



Section 5: Securing and Managing the Storage Infrastructure

This section contains two module:

- Module 14: Securing the Storage Infrastructure
- Module 15: Managing the Storage Infrastructure

Welcome to Section 5: Securing and Managing the Storage Infrastructure. This section covers two modules, namely Securing the Storage Infrastructure and Managing the Storage Infrastructure.



Module 14: Securing the Storage Infrastructure

Upon completion of this module, you should be able to:

- Describe information security framework
- Explain various storage security domains
- Discuss security implementations in SAN, NAS, and IP SAN
- Explain security in virtualized and cloud environments

This module focuses on information security framework and various storage security domains. This module also focuses on security implementation in SAN, NAS, and IP SAN. Further, this module focuses on security in virtualized and cloud environment.

Module 14: Securing the Storage Infrastructure

Lesson 1: Information Security Framework


During this lesson the following topics are covered:

- Building information security framework
- Risk triad
- Security elements
- Security controls

This lesson covers building information security framework and risk triad. This lesson also covers security elements such as assets, threats and vulnerabilities. Additionally this lesson also focuses on security controls.

Storage Security

- Process of applying information security principles and practices within the domain of storage networking technologies
- Storage security focuses on securing access to information by implementing safeguards or controls
- Storage security begins with building 'information security framework'



Valuable information, including intellectual property, personal identities, and financial transactions, is routinely processed and stored in storage arrays, which are accessed through the network. As a result, storage is now more exposed to various security threats that can potentially damage business-critical data and disrupt critical services. Securing storage infrastructure has become an integral component of the storage management process in traditional and virtualized data centers. It is an intensive and necessary task, essential to managing and protecting vital information.

Storage security is the process of applying information security principles and practices within the domain of storage networking technologies. Storage security implements various kinds of safeguards or controls, in order to lessen the risk of an exploitation or a vulnerability in the storage network which could otherwise cause a significant impact to organization's business. From this perspective, security is an ongoing process, not static and requires continuing revalidation and modification. Storage security begins with building a framework.

Information Security Framework

- A systematic way of defining security requirements
- Framework should incorporate:
 - ▶ Anticipated security attacks
 - ▶▶ Actions that compromise the security of information
 - ▶ Security measures
 - ▶▶ Control designed to protect from these security attacks
- Security framework is built to achieve four security goals:
 - ▶ Confidentiality
 - ▶ Integrity
 - ▶ Availability
 - ▶ Accountability
- Securing infrastructure begins with understanding the risk

The basic information security framework is built to achieve four security goals, confidentiality, integrity, and availability (CIA) along with accountability. This framework incorporates all security standards, procedures and controls, required to mitigate threats in the storage infrastructure environment.

Confidentiality : Provides the required secrecy of information and ensures that only authorized users have access to data. This requires authentication of users who need to access information. Data in transit (data transmitted over cables) and data at rest (data residing on a primary storage, backup media, or in the archives) can be encrypted to maintain its confidentiality. In addition to restricting unauthorized users from accessing information, confidentiality also requires to implement traffic flow protection measures as part of the security protocol. These protection measures generally include hiding source and destination addresses, frequency of data being sent, and amount of data sent.


Integrity: Ensures that the information is unaltered. Ensuring integrity requires detection and protection against unauthorized alteration or deletion of information. Ensuring integrity stipulate measures such as error detection and correction for both data and systems.

Availability: This ensures that authorized users have reliable and timely access to systems, data and applications residing on these systems. Availability requires protection against unauthorized deletion of data and denial of service. Availability also implies that sufficient resources are available to provide a service.

Accountability: Refers to accounting for all the events and operations that take place in the data center infrastructure. The accountability service maintains a log of events that can be audited or traced later for the purpose of security.

Risk Triad

- Defines risk in terms of threats, assets, and vulnerabilities



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 5: Securing and Managing the Storage Infrastructure 6

Risk triad defines risk in terms of threats, assets, and vulnerabilities. Risk arises when a threat agent (an attacker) uses an existing vulnerability to compromise the security services of an asset, for example, if a sensitive document is transmitted without any protection over an insecure channel, an attacker might get unauthorized access to the document and may violate its confidentiality and integrity. This may, in turn, result in business loss for the organization. In this scenario potential business loss is the risk, which arises because an attacker uses the vulnerability of the unprotected communication to access the document and tamper with it.

To manage risks, organizations primarily focus on vulnerabilities because they cannot eliminate threat agents that appear in various forms and sources to its assets. Organizations can enforce countermeasures to reduce the possibility of occurrence of attacks and the severity of their impact.

Risk assessment is the first step to determine the extent of potential threats and risks in an IT infrastructure. The process assesses risk and helps to identify appropriate controls to mitigate or eliminate risks. Based on value of assets, risk assessment helps to prioritize the investment and provisioning of security measures. To determine the probability of an adverse event occurring, threats to an IT system must be analyzed with the potential vulnerabilities and the existing security controls.

The severity of an adverse event is estimated by the impact that it may have on critical business activities. Based on this analysis, a relative value of criticality and sensitivity can be assigned to IT assets and resources. For example, a particular IT system component may be assigned a high-criticality value if an attack on this particular component can cause a complete termination of mission-critical services.

Assets

- “Information” – the most important asset for any organization
 - ▶ Other assets include hardware, software, and network infrastructure
- Protecting assets is the primary concern
- Security considerations
 - ▶ Must provide easy access to assets for authorized users
 - ▶ Cost of securing the assets should be a fraction of the value of the assets
 - ▶ Make it difficult for potential attackers to access and compromise the assets
 - ▶▶ Should cost heavily to a potential attacker in terms of money, effort, and time

Information is one of the most important *assets* for any organization. Other assets include hardware, software, and other infrastructure components required to access the information. To protect these assets, organizations must develop a set of parameters to ensure the availability of the resources to authorized users and trusted networks. These parameters apply to storage resources, network infrastructure, and organizational policies.

Security methods have two objectives. The first objective is to ensure that the network is easily accessible to authorized users. It should also be reliable and stable under disparate environmental conditions and volumes of usage. The second objective is to make it difficult for potential attackers to access and compromise the system.

The security methods should provide adequate protection against unauthorized access, viruses, worms, trojans, and other malicious software programs. Security measures should also include options to encrypt critical data and disable unused services to minimize the number of potential security gaps. The security method must ensure that updates to the operating system and other software are installed regularly. At the same time, it must provide adequate redundancy in the form of replication and mirroring of the production data to prevent catastrophic data loss if there is an unexpected data compromise. For the security system to function smoothly, all users are informed about the policies governing the use of the network.

The effectiveness of a storage security methodology can be measured by two key criteria. One, the cost of implementing the system should be a fraction of the value of the protected data. Two, it should cost heavily to a potential attacker, in terms of money, effort, and time.

Threats

- Potential attacks that can be carried out on an IT infrastructure
- Attacks can be classified as passive or active
 - ▶ Passive attacks
 - ▶▶ Attempt to gain unauthorized access into the system
 - ▶▶ Attempt to threaten the confidentiality of information
 - ▶ Active attacks
 - ▶▶ Attempt data modification, Denial of Service (DoS), and repudiation attacks
 - ▶▶ Attempt to threaten data integrity, availability, and accountability

Threats are the potential attacks that can be carried out on an IT infrastructure. These attacks can be classified as active or passive. Passive attacks are attempts to gain unauthorized access into the system. They pose threats to confidentiality of information. Active attacks include data modification, denial of service (DoS), and repudiation attacks. They pose threats to data integrity, availability, and accountability.

In a data modification attack, the unauthorized user attempts to modify information for malicious purposes. A modification attack can target the data at rest or the data in transit. These attacks pose a threat to data integrity.

Denial of service (DoS) attacks prevent legitimate users from accessing resources and services. These attacks generally do not involve access to or modification of information. Instead, they pose a threat to data availability. The intentional flooding of a network or website to prevent legitimate access to authorized users is one example of a DoS attack.

Repudiation is an attack against the accountability of information. It attempts to provide false information by either impersonating someone's identity or denying that an event or a transaction has taken place. For example, a repudiation attack may involve performing an action and eliminating any evidence that could prove the identity of the user (attacker) who performed that action. Repudiation attacks include circumventing the logging of security events or tampering with the security log to conceal the identity of the attacker.

Vulnerabilities

- Paths that provide access to information are vulnerable to potential attacks
- Requires implementation of “defense in depth”
- Factors to consider when assessing the extent to which an environment is vulnerable:
 - ▶ Attack surface
 - ▶ Attack vectors
 - ▶ Work factor
- Managing vulnerabilities
 - ▶ Minimize the attack surface and maximize the work factor
 - ▶ Install controls (or countermeasures)

The paths that provide access to information are vulnerable to potential attacks. Each of the paths may contain various access points, which provide different levels of access to the storage resources. It is important to implement adequate security controls at all the access points on an access path. Implementing security controls at each access point of every access path is known as *defense in depth*. Defense in depth recommends using multiple security measures to reduce the risk of security threats if one component of the protection is compromised. It is also known as a “layered approach to security”. Because there are multiple measures for security at different levels and defense in depth gives additional time to detect and respond to an attack. This can reduce the scope or impact of a security breach. *Attack surface*, *attack vector*, and *work factor* are the three factors to consider when assessing the extent to which an environment is vulnerable to security threats. *Attack surface* refers to the various entry points that an attacker can use to launch an attack. Each component of a storage network is a source of potential vulnerability. An attacker can use all the external interfaces supported by that component, such as the hardware and the management interfaces, to execute various attacks. These interfaces form the attack surface for the attacker. Even unused network services, if enabled, can become a part of the attack surface.

Cont..

An *attack vector* is a step or a series of steps necessary to complete an attack. For example, an attacker might exploit a bug in the management interface to execute a snoop attack whereby the attacker can modify the configuration of the storage device to allow the traffic to be accessed from one more host. This redirected traffic can be used to snoop the data in transit. *Work factor* refers to the amount of time and effort required to exploit an attack vector. For example, if attackers attempt to retrieve sensitive information, they consider the time and effort that would be required for executing an attack on a database. This may include determining privileged accounts, determining the database schema, and writing SQL queries. Instead, based on the work factor, they may consider a less effort-intensive way to exploit the storage array by attaching to it directly and reading from the raw disk blocks.

Security Controls

- Reduces the impact of vulnerabilities
- Any control measure should involve all the three aspects of infrastructure
 - ▶ People, process, and technology
- Controls can be technical or non-technical
 - ▶ Technical: antivirus, firewalls, and intrusion detection system
 - ▶ Non-technical: administrative policies and physical controls

Having assessed the vulnerability of the environment, organizations can deploy specific control measures. Any control measures should involve all the three aspects of infrastructure: people, process and technology, and their relationship. To secure people, first step is to establish and assure their identity. Based on their identity, selective controls can be implemented for their access to data and resources. The effectiveness of any security measure is primarily governed by the process and policies. The processes should be based on a thorough understanding of risks in the environment and recognize the relative sensitivity of different types of data, the needs of various stakeholders to access the data. Without an effective process, the deployment of technology is neither cost-effective nor aligned to organizations' priorities. And finally, the technologies or controls that are deployed should ensure compliance with the processes, policies, and people for its effectiveness. These security technologies are directed at reducing vulnerability by minimizing attack surfaces and maximizing the work factors. These controls can be technical or nontechnical. Technical controls are usually implemented through computer systems, whereas nontechnical controls are implemented through administrative and physical controls. Administrative controls include security and personnel policies or standard procedures to direct the safe execution of various operations. Physical controls include setting up physical barriers, such as security guards, fences, or locks.

Security Controls: Classification

- Controls are classified as:
 - ▶ Preventive
 - ▶▶ Attempts to prevent an attack
 - ▶ Corrective
 - ▶▶ Reduce the effect of an attack
 - ▶ Detective
 - ▶▶ Discover attacks and trigger preventive and corrective controls

Based on the roles they play, controls are categorized as preventive, detective, and corrective. The preventive control attempts to prevent an attack; the detective control detects whether an attack is in progress; and after an attack is discovered, the corrective controls are implemented. *Preventive controls* avert the vulnerabilities from being exploited and prevent an attack or reduce its impact. *Corrective controls* reduce the effect of an attack, whereas *detective controls* discover attacks and trigger preventive or corrective controls. For example, an Intrusion Detection/Intrusion Prevention System (IDS/IPS) is a detective control that determines whether an attack is underway and then attempts to stop it by terminating a network connection or invoking a firewall rule to block traffic.

