Chapter 16
Managing the Storage Infrastructure

KEY CONCEPTS

- Alerts
- Management Platform Standards
- Internal Chargeback

Proliferation of applications, complexity of business processes, and requirements of 24×7 availability of information have put increasingly higher demands on IT infrastructure. Managing storage infrastructure is a key that enables organizations to address these challenges and ensures continuity of business.

Comprehensive storage infrastructure management requires the implementation of intelligent processes and tools. This ensures availability and performance of all storage infrastructure elements, greater data protection and security, centralized auditing, and meeting compliance requirements. It also ensures the consolidation and better utilization of resources, limiting the need for excessive investment in technology, and helps to efficiently leverage existing resources.

Managing the storage infrastructure comprises various activities including availability, capacity, performance, and security management. Each of these aspects are interrelated and must work together to maximize the return on investment. Virtualization technologies have dramatically changed the storage management scenario and have simplified the storage infrastructure management.

Monitoring is one of the most important aspects that forms the basis for managing a storage infrastructure. Monitoring provides status of various storage components and information to perform essential management activities.

The management of the storage infrastructure components begins with defining products and services followed by defining service levels for these products and services. The service level defines the scope of the service, the associated security definitions, level of criticality to the business, and severity or priority levels. It also defines the response time associated with each level, hours of normal operation, promises of quality, and speed of service. These service levels are defined with metrics that can be monitored and measured.

Establishing management processes and implementing appropriate tools is the key to meeting service levels proactively. The management process establishes procedures for efficient handling of incidents, problems, and change requests to the storage infrastructure environment. The tools help in monitoring and executing management activities on the infrastructure. It is imperative to manage not just the individual components, but the storage infrastructure end-to-end due to the components' interdependency.

This chapter details the monitoring and other management activities of the storage infrastructure; it also describes emerging standards in storage resource management tools.

16.1 Monitoring the Storage Infrastructure

Monitoring helps to analyze the status and utilization of various storage infrastructure components. This analysis facilitates optimal use of resources and proactive management. Monitoring supports capacity planning, trend analysis, and root cause/impact analysis. As the business grows, monitoring helps to optimize the storage infrastructure resources. The monitoring process also includes the storage infrastructure's environmental controls and the operating environments for key components such as storage arrays and servers.

16.1.1 Parameters Monitored

Storage infrastructure components should be monitored for accessibility, capacity, performance, and security.

Accessibility refers to the availability of a component to perform a desired operation. A component is said to be accessible when it is functioning without any fault at any given point in time. Monitoring hardware components (e.g., a SAN interconnect device, a port, an HBA, or a disk drive) or software components (e.g., a database instance) for accessibility involves checking their availability status by listening to pre-determined alerts from devices. For example, a port may go down resulting in a chain of availability alerts. A storage infrastructure uses redundant components to avoid a single point of failure. Failure of a component may cause an outage that affects application availability, or it may cause serious performance degradation even though accessibility is not compromised.

For example, an HBA failure can restrict the server to a few paths for access to data devices in a multipath environment, potentially resulting in degraded performance. In a single-path environment, an HBA failure results in complete accessibility loss between the server and the storage. Continuously monitoring for expected accessibility of each component and reporting any deviations helps the administrator to identify failing components and plan corrective action to maintain SLA requirements.

Capacity refers to the amount of storage infrastructure resources available. Examples of capacity monitoring include examining the free space available on a file system or a RAID group, the mailbox quota allocated to users, or the numbers of ports available on a switch. Inadequate capacity may lead to degraded performance or affect accessibility or even application/service availability. Capacity monitoring ensures uninterrupted data availability and scalability by averting outages before they occur. For example, if a report indicates that 90 percent of the ports are utilized in a particular SAN fabric, a new switch should be added if more arrays and servers need to be installed on the same fabric. Capacity monitoring is preventive and predictive, usually leveraged with advanced analytical tools for trend analysis. These trends help to understand emerging challenges, and can provide an estimation of time needed to meet them.

Performance monitoring evaluates how efficiently different storage infrastructure components are performing and helps to identify bottlenecks. Performance monitoring usually measures and analyzes behavior in terms of response time or the ability to perform at a certain predefined level. It also deals with utilization of resources, which affects the way resources behave and respond. Performance measurement is a complex task that involves assessing various components on several interrelated parameters. The number of I/Os to disks, application response time, network utilization, and server CPU utilization are examples of performance monitoring.

Monitoring a storage infrastructure for *security* helps to track and prevent unauthorized access and login failures, whether accidental or malicious. Security monitoring also helps to tracks unauthorized configuration changes of storage infrastructure elements. For example, security monitoring tracks and reports the initial zoning configuration performed and all subsequent changes. Physical security of a storage infrastructure is also continuously monitored using badge readers, biometric scans, or video cameras.

16.1.2 Components Monitored

Hosts, networks, and storage are components within the storage environment that should be monitored for accessibility, capacity, performance, and security.

Hosts

Mission-critical application hosts should be monitored continuously. The accessibility of a host depends on the status of the hardware components and software processes running on it. For example, an application crash due to host hardware failure can cause instant unavailability of the data to the user.

Servers are used in a cluster to ensure high availability. In a server virtualization environment, multiple virtual machines share a pool of resources. These resources are dynamically reallocated, which ensures application accessibility and ease of management.

File system utilization of hosts also needs to be monitored. Monitoring helps in estimating the file system's growth rate and helps in predicting when it will reach 100 percent. Accordingly, the administrator can extend (manually or automatically) the file system's space proactively to prevent a failure resulting from a file system being full. New provisioning technologies even enable the allocation of storage on demand as the need arises. Alternatively, system administrators can enforce a quota for users, provisioning a fixed amount of space for their files. For example, a quota could be specified at a user level, restricting the maximum space to 10 GB per user, or at a file level that restricts a file to a maximum of 100 MB.

Server performance mainly depends on I/O profile, utilization of CPU and memory. For example, if a server running an application is experiencing 80 percent of CPU utilization continuously, this suggests that the server may be running out of processing power, which can lead to degraded performance and slower response time. Administrators can take several actions to correct the problem, such as upgrading or adding more processors, shifting the workload to different servers, and restricting the number of simultaneous client access. In a virtualized environment, CPU and memory may be allocated dynamically from another physical server or from the same server.

Memory utilization is measured by the amount of free memory available. Databases, applications, and file systems utilize the server's physical memory (RAM) for data manipulation. Insufficient memory leads to excessive swapping and paging on the disk, which in turn affects response time to the applications.

Security monitoring on servers involves tracking of login failures and execution of unauthorized applications or software processes. Proactive measures against unauthorized access to the servers are based on the threat identified. For example, an administrator can block access to an unauthorized user if multiple login failures are logged.

Storage Network

The storage network needs to be monitored to ensure proper communication between the server and the storage array. Uninterrupted access to data over the storage network depends on the accessibility of the physical and logical components in the storage network. The physical components of a storage network include elements such as switches, ports, cables, GBICs, and power supplies. The logical components include constructs, such as zones and fabrics. Any failure in the physical or logical components may cause data unavailability. For example, errors in zoning such as specifying the wrong WWN of a port results in failure to access that port, which potentially prevents access from a host to its storage.

Capacity monitoring in a storage network involves monitoring the availability of ports on a switch, the number of available ports in the entire fabric, the utilization of the inter-switch links, individual ports, and each interconnect device in the fabric. Capacity monitoring provides all required inputs for future planning and optimization of the fabric with additional interconnect devices.

Monitoring the performance of a storage network is useful in assessing individual component performance and helps to identify network bottlenecks. For example, monitoring port performance is done by measuring receive or transmit link utilization metrics, which indicate how busy the switch port is, based on expected maximum throughput. Heavily used ports can cause queuing delays on the server.

For IP networks, monitoring performance includes monitoring network latency, packet loss, bandwidth utilization for I/O, network errors, and collisions.

Storage network security monitoring provides information for any unauthorized change to the configuration of the fabric—for example, changes to the zone policies that can affect data security. Login failures and unauthorized access to switches for performing administrative changes should be logged and monitored continuously.

Storage

The accessibility of the storage array should be monitored for its hardware components and various processes. Storage arrays configured with redundant components do not affect accessibility in the event of an individual component failure, but failure of any process can disrupt or compromise business continuity operations. For example, the failure of a replication task affects disaster recovery capabilities. Some storage arrays also provide the capability to send a message to the vendor's support center in the event of hardware or process failures, referred to as a *call home*.

Capacity monitoring of a storage array enables the administrator to respond to storage needs as they occur. Information about fan-in or fan-out ratios and the availability of front-end ports is useful when a new server is given access to the storage array.

A storage array can be monitored by a number of performance metrics, such as utilization rates of the various storage array components, I/O response time, and cache utilization. A high utilization rate of storage array components may lead to performance degradation.

A storage array is usually a shared resource, which may be exposed to security breaches. Monitoring security helps to track unauthorized configuration of the storage array or corruption of data and ensures that only authorized users are allowed to access it.

