Cisco Storage Media Encryption Design Guide for Cisco MDS 9000 NX-OS Software Release 5.2(6)

Design Guide

June, 2012

For further information, questions and comments please contact ccbu-pricing@cisco.com

Contents

Audience 4
What Is New 4
Cisco Storage Media Encryption Overview 4
Cisco SME Terminology 6
- Terminology Specific to Cisco SME Tape 6
- Terminology Specific to Cisco SME Disk 6

Cisco SME Requirements 7
- Software Requirements 7
- Hardware Requirements 7
  - Security Engine Requirements 7
  - Cisco DCNM and KMC Workstation Requirements 8
  - Smart Card Readers 8
- License Requirements 9
- Topology Requirements 9
- Zoning Requirements 11
- FC-Redirect Requirements 11
  - Requirements Specific to Cisco SME Disk 12
  - Requirements Specific to Cisco SME Tape 12
- Configuration Requirements 12
- Requirements for Replication and Snapshot Technologies 12

Cisco SME Data Flow 12
- Write I/O Data Flow 12
- Read I/O Data Flow 13
- Single Cisco SME Switch 14

Cisco SME Solution Architecture 15
Cisco SME Security Architecture 16
- Securing Communications Among Components 16
- Securing Cisco SME Keys 17

Cisco SME Administrative Roles 18
Cisco SME Clustering 19
- Cisco SME Data Flow in a Cluster 19
- Failure Conditions 20

Cisco SME Design Considerations 21
- Host-Based Cluster Considerations (Cisco SME Disk Only) 21
- Cisco SME Clustering: High-Availability Considerations 21
- Network Topologies 22
- Topology Considerations 22
  - General Guidelines 22
  - Guidelines Specific to Multiple Clusters 22
  - Core-Edge Topology 22
  - Edge-Core-Edge Topology 24
  - FCIP Topology 25
- Cisco SME Load-Balancing Considerations 26
- Cisco KMC Considerations 26
  - Cisco KMC High Availability 26
  - Key Database High Availability 27
- Cisco DCNM Guidelines 27

Cisco SME Solution Scalability Limits ... 28
Encryption Throughput Guidelines ... 29
Disk Encryption Latency Guidelines (Cisco SME Disk Only) ... 30
Read Operations ... 30
Write Operations ... 30
Sizing Guidelines ... 30
Sizing Using Line-Rate Disk or Tape Drive Throughput ... 30
Sizing Using Backup Window Calculations ... 31
Inserting Cisco SME Tape into Existing Cisco SANs ... 31
Cisco SME Disk Deployment in Storage Environments ... 31
Managing Replication with Cisco SME ... 32
Managing Snapshots in Cisco SME ... 33
Replication and Mirroring ... 33
Copy-on-Write Snapshots ... 34
EMC Recover Point IO Journal Snapshots ... 35
Mix of Disk and Tape Devices ... 35
Inserting Cisco SME Disk into Existing Storage Environments ... 36
LUNs Containing Preexisting Data ... 36
LUNs Not Containing Preexisting Data ... 36
Maintaining Data Security ... 37
Rekey (Cisco SME Disk Only) ... 37
Master-Key Rekey ... 37
Disk Key Replication ... 37
Appendix ... 38
Cisco SME Tape Checklist ... 38
Cisco SME Disk Checklist ... 38
Generic Checklist ... 39
Sizing and Placement Deployment Examples ... 39
Example 1: Single-Switch Single-Fabric Tape Backup Environment ... 39
Example 2: Core-Edge Topology Tape Backup Environment ... 40
Example 3: Edge-Core-Edge Topology Tape Backup Environment with Restricted Zoning ... 42

Audience

This guide is for sales engineers and storage administrators who want to understand the Cisco® Storage Media Encryption (SME) service. Basic knowledge of Cisco MDS 9000 Family Fibre Channel concepts and SANs, including data storage and tape backup environments, is expected. Familiarity with the Fibre Channel Redirect (FC-Redirect) feature of the Cisco MDS 9000 NX-OS Software is desirable.

This design guide provides details about Cisco SME data flow, supported topologies, and best practices for Cisco SME deployment in data storage and tape backup environments.

What Is New

Cisco MDS 9000 NX-OS Software Release 5.2(6) provides the following new enhancements:

- Cisco SME Disk support for storage snapshots
- Cisco SME Disk support for master-key rekey
- Cisco SME Disk support for logical unit numbers (LUNs) over 2 terabytes (2 TB)

Cisco Storage Media Encryption Overview

Encryption of storage media in the data center has become a critical issue. Numerous high-profile incidents of lost or stolen tapes and disk devices have underscored the risk and exposure that companies face when sensitive information falls into the wrong hands. As these devices are retired, disposed of, sent out for service, or repurposed and moved from the data center or reused, the data on them must be kept encrypted and must not be vulnerable to recovery. Often, however, discarded devices contain clear data that can be accessed by an unauthorized user.

Furthermore, regulatory requirements arising from Sarbanes-Oxley, Health Insurance Portability and Accountability Act (HIPAA), Payment Card Industry Data Security Standard (PCI-DSS), and other laws have made encryption a top priority.

To meet these requirements, Cisco has introduced the Cisco SME solution. Cisco SME protects against such threats by encrypting data on heterogeneous disk arrays, tape drives, and virtual tape libraries (VTLs), using Advanced Encryption Standard (AES) 256-bit algorithms: specifically, AES Galois counter mode (AES-GCM) for tape drives and AES XOR-encrypt-XOR (XEX)-based tweaked-codebook mode with ciphertext stealing (AES-XTS) for disk arrays. Cisco SME is a comprehensive network-integrated encryption service with enterprise-class key management that works transparently with existing and new SANs. Figure 1 shows a high-level view of a SAN with Cisco SME service deployed.

Figure 1. Cisco Storage Media Encryption

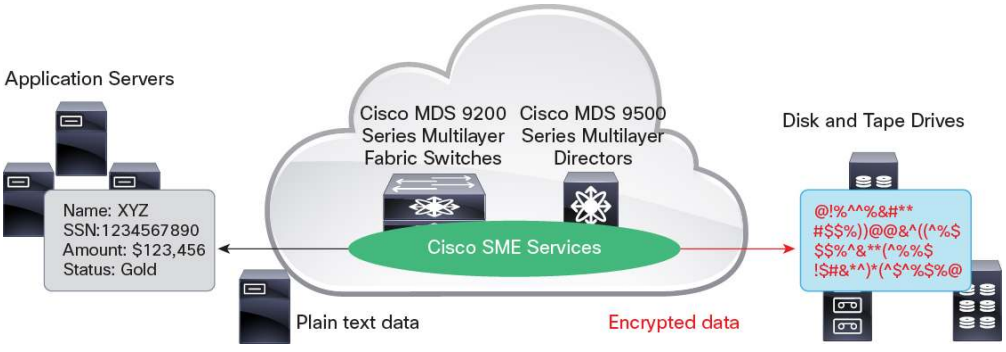

Cisco SME is a secure, integrated solution that delivers encryption as a SAN service (Figure 2). It provides intuitive provisioning, support for heterogeneous SAN devices, comprehensive key management, and role-based access control (RBAC). Using the clustering infrastructure, Cisco SME provides scalability, high availability, and automatic load balancing.

