

# Simplifying Management of Storage Area Networks

By Jon Hass, Keith R. Kammer, and Dan McConnell

Storage management allows IT administrators to realize the full capability of a storage area network. From the basic elements of administration to the advanced techniques of virtualization, storage management provides several means for maximizing storage space.

Storage management is the ability to manage all aspects of an environment's storage needs by discovering, visualizing, and supervising all the devices within the environment. This device management involves the devices' logical and physical characteristics as well as the data that resides on or passes through them.

Storage management consists of varying levels of capability that can be depicted as a pyramid (see Figure 1). The lowest level of the pyramid is storage administration, which encompasses the six basic tools to ensure the setup and continuous operation of an enterprise. The second level includes the storage operations—the capabilities to automate an environment and ensure continuing operations.

The next level is file management, which provides the ability to manage the data that resides on the storage devices. Finally, the top of the pyramid is virtualization, a process that allows dynamic allocation of storage and enables the administrator to achieve better utilization of available storage resources.

## Storage administration: The base to build a storage management environment

Storage administration comprises the fundamental tasks of storage management:

- » Installing and configuring the storage
- » Discovering and visualizing the environment
- » Providing notification and monitoring for event management

- » Centrally administering the devices
- » Reporting on basic levels of asset and inventory
- » Using diagnostic capabilities to isolate and fix problems

## Installation and configuration

Storage management software should provide a centralized location for installing and configuring an environment. This software should be automated with validation routines to verify that the correct software, device drivers, firmware, and hardware revision levels have been installed. Ideally, the software installation should push the software or agents from the central console to other servers that require the software to operate. The software also should be able to create a script of the site's installation and configuration procedures; such a script would then enable the site to be duplicated at another location.

## Discovery and visualization

Manageable devices within the environment must be identified before an administrator can centrally manage the storage. This process, known as *discovery*, creates a database repository of information about the devices. In its simplest form, discovery is performed by polling the devices, which may include storage arrays, switches, managed hubs, routers, storage area network (SAN) hosts, storage appliances, network-attached storage (NAS) devices, tape devices, subsystems, or storage domain management servers.

Polling is often effective in small environments, but it can become cumbersome and degrade system performance in larger

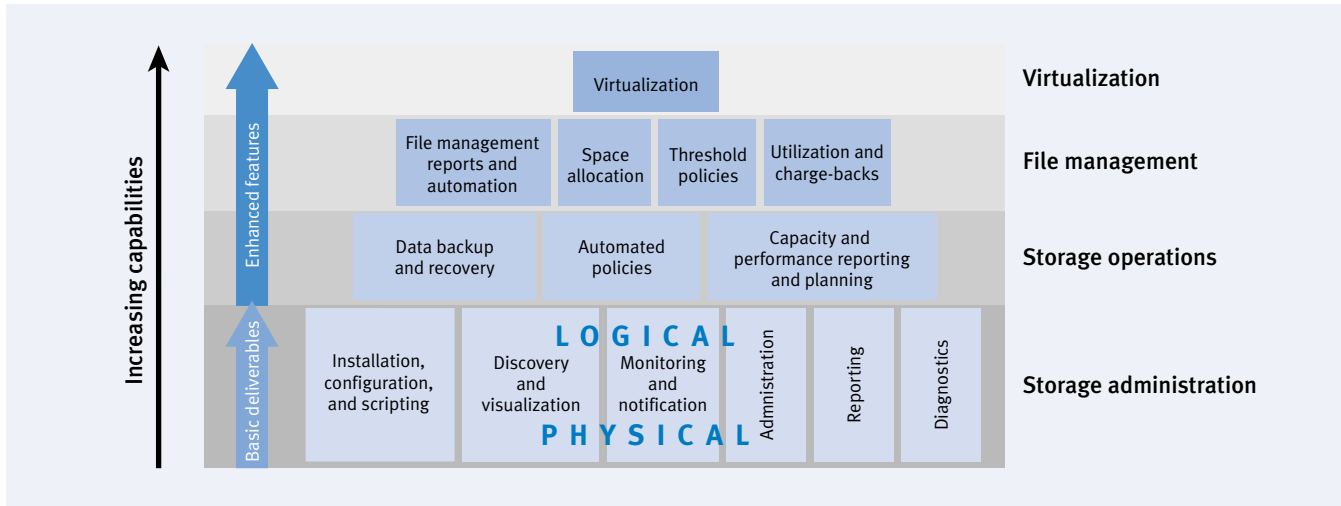


Figure 1. The tools required for storage availability

environments. To circumvent this degradation, agents can be distributed throughout the environment to gather information at a sublevel and then report it to the management station. Storage management applications that use agents to help deliver information offer more scalability compared to applications that use polling only.