16.1.3 Monitoring Examples

A storage infrastructure requires implementation of an end-to-end solution to actively monitor all the parameters of its critical components. Early detection and instant alerts ensure the protection of critical assets. In addition, the monitoring tool should be able to analyze the impact of a failure and deduce the root cause of symptoms.

Accessibility Monitoring

Failure of any component may affect the accessibility of one or more components due to their interconnections and dependencies, or it may lead to overall performance degradation. Consider an implementation in a storage infrastructure with three servers: H1, H2, and H3. All the servers are configured with two HBAs, each connected to the storage array through two switches, SW1 and SW2, as shown in Figure 16-1. The three servers share two storage ports on the storage array. Path failover software has been installed on all three servers.

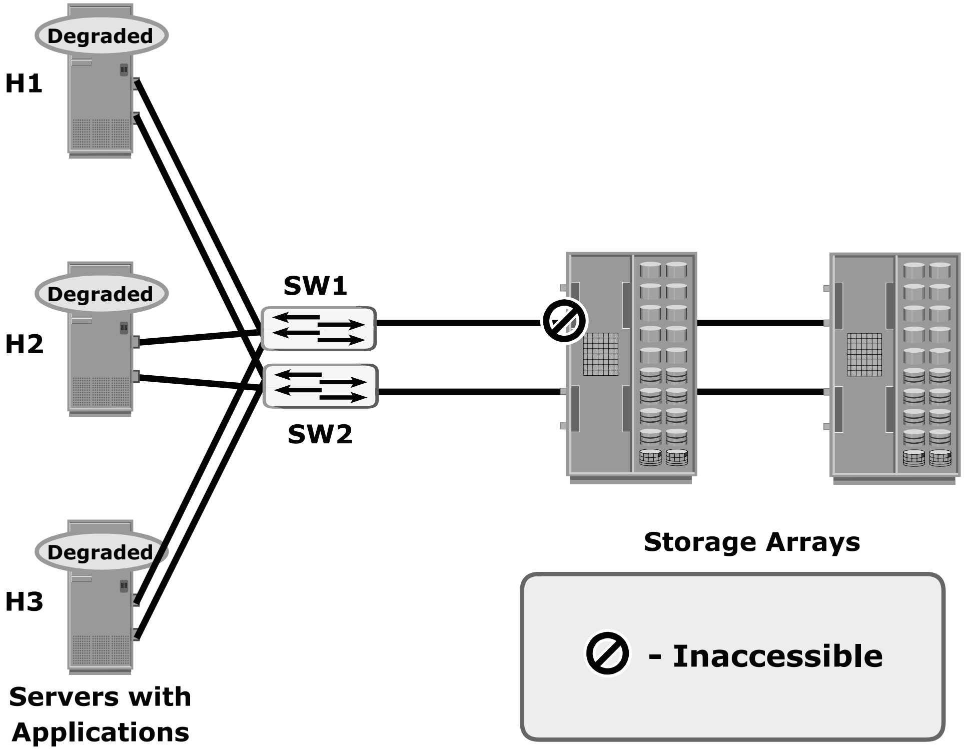

Figure 16-1: Storage array port failure in a storage infrastructure

If one of the *storage array ports* fails, all the storage volumes that were accessed through the switch connected to that port may become unavailable, depending on the type of storage array. If the storage volume becomes unavailable, path failover software initiates a path failover. However, due to redundant ports, the servers continue to access data through another switch, SW2. The servers H1, H2, and H3 may experience degraded performance due to an increased load on the path through SW2.

In the same example, if a single HBA fails on server H1, the server experiences path failure as shown in Figure 16-2. However, due to redundant HBAs, H1 can still access the storage device but it may experience degraded application response time (depends on I/O load).

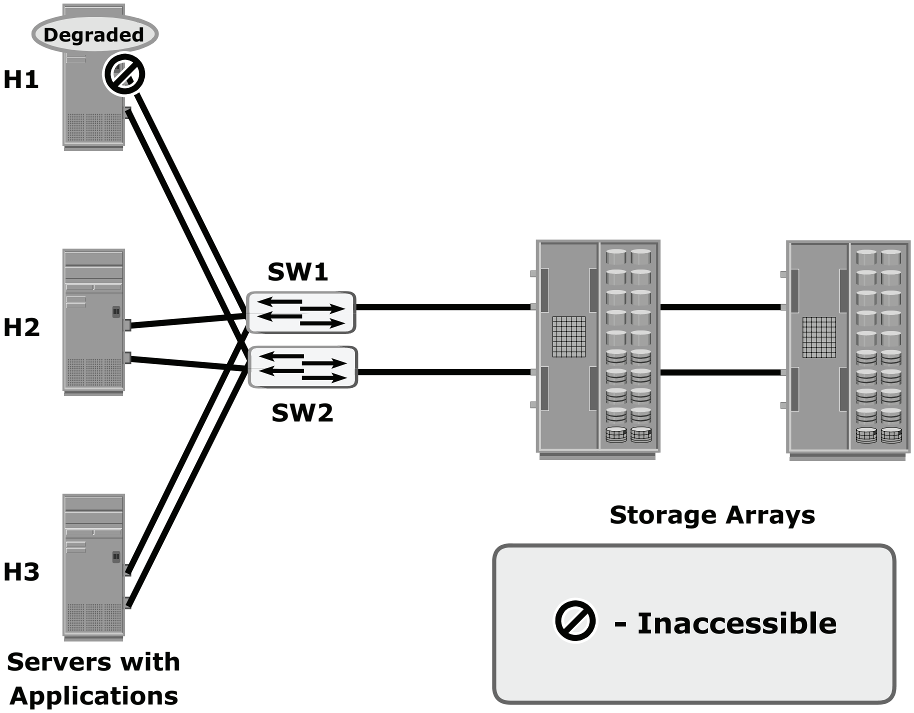

Figure 16-2: HBA failure in a storage infrastructure

Consider a scenario in which a number of servers with two HBAs each are connected to the storage array through two switches, SW1 and SW2, as shown in Figure 16-3. If SW1 fails, all the servers that were accessing the storage array through SW1 experience a path failure and data is redirected to SW2. The applications on all of the servers may experience degradation in response time depending on I/O workload. In this case, the failure of a single component has affected multiple storage infrastructure components.

Capacity Monitoring

In the scenario shown in Figure 16-4, each of the servers is allocated storage on the storage array. When a new server is deployed in this configuration, the applications on the new servers have to be given access to the storage devices from the array through switches SW1 and SW2. Monitoring the available capacity on the array helps to proactively decide whether the array can provide the required storage to the new server. Other considerations include the availability of ports on SW1 and SW2 to connect to the new server as well as the availability of storage ports to connect to the switches. Proactive monitoring also helps to identify the availability of an alternate fabric or an array to connect to the server.

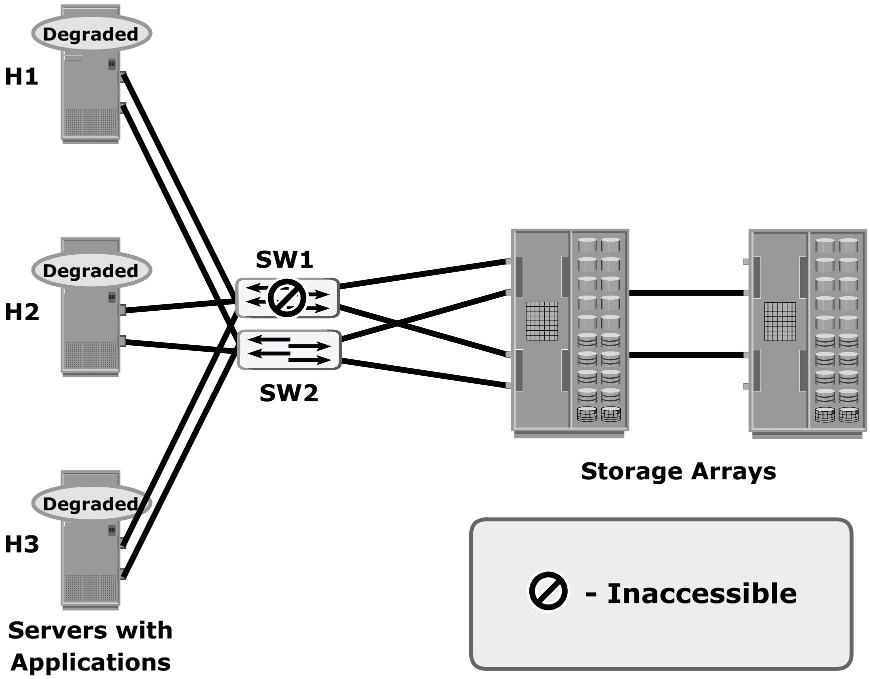

Figure 16-3: Switch failure in a storage infrastructure

The following example illustrates the importance of monitoring file system capacity on file servers. If file system capacity monitoring is not implemented, as shown in Figure 16-5 (a), and the file system is full, the application most likely will not function properly. Monitoring can be configured to issue a message when thresholds are reached on file system capacity. For example, when the file system reaches 66 percent of its capacity a warning message is issued, and a critical message when the file system reaches 80 percent of its capacity (see Figure 16-5 [b]). This enables the administrator to take action manually or automatically to extend the file system before the full condition is reached. Proactively monitoring the file system can prevent application outages caused by a lack of file system space. Applying trend analysis provides even more proactive help in dealing with such a scenario.