Figure 2. Cisco SME: Secure, Integrated Solution

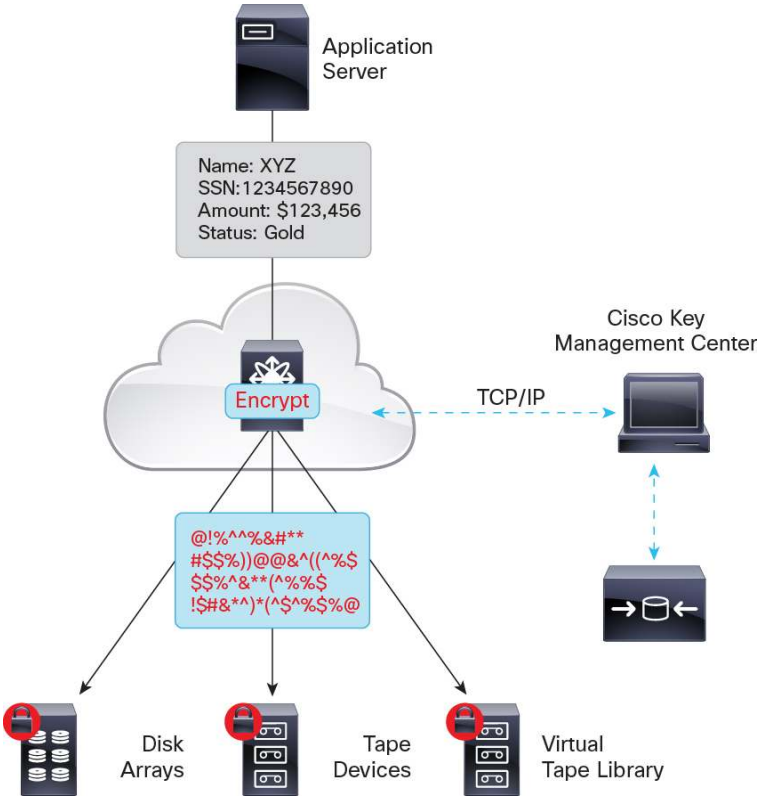

The Cisco SME solution offers numerous advantages when you are implementing encryption for data at rest:

- Cisco SME installation and provisioning are simple and nondisruptive. Cisco SME does not require rewiring or SAN reconfiguration.
- Cisco SME uses FC-Redirect infrastructure to redirect the data flows that need encryption to its engines.
- Cisco SME for disk (SME Disk) and Cisco SME for tape (SME Tape) are supported on the same Cisco MDS 9000 16-Port Storage Services Node (SSN) and can run concurrently.
- Service engines are integrated into the Cisco MDS 9000 18/4-Port Multiservice Module (MSM), Cisco MDS 9000 16-Port SSN, and Cisco MDS 9222i Multiservice Modular Switch (MMS), eliminating the need to purchase and manage additional switch ports, cables, and appliances.
- Traffic from any virtual SAN (VSAN) can be encrypted using Cisco SME, enabling flexible, automated load balancing through network traffic management across multiple SANs.
- No additional software is required for provisioning, key, and user-role management; Cisco SME is integrated into Cisco Data Center Network Manager (DCNM) for the SAN (DCNM-SAN), therefore reducing operating expenses.
- Cisco SME does not increase the size of the data set or LUN.

Cisco SME Terminology

This section explains the terminology that is required to build a Cisco SME solution.

- Cisco SME interface: The security engine in the Cisco MDS 9000 18/4-Port MSM and 16-Port SSN Cisco Service line cards or in a fixed slot of a Cisco MDS 9222i fabric switch. Cisco MSM-18/4 line card and Cisco MDS 9222i switch each have one security engine; the Cisco MDS 9000 16-Port SSN line card has four security engines
- Cisco SME cluster: A network of Cisco MDS 9000 Family switches that are configured to provide the Cisco SME function; each switch includes one or more Cisco Service line cards, and the switches in the cluster use IP connectivity through the management interface for communication
- Cisco SME cluster node: The Cisco MDS 9000 Family switch that is part of a Cisco SME cluster
- Fabric: A physical fabric topology in the SAN as seen by Cisco DCNM-SAN; the physical fabric can include multiple VSANs (logical fabrics)
- Cisco Key Management Center (KMC): A component of the Cisco DCNM server that stores the encryption keys (for details, refer to the Cisco Storage Encryption Media Key Management white paper)
- Cisco SME key hierarchy: The keys included in the Cisco SME key management system
- Smart card: A card (approximately the size of a credit card) with a built-in microprocessor and memory used for authentication; it is used to store the master key recovery shares for Cisco SME recovery officers
- Cisco SME administrator: A network administrator who configures Cisco SME
- Cisco SME recovery officer: A data security officer entrusted with smart cards and the associated personal identifier numbers (PINs):
  - Each smart card stores a share of the master key of the cluster.
  - Recovery officers must present their cards and PINs to recover the key database of an archived cluster.
  - A quorum of recovery officers is required to perform this operation.
- FC-Redirect: Capability in Cisco MDS 9000 NX-OS Software that enables traffic from any switch port to be encrypted without SAN reconfiguration or rewiring
- Cisco SME line card: A module capable of providing Cisco SME services: the Cisco MDS 9000 18/4-Port MSM or 16-Port SSN or the integrated supervisor module on the Cisco MDS 9222i switch; for simplicity, this document uses “Cisco SME line card”

Terminology Specific to Cisco SME Tape

- Tape group: A logical grouping of tape volumes that are configured for a specific use: for example, a group of tape volumes used to back up a database
- Tape device: A tape device that is configured for encryption
- Tape volume: A physical tape cartridge identified by a bar code for a given use

Terminology Specific to Cisco SME Disk

- Disk group: A logical grouping of disk devices that are configured for a specific use: for example, a group of disk devices used to store employee payroll information
- Disk device: A disk drive that is configured for encryption

- Data preparation: Operation in which existing clear-text data on a LUN is read, encrypted with the current volume key, and written back to the LUN at the same logical block address (LBA)
- Rekey: Operation in which previously encrypted data on a LUN is decrypted with the current volume key, reencrypted with a new volume key, and written back to the same LUN at the same LBA
- Mirror (or clone): Destination disk created when data for the source disk is duplicated by the disk array on another disk in the same storage system
- Replication: Operation in which data for the source disk is duplicated by the disk array on another disk in a remote storage system
- Snapshot: Point-in-time copy that can be created instantly for a source disk; if any write operations to the source disk occur after a snapshot is created, the previous data will be saved elsewhere before modification, allowing the disk array to present a specific point-in-time copy of the data of the source disk
- Disk key replication (DKR): Feature in Cisco MDS 9000 NX-OS Software that helps ensure that the appropriate key association occurs on LUN clones, LUN replicas, and snapshots; after this association, any key modifications that occur on one side of the relationship

Cisco SME Requirements

This section introduces the main requirements that must be met before configuration of the Cisco SME solution can begin.