Common methods of exchanging information, such as Simple Network Management Protocol (SNMP) or other standards-based application programming interfaces (APIs), can ease the discovery process. Devices that do not comply with these standards are usually defined through proprietary APIs on a per-device basis. Some older devices that do not have APIs or management information bases (MIBs) must be defined manually to be visible.

Following discovery, an environment should be visualized. Current storage management applications use graphical representations for both the physical and logical components of an environment to accomplish visualization. On a physical level, the graphical representations should not only display the actual devices and interconnections, but also the significant internal components, such as host bus adapters (HBAs) and hard drives.

The representation of the logical environment should include logical unit numbers (LUNs), logical volumes, RAID (redundant array of independent disks) configurations, zones, mirrored environments, and clusters, as well as the association of these items with the servers in the environment. These representations enable administrators to isolate events and identify specific problems in the environment.

### Event management: Monitoring and notification

The next task of storage administration, event management, encompasses two very important responsibilities of storage

management: monitoring devices for status changes and notifying the administrator of these changes.

Administrators can perform event monitoring by polling the managed devices for status changes or by setting up the management station to receive events from managed devices when a status change occurs. (Agents also can be used to distribute the process and assist scalability.) The IT industry is slowly standardizing methods for polling and event receipts. Currently, SNMP is the most prevalent of these standards, but newer, more secure and scalable technologies are being introduced.

Monitoring physical devices helps identify when devices may potentially fail or become critical, ensuring continuous operation of the environment. Monitoring the logical environment helps identify events associated with capacity availability and performance. Ideally, the storage management application will allow users to set threshold levels for various logical issues so that when these levels are crossed, the storage management application can notify the administrator, run a script, or perform another predetermined action.

The other aspect of event management, notification, uses the storage management application to notify the administrator of events. Notification can occur in several forms: a pop-up message on the management station screen, an e-mail, or a wireless alphanumeric page. The event notification can be used to trigger policies that automate or auto-fix specific conditions, or to even “phone home” and advise the hardware or application vendor of an impending failure.

### Device administration

Storage management applications centralize the administration of storage devices. Centralization is crucial; without it, managing

storage is burdensome and requires significantly more resources to support and administer the environment.

The storage management application should centralize the distribution of firmware and software to update devices and applications within the environment. It also should deliver asset and inventory management, security, and management of storage devices, which includes the following tasks:

- » Creation of LUNs and volumes
- » RAID configurations
- » LUN masking and mapping
- » Zoning
- » Online expansion of capacity, virtual disks, and volumes
- » Physical management of arrays, HBAs, switches, and disaster recovery solutions
- » Configuration of failover paths and load balancing
- » Management of tape devices, libraries, and tape backup

### Reporting asset inventory

The storage management application provides a collection point for information about all devices within an environment. Depending on the data collected by the application, reports can range from basic asset inventory to projections of future resource needs. Unfortunately, collecting data can consume network bandwidth and a significant number of CPU cycles, so the application's reporting mechanism should have the ability to schedule data collection and report generation as well as the ability to run the reports interactively.

### Diagnostics: Isolating and fixing problems

Perhaps the most difficult aspect of storage management is handling problems. Most storage management applications record events as they occur to identify critical and warning situations, but in a SAN environment that contains fabrics of interconnected devices, multiple events can be generated for a single problem.

Diagnostic capabilities enable administrators to identify these problems faster. The storage management application should provide a visual representation of an event that helps identify the problem. It should also include an event correlation engine to isolate problems and diagnostic modules to help troubleshoot the problematic storage devices. These diagnostic modules can perform both individual and system-wide tests to determine which devices are failing. Using these capabilities makes solving problems within a storage environment more efficient.

### Advanced levels of storage management

The next three tiers of the storage management pyramid represent increasing levels of quality of service. These levels, often called

storage resource management, consist of storage operations, file management, and virtualization capabilities.