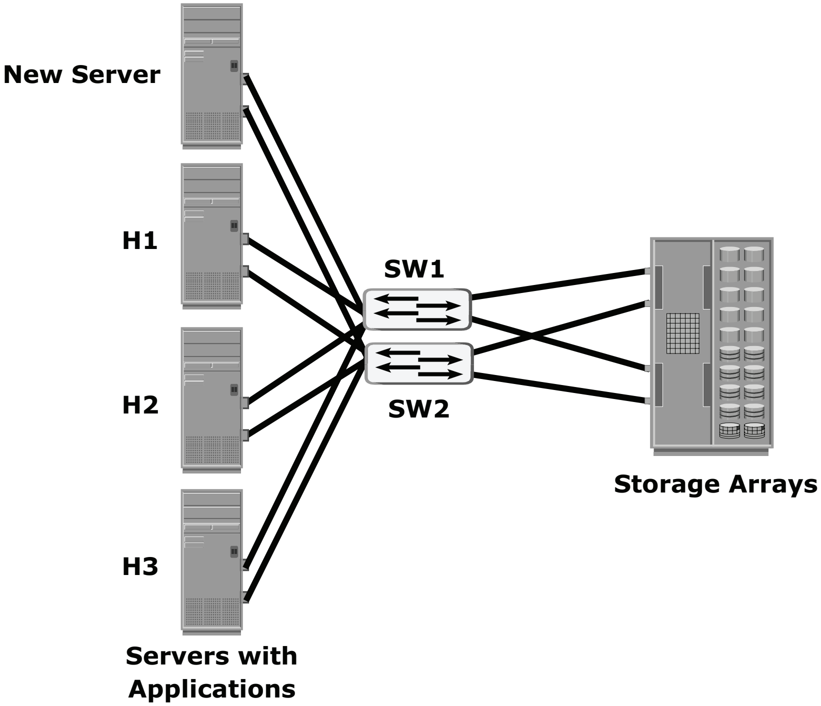

Figure 16-4: Monitoring storage array capacity

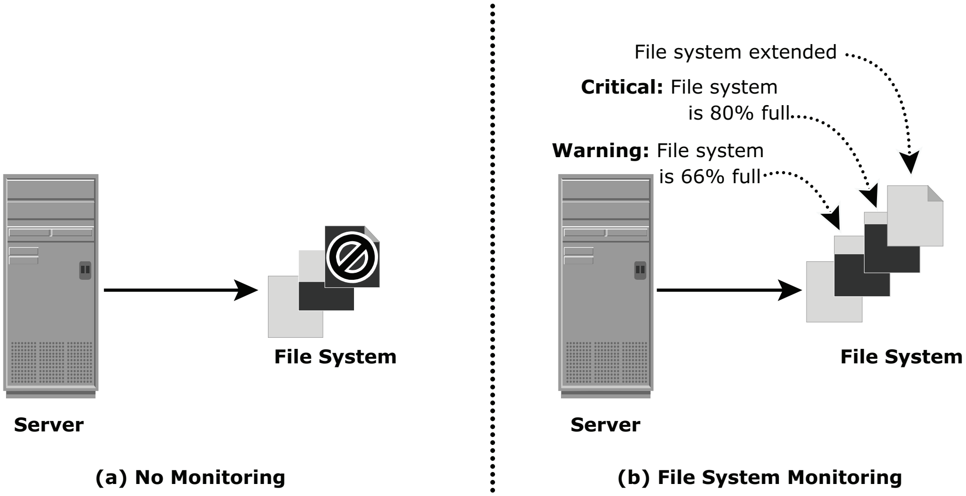
Server
File System
(a) No Monitoring
File system extended
Critical: File system is 80% full
Warning: File system is 66% full
Server
File System
(b) File System Monitoring

Figure 16-5: Monitoring server file-system space

Performance Monitoring

The example shown in Figure 16-6 illustrates the importance of monitoring performance on storage arrays. In this example, servers H1, H2, and H3 (with two HBAs each) are connected to the storage array through switches SW1 and SW2. The three servers share the same storage ports on the storage array. A new server, H4 running an application with high work load, has to be deployed to share the same storage ports as H1, H2, and H3.

Monitoring array port utilization ensures that the new server does not adversely affect the performance of the other servers. In this example, utilization for the shared ports is shown by the solid and dotted lines in the line graph for the storage ports. Notice that the port represented by a solid line is close to 100 percent utilization. If the actual utilization of both ports prior to deploying the new server is closer to the dotted line, there is room to add the new server. Otherwise, deploying the new server will affect the performance of all servers.

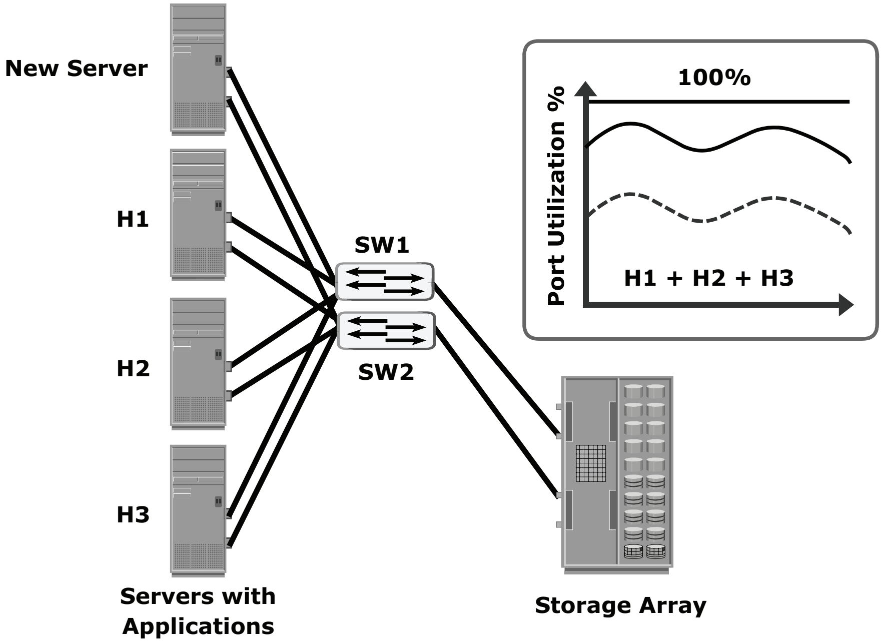

Figure 16-6: Monitoring array port utilization

Most servers offer tools that enable interactive monitoring of server CPU usage. For example, Windows Task Manager displays CPU and memory usage, as shown in Figure 16-7. These interactive tools are useful only when a few servers need to be managed. A storage infrastructure requires performance monitoring tools that are capable of monitoring many servers simultaneously. Although it is inefficient to monitor hundreds of servers continuously in real-time, this monitoring often uses polling servers at regular intervals. These monitoring tools must have the capability to send alerts whenever the CPU utilization exceeds a specified threshold.

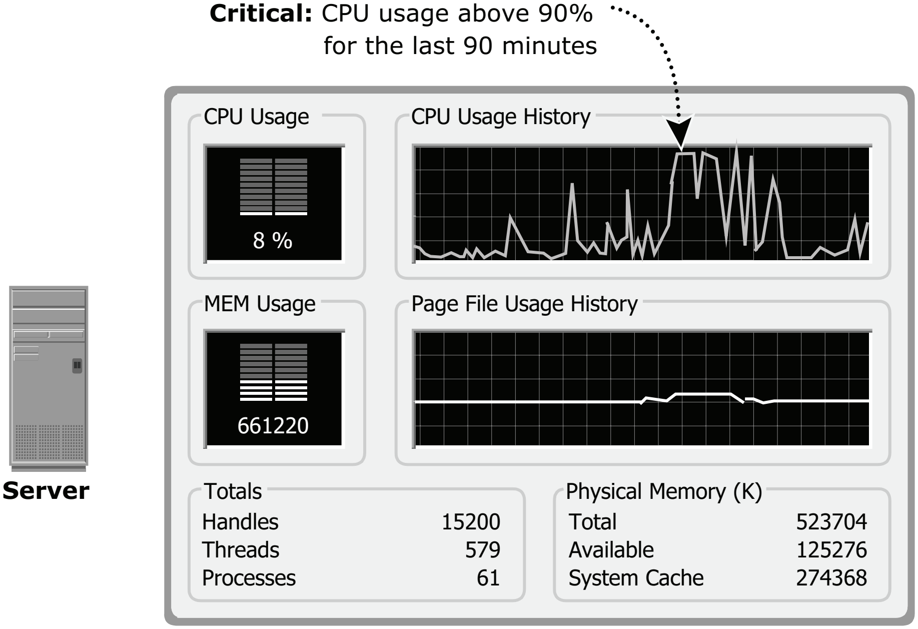

Critical: CPU usage above 90% for the last 90 minutes

Totals	
Handles	15200
Threads	579
Processes	61

Physical Memory (K)	
Total	523704
Available	125276
System Cache	274368

Figure 16-7: Monitoring the performance of servers

Security Monitoring

The example shown in Figure 16-8 illustrates the importance of monitoring security breaches in a storage array.