Software Requirements

All Cisco MDS 9000 Family switches in the Cisco SME cluster must be running the current release of Cisco DCNM-SAN and Cisco MDS 9000 NX-OS Software:

- Cisco DCNM must be running Cisco DCNM-SAN 5.2(1) or later
- Cisco MDS 9000 Family switches attached to disk devices must be running Cisco MDS 9000 NX-OS Software Release 5.2(1) or later
- Cisco MDS 9000 Family switches attached to tape devices must be running Cisco MDS 9000 NX-OS Software Release 3.4(0) or later
- All switches that include the Cisco MDS 9000 18/4-Port MSM or 16-Port SSN must be running Cisco MDS 9000 NX-OS Software Release 5.2(1) or later
- All switches that support FC-Redirect must be running Cisco MDS 9000 NX-OS Software Release 5.2(1) or later

Hardware Requirements

Security Engine Requirements

Cisco SME requires at least one encryption service engine in each cluster. The Cisco SME engines provide the transparent encryption and compression services to the hosts and storage devices. The following hardware supports Cisco SME:

- Cisco MDS 9000 18/4-Port MSM
- Cisco MDS 9000 16-Port SSN
- Integrated supervisor module on the Cisco MDS 9222i

The Cisco MDS 9000 18/4-Port MSM and Cisco MDS 9222i fixed module can host one Cisco I/O Accelerator (IOA) engine; the Cisco MDS 9000 16-Port SSN can host up to four engines.

Cisco DCNM and KMC Workstation Requirements

A separate dedicated workstation with the following configuration should be used for Cisco DCNM and KMC for Cisco SME purposes:

- CPU: 2 GHz or more
- Memory: 2 GB or more
- Disk space: 20 GB or more

A best practice is to have a separate server running Cisco KMC and DCNM.

Smart Card Readers

To deploy standard and advanced security levels, Cisco SME requires the following:

- Smart card reader for Cisco SME (DS-SCR-K9)
- Smart card for Cisco SME (DS-SC-K9)

The smart card reader is a USB device that is connected to a management workstation (Microsoft Windows 2000, XP, and 2003 platforms only). The Cisco SME web client on the management workstation is used to configure the Cisco SME cluster (Figure 3).

Figure 3. Cisco SME Smart Card and Reader

The smart card reader requires the smart card drivers that are included on the installation CD. These must be installed on the management workstation to which the reader is attached (Figure 4).

Figure 4. Smart Card Reader

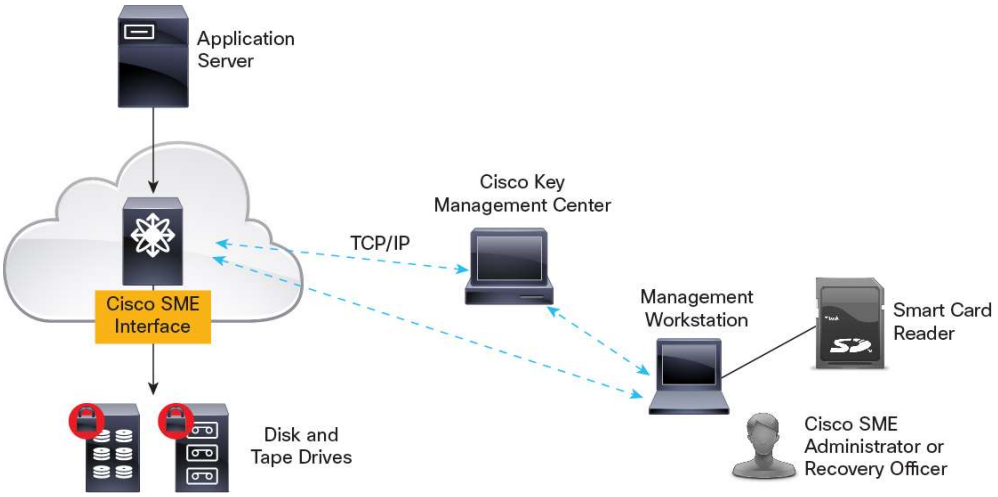

The smart card reader is required only for initial configuration and data recovery. It is not required for normal Cisco SME operations.

License Requirements

Each Cisco SME engine needs a separate Cisco SME license. For Cisco MDS 9000 18/4-Port MSM and MDS 9222i MMS, one license is required. For Cisco MDS 9000 16-Port SSN, one Cisco SME license is required for each of the four engines; for example, to run all four Cisco SME engines on the MDS 9000 16-Port SSN, four Cisco SME licenses are required.

License packages are summarized in Table 1.

Table 1. Cisco SME License Packages

Part Number	Description	Applicable Product
M9500SME1MK9	Cisco SME license for Cisco MDS 9500 MSM-18/4	Cisco MDS 9500 Series with 18/4-Port MSM
M9200SME1MK9	Cisco SME license for Cisco MDS 9200 MSM-18/4	Cisco MDS 9200 Series with 18/4-Port MSM
M9200SME1FK9	Cisco SME package for fixed slot	Cisco MDS 9222i switch only
M95SMESSNK9=	Cisco SME package for Cisco MDS 9500 SSN-16	Cisco MDS 9500 Series with 16-Port SSN
M92SMESSNK9=	Cisco SME package for Cisco MDS 9200 SSN-16	Cisco MDS 9200 Series with 16-Port SSN

The DCNM SAN Advanced license is required to run Cisco KMC. A single Cisco KMC can support multisite deployment. Two instances provide 1+1 high availability. The web client (supported by the DCNM SAN Advanced license) provides the Cisco SME provisioning wizard. At least a single instance of DCNM SAN Advanced is thus required.

Cisco DCNM license packages are summarized in Table 2.

Table 2. Cisco DCNM License Packages

Part Number	Description	Applicable Product
DCNM-SAN-M95-K9	DCNM for SAN Advanced Edition for MDS 9500	Cisco MDS 9500 Series
DCNM-SAN-M92-K9	DCNM for SAN Advanced Edition for MDS 9200	Cisco MDS 9200 Series

Note: Cisco MDS 9000 Series Switches do not need Cisco Fabric Manager Server (FMS) license packages to provision Cisco SME or to use the associated key management capabilities.

Topology Requirements

Cisco SME supports dual-fabric, Fibre Channel over IP (FCIP), and remote replication topologies. The following prerequisites for a Cisco SME Cluster are assumed:

- All switches configured in the cluster must have IP connectivity with each other and with the KMC server.
- The initiator and target pair for which the data is being encrypted must be in the same Cisco SME cluster.
- The FCIP link cannot span two separate Cisco SME clusters.

Cisco SME is fully supported in a dual-fabric, FCIP, and remote-replication SAN consisting of Cisco switches only. Cisco SME may be supported in some environments consisting of both Cisco's and other vendors' switches; such a configuration must be evaluated on a case-by-case basis.

In a dual-fabric SAN topology, one or more (up to a maximum of eight) switches capable of supporting Cisco SME form a cluster (Figure 5). Cisco SME supports two clusters for each dual fabric in a data center SAN environment.

In FCIP and remote replication topologies, a Cisco SME cluster should not span remote data centers.