### Storage operations: Ensuring secure storage management

Storage operations, the second level of the storage management pyramid, provide a more automated and reliable storage environment. Three key elements of storage operations include:

- » Data backup and recovery
- » Automated policies
- » Capacity and performance reporting and planning

### Data backup and recovery

Data backup and recovery services, vital to storage operations, can span from a simple, manually initiated tape backup to multidirectional remote mirroring. Segregating the storage environment by data types can provide guidelines for backup and recovery requirements and also substantial cost savings. If data is mission critical, then remote mirroring is a necessity for true disaster recovery. If the data represents e-mail or nonessential files, then a lower level of backup and recovery support can be used.

### Automated policies

Storage operations also incorporate automated policies, which allow event-based actions to make adjustments to the storage environment. Administrators can define and implement automated policies based on physical events so that the system will automatically recover from device failure or reroute failover paths. For example, an administrator could implement a policy that would automatically assign a hot spare to a failing array. If a monitored event reports that a hard drive failed in an array (or is about to fail), the automated policy would detect this and assign the failing array a hot spare, thereby ensuring continuous operation of the array.

Logical events can initiate other automated policies, such as adjusting available capacity based on preset or user-established thresholds. An example of these policies is *online expansion capabilities*, which enable the administrator to seamlessly expand storage without any downtime. Online expansion capabilities include:

**Online capacity expansion (OLCE).** OLCE provides the capability to add hard drives or an entirely new array to a storage subsystem and have them available without rebooting the storage subsystem.

**Online virtual disk expansion (OLVDE).** OLVDE enables a storage subsystem to add capacity to a pre-existing virtual disk and dynamically grow the size of the exported LUN. With OLVDE, the concept of a LUN is further refined to encompass virtual disks and dynamic disks. A virtual disk is made up of

1 to 32 physical disks; a dynamic disk is made up of partial or whole virtual disks. Operating system (OS) volumes may be defined on one or more dynamic disks. OLVD allows the reconfiguration of the virtual disk setup without interrupting production environment activity.

**Online volume expansion (OLVE).** OLVE allows the resizing of an OS volume without rebooting the server. While OLVD allows a storage subsystem to expand the exported LUN, OLVE allows the OS to recognize the expanded LUN and expand the file system to use the larger LUN.

To maximize storage use while minimizing management costs, automated policies such as OLCE, OLVD, and OLVE are crucial to managing the storage environment.

### Reporting on and planning for capacity and performance trends

The final element of storage operations is the ability to plan and improve the storage environment. Storage needs are growing at a tremendous rate, and for most organizations, it is no longer possible to simply order an extra drive when storage needs increase, especially since adding storage can affect system performance.

Storage management software that tracks and stores capacity and performance trends can project future storage needs and enable administrators to plan accordingly. Planning for inevitable capacity and performance needs ensures storage security and stability.

### File management: Managing data within the storage environment

File management, the next level in the storage management pyramid, represents the ability to not only manage storage devices, but also to manage the data that resides on those devices. This data management includes direct file monitoring and the monitoring and control of the space being utilized. Incorporating file management into a storage management process gives administrators an immediate and recognizable payback—the recovery of misused storage space. File management encompasses four areas:

- » File management reports and automation
- » Space allocation
- » Threshold policies
- » Utilization and charge-backs

### File management reports and automation

File reports and automated policies can help users recover a significant amount of storage space. File reports obtained through direct file monitoring provide information on the types of files that are

---

Incorporating file management into a storage management process gives administrators an immediate and recognizable payback—the recovery of misused storage space.

---

stored in the environment. File types that should be monitored include old, duplicate, oversized (80/20 rule), unauthorized, insecure, and orphaned files as well as files that have been modified but not backed up. In some cases, the management of this information can be automated through the use of policies. For example, an automated policy could, on a monthly basis, archive files that are greater than two years old and have not been updated.

A file management application also can notify administrators when storage is being misused, such as when .MPG and .MP3 files are stored on the system. Although the administrator can physically block users from storing unauthorized files, such blocking requires the use of agents to constantly

monitor the file system, which can degrade system performance for each and every I/O. With the notification process (compared to physically blocking the write), the file management application scans the file system, much the way a virus-checking application does. If unauthorized files are found, the automated policies delete them. Scheduling these scans during unused time frames allows the administrator to avoid performance degradation.

### Space allocation

Storage space always seems to be in short supply and at a premium, especially with the prevalence of graphics-intensive applications. As a result, monitoring and allocating storage space to users, groups, or applications is becoming a necessity in storage management.