In this example, the storage array is shared between two workgroups, WG1 and WG2. The data of WG1 should not be accessible by WG2. Likewise, WG2 should not be accessible by WG1. A user from WG1 may try to make a local replica of the data that belongs to WG2. Usually, available mechanisms prevent such an action. However, if this action is not monitored or recorded, it is difficult to track such a violation of security protocols. Conversely, if this action is monitored, a warning message can be sent to prompt a corrective action or at least enable discovery as part of regular auditing operations.

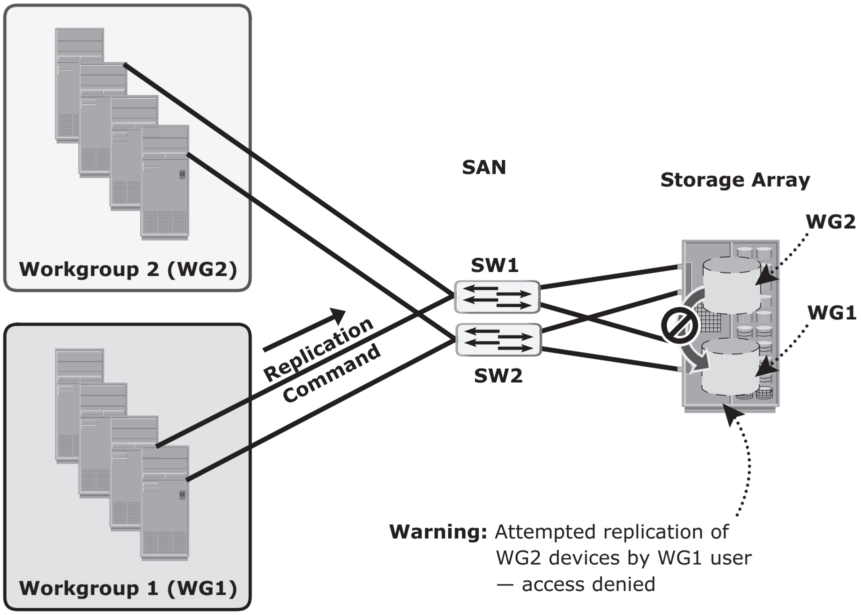

Figure 16-8: Monitoring security in a storage array

Example of host security monitoring involves login failures at the host. These login failures may be accidental (mistyping) or a deliberate attempt to access a server. Many servers usually allow two successive login failures, prohibiting additional attempts after three consecutive login failures. In most environments, login information is recorded in a system log file. In a monitored environment, three successive login failures usually triggers a message, warning of a possible security threat.

16.1.4 Alerts

Alerting of events is an integral part of monitoring. There are conditions observed by monitoring, such as failure of power, disks, memory, or switches, which may impact the availability of services that requires immediate administrative attention. Other conditions, such as a file system reaching a capacity threshold or a soft media error, are considered warning signs, and may also require administrative attention.

Monitoring tools enables administrators to assign different severity levels for different conditions in the storage infrastructure. Whenever a condition with a particular severity level occurs, an alert is sent to the administrator or triggers a script, or opens an incident ticket to initiate a corrective action. Alert classifications can range from information alerts to fatal alerts. *Information alerts* provide useful information that does not require any intervention by the administrator.

Creation of zone or LUN is an example of an information alert. *Warning alerts* require administrative attention so that the alerted condition is contained and does not affect accessibility. For example, when an alert indicates a soft media error on a disk that is approaching a predefined threshold value, the administrator can decide whether the disk needs to be replaced. *Fatal alerts* require immediate attention because the condition may affect overall performance or availability. For example, if a disk fails, the administrator must ensure that it is replaced quickly. Alerts can be assigned a severity level based on the impact of the alerted condition.

Continuous monitoring, in conjunction with automated alerting, enables administrators to respond to failures quickly and proactively. Alerting provides information to prioritize administrator's response to events.

16.2 Storage Management Activities

All the management tasks in a storage infrastructure can be broadly categorized into availability management, capacity management, performance management, security management, and reporting.

16.2.1 Availability management

The critical task in availability management is establishing a proper guideline for all configurations to ensure availability based on service levels. For example, when a server is deployed to support a critical business function, the highest availability standard is usually required. This is generally accomplished by deploying two or more HBAs, multipathing software with path failover capability, and server clustering. The server must be connected to the storage array using at least two independent fabrics and switches that have built-in redundancy. Storage devices with RAID protection are made available to the server using at least two front-end ports. In addition, these storage arrays should have built-in redundancy for various components, support backup, and local and remote replication. Virtualization technologies have significantly improved the availability management task. With virtualization in place resources can be dynamically added or removed to maintain the availability.

16.2.2 Capacity management

The goal of *capacity management* is to ensure adequate availability of resources for all services based on their service level requirements. Capacity management provides capacity analysis, comparing allocated storage to forecasted storage on a regular basis. It also provides trend analysis of actual utilization of allocated storage and rate of consumption, which must be rationalized against storage acquisition and deployment timetables.

Storage provisioning is an example of capacity management. It involves activities such as device configuration and LUN masking on the storage array and zoning configuration on the SAN and HBA components. Capacity management also takes into account the future needs of resources, and setting up monitors and analytics to gather such information.

16.2.3 Performance management

Performance management ensures the optimal operational efficiency of all components. Performance analysis is an important activity that helps to identify the performance of storage infrastructure components. This analysis provides the information — whether a component is meeting expected performance levels.

Several performance management activities are initiated for the deployment of an application or server in the existing storage infrastructure. Every component must be validated for adequate performance capabilities as defined by the service levels. For example, to optimize expected performance levels, activities on the server such as the volume configuration, designing the database, application layout configuration of multiple HBAs, and intelligent multipathing software must be fine-tuned. The performance management tasks on a SAN include designing sufficient ISLs in a multi-switch fabric with adequate bandwidth to support the required performance levels. The storage array configuration tasks include selecting the appropriate RAID type and LUN layout, front-end and back-end ports, and LUN accessibility (LUN masking) while considering the end-to-end performance.

16.2.4 Security Management

Security management prevents unauthorized access and configuration of storage infrastructure components. For example, while deploying an application or a server, the security management tasks include managing user accounts and access policies, that authorizes users to perform role-based activities. The security management tasks in the SAN environment include configuration of zoning to restrict an HBA's unauthorized access to the specific storage array ports. LUN masking prevents data corruption on the storage array by restricting host access to a defined set of logical devices.

16.2.5 Reporting

It is difficult for businesses to keep track of the resources they have in their data centers, for example, the number of storage arrays, the array vendors, how the storage arrays are being used, and by which applications. Reporting on a storage infrastructure involves keeping track and gathering information from various components/processes. This information is compiled to generate reports for trend analysis, capacity planning, chargeback, performance, and to illustrate the basic configuration of storage infrastructure components. Capacity planning reports also contain current and historic information about utilization of storage, file system, database tablespace, and ports. Configuration or asset management reports include details about device allocation, local or remote replicas, and fabric configuration; and list all equipment, with details such as their value, purchase date, lease status, and maintenance records. Chargeback reports contain information about the allocation or utilization of storage infrastructure components by various departments or user groups. Performance reports provide details about the performance of various storage infrastructure components.

16.2.6 Storage Management Examples

The discussion that follows details various storage management activities with examples.

Example 1: Storage Allocation to a New Server/Host

Consider a deployment of the new RDBMS server to the existing non-virtualized SAN environment. As a part of storage array management activities, the administrator needs to configure new volumes on the array and assign those volumes to the array front-end ports. In addition, a LUN masking configuration is performed on the storage array by assigning new servers and volumes to the storage group.

The installation and configuration of the HBA hardware (at least two to ensure redundancy) and driver has to be performed on the server before it can be physically connected to the SAN. Server reconfiguration may be required, depending on the operating system installed on the server, so it can recognize the new devices either with a *bus rescan* process or sometimes through a server reboot. Optional multipathing software can be installed on the server, which might require additional configuration.

The administrator configures the fabric's zoning policies for the new server's HBA, allowing the host to access the storage array port via the specific HBA port. This operation should probably be done at two or more fabrics to ensure redundant paths between the hosts and the storage array. The switches should have free ports available for the new server, and the array port utilization is validated against the required I/O performance of the server if the port is shared between many servers.

The volume management tasks involve the creation of volume groups, logical volumes, and file systems. The number of logical volumes or file systems to create depends on how the database or the application is expected to use the storage.

On the application side, whether it is a database or any other type of application, administrator tasks include installation of the database or the application on the logical volumes or file systems that were created. Other required activities to perform include the implementation of procedures to start the database or application. Figure 16-9 illustrates the individual tasks on the server, the SAN, and the storage arrays for this new allocation. It is a new trend in virtualization, where the application is already installed and sometimes already configured.