Module 14: Securing the Storage Infrastructure


Lesson 2: Storage Security Domains

During this lesson the following topics are covered:

- Storage security domains
- Security threats in each domain
- Controls applied to reduce the risk in each domain

This lesson covers various storage security domains such as application access, management access, and back, replication, and archive. This lesson also covers security threats and control in each domain.

Storage Security Domains



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 5: Securing and Managing the Storage Infrastructure 14

Storage devices connected to a network raises the risk level and more exposed to security threats via networks. However, with increasing use of networking in storage environments, storage devices are becoming highly exposed to security threats from a variety of sources. Specific controls must be implemented to secure a storage networking environment. This requires a closer look at storage networking security and a clear understanding of the access paths leading to storage resources. If a particular path is unauthorized and needs to be prohibited by technical controls, ensure that these controls are not compromised. If each component within the storage network is considered a potential access point, the attack surface of all these access points must be analyzed to identify the associated vulnerabilities. To identify the threats that apply to a storage network, access paths to data storage can be categorized into three security domains: *application access*, *management access*, and *backup, replication, and archive*. Figure on the slide depicts the three security domains of a storage system environment. The first security domain involves application access to the stored data through the storage network. The second security domain includes management access to storage and interconnect devices and to the data residing on those devices. This domain is primarily accessed by storage administrators who configure and manage the environment. The third domain consists of backup, replication, and archive access. Along with the access points in this domain, the backup media also needs to be secured. To secure the storage networking environment, identify the existing threats within each of the security domains and classify the threats based on the type of security services—availability, confidentiality, integrity, and accountability. The next step is to select and implement various controls as countermeasures to the threats.

Securing the Application Access Domain

- Protect data and access to the data

Common Threats	Available Controls	Examples
• Spoofing user or host identity • Elevation of privileges • Tampering with data in-flight and at rest • Network snooping • Denial of service • Media theft	• Strong user and host authentication and authorization • Access control to storage objects • Data encryption • Storage network encryption	• Multi-factor authentication • RBAC, DH-CHAP • Zoning, LUN masking • Storage encryption • IP-Sec, FC security protocol • Antivirus • Controlling physical access to data center

Access control services regulate user and host access to data. These services mitigate the threats of spoofing user identity and elevating their privileges. Both these threats affect data integrity and confidentiality. Access control mechanisms used in application access domain are user and host authentication (technical control) and authorization (administrative control). These mechanisms may lie outside the boundaries of the storage network and require various systems to interconnect with other enterprise identity management and authentication systems. NAS devices support the creation of *access control lists* that regulates user access to specific files. The Enterprise Content Management application enforces access to data by using Information Rights Management (IRM) that specifies which users have what rights to a document.

Restricting access at the host level starts with authenticating a node when it tries to connect to a network. Different storage networking technologies, such as iSCSI, FC, and IP-based storage, use various authentication mechanisms, such as Challenge-Handshake Authentication Protocol (CHAP), Fibre Channel Security Protocol (FC-SP), and IPSec, respectively, to authenticate host access. After a host has been authenticated, the next step is to specify security controls for the storage resources, such as ports, volumes, or storage pools, that the host is authorized to access. *Zoning* is a control mechanism on the switches that segments the network into specific paths to be used for data traffic; *LUN masking* determines which hosts can access which storage devices.

Cont..

It is also important to ensure that administrative controls are implemented. Regular auditing is required to ensure proper functioning of administrative controls. This is enabled by logging significant events on all participating devices. Event logs should also be protected from unauthorized access because they may fail to achieve their goals if the logged content is exposed to unauthorized modifications by an attacker.

Security controls for protecting the storage infrastructure address the threats of unauthorized tampering of data in transit that leads to a loss of data integrity, denial of service that compromises availability, and network snooping that may result in loss of confidentiality.

The security controls for protecting the network fall into two general categories: network infrastructure integrity and storage network encryption. Controls for ensuring the infrastructure integrity include a fabric switch function that ensures fabric integrity. This is achieved by preventing a host from being added to the SAN fabric without proper authorization. Storage network encryption methods include the use of IPSec for protecting IP-based storage networks, and FC-SP for protecting FC networks. In secure storage environments, root or administrator privileges for a specific device are not granted to every user. Instead, *role-based access control* (RBAC) is deployed to assign necessary privileges to users, enabling them to perform their roles. A role may represent a job function, for example, an administrator. Privileges are associated with the roles and users acquire these privileges based upon their roles. It is also advisable to consider administrative controls, such as “separation of duties,” when defining data center procedures. Management networks for storage systems should be logically separate from other enterprise networks. Finally, physical access to the device console and the cabling of FC switches must be controlled to ensure protection of the storage infrastructure.

The most important aspect of securing data is protecting data held inside the storage arrays. Threats at this level include tampering with data, which violates data integrity, and media theft, which compromises data availability and confidentiality. To protect against these threats, encrypt the data held on the storage media or encrypt the data prior to being transferred to the disk. Data should be encrypted as close to its origin as possible. If it is not possible to perform encryption on the host device, an encryption appliance can be used for encrypting data at the point of entry into the storage network. It is also critical to decide upon a method for ensuring that data deleted at the end of its lifecycle has been completely erased from the disks and cannot be reconstructed for malicious purposes. On NAS devices, adding antivirus checks and file extension controls can further enhance data integrity. In the case of CAS, use of MD5 or SHA-256 cryptographic algorithms guarantees data integrity by detecting any change in content bit patterns. In addition, the data erasure service ensures that the data has been completely over written by bit sequence before the disk is discarded.

Securing the Management Access Domain

- Involves protecting administrative access and management infrastructure
- Common threats
 - ▶ Spoofing administrator's identity
 - ▶ Elevating administrative privileges
 - ▶ Network snooping and DoS
- Available controls
 - ▶ Authentication, authorization, and management access control
 - ▶ Private management network
 - ▶ Disable unnecessary network services
 - ▶ Encryption of management traffic

Management access, whether monitoring, provisioning, or managing storage resources, is associated with every device within the storage network. Implementing appropriate controls for securing storage management applications is important because the damage that can be caused by using these applications can be far more extensive.

Controlling administrative access to storage aims to safeguard against the threats of an attacker spoofing an administrator's identity or elevating privileges to gain administrative access. To protect against these threats, administrative access regulation and various auditing techniques are used to enforce accountability of users and processes. Access control should be enforced for each storage component. In some storage environments, it may be necessary to integrate storage devices with third-party authentication directories, such as Lightweight Directory Access Protocol (LDAP) or Active Directory. Security best practices stipulate that no single user should have ultimate control over all aspects of the system. It is better to assign various administrative functions by using RBAC. Auditing logged events is a critical control measure to track the activities of an administrator. However, access to administrative log files and their content must be protected. In addition, having a Security Information Management (SIM) solution supports effective analysis of the event log files.

Mechanisms to protect the management network infrastructure include encrypting management traffic, enforcing management access controls, and applying IP network security best practices. Restricting network activity and access to a limited set of hosts minimizes the threat of an unauthorized device attaching to the network and gaining access to the management interfaces. Access controls need to be enforced at the storage-array level to specify which host has management access to which array. A separate private management network is highly recommended for the management traffic. If possible, management traffic should not be mixed with either production data traffic or other LAN traffic used in the enterprise. Unused network services must be disabled on every device within the storage network. This decreases the attack surface for that device by minimizing the number of interfaces through which the device can be accessed.

Securing Backup, Replication, and Archive Domain

- Involves protecting backup, replication, and archive infrastructure
- Common threats
 - ▶ Spoofing DR site identity
 - ▶ Tampering with data in-flight and at rest
 - ▶ Network snooping
- Available controls
 - ▶ Access control – primary to secondary storage
 - ▶ Backup encryption
 - ▶ Replication network encryption

Backup, replication, and archive is the third domain that needs to be secured against an attack. A backup involves copying the data from a storage array to backup media, such as tapes or disks. Securing backup is complex and is based on the backup software that accesses the storage arrays. It also depends on the configuration of the storage environments at the primary and secondary sites, especially with remote backup solutions performed directly on a remote tape device or using array-based remote replication.

Organizations must ensure that the disaster recovery (DR) site maintains the same level of security for the backed up data. Protecting the backup, replication, and archive infrastructure requires addressing several threats, including spoofing the legitimate identity of a DR site, tampering with data, network snooping, DoS attacks, and media theft. Such threats represent potential violations of integrity, confidentiality, and availability. In a remote backup solution where the storage components are separated by a network, the threats at the transmission layer need to be countered. Otherwise, an attacker can spoof the identity of the backup server and request the host to send its data. The unauthorized host claiming to be the backup server may lead to a remote backup being performed to an unauthorized and unknown site. In addition, attackers can use the DR network connection to tamper with data, snoop the network, and create a DoS attack against the storage devices.

The physical threat of a backup tape being lost, stolen, or misplaced, especially if the tapes contain highly confidential information, is another type of threat. Backup-to-tape applications are vulnerable to severe security implications if they do not encrypt data while backing it up.

Module 14: Securing the Storage Infrastructure

Lesson 3: Security Implementations in Storage Networking

During this lesson the following topics are covered:

- SAN security implementations
- NAS security implementations
- IP SAN security implementations

This lesson covers various security implementation in SAN, NAS and IP SAN environment.

Security Implementation in SAN

- Common SAN security mechanisms are:
 - ▶ LUN masking and zoning
 - ▶ Securing FC switch ports
 - ▶ Switch-wide and fabric-wide access control
 - ▶ Logical partitioning of a fabric: VSAN

Traditional FC SANs enjoy an inherent security advantage over IP-based networks. An FC SAN is configured as an isolated private environment with fewer nodes than an IP network. Consequently, FC SANs impose fewer security threats. However, this scenario has changed with converged network, storage consolidation, driving rapid growth and necessitating designs for large, complex SANs that span multiple sites across the enterprise. Today, no single comprehensive security solution is available for FC SANs. Many FC SAN security mechanisms have evolved from their counterpart in IP networking, thereby bringing in matured security solutions.

Fibre Channel Security Protocol (FC-SP) standards (T11 standards), published in 2006, align security mechanisms and algorithms between IP and FC interconnects. These standards describe protocols used to implement security measures in a FC fabric, among fabric elements and N_Ports within the fabric. They also include guidelines for authenticating FC entities, setting up session keys, negotiating the parameters required to ensure frame-by-frame integrity and confidentiality, and establishing and distributing policies across an FC fabric.

LUN masking and zoning, security in FC switch port, switch-wide and fabric-wide access control, and logical partitioning of a fabric (Virtual SAN) are the most commonly used SAN security methods. A stronger variant of LUN masking may sometimes be offered whereby masking can be done on the basis of source FC address. It offers a mechanism to lock down the FC address of a given node port to its WWN.

Securing FC Switch Ports

- Port binding
 - ▶ Restricts devices that can attach to a particular switch port
 - ▶▶ Allows only the corresponding switch port to connect to a node for fabric access
- Port lockdown and port lockout
 - ▶ Restricts a switch port's type of initialization
- Persistent port disable
 - ▶ Prevents a switch port from being enabled even after a switch reboot

Apart from zoning and LUN masking, additional security mechanisms, such as port binding, port lockdown, port lockout, and persistent port disable, can be implemented on switch ports.

Port binding : Limits the devices that can attach to a particular switch port and allows only the corresponding switch port to connect to a node for fabric access. Port binding mitigates but does not eliminate WWPN spoofing.

Port lockdown and port lockout: Restrict a switch port's type of initialization. Typical variants of port lockout ensure that the switch port cannot function as an E-Port and cannot be used to create an ISL, such as a rogue switch. Some variants ensure that the port role is restricted to only F-Port, E-Port, or a combination of these.

Persistent port disable: Prevents a switch port from being enabled even after a switch reboot.

Switch-wide and Fabric-wide Access Control

- Access control lists (ACLs)
 - ▶ Include device connection and switch connection control policies
 - ▶▶ Device connection control policy specifies which HBAs, storage ports can be connected to a particular switch
 - ▶▶ Switch connection control policy prevents unauthorized switches to join a particular switch
- Fabric Binding
 - ▶ Prevents unauthorized switch from joining a fabric
- Role-based access control (RBAC)
 - ▶ Enables assigning roles to users that explicitly specify access rights

As organizations grow their SANs locally or over longer distances, there is a greater need to effectively manage SAN security. Network security can be configured on the FC switch by using access control lists (ACLs) and on the fabric by using fabric binding.


