04/tech/corbett.html>

<http://h18004.www1.hp.com/products/servers/proliantstorage/arraycontrollers/adg/index.html>

How Much Information Produced in 2003?

<http://www.sims.berkeley.edu/research/projects/how-much-info-2003/execsum.htm>

Backup Basics

<http://www.adic.com/us/collateral/WPBackup.pdf>

Gartner Research – Storage Systems

http://www4.gartner.com/resources/116400/116480/2003_disk_stora.pdf