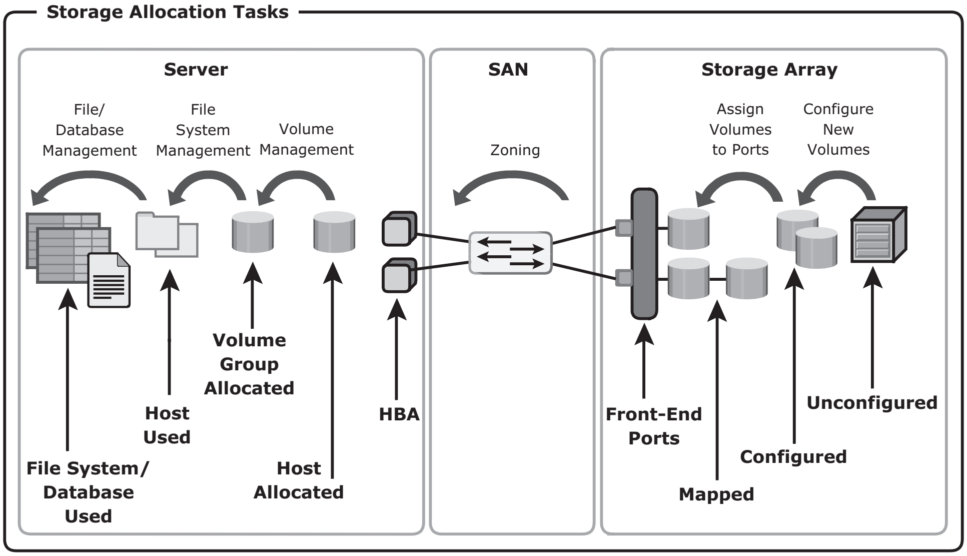

Figure 16-9: Storage allocation tasks

Example 2: File System Space Management

To prevent a file system from running out of space, administrators need to perform tasks to offload data from the existing file system. This includes deleting unwanted files and offloading files to backup media that have not been accessed for a long time.

Alternatively, an administrator can extend the file system to increase its size and avoid an application outage. The dynamic extension of file systems or a logical volume is dependent on the specific operating system or a logical volume manager (LVM) in use, and the volume management tasks detailed in the previous example may need to be commenced.

The steps and considerations for the extension of file systems are illustrated in the flow chart shown in Figure 16-10.

While extending the file system also consider whether the volume is replicated or not. If the application uses remote or local replication for business continuity operations and a new device is added to the volume group, it must be ensured that the new device is replicated as well.

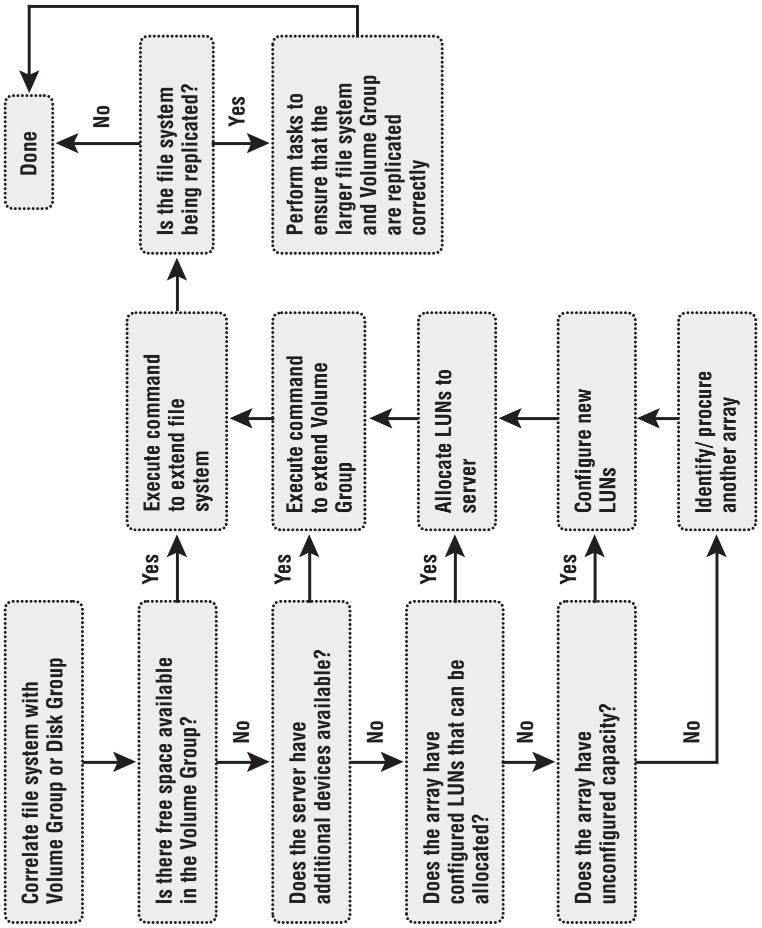
Correlate file system with Volume Group or Disk Group
Is there free space available in the Volume Group?
No
Does the server have additional devices available?
No
Does the array have configured LUNs that can be allocated?
No
Does the array have unconfigured capacity?
No
Identify/ procure another array
Yes
Yes
Yes
Yes
Configure new LUNs
Allocate LUNs to server
Execute command to extend Volume Group
Execute command to extend file system
Is the file system being replicated?
No
Done
Yes
Perform tasks to ensure that the larger file system and Volume Group are replicated correctly

Figure 16-10: Extending a file system

Example 3: Chargeback Report

This example explores the storage infrastructure management tasks necessary to create a specific report.

Figure 16-11 shows a configuration deployed in a storage infrastructure. Three servers with two HBA each are connected to a storage array via two switches, SW1 and SW2. Individual departmental applications are running on each of the servers, and array replication technology is used to create local and remote replicas. The production volume is represented as A, local replica volume as B and the remote replica volume as C.

A report documenting the exact amount of storage resources used by each application is created using a chargeback analysis for each department. If the unit for billing is based on the amount of raw storage (usable capacity plus protection provided) configured for an application used by a department, the exact amount of raw space configured must be reported for each application. A sample report is shown at the bottom of Figure 16-11. The report shows the information for two applications, Payroll_1 and Engineering_1.

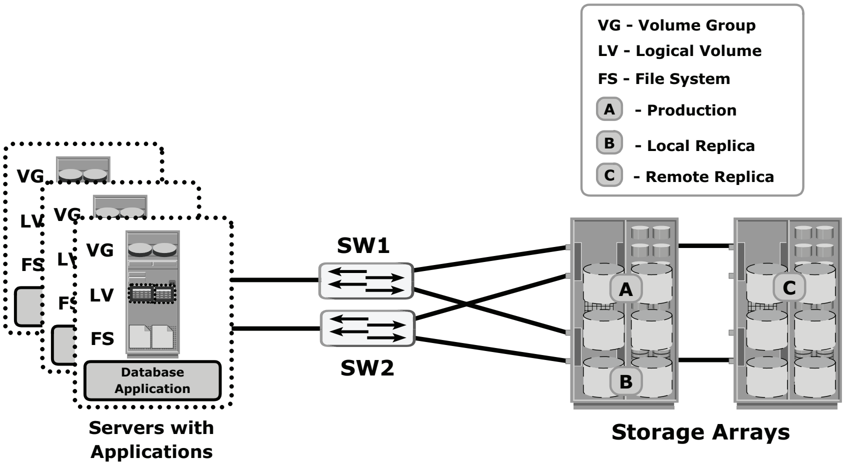

Application	Storage (GB)	Production Storage Raw (GB)	Local Replica Storage Raw (GB)	Remote Replica Storage Raw (GB)	Total Storage Raw (GB)	Chargeback Cost $ 0.25/Raw (GB)
Payroll_1	100	200	100	125	425	$ 106.25
Engineering_1	200	250	200	250	700	$ 175.00

Figure 16-11: Chargeback report

The first step to determine chargeback costs associated with an application is to correlate the application with the exact amount of raw storage configured for that application.

As indicated in Figure 16-12, the Payroll_1 application storage space is traced from file systems to logical volumes to volume groups to the LUNs on the array. When the applications are being replicated, the storage space used for local replication and remote replication is also identified. In the example shown, the application is using "Source Vol 1 and Vol 2" (in Array 1). The replication volumes are "Local Replica Vol 1 and Vol 2" (in Array 1) and "Remote Replica Vol 1 and Vol 2" (in the Remote Array). As the application grows, more file systems and storage space may be used; therefore, configuration changes are inevitable. Before a new report is generated, a correlation of the application to the array LUNs should be done to ensure that the most current information is used.