Figure 5. Cisco SME Cluster in a Dual-Fabric SAN

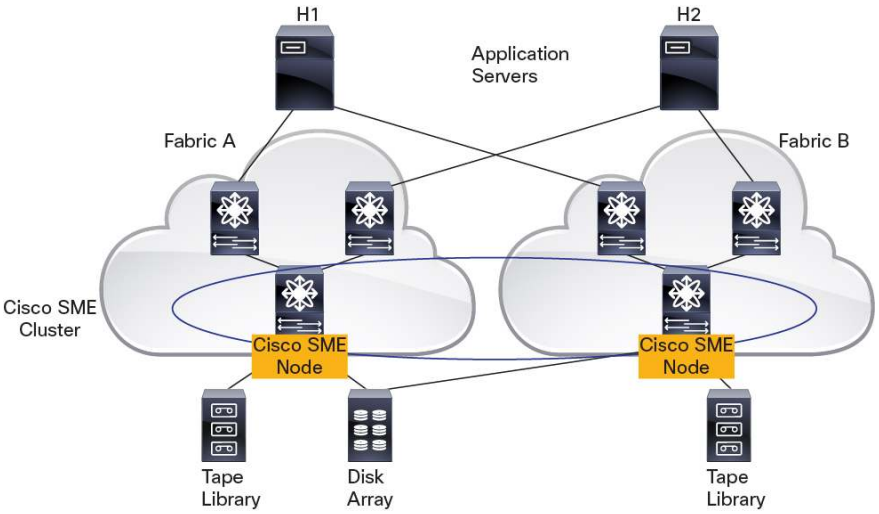

A Cisco SME cluster can span multiple VSANs in a fabric. In Figure 6, the same Cisco SME cluster encrypts traffic in multiple VSANs.

Figure 6. Cisco SME Cluster Spanning Multiple VSANs

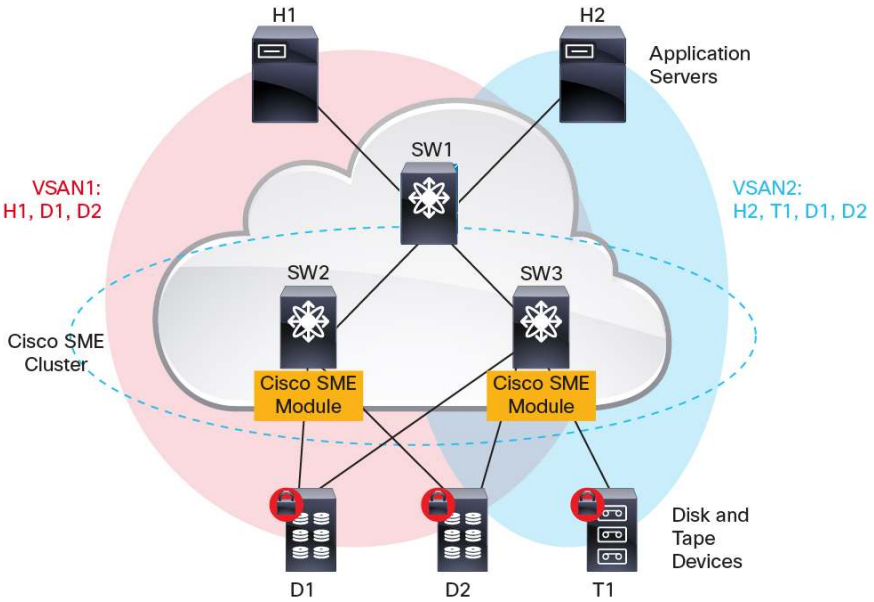

Note: Figure 6 is for illustration purposes only. Heterogeneous storage environments are not supported within a single Cisco SME cluster. Separate Cisco SME clusters must be created to support disks and tapes.

For more guidelines on the topology design, refer to “Network Topologies” later in this document.

Zoning Requirements

Cisco SME creates internal virtual N-ports in the default zone. The default zone must be set to Deny, and these virtual N-ports must not be zoned with any other host or target.

FC-Redirect Requirements

FC-Redirect is a generic flow-redirection technology that enables redirection of a flow to a specific service engine in the fabric to provide intelligent services such as Cisco SME, IOA, and Data Mobility Manager (DMM). The target device that needs the service must connect to a switch that supports FC-Redirect for the flows to be redirected to the correct service engine. FC-Redirect switches use Cisco Fabric Services to remain synchronized with peer switches.

Cisco SME requires that each target switch be capable of FC-Redirect, and FC-Redirect Version 2 (v2) should be enabled on all switches in the network. FC-Redirect is supported on the following Cisco switches:

- Cisco MDS 9500 Series
- Cisco MDS 9222i
- Cisco MDS 9216i and 9216A Multilayer Fabric Switches
- Cisco MDS 9124 Multilayer Fabric Switch (24 ports; 4 Gbps)
- Cisco MDS 9148 Multilayer Fabric Switch (48 ports; 8 Gbps)

FC-Redirect is a fundamental technology needed for Cisco SME. Hence, the following considerations must be addressed:

- The target devices must be connected to a FC-Redirect-capable switch running Cisco MDS SAN-OS Software Release 3.3(1c) or later.
- FC-Redirect is limited to 32 target devices per switch if the devices are connected to generation-1 switching modules or if the switch contains Inter-Switch Links (ISLs) from a generation-1 switching module. To avoid any such limitation, you should use generation-2 (or later) switching modules for Cisco SME and ISLs targets.
- FC-Redirect v2 supports up to 128 hosts per target.
- Advanced zoning capabilities such as quality of service (QoS), LUN zoning, and read-only LUNs must not be used for FC-Redirect hosts and targets.
- Cisco Fabric Services must not be disabled on any of the required switches for FC-Redirect.
- Scalability of FC-Redirect depends on the number of switches in the fabric since it uses Cisco Fabric Services to synchronize with other switches in the fabric. You should limit the number of switches to 34. If more than 34 switches are needed, you should use Cisco Fabric Services regions. Cisco Fabric Services regions provide a way to segregate the switch to limit the scope of the synchronization domain.
- FC-Redirect and Fibre Channel over Ethernet (FCoE) cannot coexist in the same fabric as of today. Therefore, Cisco SME cannot be used to encrypt FCoE flows.
- Cisco SME must not be used in conjunction with SAN device virtualization (SDV) or Cisco DMM.

Requirements Specific to Cisco SME Disk

- FC-Redirect v2 must be enabled on all fabric switches.
- The Cisco MDS 9000 Family switch with a Cisco SME line card must be running Cisco 9000 NX-OS Software Release 5.2(1) or later.
- Servers and disk or tape devices using Cisco SME cannot be part of an inter-VSAN routing (IVR) zone set. Cisco SME Disk is not supported with IVR.

Requirements Specific to Cisco SME Tape

- Either FC-Redirect v1 or v2 can be configured.
- The target devices must be connected to a Cisco MDS 9500 Series or 9222i switch running Cisco MDS 9000 NX-OS Software Release 4.2(1) or later.
- Cisco SME Tape is supported with IVR by Cisco MDS 9000 NX-OS Software Release 5.0(1a) or later.

Configuration Requirements

- On a Cisco SME interface, either Cisco SME Disk or Cisco SME Tape can be configured. Cisco SME Disk and Cisco SME Tape cannot both be configured on the same encryption interface.
- On a Cisco SME cluster, either Cisco SME Disk or Cisco SME Tape can be configured. Cisco SME Disk and Cisco SME Tape cannot both be configured on the same Cisco SME cluster.
- Cisco SME cannot coexist with any other IP services such as Small Computer System Interface over IP (iSCSI) or FCIP.
- Host and target devices using Cisco SME cannot be part of an IVR zone set.
- FCIP write acceleration and FCIP tape acceleration must not be configured on the Cisco SME data flow (that is, Cisco SME traffic between the host and the target must not pass through FCIP tunnels with FCIP write acceleration or FCIP tape acceleration enabled).