File management applications can limit the space used by individuals, groups, or applications with soft or hard quotas. A soft quota provides the administrator and user with a notification or a report indicating that the allocated space has been exceeded. A periodic sweep of the file system determines whether the soft quota has been exceeded, and the sweep can be conducted when the activity will not affect performance. Hard quotas, on the other hand, actively monitor the allocated space. When the space is exceeded, it notifies users or actually prevents them from saving additional files. Unfortunately, hard quotas may also negatively affect system I/O performance.

### Threshold policies

Administrators can set threshold policies to define limits on various file management issues and then receive notification when these limits are exceeded. For example, an administrator can set a threshold on the space available on a certain volume. When the free space drops below the threshold, the management software can either notify the administrator or proactively grow the volume.

Setting thresholds can be a complicated task depending on the file management application used. In some applications, wizards or graphical user interfaces (GUIs) simplify the development of threshold policies. In other applications, these policies are created in a complicated scripting environment. Scripting enables more complex policies to be implemented but requires that the administrator learn and understand the scripting language.

Virtualization is a major step toward delivering the true promise of SAN technology.

**Utilization and charge-backs**

Because most IT departments are considered cost centers (not providing any revenue, only costs), monitoring storage becomes a financial issue. A file management application enables administrators to monitor storage utilization per user or department and charge back the storage costs to other departments within an organization. This ability allows the IT department to associate the cost of storage with the profit centers within the business, facilitating better management of the business as a whole. File management applications that monitor space utilization by owner can provide this function.

**Virtualization: The culmination of storage management capabilities**

Virtualization capabilities represent the ability to fully use all available storage by implementing *storage pooling*, a process promoted by SANs and other storage environments. Storage pooling consolidates an enterprise’s storage assets so that they can be allocated at any particular time to meet the needs of the servers and applications that require the storage.

For example, an accounting program requires significant storage space at the end of each accounting period, but between accounting periods, the space is open and available. Storage pooling allows the accounting program to expand its storage space by acquiring the excess storage from a pool of unused storage for the necessary amount of time. Once the space is no longer needed, it is returned to the pool.

Virtualization is a major step toward delivering the true promise of SAN technology. It allows administrators to redefine and parcel out various storage devices from a single console. This type of virtual storage makes physical LUN management obsolete, replacing it with attribute-based virtual disk controls, which simplify storage management. However, to use virtualization and maintain control of it, the lower levels of the storage management pyramid must also be employed. Without storage administration, storage operations, and file management, virtualization of an environment becomes an unmanageable task.

**Implementing the storage management pyramid**

When developing a storage management plan, it is important to consider two factors: the level of automation and redundancy necessary to secure the storage, and the available resources to manage the storage. Storage management applications also should be modular to allow for growth. Once a storage management application is implemented, converting to a different application is difficult.

Not every enterprise requires the implementation of all capabilities illustrated in the storage management pyramid. In some situations, an enterprise may implement some capabilities in one area and different capabilities in another area. The level of storage management required is dependent on the business needs of the storage environment, the level of reliability that the data dictates, and the budget to maintain it.

*Jon Hass (jon\_hass@dell.com) is a storage management technologist with the System Storage Group at Dell. Jon holds a B.S. in Computer Science from Texas A&M University. He has worked as a software engineer and systems architect for Data General, IBM, and most recently, Healthway Interactive, a provider of interactive television systems for the hospital industry. Jon is currently involved in the Storage Networking Industry Association (SNIA) and other organizations working to standardize management interfaces for the storage industry.*

*Keith R. Kammer (keith\_kammer@dell.com) is a senior product marketing consultant with the System Storage Group at Dell. He has more than 20 years of experience in the storage industry working with Burroughs, Unisys, and Dell. Keith holds a B.B.A. in Finance from Walsh College of Business and Accountancy in Troy, Michigan, and an M.S. in Administration from Central Michigan University.*

*Dan McConnell (dan\_mcconnell@dell.com) is a systems engineer in Storage Architecture and Technology Evaluation at Dell. He is currently responsible for evaluation of future technologies and system prototyping of storage products. He has worked as a systems analyst, software engineer, and systems architect for Southern Company, EnerLink, and Dell. Dan holds a B.S. in Computer Engineering from the Georgia Institute of Technology with a specialization in Computer Systems and Architectures.*

**FOR MORE INFORMATION**

For more information, visit <http://www.dell.com>