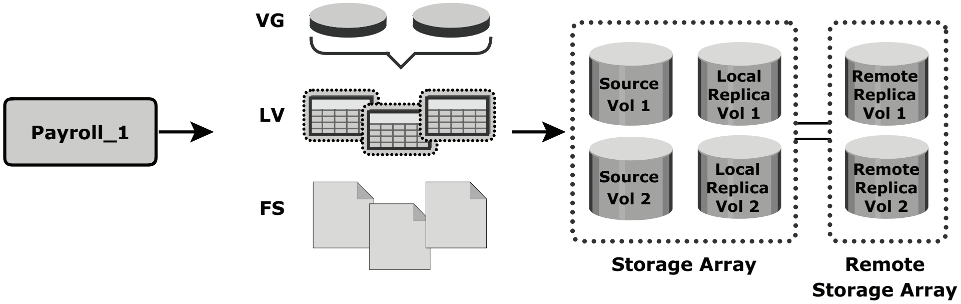

Figure 16-12: Correlation of capacity configured for an application

The amount of storage allocated to the application can be easily computed after the array devices are identified. In this example, if we assume that Source Vol 1 and Vol 2 are each 50 GB in size, the storage allocated to the application is 100 GB (50 + 50). The allocated storage for replication is 100 GB for local replication and 100 GB for remote replication. From the allocated storage, the raw storage configured for the application is determined based on the RAID protection that is used for various array devices. If the Payroll_1 application production volumes are RAID 1 protected, the raw space used by the production volumes is 200 GB. If we assume that the local replicas are on unprotected volumes and the remote replicas are protected with a RAID 5 configuration, the raw space used by local replicas is 100 GB, and 125 GB by the remote replicas. Therefore, the total raw capacity used by the Payroll_1 application is 425 GB. If provisioned capacity is equal to the consumed capacity, the total cost of storage will be $106.25 (assume cost per GB of storage is $0.25). This exercise must be repeated for each application in the enterprise to generate the required report.

Chargeback reports can be extended to include a preestablished cost for other resources, such as the number of switch ports, HBAs, and array ports in the configuration. Chargeback reports are used by data center administrators to ensure that storage consumers are well aware of the costs of the service levels they have requested.

16.3 Storage Infrastructure Management Challenges

Monitoring and managing today's complex storage infrastructure environment has become very challenging due to the number and variety of storage arrays, networks, servers, databases, and applications. There is a variety of storage devices varying in capacity, performance, and protection methodologies. Storage infrastructures deploy both SAN and IP networks and servers with different operating systems such as UNIX, LINUX, Windows, or mainframe. These products and services from multiple vendors may have interoperability issues which add complexity in managing storage infrastructure.

All of these components are provided with vendor-specific tools to manage and monitor them. However, in an environment where multiple tools are in use, it is difficult to understand the overall status of all the components and to be able to provide cross-component failure, impact, and behavior analysis. Ideally, monitoring tools should be able to correlate information from all components in one place. That way, analysis and actions are taken based on a holistic, end-to-end view of the environment and corrective measures can be taken proactively.

16.4 Developing an Ideal Solution

An ideal solution offers meaningful insight into the accessibility and status of the overall infrastructure and provides remedial solutions to each failure based on interdependencies and priorities, as discussed in the preceding section. There is value in building a central monitoring and management system that can work in a multi-vendor storage environment and is able to create an end-to-end view that includes various technology stacks and different deployment configurations. The other benefit of end-to-end monitoring is the ability to correlate one component's behavior to others. This will be helpful to debug or analyze a problem, where looking into each component individually might not be enough. The infrastructure management system should be able to gather information from all of the components and manage them through a single-user interface. It should also be able to perform root cause analysis and indicate the impact of individual component failure on various business applications/processes. In addition, it must provide a mechanism to notify administrators about various events using methods such as e-mail and SNMP traps, and generate monitoring reports or run automated scripts for task automation.

The ideal solution must be based on industry standards, leveraging common APIs, data model terminology, and taxonomy. This enables the implementation of policy-based management across heterogeneous classes of devices, services, applications, storage infrastructure, and deployed topologies.

The *SNMP* protocol was the standard used to manage multi-vendor SAN environments. However, SNMP was primarily a network management protocol and was inadequate for providing the detailed information and functionality required to manage the SAN environment. The unavailability of automatic discovery functions, weak modeling constructs, and lack of transactional support are some inadequacies of SNMP in a SAN environment. Even with these limitations, SNMP still holds a predominant role in SAN management, although newer open storage SAN management standards have emerged to monitor and manage these environments more effectively.

16.4.1 Storage Management Initiative

The Storage Networking Industry Association (SNIA) has been engaged in an initiative to develop a common, open storage, and SAN management interface. SMI-S is based on Web-Based Enterprise Management (WBEM) technology and the DMTF's Common Information Model (CIM). The initiative was formally created to enable broad interoperability among heterogeneous storage vendor systems and to enable better management solutions that span these environments. This initiative is known as the Storage Management Initiative (SMI). For more information, see `www.snia.org`.

The SMI Specification, known as SMI-S, offers substantial benefits to users and vendors. It forms a normalized, abstracted model to which a storage infrastructure's physical and logical components can be mapped, and which can be used by management applications such as storage resource management, device management, and data management for standardized, effective, end-to-end control of storage resources (see Figure 16-13).

Using SMI-S, the storage software developers have a single normalized and unified object model comprising the detailed document that contains information about managing the breadth of SAN components. Moreover, SMI-S eliminates the need for development of vendor-proprietary management interfaces, enabling vendors to focus on added value functions and offering solutions in a way that will support new devices as long as they adhere to the standard. Using SMI-S, device vendors can build new features and functions to manage storage subsystems and expose them via SMI-S. The SMI-S-compliant products lead to easier, faster deployment, and accelerated adoption of policy-based storage management frameworks.

The information required to perform management tasks is better organized or structured in a way that enables disparate groups of people to use it. This can be accomplished by developing a model or representation of the details required by users working within a particular domain. Such an approach is referred to as an *information model*. An information model requires a set of legal statements or syntax to capture the representation and expressions necessary to manage common aspects of that domain.

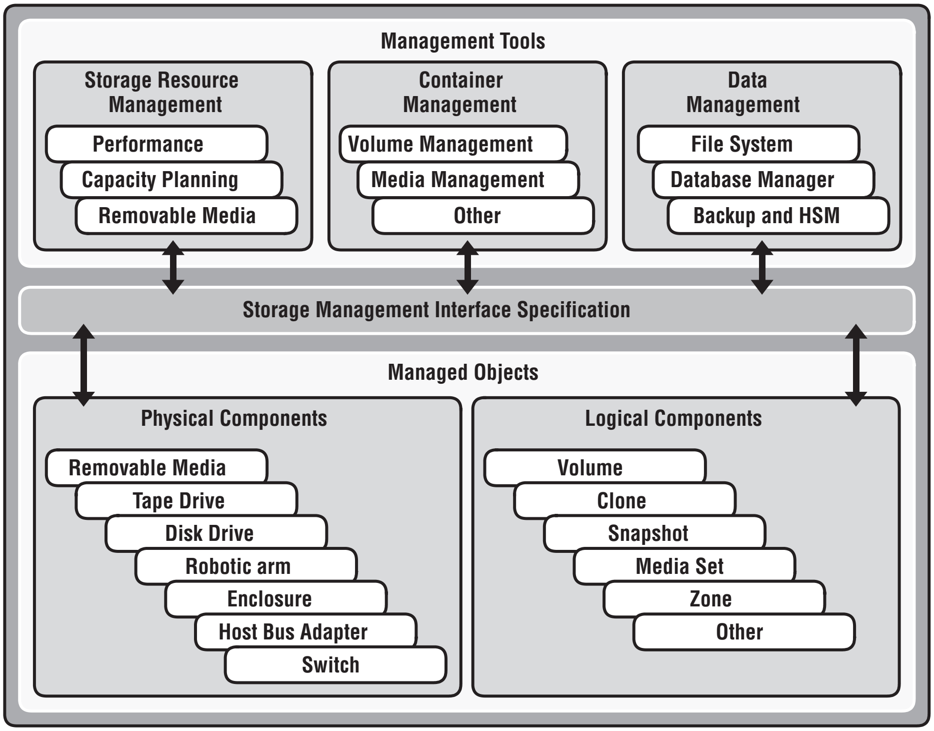

Figure 16-13: Storage Management Initiative Specification

The CIM is a language and methodology for describing management elements. A CIM schema includes models for systems, applications, networks, and devices. This schema also enables applications from different vendors working on different platforms to describe the management data in a standard format so that it can be shared among a variety of management applications.

WBEM is an initiative in the DMTF. It is a set of technologies that enables interoperable management of an enterprise. WBEM consists of CIM, an XML DTD defining the tags (XML encodings) to describe the CIM Schema and its data, and a set of HTTP operations for exchanging the XML-based information. The main objective of developing a WBEM is to unify the management of enterprise-computing environments that were traditionally administered through the management stacks, such as SNMP and Common Management Information Protocol (CMIP).

WBEM includes the CIM standard, the xmlCIM encoding specification, and a transport mechanism — CIM operation over HTTP.