- IP Security (IPsec) is not supported on modules running Cisco SME.
- FCoE modules are not supported on the fabric running Cisco SME.

Requirements for Replication and Snapshot Technologies

This section applies to the Cisco SME Disk feature only.

Cisco SME Disk uses the disk LUN World Wide Name (WWN) to identify the encryption status and encryption key on the key manager.

Certain disk mirror and snapshot technologies, such as EMC RecoverPoint I/O journal snapshots and copy-on-write snapshots, lack a LUN WWN, and for this reason they are not supported in Cisco SME rekey and data preparation operations.

Best practices recommend destroying copy-on-write snapshots mapped to encrypted LUNs after a rekey or data preparation operation completes successfully.

Cisco SME Data Flow

This section describes the handling of I/O traffic for Cisco SME.

Write I/O Data Flow

Figure 7 shows the write I/O data flow.

Figure 7. Write I/O

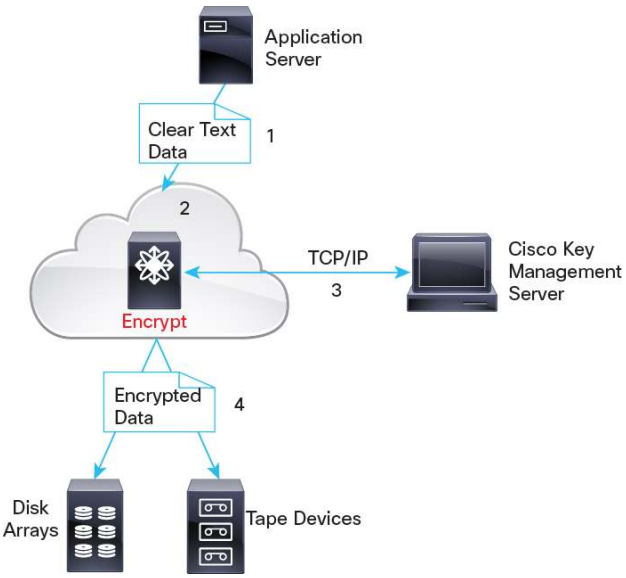

1. The application server sends a SCSI write I/O command.
2. The infrastructure uses FC-Redirect to redirect the I/O based on the initiator and target pair and its associated cryptographic node in the fabric.
3. The cryptographic node communicates with the Cisco KMC (if required) to identify the key to be used to encrypt the data and then processes it.
4. The encrypted data is then written to the disk array or tape device.

Read I/O Data Flow

Figure 8 shows the read I/O data flow.

Figure 8. Read I/O

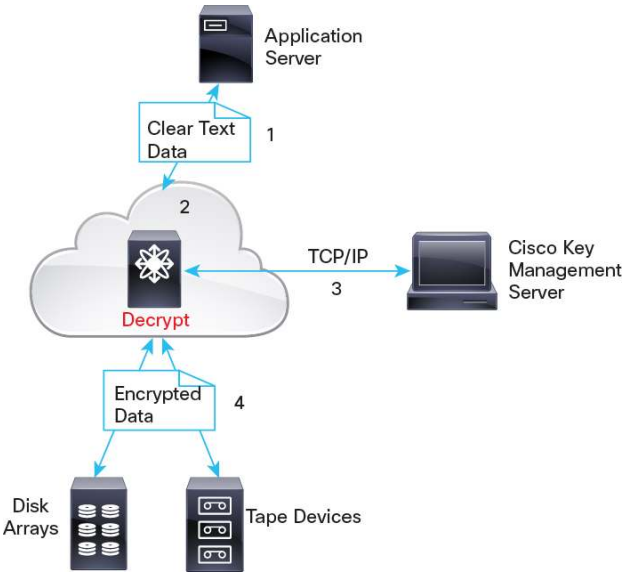

1. The application server sends a SCSI read I/O command.
2. The network infrastructure uses FC-Redirect to redirect the I/O based on the initiator and target pair and its associated cryptographic node in the fabric.
3. The cryptographic node communicates with the Cisco KMC (if required) to identify the key to be used to decrypt the data.
4. The cryptographic node sends SCSI read command to the disk array or tape library, decrypts the data, and returns the decrypted data to the server.

Caution: All hosts that access the same encrypted LUN must be provisioned for encryption to avoid data corruption. If any host writes clear text to an encrypted LUN, the data on the LUN will be lost.

Note: Cisco SME Tape can optionally compress the data.

Single Cisco SME Switch

Figure 9 shows a single-fabric topology with a Cisco SME line card on one switch. In this case, the data from server H2 (I/O marked in blue) is compressed (tape only) and encrypted by Cisco SME. Cisco SME does not process data from server H1 (I/O marked in black).

Figure 9. Cisco SME Data Flow: Single Cisco SME Switch

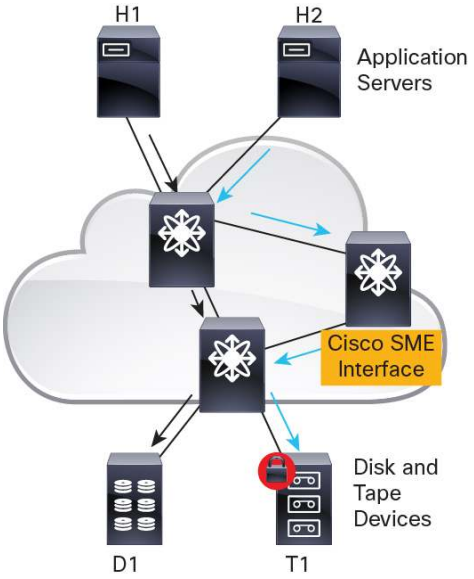

Cisco employs an FC-Redirect scheme that automatically redirects the traffic flow for the desired initiator-target pair to an appropriate Cisco SME line card in the fabric. There is no appliance inline in the data path, and there is no SAN rewiring or reconfiguration. Encryption and compression services are transparent to the hosts and storage devices. These services are available for devices in any VSAN in a physical fabric and can be used without rezoning.

Cisco SME Solution Architecture

Figure 10 shows the components of the Cisco SME solution.