ACLs incorporate the device connection control and switch connection control policies. The device connection control policy specifies which HBAs and storage ports can be connected to a particular switch, preventing unauthorized devices from accessing it. Similarly, the switch connection control policy specifies which switches are allowed to be connected to a particular switch, preventing unauthorized switches from joining it.

Fabric binding prevents an unauthorized switch from joining any existing switch in the fabric. It ensures that authorized membership data exists on every switch and any attempt to connect any switch in the fabric by using an ISL causes the fabric to segment.

Role-based access control provides additional security to a SAN by preventing unauthorized activity on the fabric for management operations. It enables the security administrator to assign roles to users that explicitly specify privileges or access rights after logging into the fabric. For example, the zone admin role can modify the zones on the fabric, whereas a basic user may view only fabric-related information, such as port types and logged-in nodes.

Logical Partitioning of a Fabric: VSAN

- Enables the creation of multiple logical SANs over a common physical SAN
- Fabric events in one VSAN are not propagated to the others
- Zoning should be configured for each VSAN



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.


Section 5: Securing and Managing the Storage Infrastructure 24

VSANs enable the creation of multiple logical SANs over a common physical SAN. They provide the capability to build larger consolidated fabrics and still maintain the required security and isolation between them. Figure on the slide depicts logical partitioning of a fabric using VSAN.

The SAN administrator can create distinct VSANs by populating each of them with switch ports. In the example, the switch ports are distributed over two VSANs: 10 and 20—for the Engineering and HR divisions, respectively. Although they share physical switching gear with other divisions, they can be managed individually as standalone fabrics. Zoning should be done for each VSAN to secure the entire physical SAN. Each managed VSAN can have only one active zone set at a time.

VSANs minimize the impact of fabric wide disruptive events because management and control traffic on the SAN—which may include RSCNs, zone set activation events, and more—does not traverse VSAN boundaries. Therefore, VSANs are a cost-effective alternative for building isolated physical fabrics. They contribute to information availability and security by isolating fabric events and providing authorization control within a single fabric.

SAN Security Architecture: Defense-in-Depth



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 5: Securing and Managing the Storage Infrastructure 25

Storage networking environments are a potential target for unauthorized access, theft, and misuse because of the vastness and complexity of these environments. Therefore, security strategies are based on the *defense in depth* concept, which recommends multiple integrated layers of security. This ensures that the failure of one security control will not compromise the assets under protection. Figure on the slide illustrates various levels (zones) of a storage networking environment that must be secured. FC SANs not only suffer from certain risks and vulnerabilities that are unique, but also share common security problems associated with physical security and remote administrative access. In addition to implementing SAN-specific security measures, organizations must simultaneously leverage other security implementations in the enterprise.

Comprehensive list of protection strategies that must be implemented in various security zones are listed below:

Zone A (Authentication at the Management Console):

(a) Restrict management LAN access to authorized users (lock down MAC addresses) (b) implement VPN tunneling for secure remote access to the management LAN and (c) use two-factor authentication for network access

Zone B (Firewall)

Block inappropriate traffic by (a) filtering out addresses that should not be allowed on your LAN, and (b) screening for allowable protocols block ports that are not in use

Cont..

Zone C (Access Control-Switch)

Authenticate users/administrators of FC switches using Remote Authentication Dial In User Service (RADIUS) and DH-CHAP (Diffie-Hellman Challenge Handshake Authentication Protocol)

Zone D (Host to switch)

Restrict Fabric access to legitimate hosts by implementing (a) ACLs: Known HBAs can connect on specific switch ports only; and (b) a secure zoning method, such as port zoning (also known as hard zoning)

Zone E (Switch to Switch/Switch to Router)

Protect traffic on fabric by (a) using E_Port authentication; (b) encrypting the traffic in transit; and (c) implementing FC switch controls and port controls

Zone F (Distance Extension)

Implement encryption for in-flight data (a) FC-SP for long-distance FC extension, and (b) IPSec for SAN extension via FCIP

Zone G (Switch to Storage)

Protect the storage arrays on your SAN via (a) WWPN-based LUN masking and (b) S_ID locking: masking based on source Fibre Channel address

Security Implementation in NAS

- Permissions and ACLs
 - ▶ Protection to NAS resources by restricting access
- Other authentication and authorization mechanisms
 - ▶ Kerberos and Directory services
 - ▶▶ Implemented to verify the identity of network users and define their privileges
 - ▶ Firewalls
 - ▶▶ To protect the storage infrastructure from unauthorized access and malicious attacks

NAS is open to multiple exploits, including viruses, worms, unauthorized access, snooping, and data tampering. Various security mechanisms are implemented in NAS to secure data and the storage networking infrastructure.

Permissions and ACLs form the first level of protection to NAS resources by restricting accessibility and sharing. These permissions are deployed over and above the default behaviors and attributes associated with files and folders. In addition, various other authentication and authorization mechanisms, such as Kerberos and directory services, are implemented to verify the identity of network users and define their privileges. Similarly, firewalls protect the storage infrastructure from unauthorized access and malicious attacks.

NAS File Sharing: Windows ACLs

- Types of ACLs
 - ▶ Discretionary access control lists (DACL)
 - ▶▶ Commonly referred to as ACL and used to determine access control
 - ▶ System access control lists (SACL)
 - ▶▶ Determine what access needs to be audited if auditing is enabled
- Object Ownership
 - ▶ Object owner has hard-coded rights to that object
 - ▶ Child objects within a parent object automatically inherit the ACLs of parent object
- Security identifiers (SIDs)
 - ▶ SIDs uniquely identify a user or a user group
 - ▶ ACLs use SIDs to control access to the objects

Windows supports two types of ACLs: *discretionary access control lists* (DACLs) and *system access control lists* (SACLs). The DACL, commonly referred to as the ACL, that determines access control. The SACL determines what accesses need to be audited if auditing is enabled.

In addition to these ACLs, Windows also supports the concept of object ownership. The owner of an object has hard-coded rights to that object, and these rights do not need to be explicitly granted in the SACL. The owner, SACL, and DACL are all statically held as attributes of each object. Windows also offers the functionality to inherit permissions, which allows the child objects existing within a parent object to automatically inherit the ACLs of the parent object.

ACLs are also applied to directory objects known as security identifiers (SIDs). These are automatically generated by a Windows server or domain when a user or group is created, and they are abstracted from the user. In this way, though a user may identify his login ID as "User1," it is simply a textual representation of the true SID, which is used by the underlying operating system. Internal processes in Windows refer to an account's SID rather than the account's username or group name while granting access to an object. ACLs are set by using the standard Windows Explorer GUI but can also be configured with CLI commands or other third-party tools.

NAS File Sharing: UNIX Permissions


- UNIX permissions specify what can be done to a file and by whom
 - ▶ Common permissions: Read/Write/Execute
- Every file and directory (folder) has three ownership relations:
 - ▶ Rights for the file owner
 - ▶ Rights for the group the user belong to
 - ▶ Rights for all other users

For the UNIX operating system, a *user* is an abstraction that denotes a logical entity for assignment of ownership and operation privileges for the system. A user can be either a person or a system operation. A UNIX system is only aware of the privileges of the user to perform specific operations on the system and identifies each user by a user ID (UID) and a username, regardless of whether it is a person, a system operation, or a device.

In UNIX, users can be organized into one or more groups. The concept of group serves the purpose to assign sets of privileges for a given resource and sharing them among many users that need them. For example, a group of people working on one project may need the same permissions for a set of files.

UNIX permissions specify the operations that can be performed by any ownership relation with respect to a file. In simpler terms, these permissions specify what the owner can do, what the owner group can do, and what everyone else can do with the file. For any given ownership relation, three bits are used to specify access permissions. The first bit denotes read (r) access, the second bit denotes write (w) access, and the third bit denotes execute (x) access. Because UNIX defines three ownership relations (Owner, Group, and All), a triplet (defining the access permission) is required for each ownership relationship, resulting in nine bits. Each bit can be either set or clear. When displayed, a set bit is marked by its corresponding operation letter (r, w, or x), a clear bit is denoted by a dash (-), and all are put in a row, such as `rw-r-xr-x`. In this example, the owner can do anything with the file, but group owners and the rest of the world can read or execute only. When displayed, a character denoting the mode of the file may precede this nine-bit pattern. For example, if the file is a directory, it is denoted as `"d"`; and if it is a link, it is denoted as `"l"`.

Authentication and Authorization in a NAS Environment



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 5: Securing and Managing the Storage Infrastructure 30

In a file-sharing environment, NAS devices use standard file-sharing protocols, NFS and CIFS. Therefore, authentication and authorization are implemented and supported on NAS devices in the same way as in a UNIX or Windows file-sharing environment.

Authentication requires verifying the identity of a network user and therefore involves a login credential lookup on a Network Information System (NIS) server in a UNIX environment. Similarly, a Windows client is authenticated by a Windows domain controller that houses the Active Directory. The Active Directory uses LDAP to access information about network objects in the directory and Kerberos for network security. NAS devices use the same authentication techniques to validate network user credentials. Figure on the slide depicts the authentication process in a NAS environment.

Authorization defines user privileges in a network. The authorization techniques for UNIX users and Windows users are quite different. UNIX files use mode bits to define access rights granted to owners, groups, and other users, whereas Windows uses an ACL to allow or deny specific rights to a particular user for a particular file.

Although NAS devices support both of these methodologies for UNIX and Windows users, complexities arise when UNIX and Windows users access and share the same data. If the NAS device supports multiple protocols, the integrity of both permission methodologies must be maintained. NAS device vendors provide a method of mapping UNIX permissions to Windows and vice versa, so a multiprotocol environment can be supported.


Kerberos – Network Authentication Protocol

- Uses secret-key cryptography
- A client can prove its identity to a server (and vice versa) across an insecure network connection
- Kerberos client
 - ▶ An entity that gets a service ticket for a Kerberos service
- Kerberos server
 - ▶ Refers to the Key Distribution Center (KDC)
 - ▶ Implements the Authentication Service (AS) and the Ticket Granting Service (TGS)

Kerberos is a network authentication protocol, which is designed to provide strong authentication for client/server applications by using secret-key cryptography. It uses cryptography so that a client and server can prove their identity to each other across an insecure network connection. After the client and server have proven their identities, they can choose to encrypt all their communications to ensure privacy and data integrity.

In Kerberos, authentications occur between clients and servers. The client gets a ticket for a service and the server decrypts this ticket by using its secret key. Any entity, user, or host that gets a service ticket for a Kerberos service is called a *Kerberos client*. The term *Kerberos server* generally refers to the Key Distribution Center (KDC). The KDC implements the Authentication Service (AS) and the Ticket Granting Service (TGS). The KDC has a copy of every password associated with every principal, so it is absolutely vital that the KDC remain secure. In Kerberos, users and servers for which a secret key is stored in the KDC database are known as *principals*.

Kerberos Authentication



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.


Section 5: Securing and Managing the Storage Infrastructure 32

The Kerberos authentication process shown in figure on the slide includes the following steps:

1. The user logs on to the workstation in the Active Directory domain (or forest) using an ID and a password. The client computer sends a request to the AS running on the KDC for a Kerberos ticket. The KDC verifies the user's login information from Active Directory.
2. The KDC responds with an encrypted Ticket Granting Ticket (TGT) and an encrypted session key. TGT has a limited validity period. TGT can be decrypted only by the KDC, and the client can decrypt only the session key.
3. When the client requests a service from a server, it sends a request, consisting of the previously generated TGT, encrypted with the session key and the resource information to the KDC.
4. The KDC checks the permissions in Active Directory and ensures that the user is authorized to use that service.
5. The KDC returns a service ticket to the client. This service ticket contains fields addressed to the client and to the server hosting the service.
6. The client then sends the service ticket to the server that houses the required resources.
7. The server, in this case the NAS device, decrypts the server portion of the ticket and stores the information in a keytab file. As long as the client's Kerberos ticket is valid, this authorization process does not need to be repeated. The server automatically allows the client to access the appropriate resources.
8. A client-server session is now established. The server returns a session ID to the client, which tracks the client activity, such as file locking, as long as the session is active.