The following features of SMI-S simplify SAN management:

- **Common data model:** SMI-S agents interact with an SMI-S-enabled device, such as a switch, a server, or a storage array, to extract relevant management data. They can also interact at the management layer to exchange information between one management application and another. They then provide this information to the requester in a consistent syntax and format.
- **Interconnect independence:** SMI-S eliminates the need to redesign the management transport and enables components to be managed by using out-of-band communications. In addition, SMI-S offers the advantages of specifying the CMI-XML over the HTTP protocol stack and utilizing the lower layers of the TCP/IP stack, both of which are ubiquitous in today's networking world.
- **Multilayer management:** SMI-S can be used in a multilayer and cross-domain environment—for example, server-based volume managers and network storage appliances. Many storage deployment environments currently employ this combination.
- **Legacy system accommodation:** SMI-S can be used to manage legacy systems by using a proxy agent or can be directly supported by the device itself. SMI-S can coexist with proprietary APIs and agents as well as providing integration framework for such mechanisms.
- **Policy-based management:** SMI-S includes object models applicable across all classes of devices, enabling a SAN administrator to implement policy-based management for entire storage networks.

A Common Information Model eases the huge tasks of interpreting relations and correlations of information across technology stacks and vendor devices. A CIM makes it easier for management applications as well as device/solution developers, as it creates normalized abstracted targets when both management and managed elements interact.

16.4.2 Enterprise Management Platforms

Enterprise management platforms (EMPs) are complex applications or suites of applications that provide an integrated solution for managing and monitoring an enterprise storage infrastructure. These applications proactively monitor many storage infrastructure components and are designed to alert users about any problems or the approach of a predefined threshold in monitored components.

These alerts are either shown on a console depicting the faulty component in a different color, or they can be configured to send the alert by e-mail or other means. In addition to monitoring functionality, the EMPs can provide the necessary management functionality, which can be natively implemented by code embedded into the EMP or may involve launching the proprietary management utility supplied by the target component manufacturer.

Other functionality enables easy scheduling of operations that must be performed regularly, such as the provisioning of resources, configuration management, and fault investigation. These platforms also provide extensive analytical, remedial, and reporting capabilities to ease storage infrastructure management. ControlCenter, described in this chapter, is an example of an EMP implementation.

16.5 Concepts in Practice: EMC ControlCenter

Businesses today are facing the challenge of managing storage resources. Management of storage environments has become complex due to the large number of physical and logical components from multiple vendors and complex technology stacks that involve physical, logical, structural, and virtualization concepts. Therefore organizations must need an integrated solution to effectively manage their storage resources from end to end.

EMC ControlCenter storage management suite provides an end-to-end integrated approach for dealing with multi-vendor storage reporting, monitoring, configuration, and control tasks. Using ControlCenter, administrators can see an end-to-end view of the storage infrastructure, understand how the infrastructure is performing, and do what is necessary to ensure that service levels are met. This results in better performance, improved productivity, and reduced costs.

EMC Navisphere Manager is a suite of management tools for managing EMC CLARiiON storage arrays (which is discussed briefly in Chapter 4). It enables access and management of all CLARiiON's advanced software functionality including Navisphere Quality of Service Manager, Navisphere Analyzer, SnapView, SAN Copy, and MirrorView. It can be launched from EMC Control Center. Visit `http://education.EMC.com/ismbook` for the latest information.

16.5.1 ControlCenter Features and Functionality

ControlCenter provides many features and functions to monitor and manage storage resources effectively. By supporting interoperability in a heterogeneous virtualized storage environment, it enables users to manage hardware, software and integrate resources, in multiple sites under a single management umbrella.

In addition, ControlCenter enables planning and provisioning of storage resources. It also provides customized monitoring and reporting using rule-based event alerting. Furthermore, ControlCenter provides enhanced security by enabling policy-based permission and authorization to grant controlled access to systems, operations, and information. In addition to helping to analyze and resolve existing failures, ControlCenter predicts failure scenarios and optimizes storage resource performance.

16.5.2 ControlCenter Architecture

ControlCenter is an n-tier, distributed application that consists of a user interface tier and a consolidation and analytics (infrastructure) tier as shown in Figure 16-14. The user interface tier is where user interactions are initiated and the results are reviewed for monitoring and managing the storage environment. The infrastructure tier is the data processing tier, consisting of the ControlCenter server, repository, store, and agents. Agents are software programs running on the host to collect and monitor host information and send it to the store in the infrastructure tier. In addition, the agent runs the user commands to configure and monitor activities of the storage environment components.

User Interface Tier

ControlCenter supports different types of interfaces to interact with users or administrators. A *console* is the primary interface to view, manage, configure, and handle reporting of various components (managed objects) and other required views. It is a Java-based application installed through a Web browser and launched from a desktop icon. For an object (storage infrastructure entity) to be displayed on the console, a ControlCenter agent must first discover it. Any command entered at the console is passed from the console to a ControlCenter server and then forwarded to the appropriate agent. The entities monitored and managed by various agents that appear on the console are organized into groups such as storage systems, hosts, and connectivity or may be manually managed in user-defined groups. Information about the objects is retrieved by a console from a repository or in real time directly from the agent.

The *Web console* is a web-based interface that provides support for remote or high latency local networks. It provides a portable solution because it does not require local installation.

The *StorageScope console* helps to view configuration, status, and usage information for individual objects, user-defined groups, or the entire enterprise storage environment. It is an interface to the StorageScope applications that monitor and report on all storage assets and their usage. StorageScope is bundled with ControlCenter as a Storage Resource Management (SRM) Tool. It enables users to collect high-level and detailed data about the storage components. It helps to view point-in-time snapshots of high-importance areas of your storage environment on StorageScope's customizable Dashboard page.

Infrastructure Tier

EMC ControlCenter's infrastructure tier has the following main components: ControlCenter server, Performance Manager server, Repository, Store, and StorageScope repository (refer to Figure 16-14). A *ControlCenter server* is the central component in the ControlCenter infrastructure and provides the primary interface for most components including the console. It retrieves data from the repository for the console to display and handles user-initiated requests for real-time data.

The *Repository* is a licensed, embedded Oracle 10g R2 database that holds current and historical information about both the storage environment and ControlCenter itself with the exception of performance data. This data includes configuration details, statistical data, alerts, and status information about any given device. It also contains general information about links, groups, metadata, alert definitions, components, and the data dictionary. The ControlCenter Store populates the repository with data sent by the agents. The ControlCenter Server processes transactions from the console for repository data, such as checking user group permissions. The repository has restricted access and can be updated only by the ControlCenter Server.

The *Store* is a process that populates the Repository with persistent data from the agents. It provides a store-and-retrieve interface between the agents and the Repository. ControlCenter can contain more than one store for load balancing and failover in large environments.

The *StorageScope Repository* is built on an Oracle 10g R2 and is populated with information that helps the business to make decisions about storage utilization and configuration. It provides managerial summaries and responses to online queries, and contains historical information that enables performance and utilization analysis over time. Moreover, it can be effectively used to reclaim storage resources by identifying unused or underutilized storage, as well as duplicate, rarely accessed, or nonbusiness files. It also facilitates chargeback and billing operations. StorageScope analytical modules determine future storage needs based on hierarchical usage and trends analysis.

The StorageScope Repository tables are populated through the extract, transform, and load (ETL) processes from the ControlCenter Repository.

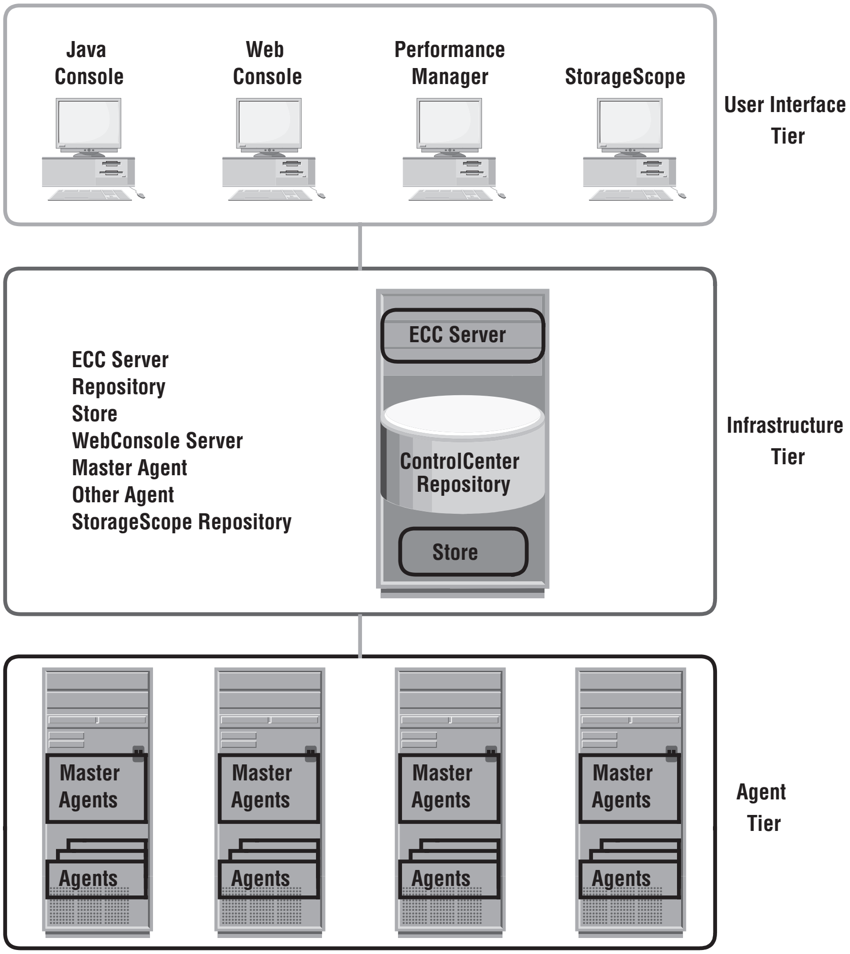

Figure 16-14: ControlCenter architecture

Agents

EMC ControlCenter uses agents installed on hosts to monitor and control the storage infrastructure components, or "managed objects" in ControlCenter. There are two type of agents; master agent and other agents. The master agent controls all other ControlCenter agents on a host by starting and stopping the agents and facilitates their remote installation and upgrade. In ControlCenter, only one master agent is required per host (object). The other agents run on the host to collect data and monitor the accessibility of storage infrastructure components. Multiple ControlCenter agents can exist on a host to monitor events and organize, analyze, and interpret collected information. In addition, the agents generate alerts when monitored events occur, passing information to the ControlCenter Store and the Server.