Figure 10. Cisco SME Solution Components

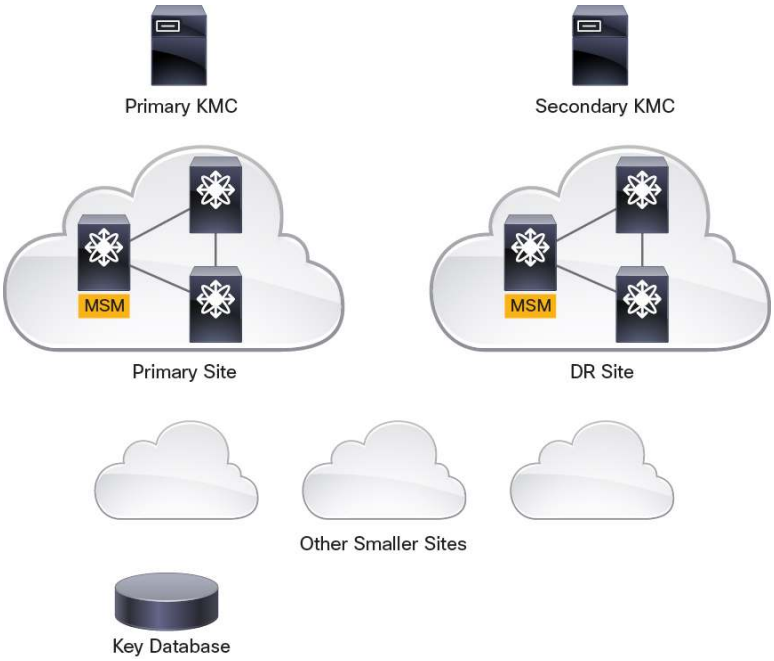

The main components of Cisco SME are:

- Primary Cisco KMC: This is the primary (or active) Cisco KMC for the Cisco SME solution. All key changes and updates to the key management database are handled by this Cisco KMC.
- Secondary Cisco KMC: This is the secondary (or passive) Cisco KMC. It is kept synchronized with the primary Cisco KMC by exchanging all updates to maintain a 1+1 configuration for high availability.
- Primary site: This is the primary site servicing all data requests. A best practice is to configure the primary Cisco KMC at this site.
- Disaster recovery site: This site handles all the redundant data and can be used in the event of failures on the primary site. A best practice is to configure the secondary Cisco KMC at this site.
- Cisco DCNM: Each site can have a separate Cisco DCNM server for regular SAN provisioning and monitoring. Cisco SME can be provisioned using a Cisco DCNM instance solely for Cisco KMC (or a pair of them for a high-availability solution). A separate Cisco DCNM instance is then used to manage the fabric and the Cisco SME solution. This configuration is more robust and allows better access control for security purposes.
- Key database: This database stores all the keys that Cisco SME is using.

Cisco SME Security Architecture

This section discusses how Cisco SME helps ensure security in the communication of each of its components.

Securing Communications Among Components

Cisco SME provisioning is accomplished using a web client to the Cisco DCNM Server. All data is exchanged using HTTPS. Further, the Cisco DCNM Server connects to the master switch of the Cisco SME cluster using the Secure Shell (SSH) Protocol (Figure 11).

Figure 11. Cisco SME Provisioning

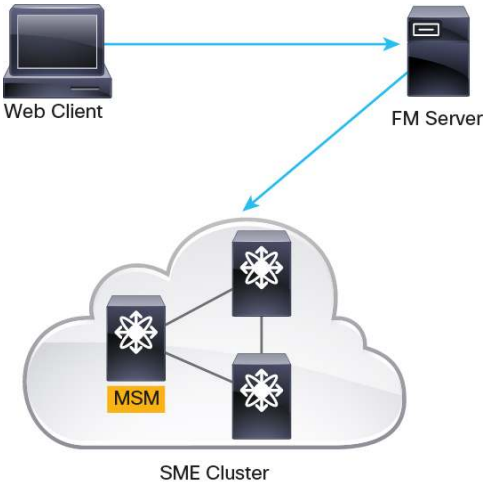

Interswitch communication within the Cisco SME cluster is protected using SSH (Figure 12).

Figure 12. Inter-Switch Communication

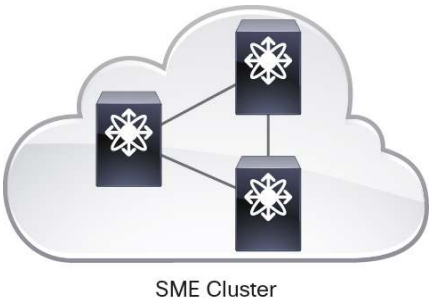

Communication between the Cisco SME switches the Cisco KMC is protected by SSL (Figure 13). For more information about SSL configuration, please refer to the Cisco SME Configuration Guide.

Figure 13. Cisco KMC Communication

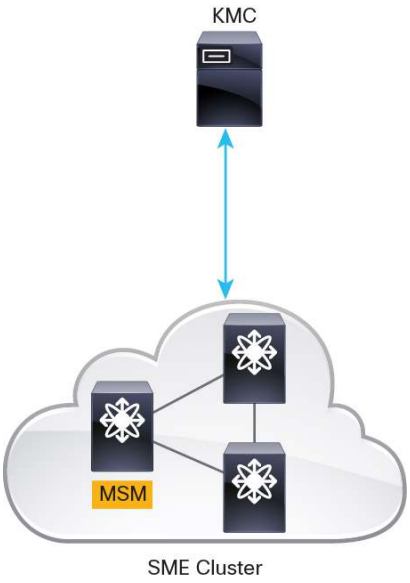

Securing Cisco SME Keys

Cisco SME Disk uses a two-level key hierarchy, and Cisco SME Tape uses a three-level key hierarchy. Cisco SME Disk uses a master key and volume keys; Cisco SME Tape uses a master key, volume group keys, and volume keys:

- Master key: This encryption key is generated when a Cisco SME cluster is created. Each cluster has a unique master key, and this key is shared across all members of the cluster. The master key is used to wrap the tape volume group keys or the disk keys.
- Volume group key (only for Cisco SME Tape): This encryption key is used to encrypt and authenticate the tape volume keys: the keys that encrypt all tapes belonging to the same tape volume group. A tape volume group can be created on the basis of a bar code range for a set of backup tapes, or it can be associated with a specific backup application.
- Volume key: This key is used to encrypt and authenticate the data on the tapes or on the disk.
  - Cisco SME Disk: This key is unique for each LUN, and it is stored in the Cisco KMC in encrypted format.
  - Cisco SME Tape: In unique key mode, the tape volume keys are unique for each physical tape. In shared key mode, one tape volume key is used to encrypt all volumes in a volume group. This key can be stored in the Cisco KMC or on the physical tape itself. The Cisco KMC database does not need to store the tape volume key if the key is on the tape, dramatically reducing the number of keys to be stored in the Cisco KMC. If the user knows the tape volume group key and if the tape volume key is on the tape, the tape can be decrypted; as a consequence, this option poses some limits to virtual shredding of an individual tape (Figure 14).

Figure 14. Cisco SME Disk Key Hierarchy

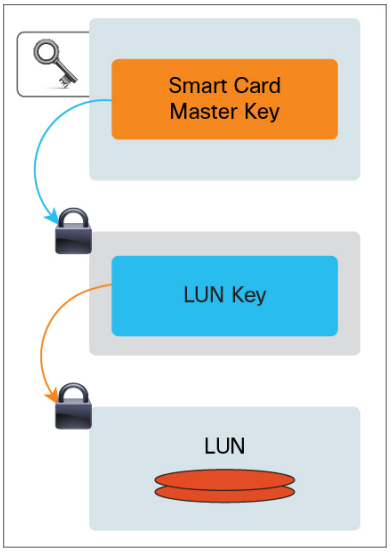

To recover encrypted data at rest, you need access to the keys that are created for the specific disk array or tape cartridge. Because the master key is used to protect all other keys, Cisco SME provides three master key security modes to protect the master key: Basic, Standard, and Advanced. During cluster configuration, you designate the level of security for the master key. Basic security writes the encrypted master key to a disk. To unlock the master key, you need access to the file. The file is encrypted and requires a password to retrieve the master key. The Standard and Advanced security modes require the use of smart cards to access the master key. For more information about Cisco SME keys, refer to the Cisco Storage Encryption Media Configuration Guide.