Network Layer Firewalls

- Firewalls are implemented in NAS environments
 - ▶ To protect against security threats in IP network
 - ▶ To examine network packets and compare them to a set of configured security rules
 - ▶▶ Packets that are not authorized by a security rule are dropped
- Demilitarized Zone (DMZ)
 - ▶ To secure internal assets while allowing Internet-based access to various resources



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.


Section 5: Securing and Managing the Storage Infrastructure 33

Because NAS devices utilize the IP protocol stack, they are vulnerable to various attacks initiated through the public IP network. Network layer firewalls are implemented in NAS environments to protect the NAS devices from these security threats. These network-layer firewalls can examine network packets and compare them to a set of configured security rules. Packets that are not authorized by a security rule are dropped and not allowed to continue to the destination. Rules can be established based on a source address (network or host), a destination address (network or host), a port, or a combination of those factors (source IP, destination IP, and port number). The effectiveness of a firewall depends on how robust and extensive the security rules are. A loosely defined rule set can increase the probability of a security breach.

A demilitarized zone (DMZ) is commonly used in networking environments. A DMZ provides a means to secure internal assets while allowing Internet-based access to various resources. In a DMZ environment, servers that need to be accessed through the Internet are placed between two sets of firewalls. Application-specific ports, such as HTTP or FTP, are allowed through the firewall to the DMZ servers. However, no Internet-based traffic is allowed to penetrate the second set of firewalls and gain access to the internal network. The servers in the DMZ may or may not be allowed to communicate with internal resources. In such a setup, the server in the DMZ is an Internet-facing web application accessing data stored on a NAS device, which may be located on the internal private network. A secure design would serve only data to internal and external applications through the DMZ.

Security Implementation in IP SAN: CHAP

- Challenge-Handshake Authentication Protocol (CHAP)
 - ▶ Provides a method for initiators and targets to authenticate each other by utilizing a secret code



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.


Section 5: Securing and Managing the Storage Infrastructure 34

The *Challenge-Handshake Authentication Protocol* (CHAP) is a basic authentication mechanism that has been widely adopted by network devices and hosts. CHAP provides a method for initiators and targets to authenticate each other by utilizing a secret code or password. CHAP secrets are usually random secrets of 12 to 128 characters. The secret is never exchanged directly over the communication channel; rather, a one-way hash function converts it into a hash value, which is then exchanged. A hash function, using the MD5 algorithm, transforms data in such a way that the result is unique and cannot be changed back to its original form. Figure on the slide depicts the CHAP authentication process.

If the initiator requires reverse CHAP authentication, the initiator authenticates the target by using the same procedure. The CHAP secret must be configured on the initiator and the target. A CHAP entry, composed of the name of a node and the secret associated with the node, is maintained by the target and the initiator.

The same steps are executed in a two-way CHAP authentication scenario. After these steps are completed, the initiator authenticates the target. If both authentication steps succeed, then data access is allowed. CHAP is often used because it is a fairly simple protocol to implement and can be implemented across a number of disparate systems.

Securing IPSAN with iSNS Discovery Domains



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 5: Securing and Managing the Storage Infrastructure 35

iSNS discovery domains function in the same way as FC zones. Discovery domains provide functional groupings of devices in an IP-SAN. For devices to communicate with one another, they must be configured in the same discovery domain. State change notifications (SCNs) inform the iSNS server when devices are added to or removed from a discovery domain.

Module 14: Securing the Storage Infrastructure

Lesson 4: Security in Virtualized and Cloud Environments

During this lesson the following topics are covered:

- Security concerns
- Security measures

This lesson covers an overview of security in virtualized and cloud environment. This lesson also covers security concerns and measures in virtualized and cloud environment.

Security in Virtualized and Cloud Environments

- These environments have additional threats due to multitenancy and lack of control over the cloud resources
- Virtualization-specific security concerns are common for all cloud models
- In public clouds, there are additional security concerns, which demand specific countermeasures
 - ▶ Clients have less control to enforce security measures in public clouds
 - ▶ Difficult for cloud service provider(CSP) to meet the security needs of all the clients

This module, so far, focused only on the security threats and measures in a traditional data center. These threats and measures are also applicable to information storage in virtualized and cloud environments. However, virtualized and cloud computing environments pose additional threats to an organization's data due to multitenancy and lack of control over the cloud resources. A public cloud has more security concerns compared to a private cloud and demands additional counter measures. This is because in a public cloud, cloud users (consumers) usually have limited control over resources, and therefore, enforcement of security mechanisms for consumers is comparatively difficult. From a security perspective, both consumers and cloud service providers (CSP) have several security concerns and face multiple threats.

Security Concerns

- Multitenancy
 - ▶ Enables multiple independent tenants to be serviced using the same set of storage resources
 - ▶▶ Co-location of multiple VMs in a single server and sharing the same resources increase the attack surface
- Velocity of attack
 - ▶ Any existing security threat in the cloud spreads more rapidly and has larger impact than that in the traditional data center
- Information assurance and data privacy

Organizations are rapidly adopting virtualization and cloud computing, however they have some security concerns. These key security concerns are multitenancy, velocity of attack, information assurance, and data privacy.

Multitenancy, by virtue of virtualization, enables multiple independent tenants to be serviced using the same set of storage resources. In spite of the benefits offered by multitenancy, it is still a key security concern for users and service providers. Colocation of multiple VMs in a single server and sharing the same resources increase the attack surface. It may happen that business critical data of one tenant is accessed by other competing tenants who run applications using the same resources.

Velocity-of-attack refers to a situation in which any existing security threat in the cloud spreads more rapidly and has a larger impact than that in the traditional data center environments. *Information assurance* for users ensures confidentiality, integrity, and availability of data in the cloud. Also the cloud user needs assurance that all the users operating on the cloud are genuine and access the data only with legitimate rights and scope.

Data privacy is also a major concern in a virtualized and cloud environment. A CSP needs to ensure that Personally Identifiable Information (PII) about its clients is legally protected from any unauthorized disclosure.

Security Measures

- Securing compute
 - ▶ Securing physical server, VMs, and hypervisor
- Securing network
 - ▶ Virtual firewall
 - ▶▶ Provides packet filtering and monitoring of the VM-to-VM traffic
 - ▶ DMZ and data encryption
- Securing storage
 - ▶ Access control and data encryption
 - ▶ Use separate LUNs for VM configuration files and VM data
 - ▶ Segregate VM traffic from management traffic

Major threats to storage systems in virtualized and cloud environments arise due to compromises at compute, network, and physical security levels. This is because access to storage systems is provided by using compute and network infrastructure. Therefore, adequate security measures should be in place at the compute and network levels to ensure storage security.

Securing a compute infrastructure includes enforcing the security of the physical server, hypervisor, VM, and guest OS (OS running within a virtual machine). Physical server security involves implementing user authentication and authorization mechanisms. These mechanisms identify users and provide access privileges on the server. To minimize the attack surface on the server, unused hardware components, such as NICs, USB ports, or drives, should be removed or disabled.

A *hypervisor* is a single point of security failure for all the VMs running on it. Rootkits and malware installed on a hypervisor make detection difficult for the antivirus software installed on the guest OS. To protect against attacks, security-critical hypervisor updates should be installed regularly. Further, the hypervisor management system must also be protected. Malicious attacks and infiltration to the management system can impact all the existing VMs and allow attackers to create new VMs. Access to the management system should be restricted to authorized administrators. Furthermore, there must be a separate firewall installed between the management system and the rest of the network.

Cont..

VM isolation and *hardening* are some of the common security mechanisms to effectively safeguard a VM from an attack. VM isolation helps to prevent a compromised guest OS from impacting other guest OSs. VM isolation is implemented at the hypervisor level. Apart from isolation, VMs should be hardened against security threats. Hardening is a process to change the default configuration to achieve greater security.

The key security measures that minimize vulnerabilities at the network layer are firewall, intrusion detection, demilitarized zone (DMZ), and encryption of data-in-flight. In a virtualized and cloud environment, a firewall can also protect hypervisors and VMs. For example, if remote administration is enabled on a hypervisor, access to all the remote administration interfaces should be restricted by a firewall. A firewall also secures VM-to-VM traffic. This firewall service can be provided using a *Virtual Firewall* (VF) running on the hypervisor. VF gives visibility and control over the VM traffic and enforces policies at the VM level. DMZ and data encryption are also deployed as security measures in the virtualized and cloud environments. However, these deployments work in the same way as in the traditional data center.

Common security mechanisms that protect storage include the following:

- Access control methods to regulate which users and processes access the data on the storage systems
- Zoning and LUN masking
- Encryption of data-at-rest (on the storage system) and data-in-transit. Data encryption should also include encrypting backups and storing encryption keys separately from the data.
- Data shredding that removes the traces of the deleted data

Apart from these mechanisms, isolation of different types of traffic using VSANs further enhances the security of storage systems. In the case of storage utilized by hypervisors, additional security steps are required to protect the storage. Storage for hypervisors using clustered file systems supporting multiple VMs may require separate LUNs for VM components and VM data.

Module 14: Securing the Storage Infrastructure

Concept in Practice

- RSA security products
- VMware vShield

The concepts in practice section covers various security products of RSA and VMware. The product includes RSA SecureID, RSA Identity and Access Management, RSA Data Protection Manager, and VMware vShield.

RSA Security Products

- RSA SecureID
 - ▶ Provides two-factor authentication
 - ▶▶ Based on something a user knows (a password or PIN) and something a user has (an authenticator device)
 - ▶▶ Authenticator device automatically changes passwords every 60 seconds
- RSA Identity and Access Management
 - ▶ Provides identity, security, and access-control management for physical, virtual, and cloud-based environments
- RSA Data Protection Manager
 - ▶ Enables deployment of encryption, tokenization, and enterprise key management

EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 5: Securing and Managing the Storage Infrastructure 42

RSA SecurID two-factor authentication provides an added layer of security to ensure that only valid users have access to systems and data. RSA SecurID is based on something a user knows (a password or PIN) and something a user has (an authenticator device). It generates a new one-time password code every 60 seconds, making it difficult for anyone other than the genuine user to input the correct token code at any given time. To access their resources, users combine their secret Personal Identification Number (PIN) with the token code that appears on their SecurID authenticator display at that given time. The result is a unique, one-time password used to assure a user's identity.

The RSA Identity and Access Management product provides identity, security, and access-control management for physical, virtual, and cloud-based environments through access management. It enables trusted identities to freely and securely interact with systems and access. The RSA Identity and Access Management family has two products: *RSA Access Manager* and *RSA Federated Identity Manager*. *RSA Access Manager* enables organizations to centrally manage authentication and authorization policies for a large number of users, online web portals, and application resources. *RSA Federated Identity Manager* enables end users to collaborate with business partners, outsourced service providers, and supply-chain partners or across multiple offices or agencies all with a single identity and login.

RSA Data Protection Manager enables deployment of encryption, tokenization, and enterprise key management simply and affordably. *RSA Data Protection Manager* family is composed of two products: *Application Encryption and Tokenization* and *Enterprise Key Management*. *Application Encryption and Tokenization* with *RSA Data Protection Manager* helps to achieve compliance with regulations related to PII. *Enterprise key management* is an easy-to-use management tool for encrypting keys at the database, file server and storage layers.

VMware vShield

- VMware vShield family includes three products
 - ▶ vShield App
 - ▶▶ Hypervisor-based application-aware firewall solution
 - ▶▶ Observes network activity between virtual machines
 - ▶ vShield Edge
 - ▶▶ Provides comprehensive perimeter network security
 - ▶▶ Deployed as a virtual appliance and serves as a network security gateway for all the hosts
 - ▶▶ Provides many services including firewall, VPN, and DHCP
 - ▶ vShield Endpoint
 - ▶▶ Consists of a hardened special security VM with a third party antivirus software

The VMware vShield family includes three products: *vShield App*, *vShield Edge*, and *vShield Endpoint*.

VMware vShield App is a hypervisor-based application-aware firewall solution. It protects applications in a virtualized environment from network-based threats by providing visibility into network communications and enforcing granular policies with security groups. VMware vShield App observes network activity between virtual machines to define and refine firewall policies and secure business processes through detailed reporting of application traffic.

VMware vShield Edge provides comprehensive perimeter network security for a virtualized environment. It is deployed as a virtual appliance and serves as a network security gateway for all the hosts within the virtualized environment. It provides many services including firewall, VPN, and Dynamic Host Configuration Protocol (DHCP) services.