Different agent types manage different objects, such as storage agents for storage arrays, host agents for operating systems, and database agents for database applications. The agents collect data from the objects and analyze commands sent from the server. As defined by the agent policy, the store periodically polls every agent to retrieve updated information.

Each agent type has a predefined set of actions called a Data Collection Policy (DCP) specific to the type of objects it manages. New policies can be defined, or existing policies can be modified based on available templates. From these policies, an agent can collect five types of information about managed objects or components: discovery, configuration, status, performance, and system log data. The discovery data contains the discovered objects' names and types. Configuration data refers to an object's configuration, hierarchy, and subcomponents. The status data contains information about the state of objects and their components, with an additional explanation of their status. The status can be ok, offline, or error. Performance statistics indicate bandwidth, disk space, memory, and CPU performance.

The *Performance Manager* is packaged as part of the SRM monitoring and reporting offering. It is the ControlCenter performance analysis tool that provides the capability to quickly analyze and report on the performance and configuration data collected by the agents. Agents use data-collection policies to obtain historical, operational, or performance-related data. The Performance Manager function has an automated report generation and distribution feature supported by a job scheduler that can be displayed at the console. It is a Windows-based, post-processing tool that is invoked after data collection and processing is complete. Performance Manager helps to create data views of system performance and configuration. *Workload Analyzer Archiver*, a part of Performance Manager processes and stores the data collected by ControlCenter agents as performance archives, revolving collections, and analyst collections. The performance archives and collections are then available for viewing through Performance Manager Console and Performance Manager Automated reports. Figure 16-15 shows Window-based Performance Manager interface of ControlCenter.

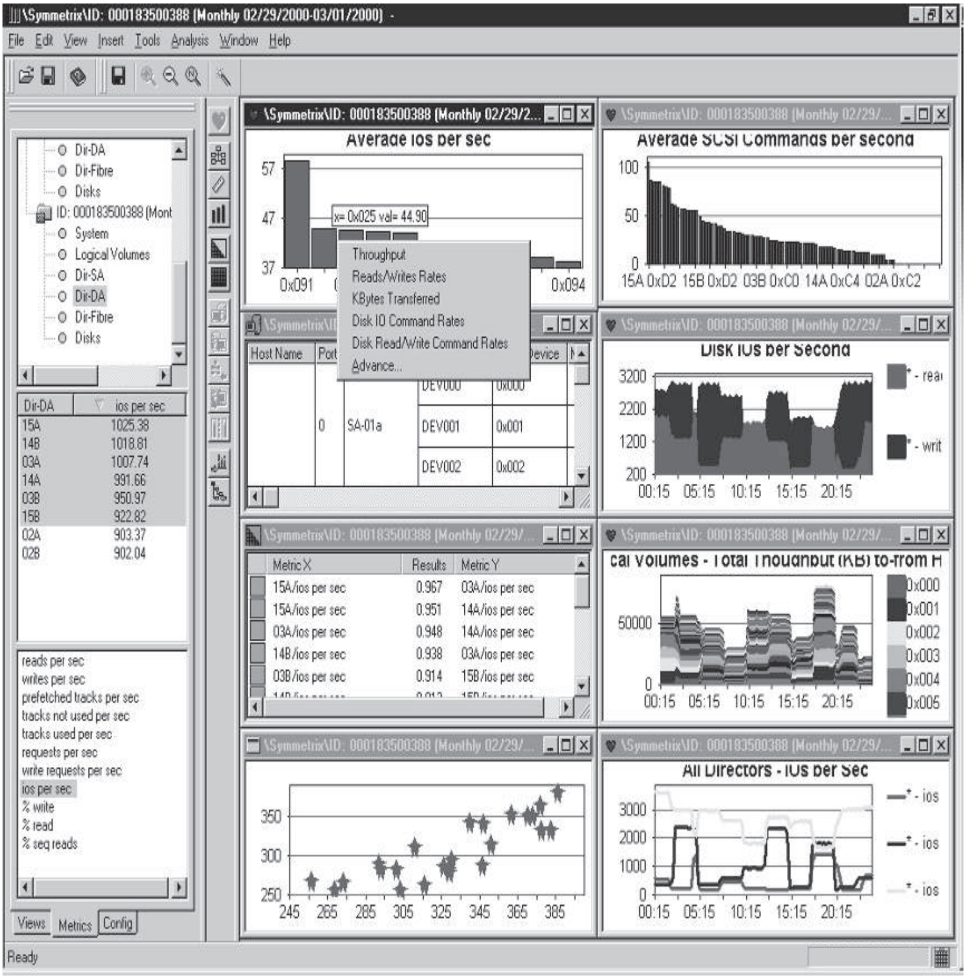

Dir-DA	ios per sec
15A	1025.38
14B	1018.81
03A	1007.74
14A	991.66
03B	950.97
15B	922.82
02A	903.37
02B	902.04

Metric X	Results	Metric Y
15A/ios per sec	0.967	03A/ios per sec
15A/ios per sec	0.951	14A/ios per sec
03A/ios per sec	0.948	14A/ios per sec
14B/ios per sec	0.938	03A/ios per sec
03B/ios per sec	0.914	15B/ios per sec

Figure 16-15: ControlCenter Performance Manager

ControlCenter Monitoring through Alerts

Each EMC ControlCenter agent can trigger predefined alerts based on monitoring various aspects of the distributed storage environment. These alerts enable the management of data availability and the proactive maintenance of various components. Alerts can be categorized as health, capacity, or performance alerts. Health alerts indicate the availability and status of ControlCenter components, capacity alerts specify the space availability of a volume or disk, and performance alerts indicate performance and usage of subsystems.

Agents trigger alerts based on a predefined alert matrix or conditions such as space availability or status of components. The alerts can be displayed on the console, or sent to the administrator or a user via e-mail, or directed to the management framework using the integration gateway or SNMP.

ControlCenter Add-on Products

ControlCenter can be effectively deployed to automatically build an accurate topology of different pieces of a multivendor SAN environment. Therefore, storage infrastructure administrators need to understand how these elements are interconnected and the nature of these interrelationships. They can monitor these elements and receive notifications when alert conditions are satisfied. ControlCenter helps to consistently manage multivendor SANs. Other products in the ControlCenter family, such as SAN Manager and SAN Advisor, provide additional functionality that facilitates administrators to enhance the overall monitoring and management processes. Some of this functionality includes LUN masking for EMC ControlCenter and other vendor storage and switch-zoning configuration setup capabilities for each major switch manufacturer. It also provides the capability to explore "what-if" scenarios, which administrators can use to quickly build an agile infrastructure that meets the needs of their business.

The provisioning of storage resources can be automated with the automated Resource Manager product. With many separate tasks being automated, the storage infrastructure can meet required service-level agreements while reducing management costs and the risk of error.

Summary

This chapter detailed the required proactive monitoring and management of storage infrastructure components that ensures established service levels are met for accessibility, capacity, performance, and security. The growing demand for storage coupled with heterogeneous operating environments has increased the complexity of monitoring and managing the storage environment. This chapter provided examples demonstrating the need for monitoring and management functionalities, and detailed industry standards, including SMI-S, that have emerged to minimize storage management interoperability concerns in the enterprise.

Visit `http://education.emc.com/ismbook` for additional reading materials.

EXERCISES

1. This chapter described how a storage array initiates a “call home,” an automatic alert sent to the vendor’s support team that enables the team to access the array to remotely fix an error. Discuss and detail the security implications of this procedure, and develop a procedure to mitigate any security threats that you may identify in providing access to remote support teams.
2. A performance problem has been reported on a database. Monitoring confirms that at 12:00 AM, a problem surfaced, and access to the database is severely affected until 3:00 PM every day. This time slot is critical for business operations and an investigation has been launched. A reporting process that starts at 12:00 PM contends for database resources and constrains the environment. What monitoring and management procedures, tools, and alerts would you establish to ensure accessibility, capacity, performance, and security in this environment?
3. Research SMI-S and write a technical paper on different vendor implementations of storage management solutions that comply with SMI-S.