Cisco SME Administrative Roles

The following user roles are supported by Cisco SME:

- Cisco SME storage administrator: Responsible for provisioning Cisco SME, including Cisco SME cluster creation and Cisco SME Tape device configuration
- Cisco SME recovery officer: Responsible for reconstructing the master key in disaster scenarios; holds the smart cards
- Cisco SME key management administrator: Responsible for key management operations in the Cisco KMC
- Network administrator: Super-user role for the Cisco MDS 9000 Family that can perform any function, including all the functions of the other roles

For more information about configuring these user roles, please refer to the Cisco Storage Encryption Media Configuration Guide.

Cisco SME Clustering

Cluster technology provides scalability, reliability, and availability; automated load balancing; failover capabilities; and a single point of management.

A Cisco SME cluster consists of all switches enabled for Cisco SME in a dual-fabric (Figure 15). A Cisco SME cluster can consist of up to eight switches enabled for Cisco SME. Scalability can be easily achieved by adding more Cisco SME line cards to switches enabled for Cisco SME. Additionally, the load can be rebalanced to take advantage of the new hardware. Each switch enabled for Cisco SME can have multiple Cisco SME line cards. Each switch can be part of a maximum of two clusters (however, each Cisco SME interface can be part of only one cluster).

With multiple Cisco SME line cards in a Cisco SME cluster, the traffic is automatically load balanced across these modules. If a Cisco SME cryptographic node or a Cisco MDS 9000 Family switch fails, the traffic automatically fails over to another Cisco SME cryptographic node in the cluster within the same fabric.

The entire Cisco SME cluster can be managed through a single point using Cisco DCNM.

Figure 15. Cisco SME Clustering

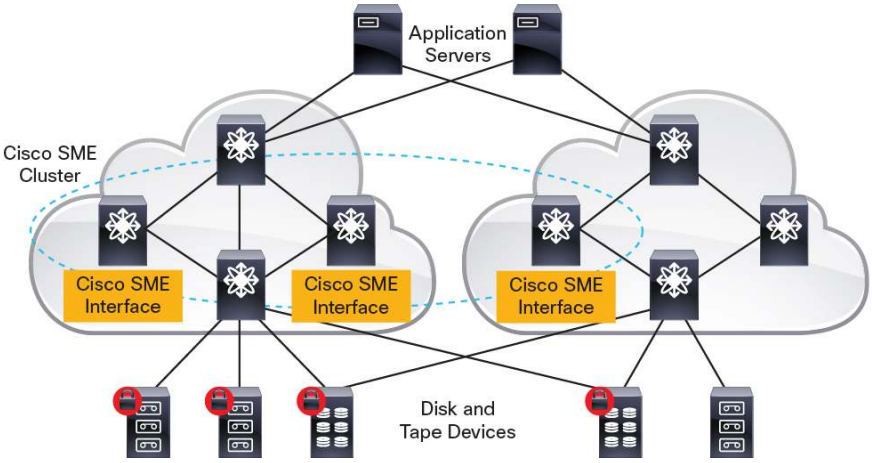

The Cisco SME cluster infrastructure uses the management interface to communicate with other switches in the cluster. A cluster view is defined as the set of switches that are part of the operation cluster. Only switches that are part of a cluster view participate in Cisco SME operations. A cluster requires a quorum of switches to be present. Refer to the Cisco Storage Encryption Media Configuration Guide for details.

Cisco SME Data Flow in a Cluster

After a Cisco SME cluster has been created and provisioned, the data is forwarded from the host to a Cisco SME module in the same fabric using FC-Redirect. The data is compressed (tape only) and encrypted and then sent to the target device. When the data is read, it follows the reverse path. Only the traffic from configured initiator and target pairs is redirected to a Cisco SME module. All other traffic is unaffected.

Each initiator and target pair is bound to a specific Cisco SME interface. When multiple Cisco SME modules are present in a Cisco SME cluster, Cisco SME uses target-based load balancing. All initiator and target pairs for a given target are always bound to the same Cisco SME interface. Initiator and target pairs for different targets are load balanced in the fabric across all available Cisco SME modules in the cluster. These Cisco SME modules can

be on any switch in the cluster that supports Cisco SME (multiple Cisco SME line cards on one switch are allowed).

In Figure 16, encryption traffic to target T1 (from hosts H1 and H2) flows through the Cisco SME module on switch SW1, and the encryption traffic to target T2 (from host H2) flows through the Cisco SME module on switch SW2, all in fabric A. Traffic from host H2 to target T4 flows through switch SW3 in fabric B. Nonencrypted data flow from host H1 to target T3 does not go through the Cisco SME modules.

Figure 16. Cisco SME Data Flow in a Cluster

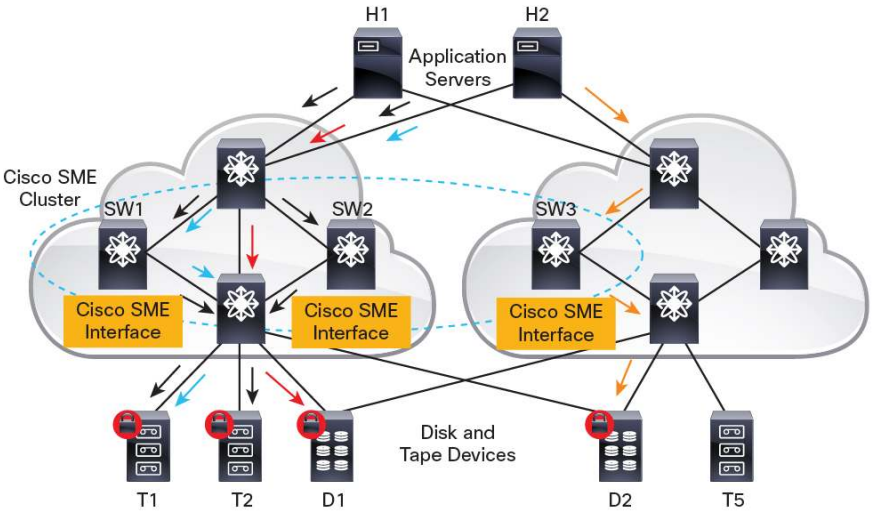

Failure Conditions

If the Cisco SME interface on switch SW2 fails (or if the entire switch SW2 fails), the traffic flow for initiator and target pairs bound to the corresponding Cisco SME interfaces will be briefly interrupted (for example, traffic from host H2 to target T2) until the affected initiator and target pairs are reassigned to other available Cisco SME interfaces in the cluster (Figure 17). The failure can cause some backup applications to stop backup jobs, and these may have to be restarted.

Note: Only Cisco SME interfaces within the same fabric are used for failover. For high availability, you should provision multiple Cisco SME interfaces in each fabric.