VMware vShield Endpoint consists of a hardened special security VM with a third party antivirus software. VMware vShield Endpoint streamlines and accelerates antivirus and antimalware deployment because antivirus engine and signature files are updated only within the special security VM. VMware vShield Endpoint improves VM performance by offloading file scanning and other tasks from VMs to the security VM. It prevents antivirus storms and bottlenecks associated with multiple simultaneous antivirus and antimalware scans and updates. It also satisfies audit requirements with detailed logging of antivirus and antimalware activities.

Module 14: Summary

Key points covered in this module:

- Information security framework
- Storage security domains
- Controls that can be deployed against identified threats in each domain
- SAN security architecture
- Protection mechanisms in SAN, NAS, and IP SAN environments
- Security in virtualized and cloud environments

This module covered information security framework and various storage security domains. This module also covered security implementation in SAN, NAS, and IP SAN. Further this module covered security in virtualized and cloud environment.

The basic security framework for storage is built around the four primary services of security: confidentiality, integrity, availability, and accountability.

To identify the threats that apply to a storage network, access paths to data storage can be categorized into three security domains: *application access*, *management access*, and *backup, replication, and archive*.

Storage networking environments are potential target for unauthorized access, theft, and misuse because of the vastness and complexity of these environments. Therefore, security strategies are based on the *defense-in-depth* concept, which recommends multiple integrated layers of security.

LUN masking and zoning, security in FC switch port, switch-wide and fabric-wide access control, and logical partitioning of a fabric (Virtual SAN) are the most commonly used SAN security methods. Permissions and ACLs, Kerberos and directory services, and firewalls are the common security mechanisms implemented in NAS. CHAP and iSNS discovery domain are the security measures implemented in IP SAN environment.

Virtualized and cloud computing environments pose additional threats to an organization's data due to multitenancy and lack of control over the cloud resources.

Check Your Knowledge – 1

- Which attack involves performing an action and eliminating evidence that could prove the identity of the attacker?
 - A. Denial of service
 - B. Eavesdropping
 - C. Repudiation
 - D. Snooping
- What is a role of Active Directory in Kerberos authentication?
 - A. Implements the authentication service and ticket granting service
 - B. Verifies the session ID when the client-server session is established
 - C. Verifies the user's login information
 - D. Maintains the security key for the servers

Check Your Knowledge – 2

- How can vulnerabilities be reduced in an IT environment?
 - A. By maximizing attack surface and minimizing work factor
 - B. By minimizing attack surface and maximizing work factor
 - C. By maximizing both attack surface and work factor
 - D. By minimizing both attack surface and work factor
- Which SAN security mechanism prevents a switch port from being enabled, even after a switch reboot?
 - A. Port lockdown
 - B. Persistent port disable
 - C. Port binding
 - D. Port zoning

Check Your Knowledge – 3

- Which security mechanism ensures that a port can only be initialized with a specific port type?
 - A. Port lockdown
 - B. Persistent port disable
 - C. Port binding
 - D. Port zoning

Module – 15

Managing the Storage Infrastructure



Module 15: Managing the Storage Infrastructure

Upon completion of this module, you should be able to:

- List the key storage infrastructure components that are monitored
- List the key monitoring parameters
- Describe the storage management activities
- Describe the storage infrastructure management challenges and their solutions
- Describe the information lifecycle management (ILM) strategy

This module focuses on management of storage infrastructure. This module lists the key storage infrastructure components that are monitored. The module also lists the monitoring parameters. Further, the module describes the storage management activities. It also describes the storage infrastructure management challenges and their solutions. Finally, this module describes the information lifecycle management strategy.

Module 15: Managing the Storage Infrastructure

Lesson 1: Monitoring the Storage Infrastructure

During this lesson the following topics are covered:

- Key storage infrastructure components that are monitored
- Monitoring parameters
- Types of alerts

This lesson covers the storage infrastructure components that are monitored. It also covers the monitoring parameters for storage infrastructure components. Further, it details the three types of alerts. Finally, the module provides monitoring examples.

Storage Infrastructure Management

- Managing storage infrastructure is key to ensure continuity of business
 - ▶ Establishing management processes and implementing appropriate tools are essential to meet service levels
 - ▶ Virtualization have changed the storage infrastructure management paradigm
- Monitoring is the most important aspect that forms the basis for storage management

Unprecedented growth of information, proliferation of applications, complexity of business processes, and requirements of 24×7 availability of information have put increasingly higher demands on the IT infrastructure.


Managing storage infrastructure efficiently is a key that enables organizations to address these challenges and ensures continuity of business.

Comprehensive storage infrastructure management requires the implementation of intelligent tools and robust processes to meet the required service levels. These tools enable performance tuning, data protection, access control, centralized auditing, and meeting compliance requirements. They also ensure the consolidation and better utilization of existing resources, thereby limiting the need for excessive ongoing investment on infrastructure. The management process defines procedures for efficient handling of various operations, such as incident, problem, and change requests. It is imperative to manage not just the individual components, but also the infrastructure end to end due to the components' interdependency.

Storage infrastructure management is also composed of strategies, such as *information lifecycle management* (ILM) that optimizes the storage cost while meeting the service levels. ILM helps to manage information based on its value to the business.

Managing the storage infrastructure requires performing various activities, including accessibility, capacity, performance, and security management. All of these activities are interrelated and should be considered to maximize the return on investment. Virtualization technologies have dramatically changed the storage infrastructure management paradigm.

Monitoring Storage Infrastructure



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 5: Securing and Managing the Storage Infrastructure 52

Monitoring is one of the most important aspects that forms the basis for managing storage infrastructure resources. Monitoring provides the performance and accessibility status of various components. It also enables administrators to perform essential management activities. Monitoring also helps to analyze the utilization and consumption of various storage infrastructure resources. This analysis facilitates capacity planning, forecasting, and optimal use of these resources. Storage infrastructure environment parameters such as heating, ventilating and air-conditioning (HVAC) are also monitored.

The key storage infrastructure components that should be monitored are:

- Servers
- Network
- Storage arrays

These components could be physical or virtualized. Each of these components should be monitored for accessibility, capacity, performance and security.

Monitoring Parameters

- Accessibility
 - ▶ To identify failure of any component that may lead to service unavailability or degraded performance
- Capacity
 - ▶ To ensure availability of adequate amount of resources and prevent service unavailability or degraded performance
- Performance
 - ▶ To evaluate efficiency and utilization of components and identify bottlenecks
- Security
 - ▶ To ensure confidentiality, integrity, and availability of storage infrastructure

Accessibility refers to the availability of a component to perform its desired operation during a specified time period. Monitoring the accessibility of hardware components (for example, a port, an HBA, or a disk drive) or software component (for example, a database instance) involves checking their availability status by reviewing the alerts generated from the system. For example, a port failure might result in a chain of availability alerts.

A storage infrastructure uses redundant components to avoid a single point of failure. Failure of a component might cause an outage that affects application availability, or it might cause performance degradation even though accessibility is not compromised. Continuously monitoring for expected accessibility of each component and reporting any deviation helps the administrator to identify failing components and plan corrective action to maintain SLA requirements.

Cont..

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 leads to degraded performance or even application/service unavailability. Capacity monitoring ensures uninterrupted data availability and scalability by averting outages before they occur. For example, if 90 percent of the ports are utilized in a particular SAN fabric, this could indicate that a new switch might be required if more arrays and servers need to be installed on the same fabric. Capacity monitoring usually leverages analytical tools to perform trend analysis. These trends help to understand future resource requirements and provide an estimation of the time required to deploy them.

Performance monitoring evaluates how efficiently different storage infrastructure components are performing and helps to identify bottlenecks. Performance monitoring measures and analyzes behavior in terms of response time or the capability to perform at a certain predefined level. It also deals with the 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 performed by a disk, application response time, network utilization, and server-CPU utilization are examples of performance parameters that are monitored.

Monitoring a storage infrastructure for security helps to track and prevent unauthorized access, whether accidental or malicious. *Security monitoring* helps to track unauthorized configuration changes of storage infrastructure elements. For example, security monitoring tracks and reports the initial zoning configuration performed and all the subsequent changes. Security monitoring also detects unavailability of information to authorized users due to security breach. Physical security of a storage infrastructure can also be continuously monitored using badge readers, biometric scans, or video cameras.

Monitoring Host: Accessibility and Capacity

- Accessibility
 - ▶ Hardware components such as HBAs, NICs, and internal disks
 - ▶ Status of various processes/applications
- Capacity
 - ▶ File system utilization
 - ▶ Database: table space/log space utilization
 - ▶ User quota

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

The accessibility of a host depends on the availability status of the hardware components and the software processes running on it. For example, a host's NIC failure might cause inaccessibility of the host to its user. Server clustering is a mechanism that provides high availability if a server failure occurs. In a server virtualization environment, multiple virtual machines (VMs) share a pool of resources. These resources are dynamically reallocated, which ensures better accessibility and utilization of the resources.

Monitoring the file system utilization of a host is important to ensure that sufficient storage capacity is available to the applications. Running out of file system space disrupts application availability. Monitoring helps estimate the file system's growth rate and predict when it will reach 100 percent. Accordingly, the administrator can extend (manually or automatically) the file system's space proactively to prevent application outage. Use of virtual provisioning technology enables efficient management of storage capacity requirements but is highly dependent on capacity monitoring.

Monitoring Host: Performance and Security

- Performance
 - ▶ CPU and memory utilization
 - ▶ Transaction response time
- Security
 - ▶ Authentication and authorization
 - ▶ Physical security (data center access)

Host performance monitoring mainly involves a status check on the utilization of various server resources, such as CPU and memory. For example, if a server running an application is experiencing 80 percent of CPU utilization continuously, it 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, additional CPU and memory may be allocated dynamically from the pool, if available, to meet performance requirements.

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 user access if multiple login failures are logged.

Monitoring Network: Accessibility and Capacity

- Accessibility
 - ▶ Physical components such as switches and ports
 - ▶ Logical components such as zones
- Capacity
 - ▶ Interswitch links and port utilization

Storage networks need 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 switches, ports, cables, and power supplies. The logical components include constructs, such as zones. Any failure in the physical or logical components causes data unavailability. For example, errors in zoning, such as specifying the wrong WWN of a port, result 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 interswitch links, individual ports, and each interconnect device in the fabric. Capacity monitoring provides all the required inputs for future planning and optimization of fabric resources.

Monitoring Network: Performance and Security

- Performance
 - ▶ Assess individual component performance and help to identify network bottlenecks
 - ▶▶ Monitoring port performance and link utilization
- Security
 - ▶ Unauthorized changes to the fabric
 - ▶ Login failures

Monitoring the performance of the storage network is useful in assessing individual component performance and helps to identify network bottlenecks. For example, monitoring port performance involves measuring the receive or transmit link utilization metrics, which indicates how busy the switch port is. Heavily used ports can cause queuing of I/Os on the server, which results in poor performance.

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

Storage network security monitoring provides information about 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.

Components Monitored – Storage

- Accessibility
 - ▶ Hardware components such as controllers and disks
 - ▶ Software processes such as replication processes
- Capacity
 - ▶ Capacity utilization and consumption trend
- Performance
 - ▶ Utilization rates of storage array components
 - ▶ I/O response time, cache utilization
- Security
 - ▶ Access control and physical security (data center access)


The accessibility of the storage array should be monitored for its hardware components and various processes. Storage arrays are typically configured with redundant components and therefore individual component failure does not usually affect their accessibility. However, failure of any process in the storage array can disrupt or compromise business continuity operations. For example, the failure of a replication task affects disaster recovery capabilities. Some storage arrays provide the capability to send messages to the vendor's support center if hardware or process failures occur, referred to as a *call home*.

Capacity monitoring of a storage array enables the administrator to respond to storage needs preemptively based on capacity utilization and consumption trends. Information about unconfigured and unallocated storage space is also required to decide whether a new server can be allocated storage capacity from the array.

A storage array can be monitored using a number of performance metrics, such as utilization rates of the various storage array components, I/O response time, and cache utilization. For example, an over utilized storage array component might lead to performance degradation.

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

Accessibility Monitoring Example: Switch Failure



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.


Section 5: Securing and Managing the Storage Infrastructure 60

A storage infrastructure requires implementation of an end-to-end solution to actively monitor all the parameters of its critical components. Early detection and preemptive alerting ensure uninterrupted services from critical assets. In addition, the monitoring tool should analyze the impact of a failure and deduce the root cause of symptoms. Next few slides illustrates examples of monitoring storage infrastructure components for accessibility, capacity, performance, and security.

Failure of any component might affect the accessibility of one or more components due to their interconnections and dependencies. 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 production storage array through two switches, SW1 and SW2, as shown on slide. All the servers share two storage ports on the storage array. Multipathing software has also been installed on all the three servers.


If one of the switches, SW1 fails, the multipathing software initiates a path failover, and all the servers continue to access data through the other switch, SW2. However, due to absence of redundant switch, a second switch failure could result in inaccessibility of the array. Monitoring for accessibility enables detecting the switch failure and helps administrator to take corrective action before another failure occurs.

In most cases, the administrator receives symptom alerts for a failing component and can initiate actions before the component fails.



This example illustrates the importance of monitoring the file system capacity on file servers. Figure on the slide illustrates the environment of a file system when full and that results in application outage when no capacity monitoring is implemented. Monitoring can be configured to issue a message when thresholds are reached on the file system capacity. For example, when the file system reaches 66 percent of its capacity, a warning message is issued, and a critical message is issued when the file system reaches 80 percent of its capacity. This enables the administrator to take action to extend the file system before it runs out of capacity. Proactively monitoring the file system can prevent application outages caused due to lack of file system space.

Performance Monitoring Example: Array Port Utilization




EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 5: Securing and Managing the Storage Infrastructure 62

This example 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 switch SW1. The three servers share the same storage ports on the storage array to access LUNs. A new server running an application with a high work load must be deployed to share the same storage port 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 of the shared storage port is shown by the solid and dotted lines in the graph. If the port utilization prior to deploying the new server is high, then deploying the new server is not recommended because it might impact the performance of the other servers. However, if the utilization of the port prior to deploying the new server is closer to the dotted line, then there is room to add a new server.

Security Monitoring Example: Storage Array



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 5: Securing and Managing the Storage Infrastructure 63

This slide illustrates the importance of monitoring security 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 and vice versa. A user from WG1 might try to make a local replica of the data that belongs to WG2. 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.

Alerts

- It is an integral part of monitoring
- It keeps administrators informed on the status of components and processes
- Monitoring tools enable administrators to assign different severity levels for different events
 - ▶ Classified as information, warning, and fatal alerts

Type of alerts	Description	Example
Information	• Provide useful information • Does not require administrator intervention	• Creation of zone or LUN
Warning	• Require administrative attention	• FS becoming full • Soft media errors
Fatal	• Require immediate attention	• Component failure • Power failure

Alerting of events is an integral part of monitoring. Alerting keeps administrators informed about the status of various components and processes—for example, conditions such as failure of power, disks, memory, or switches, which can impact the availability of services and require immediate administrative attention. Other conditions, such as a file system reaching a capacity threshold or a soft media error on disks, are considered warning signs and may also require administrative attention.

Monitoring tools enable administrators to assign different severity levels for different conditions based on the impact of the alerted condition. Whenever a condition with a particular severity level occurs, an alert is sent to the administrator, a script is triggered, or an incident ticket is opened to initiate a corrective action. Alert classifications can range from information alerts to fatal alerts. *Information alerts* provide useful information but do not require any intervention by the administrator. The creation of a 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, if an alert indicates that the number of soft media errors on a disk 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 might affect overall performance or availability. For example, if a disk fails, the administrator must ensure that it is replaced quickly.

Continuous monitoring, with automated alerting, enables administrators to respond to failures quickly and proactively. Alerting provides information that helps administrators prioritize their response to events.

Module 15: Managing the Storage Infrastructure


Lesson 2: Managing the Storage Infrastructure

During this lesson the following topics are covered:

- Storage infrastructure management activities
- Storage infrastructure management in virtualized environment
- Storage infrastructure management challenges
- Ideal solution for storage infrastructure management

This lesson details storage infrastructure management activities, and management in virtualized environment. It covers storage management examples and challenges. Finally, it describes the ideal solution for storage infrastructure management.

Storage Infrastructure Management Activities



The pace of information growth, proliferation of applications, heterogeneous infrastructure, and stringent service-level requirements have resulted in increased complexity of managing storage infrastructures. However, the emergence of storage virtualization and other technologies, such as data deduplication and compression, thin provisioning, federated storage access, and storage tiering, have enabled administrators efficiently manage storage resources.

The key storage infrastructure management activities performed in a data center can be broadly categorized into availability management, capacity management, performance management, security management, and reporting.

Availability Management

Availability Management

The goal of availability management is to ensure that the availability requirements of all the components and services are constantly met.

- Key activity is to provision redundancy at all levels
- Example of availability management tasks are:
 - ▶ Installing two or more HBAs per server
 - ▶ Installing multipathing software
 - ▶ Deploying clustered server
 - ▶ Configuring RAID
 - ▶ Deploying redundant fabric
 - ▶ Configuring data backup and replication

A critical task in availability management is establishing a proper guideline based on defined service levels to ensure availability. *Availability management* involves all availability-related issues for components or services to ensure that service levels are met. A key activity in availability management is to provision redundancy at all levels, including components, data, or even site levels. For example, when a server is deployed to support a critical business function, it requires high availability. This is generally accomplished by deploying two or more HBAs, multipathing software, 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. Provision RAID-protected LUNs to the server using at least two front-end ports. In addition, the storage arrays should have built-in redundancy for various components and should support local and remote replication.

Capacity Management

Capacity Management

The goal of capacity management is to ensure adequate availability of resources, based on their service level requirements.

- Example of capacity management activities are:
 - ▶ Storage provisioning
 - ▶ Enforcing capacity quota for users
 - ▶ Capacity consumption trend analysis
- Technologies such as data deduplication, compression, and virtual provisioning help to manage storage capacity efficiently

The goal of *capacity management* is to ensure adequate availability of resources based on their service level requirements. Capacity management also involves optimization of capacity based on the cost and future needs. Capacity management provides capacity analysis that compares allocated storage to forecasted storage on a regular basis. It also provides trend analysis based on the 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 creating RAID sets and LUNs, and allocating them to the host. Enforcing capacity quotas for users is another example of capacity management. Provisioning a fixed amount of space for their files restricts users from exceeding the allocated capacity.

Technologies, such as data deduplication and compression, have reduced the amount of data to be backed up and thereby reduced the amount of storage capacity to be managed.

Performance Management

Performance Management

The goal of performance management is to ensure the optimal operational efficiency of all components.

- Key activities are:
 - ▶ Fine tuning for performance enhancement
 - ▶ Identifying performance bottlenecks
- Example of performance management activities are:
 - ▶ Configuring multiple paths
 - ▶ Choosing appropriate RAID type and cache configuration

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 information on whether a component meets expected performance levels.

Several performance management activities need to be performed when deploying a new 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 the expected performance levels, activities on the server, such as the volume configuration, database design or application layout, configuration of multiple HBAs, and intelligent multipathing software, must be fine-tuned. The performance management tasks on a SAN include designing and implementing sufficient ISLs in a multiswitch fabric with adequate bandwidth to support the required performance levels. The storage array configuration tasks include selecting the appropriate RAID type, LUN layout, front-end ports, back-end ports, and cache configuration, when considering the end-to-end performance.

Security Management

Security Management

The goal of security management is to ensure confidentiality, integrity, and availability of information.

- It prevents unauthorized access and configuration of storage infrastructure components
- Examples of security management tasks are:
 - ▶ Managing user accounts
 - ▶ Configuring zoning and LUN masking
 - ▶ Configuring encryption services
 - ▶ Installing antivirus and firewalls
 - ▶ Auditing of event logs

The key objective of the *security management* activity is to ensure confidentiality, integrity, and availability of information in both virtualized and nonvirtualized environments. 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 the user accounts and access policies that authorize users to perform role-based activities. The security management tasks in a SAN environment include configuration of zoning to restrict an unauthorized HBA from accessing specific storage array ports. Similarly, the security management task on a storage array includes LUN masking that restricts a host from accessing a defined set of LUNs.

Reporting

- Involves gathering information from various components or processes and generating reports
- Commonly used reports are:
 - ▶ Capacity planning report
 - ▶ Configuration and asset management reports
 - ▶ Performance report
 - ▶ Chargeback report

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, and performance. Capacity planning reports contain current and historic information about the utilization of storage, file systems, database tablespace, ports, and so on. Configuration and asset management reports include details about device allocation, local or remote replicas, and fabric configuration. This report also lists all the equipment, with details, such as their 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.

Storage Infrastructure Management in Virtualized Environment

- Virtualization technology has enabled managing storage infrastructure efficiently
 - ▶ Virtualization at storage layer
 - ▶▶ Example: virtual provisioning of LUN
 - ▶ Virtualization at network layer
 - ▶▶ Example: VLAN and VSAN
 - ▶ Virtualization at compute layer
 - ▶▶ Example: Virtual machines, memory virtualization

Virtualization technology has dramatically changed the complexity of storage infrastructure management. In fact, flexibility and ease of management is one of the key drivers for wide adoption of virtualization at all layers of the IT infrastructure.

At the storage layer, storage virtualization has enabled dynamic migration of data and extension of storage volumes. Due to dynamic extension, storage volumes can be expanded nondisruptively to meet both capacity and performance requirements. Because virtualization breaks the bond between the storage volumes presented to the host and its physical storage, data can be migrated both within and across data centers without any downtime. This has made the administrators' tasks easier while reconfiguring the physical environment. Virtual storage provisioning is another tool that has changed the infrastructure management cost and complexity scenario. In conventional provisioning, storage capacity is provisioned upfront in anticipation of future growth. Because growth is uneven, some users or applications find themselves running out of capacity, whereas others have excess capacity that remains underutilized. Use of virtual provisioning can address this challenge and make capacity management less challenging. In virtual provisioning, storage is allocated from the shared pool to hosts on-demand. This improves the available capacity utilization, and thereby reduces capacity management complexities. Virtualization has also contributed to network management efficiency. VSANs and VLANs made the administrators' job easier, by isolating different networks logically using management tools rather than physically separating them. Disparate virtual networks can be created on a single physical network, and reconfiguration of nodes can be done quickly without any physical changes. It has also addressed some of the security issues that might exist in a conventional environment. On the host side, compute virtualization has made host deployment, reconfiguration, and migration easier than physical environment. Compute, application, and memory virtualization have not only improved the provisioning, but also contributed to the high availability of resources.

Storage Multitenancy


- Enables multiple tenants to share the same storage resources provided by a single landlord (resource provider)
- Security and service level assurance are key concerns in a multitenant storage environment
- Secure storage multitenant environment should follow the four pillars of multitenancy
 - ▶ Secure separation
 - ▶ Service assurance
 - ▶ Availability
 - ▶ Management

Multiple tenants sharing the same resources provided by a single landlord (resource provider) is called multitenancy. Two common examples of multitenancy are multiple virtual machines sharing the same server hardware through the use of hypervisor running on the server, and multiple user applications using the same storage platform. Multitenancy is not a new concept; however, it has become a topic of much discussion due to the rise in popularity of cloud deployments as shared infrastructure is a core component of any cloud strategy.

As with any shared services, security and service level assurance are key concern in a multitenant storage environment. Secure multitenancy means that no tenant can access another tenant's data. To achieve this, any storage deployment should follow the four pillars of multitenancy:

- **Secure Separation:** This enables data path separation across various tenants in a multitenant environment. At the storage layer, this pillar can be divided into four basic requirements: separation of data at rest, address space separation, authentication and name service separation, and separation of data access.
- **Service Assurance:** Consistent and reliable service levels are integral to storage multitenancy. Service assurance plays an important role in providing service levels that can be unique to each tenant.
- **Availability:** High availability ensures a resilient architecture that provides fault tolerance and redundancy. This is even more critical when IT infrastructure is shared by multiple tenants – as the impact of any outage is magnified.
- **Management:** This includes provisions that allow a landlord to manage basic infrastructure while delegating management responsibilities to tenants for the resources that they interact with day to day. This concept is known as balancing the provider (landlord) in- control with the tenant in-control capabilities.

Storage Management Example 1 – Storage Allocation to a New Server



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 5: Securing and Managing the Storage Infrastructure 74

Figure on the slide illustrates the activities performed on server, SAN, and storage array while allocating storage to a new server.

Consider the deployment of a new RDBMS server to the existing nonvirtualized storage infrastructure. As a part of storage management activities, first, the administrator needs to install and configure the HBAs and device drivers on the server before it is physically connected to the SAN. Optionally, multipathing software can be installed on the server, which might require additional configuration. Further, storage array ports should be connected to the SAN.

As the next step, the administrator needs to perform zoning on the SAN switches to allow the new server access to the storage array ports via its HBAs. To ensure redundant paths between the server and the storage array, the HBAs of the new server should be connected to different switches and zoned with different array ports.

Further, the administrator needs to configure LUNs on the array and assign these LUNs to the storage array front-end ports. In addition, LUN masking configuration is performed on the storage array, which restricts access to LUNs by a specific server.

Cont..

The server then discovers the LUNs assigned to it by either a *bus rescan* process or sometimes through a server reboot, depending upon the operating system installed. A volume manager may be used to configure the logical volumes and file systems on the host. The number of logical volumes or file systems to be created depends on how a database or an application is expected to use the storage.

On the application side, the administrator's task includes installation of a database or an application on the logical volumes or file systems that were created. The last step is to make the database or application capable of using the new file system space.


In a virtualized environment, provisioning storage to a VM that runs an RDBMS requires different administrative tasks.

Similar to a nonvirtualized environment, a physical connection must be established between the physical server, which hosts the VMs, and the storage array through the SAN. At the SAN level, a VSAN can be configured to transfer data between the physical server and the storage array. The VSAN isolates this storage traffic from any other traffic in the SAN. Further, the administrator can configure zoning within the VSAN.

At the storage side, administrators need to create thin LUNs from the shared storage pool and assign these thin LUNs to the storage array front-end ports. Similar to a physical environment, LUN masking needs to be carried out on the storage array.

At the physical server side, the hypervisor discovers the assigned LUNs. The hypervisor creates a logical volume and file system to store and manage VM files. Then, the administrator creates a VM and installs the OS and RDBMS on the VM. While creating the VM, the hypervisor creates a virtual disk file and other VM files in the hypervisor file system. The virtual disk file appears to the VM as a SCSI disk and is used to store the RDBMS data. Alternatively, the hypervisor enables virtual provisioning to create a thin virtual disk and assigns it to the VM. Hypervisors usually have native multipathing capabilities. Optionally, a third-party multipathing software may be installed on the hypervisor.

Storage Management Example 2 – Chargeback Report



Application	Storage (GB)	Production Storage Raw (GB)	Local Replica Storage Raw (GB)	Remote Replica Storage Raw (GB)	Total Storage RAW (GB)	Chargeback Cost \$/Raw (GB)
Payroll_1	100	200	100	125	425	2125
Engineering_1	200	250	200	250	700	3500

EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 5: Securing and Managing the Storage Infrastructure 76

This example explores the storage infrastructure management tasks necessary to create a chargeback report. Figure on the slide shows a configuration deployed in a storage infrastructure. Three servers with two HBAs each connect to a storage array via two switches, SW1 and SW2. Individual departmental applications run on each of the servers. Array replication technology is used to create local and remote replicas. The production volume is represented as A, the 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. Slide shows a sample report. The report shows the information for two applications, Payroll_1 and Engineering_1.


The first step to determine chargeback costs is to correlate the application with the exact amount of raw storage configured for that application. The Payroll_1 application storage space is traced from file systems to logical volumes to volume groups and to the LUNs on the array. When the applications are 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 the production array). The replication volumes are Local Replica Vol 1 and Vol 2 (in the production array) and Remote Replica Vol 1 and Vol 2 (in the remote array).

Cont..

The amount of storage allocated to the application can be easily computed after the array devices are identified. In this example, consider 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's production volumes are RAID 1-protected, the raw space used by the production volumes is 200 GB. Assume that the local replicas are on unprotected volumes, and the remote replicas are protected with a RAID 5 configuration, then 100 GB of raw space is used by the local replica and 125 GB by the remote replica. Therefore, the total raw capacity used by the Payroll_1 application is 425 GB. The total cost of storage provisioned for Payroll_1 application will be \$2,125 (assume cost per GB of storage is \$5). This exercise must be repeated for each application in the enterprise to generate the chargeback report.


Chargeback reports can be extended to include a pre-established cost of 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.



Monitoring and managing today's complex storage infrastructure is challenging. This is due to the heterogeneity of storage arrays, networks, servers, databases, and applications in the environment. For example, heterogeneous storage arrays vary in their capacity, performance, protection, and architectures.

Each of the components in a data center typically comes with vendor-specific tools for management. An environment with multiple tools makes understanding the overall status of the environment challenging because the tools might not be interoperable. Ideally, management tools should correlate information from all components in one place. Such tools provide an end-to-end view of the environment, and a quicker root cause analysis for faster resolution to alerts.

Developing an Ideal Solution



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 5: Securing and Managing the Storage Infrastructure 79

An ideal solution should offer meaningful insight into the status of the overall infrastructure and provide root cause analysis for each failure. This solution should also provide central monitoring and management in a multivendor storage environment and create an end-to-end view of the storage infrastructure.

The benefit of end-to-end monitoring is the ability to correlate one component's behavior with the other. This is helpful to debug or analyze a problem, when looking at each component individually might not be sufficient to reveal the actual cause of the problem. The central monitoring and management system should gather information from all the components and manage them through a single-user interface. In addition, it must provide a mechanism to notify administrators about various events using methods, such as e-mail and Simple Network Management Protocol (SNMP) traps. It should also have the capability to generate monitoring reports and run automated scripts for task automation.

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

Traditionally, SNMP protocol was the standard used to manage multivendor SAN environments. However, SNMP was primarily a network management protocol and was inadequate for providing the detailed information required to manage the SAN environment. The unavailability of automatic discovery functions and weak modeling constructs 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 storage environments more effectively.

Storage Management Initiative (SMI)

- SNIA has been engaged in an initiative to develop a common storage management interface
- SNIA developed a specification called Storage Management Initiative-Specification (SMI-S)
 - ▶ It forms a normalized, abstracted model to which a storage infrastructure components can be mapped
 - ▶ Enables standardized and end-to-end control of resources
 - ▶ SMI-S compliant products are easier and faster to deploy
 - ▶ Eliminates the need for development of vendor-proprietary management interfaces

The Storage Networking Industry Association (SNIA) has been engaged in an initiative to develop a common storage management interface. SNIA has developed a specification called Storage Management Initiative-Specification (SMI-S). This specification is based on the Web-Based Enterprise Management (WBEM) technology, and Distributed Management Task Force's (DMTF) Common Information Model (CIM). The initiative was formally created to enable broad interoperability and management among heterogeneous storage and SAN components. For more information, see www.snia.org.

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. This model is used by management applications, such as storage resource management, device management, and data management, for standardized, end-to-end control of storage resources.

Using SMI-S, device software developers have a unified object model with details about managing the breadth of storage and SAN components. SMI-S-compliant products lead to easier, faster deployment and accelerated adoption of policy-based storage management frameworks. Moreover, SMI-S eliminates the need for the development of vendor-proprietary management interfaces and enables vendors to focus on value-added features.

Enterprise Management Platform

- Suite of applications for managing and monitoring storage infrastructure
 - ▶ Provides end-to-end management of physical and virtual resources
 - ▶ Generates alerts to inform the status of components and processes
 - ▶ Schedules operations such as provisioning and configuration management

An enterprise management platform (EMP) is a suite of applications that provides an integrated solution for managing and monitoring an enterprise storage infrastructure. These applications have powerful, flexible, unified frameworks that provide end-to-end management of both physical and virtual resources. EMP provides a centrally managed, single point of control for resources throughout the storage environment.

These applications can proactively monitor storage infrastructure components and alert users about events. 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. In addition to monitoring, an EMP provides the necessary management functionality, which can be natively implemented into the EMP or can launch the proprietary management utility supplied by the component manufacturer.

An EMP also 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. EMC ControlCenter and EMC Prosphere described in this module are examples of an EMP.

Module 15: Managing the Storage Infrastructure

Lesson 3: Information Lifecycle Management

During this lesson the following topics are covered:

- Challenges of managing information
- Information lifecycle management
- Storage tiering

This lesson describes the challenges of managing information. It also describes information lifecycle management strategy. Finally, the lesson explains storage tiering.


Challenges in Managing Information

- Exploding digital universe
 - ▶ Multifold increase of information
- Increasing dependency on information
 - ▶ Strategic use of information plays an important role in determining the success of a business
- Changing value of information
 - ▶ Information that is valuable today may become less important in future

In both traditional data center and virtualized environments, managing information can be expensive if not managed appropriately. Along with the tools, an effective management strategy is also required to manage information efficiently. This strategy should address the following key challenges that exist in today's data centers:

- Exploding digital universe: The rate of information growth is increasing exponentially. Creating copies of data to ensure high availability and repurposing has contributed to the multifold increase of information growth.
- Increasing dependency on information: The strategic use of information plays an important role in determining the success of a business and provides competitive advantages in the marketplace.
- Changing value of information: Information that is valuable today might become less important tomorrow. The value of information often changes over time.

Information Lifecycle Management



A proactive strategy that enables an IT organization to effectively manage the information throughout its lifecycle

EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 5: Securing and Managing the Storage Infrastructure 84

Framing a strategy to meet these challenges involves understanding the value of information over its life cycle. When information is first created, it often has the highest value and is accessed frequently. As the information ages, it is accessed less frequently and is of less value to the organization. Understanding the value of information helps to deploy the appropriate infrastructure according to the changing value of information.

For example, in a sales order application, the value of the information (customer data) changes from the time the order is placed until the time that the warranty becomes void. The value of the information is highest when a company receives a new sales order and processes it to deliver the product. After the order fulfillment, the customer data does not need to be available for real-time access. The company can transfer this data to less expensive secondary storage with lower performance until a warranty claim or another event triggers its need. After the warranty becomes void, the company can dispose of the information.

Information lifecycle management (ILM) is a proactive strategy that enables an IT organization to effectively manage information throughout its life cycle based on predefined business policies. From data creation to data deletion, ILM aligns the business requirements and processes with service levels in an automated fashion. This allows an IT organization to optimize the storage infrastructure for maximum return on investment.

Benefits of ILM

- Provides lower total cost of ownership
 - ▶ By aligning the infrastructure and management costs with changing value of information
- Provides simplified management
 - ▶ By automating the processes
- Maintains compliance
 - ▶ Knowledge of what data needs to be protected for what length of time
- Optimized utilization
 - ▶ By deploying storage tiering

Implementing an ILM strategy has the following key benefits that directly address the challenges of information management:

- **Simplified management:** By integrating process steps and interfaces with individual tools and by increasing automation.
- **Maintaining compliance:** By knowing what data needs to be protected for what length of time.
- **Lower Total Cost of Ownership (TCO):** By aligning the infrastructure and management costs with information value. As a result, resources are not wasted, and complexity is not introduced by managing low-value data at the expense of high-value data.
- **Optimized utilization:** By deploying storage tiering.

Storage Tiering

Storage Tiering

It is a technique of establishing a hierarchy of storage types and identifying the candidate data to relocate to the appropriate storage type to meet service level requirements at a minimal cost.

- Each tier has different levels of protection, performance, and cost
- Efficient storage tiering requires defining tiering policies

Storage tiering is a technique of establishing a hierarchy of different storage types (tiers). This enables storing the right data to the right tier, based on service level requirements, at a minimal cost. Each tier has different levels of protection, performance, and cost. For example, high performance solid-state drives (SSDs) or FC drives can be configured as tier 1 storage to keep frequently accessed data and low cost SATA drives as tier 2 storage to keep the less frequently accessed data. Keeping frequently used data in SSD or FC improves application performance. Moving less-frequently accessed data to SATA can free up storage capacity in high performance drives and reduce the cost of storage. This movement of data happens based on defined tiering policies. The tiering policy might be based on parameters, such as file type, size, frequency of access, and so on. For example, if a policy states “Move the files that are not accessed for the last 30 days to the lower tier,” then all the files matching this condition are moved to the lower tier.

Storage Tiering Implementation


- Storage tiering implementations are:
 - ▶ Manual storage tiering
 - ▶ Automated storage tiering
- Data movement occurs between tiers
 - ▶ Within a storage array (Intra-array)
 - ▶ Between storage arrays (Inter-array)

Storage tiering can be implemented as a manual or an automated process. *Manual storage tiering* is the traditional method where the storage administrator monitors the storage workloads periodically and moves the data between the tiers. Manual storage tiering is complex and time-consuming. *Automated storage tiering* automates the storage tiering process, in which data movement between the tiers is performed nondisruptively. In automated storage tiering, the application workload is proactively monitored; the active data is automatically moved to a higher performance tier and inactive data to higher capacity, lower performance tier. Data movements between various tiers can happen within (intra-array) or between (inter-array) storage arrays.

Intra-array Storage Tiering


- LUN tiering

- ▶ Moves entire LUN from one tier to another
- ▶ Does not give effective cost and performance benefits



- Sub-LUN tiering


- ▶ A LUN is broken down into smaller segments and tiered at that level
- ▶ Provides effective cost and performance benefits



The process of storage tiering within a storage array is called *intra-array storage tiering*. It enables the efficient use of SSD, FC, and SATA drives within an array and provides performance and cost optimization. The goal is to keep the SSDs busy by storing the most frequently accessed data on them, while moving out the less frequently accessed data to the SATA drives. Data movements executed between tiers can be performed at the LUN level or at the sub-LUN level. The performance can be further improved by implementing tiered cache.

Traditionally, storage tiering is operated at the LUN level that moves an entire LUN from one tier of storage to another. This movement includes both active and inactive data in that LUN. This method does not give effective cost and performance benefits. Today, storage tiering can be implemented at the sub-LUN level. In sub-LUN level tiering, a LUN is broken down into smaller segments and tiered at that level. Movement of data with much finer granularity, for example 8 MB, greatly enhances the value proposition of automated storage tiering. Tiering at the sub-LUN level effectively moves active data to faster drives and less active data to slower drives.

Inter-array Storage Tiering




EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Section 5: Securing and Managing the Storage Infrastructure 89

The process of storage tiering between storage arrays is called *inter-array storage tiering*. Inter-array storage tiering automates the identification of active or inactive data to relocate them to different performance or capacity tiers between the arrays. Figure on the slide illustrates an example of a two-tiered storage environment. This environment optimizes the primary storage for performance and the secondary storage for capacity and cost. The policy engine, which can be software or hardware where policies are configured, facilitates moving inactive or infrequently accessed data from the primary to the secondary storage. Some prevalent reasons to tier data across arrays is archival or to meet compliance requirements. As an example, the policy engine might be configured to relocate all the files in the primary storage that have not been accessed in one month and archive those files to the secondary storage. For each archived file, the policy engine creates a small space-saving stub file in the primary storage that points to the data on the secondary storage. When a user tries to access the file at its original location on the primary storage, the user is transparently provided with the actual file from the secondary storage.

Cache Tiering

- Enables creation of a large capacity secondary cache using SSDs
- Enables tiering between DRAM cache and SSDs (secondary cache)
- Most reads are served directly from high performance tiered cache



Benefits

- Enhances performance during peak workload
- Non-disruptive and transparent to applications

Tiering is also implemented at the cache level. A large cache in a storage array improves performance by retaining large amount of frequently accessed data in a cache, so most reads are served directly from the cache. However, configuring a large cache in the storage array involves more cost. An alternative way to increase the size of the cache is by utilizing the SSDs on the storage array. In cache tiering, SSDs are used to create a large capacity secondary cache and to enable tiering between DRAM (primary cache) and SSDs (secondary cache). Server flash-caching is another tier of cache in which flash-cache card is installed in the server to further enhance the application performance.

Module 15: Managing the Storage Infrastructure

Concept in Practice

- EMC ControlCenter
- EMC Prosphere
- EMC Unisphere
- EMC Unified Infrastructure Management (UIM)

The Concept in Practice covers the product example of storage infrastructure management software. It covers four products: EMC ControlCenter, EMC Prosphere, EMC Unisphere, and EMC Unified Infrastructure Management.

EMC ControlCenter

- An enterprise management platform
 - ▶ Contains storage resource management applications to manage a multi-vendor storage infrastructure
- Provides an end-to-end view of the entire networked storage infrastructure including virtualized environment
- Enables implementation of ILM strategy by providing management of tiered storage infrastructure
- Has built-in security features that provide access control, data confidentiality, data integrity, logging, and auditing

EMC ControlCenter is a family of storage resource management (SRM) applications that provide a unified solution to manage a multivendor storage infrastructure. It helps address the challenges to manage a large, complex storage environment that includes hosts, storage networks, storage, and virtualization across all the layers. ControlCenter provides capabilities, such as storage planning, provisioning, monitoring, and reporting. It enables implementing an ILM strategy by providing comprehensive management of tiered storage infrastructure. It also provides an end-to-end view of the entire networked storage infrastructure that includes SAN, NAS, and host storage resources, including virtualized environment. It provides a central administrative console, discovery of new components, quota management, event management, root cause analysis, and chargeback. ControlCenter comes with built-in security features that provide access control, data confidentiality, data integrity, logging, and auditing. It offers an intuitive, easy-to-use interface that provides insight into the complex relationships of the environment. ControlCenter uses an agent to discover the components in the environment.

EMC ProSphere

- It is a storage resource management software, built to meet the demands of the new cloud computing era
- Following are the key capabilities:
 - ▶ Provides end-to-end visibility of all objects
 - ▶ Enables multi-site management from a single console
 - ▶ Improves productivity in virtualized environments with “Smart Groups”
 - ▶ Enables fast, easy, and efficient deployment
 - ▶▶ Agent-less discovery of objects
 - ▶▶ ProSphere is packaged as virtual appliance

EMC ProSphere is also storage resource management software built to meet the demands of the new cloud computing era. EMC ProSphere improves productivity and service levels in the virtualized and cloud environment. ProSphere includes the following key capabilities:

- **End-to-end visibility:** It offers an intuitive, easy-to-use interface that provides insight into the complex relationships between objects in large, virtualized environments.
- **Multi-site management:** From a single console, ProSphere’s federated architecture aggregates information from across sites and simplifies information management between data centers. ProSphere is managed from a web browser to allow easy access over the Internet for remote management.
- **Improved productivity in growing virtualized environments:** ProSphere introduces an innovative technology called Smart Groups, which groups objects with similar characteristics into a user-defined group for performing management tasks. This enables IT to take a policy-based approach to manage objects or to set data collection policies.
- **Fast, easy, and efficient deployment:** Agent-less discovery eliminates the burden of deploying and managing host agents. ProSphere is packaged as a virtual appliance that can be installed in a short time.
- **Delivery of IT as a service:** With ProSphere, service levels can now be monitored from host-to-storage layers. This allows organizations to maintain consistent service levels at an optimal price-performance ratio to meet business objectives to delivering IT-as-a-service.

EMC Unisphere

- It is a unified storage management platform for managing:
 - ▶ EMC VNX and VNXe
 - ▶ EMC RecoverPoint/SE
- Some of the key capabilities offered by Unisphere are:
 - ▶ Provides unified management for file-based, block-based, and object-based storage
 - ▶ Supports automated storage tiering
 - ▶ Provides management of both physical and virtual components

EMC Unisphere is a unified storage management platform that provides intuitive user interfaces for managing EMC RecoverPoint SE, and EMC VNX and VNXe storage arrays. Unisphere is web-enabled and supports remote management of storage arrays. Some of the key capabilities offered by Unisphere follow:

- Provides unified management for file, block, and object storage
- Provides single sign-on for all devices in a management domain
- Supports automated storage tiering and ensures that data is stored in the correct tier to meet performance and cost
- Provides management of both physical and virtual components

EMC Unified Infrastructure Manager (UIM)

- Unified management solution for Vblocks
- Enables configuring Vblock resources and activating services
- Provides a dashboard showing Vblock infrastructure configuration and resource utilization
- Provides a topology view of Vblock infrastructure
- Provides an alerts console that lists alerts against adversely affected resources and services
- Performs compliance check during resource configuration

EMC Unified Infrastructure Manager is a unified management solution for Vblocks. (Vblock is covered in module 13.) It enables configuring the Vblock infrastructure resources and activating cloud services. It provides a single user interface to manage multiple Vblocks and eliminates the need for configuring compute, network, and storage separately using different virtual infrastructure management tools.

UIM provides a dashboard that shows how the Vblock infrastructure is configured and how the resources are used. This enables an administrator to monitor the configuration and utilization of the Vblock infrastructure resources and to plan for capacity requirements. UIM also provides a topology or a map view of the Vblock infrastructure, which enables an administrator to quickly locate and understand the interconnections of the Vblock infrastructure components and services. It provides an alerts console, which allows an administrator to see the alerts against the Vblock infrastructure resources and the associated services affected by problems. UIM performs a compliance check during resource configuration. It validates compliance with configuration best practices. It also prevents conflicting resource identity assignments, for example, accidentally assigning a MAC address to more than one virtual NIC.

Module 15: Summary

Key points covered in this module:

- Key storage infrastructure components that are monitored
- Key monitoring parameters
- Storage management activities
- Storage infrastructure management challenges
- Enterprise management platform
- Information lifecycle management
- Storage tiering

This module covered the key storage infrastructure management components that should be monitored such as servers, network, storage arrays, and environmental controls.

These key components should be monitored for accessibility, capacity, performance, and security.

This module also covered the key management activities such as availability management, capacity management, performance management, security management, and reporting.

This module also covered the key infrastructure management challenges and the ideal solution.

Further, this module covered information lifecycle management (ILM) which is a proactive strategy that enables an IT organization to effectively manage information throughout its lifecycle, based on predefined business policies.

Finally, this module detailed storage tiering techniques which identifies the candidate data and relocate them to the appropriate storage type to meet service level requirements at a minimal cost.

Check Your Knowledge – 1

- Which type of alert is generated if soft media errors on a disk drive approach its pre-defined threshold value?
 - A. Fatal
 - B. Warning
 - C. Information
 - D. Watermark
- What is a purpose of chargeback report?
 - A. Reports investment in managing infrastructure
 - B. Reports cost of decommissioning infrastructure components
 - C. Reports utilization of infrastructure components by various users
 - D. Reports charges for SLA breach

Check Your Knowledge – 2

- Which monitoring parameter helps ensuring availability of adequate amount of resources and preventing service unavailability?
 - A. Availability
 - B. Capacity
 - C. Performance
 - D. Security
- Which pillar of multitenancy ensures consistent and reliable service levels in a multitenant storage environment?
 - A. Secure separation
 - B. Service assurance
 - C. Service availability
 - D. Storage tiering

Check Your Knowledge – 3

- What best describes the SMI specification?
 - A. Restricts vendors to build new features and functions to manage storage subsystems
 - B. Eliminates the need for development of vendor-proprietary management interfaces
 - C. Prevents interoperability among multi-vendor resources
 - D. Enable deploying ILM supported infrastructure

Section 5: Summary

Key points covered in this section:

- Information security framework
- Storage security domains
- Protection mechanisms in SAN, NAS, and IP SAN environments
- Security in virtualized and cloud environments
- Key storage infrastructure components that are monitored
- Storage management activities
- Information lifecycle management and Storage tiering

This section covered information security framework and various storage security domains. It also covered security implementation in SAN, NAS, and IP SAN environments. Further, this section covered security in virtualized and cloud environment. It also covered the managing and monitoring of key storage infrastructure components.

Finally, this section covered information lifecycle management (ILM) and storage tiering.

Course Summary

Key points covered in this course:

- Key elements of data center environment
- Key components of intelligent storage systems and RAID technology
- Storage networking technologies
- Business continuity solutions
- Cloud computing technology
- Security and management of storage infrastructure

This course covered the importance of information in our daily lives. The increasing dependency of information to the businesses has amplified the challenges in storing, protecting, and managing data.

The course also detailed the five core elements of data center that are essential for its functionality such as application, database management system, host or compute, network, and storage.

The course also covered RAID which is a technique of combining multiple disk drives into a logical unit and provide protection, performance, or both.

This course also covered intelligent storage systems which are feature-rich RAID arrays that provide highly-optimized I/O processing capabilities.

Further, this course detailed storage networking technologies such as FC SAN, IP SAN, and FCoE that provide block access to storage. The course also covered NAS which is a dedicated, high-performance file sharing and storage device that enables its clients to share files over an IP network. This module covered object-based and unified storage.

Besides this, the course also covered business continuity solutions such as backup and replication. It also covered data archiving solutions that enable to meet compliance.

It also covered cloud computing which is a next generation style of computing that provides highly scalable and flexible computing that is available on demand.

Finally, this module covered security of storage infrastructure that is essential to ensure confidentiality, integrity, and availability. It also covered the management activities in storage environment.

Thank You!

This concludes the training. Thank you for your participation.