Figure 17. Failure Conditions

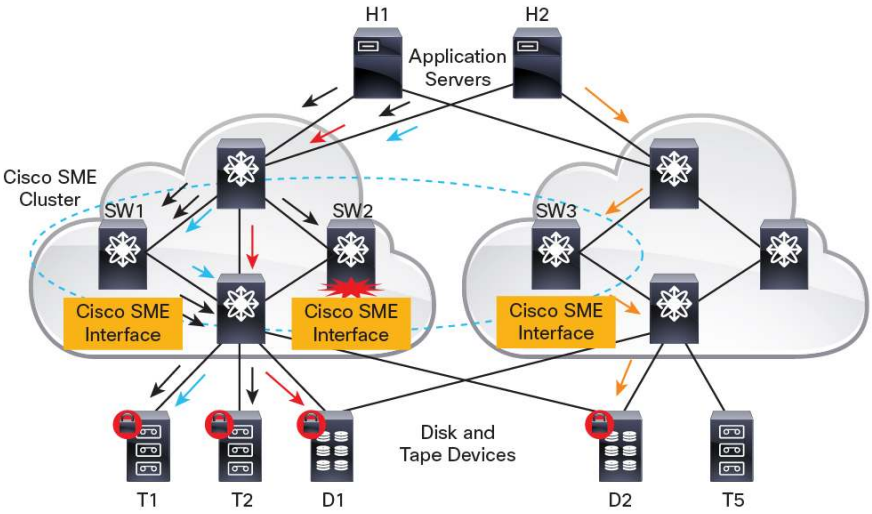

Cisco SME Design Considerations

This section discusses the guidelines you should follow when configuring Cisco SME.

Host-Based Cluster Considerations (Cisco SME Disk Only)

The quorum disk contains the host cluster state and timestamp information. It also allows continuous cluster operation without user intervention by configuring arbitrary heuristics that give each member of the cluster flexibility in handling a network-partition failure or in processing when a majority of the cluster members fails. This meta-information is required for the host cluster to function, and the quorum disk **must not** be part of the Cisco SME configuration; otherwise, a failure in the Cisco SME cluster would take down the host cluster.

Cisco SME Disk allows the use of host multipathing software such as Microsoft Multipath I/O (MPIO) and EMC PowerPath. Cisco SME Disk automatically discovers all the paths to a disk based on the unique disk WWN. In the event of incorrect configuration, SME Disk notifies the user about missed paths to prevent encrypted and clear-text data from being written on the same disk (resulting in data corruption).

Cisco SME Clustering: High-Availability Considerations

Cisco SME cluster operations require successful communication among the switches in the Cisco SME cluster using the management interface. Each switch in the cluster sends a heartbeat in 5-second intervals. Failure of this communication channel can affect the cluster membership of that switch. If a member switch loses communication with other members for more than three intervals (15 seconds) and so is no longer part of the cluster view, Cisco SME service is stopped on that switch. All the traffic fails over to other switches in the cluster.

Additionally, each switch in the cluster must have IP connectivity to Cisco KMC. In addition, for stable operations, you should use the same software revision level on all switches.

For a Cisco SME cluster to be operational, it must include more than half the number of configured switches in the cluster view. Thus, in an N-node cluster, N/2 + 1 nodes form a cluster quorum. If N is even, the cluster quorum requires N/2 nodes and also the presence of the switch with the lowest node ID. The quorum logic helps ensure that in the event of cluster partitions, at most one partition can be operational. All other switches are nonoperational. This behavior helps ensure the consistency of the cluster.

If a Cisco SME cluster is configured with two switches, a quorum requires the presence of the switch with the lowest node ID (usually the master switch). If this switch fails, the entire Cisco SME cluster becomes nonoperational due to lack of a quorum. With this scenario, all initiator and target pairs would enter the failure state, and no I/O would be possible. To avoid such a scenario, you should create a Cisco SME cluster with at least three switches. For more information about Cisco SME clustering and about quorum failures in two-, three-, and four-switch clusters, refer to the cluster quorum overview in the Cisco Storage Encryption Media Configuration Guide.

By default, all Cisco SME clusters are created to support tape. If disk support is required, you need to explicitly configure it:

1. From the command-line interface (CLI), enable the cluster-capability disk command before adding the first Cisco SME interface.
2. From Cisco DCNM, select the Include Disk Support check box during cluster creation.

Network Topologies

Cisco SME is fully supported in fabrics that consist of only Cisco MDS 9000 Family products. Cisco SME may be supported in some environments consisting of both Cisco and other vendors' switches; such configurations must be evaluated on a case-by-case basis.

Cisco SME supports single-fabric, dual-fabric, FCIP, and remote replication topologies. In a dual-fabric SAN, one or more (up to eight) switches that support Cisco SME form a cluster. Dual-fabric topologies support single-path and multipath configurations.

In a single-path configuration, a cluster includes only one path represented as an initiator-target path. In a multipath configuration, a cluster includes all paths, which are represented as multiple initiator-target paths.

Topology Considerations

This section presents guidelines and topologies for Cisco SME. The appendix presents additional examples for dedicated disk and tape SANs.

General Guidelines

- A Cisco SME Disk cluster is not compatible with IP compression or the combined configuration of IPsec and write acceleration.
- The default zone policy should be set to Deny.
- Use FC-Redirect v2 across switches for Cisco SME Disk clusters.
- For disk replication, use the same Cisco KMC for source and destination disks managed under the disk replication context.

Guidelines Specific to Multiple Clusters

- Do not use the same target LUN in two different Cisco SME Disk clusters.
- Do not use the same target ports in different clusters.
- Do not add the same Cisco SME interfaces in two different Cisco SME Disk clusters.
- Use different clusters for Cisco SME Disk and Cisco SME Tape.

Core-Edge Topology

In a core-edge topology, application or backup servers are at the edge of the network, and disk or tape libraries are at the core.

If the target devices that require Cisco SME services are connected to only one switch in each fabric in the core (Figure 18), use Cisco SME line cards and provision Cisco SME on that switch in each fabric. The number of Cisco SME line cards depends on the throughput requirements (see “Sizing Guidelines” later in this document).

Figure 18. Core-Edge Topology: Target Devices on a Single-Core Switch in Each Fabric

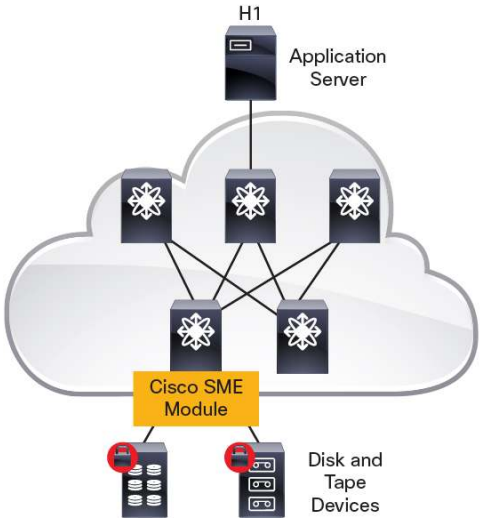

If the target devices that require Cisco SME services are connected to multiple core switches (Figure 19), connect Cisco SME line cards and provision Cisco SME on all these switches. On the basis of the throughput requirements, derive the total number of Cisco SME line cards and spread them (in proportion to the expected traffic) across the switches to which the target devices are connected. Whenever Cisco SME interfaces are available on the switch connected to a target device, those Cisco SME interfaces are used for the encryption service, thus eliminating unnecessary traffic on ISLs. To handle failure of a Cisco SME interface, each Cisco SME switch should have more than one Cisco SME interface configured (see “Sizing Guidelines” later in this document for details).

Figure 19. Core-Edge Topology: Targets on Multiple-Core Switches in Each Fabric

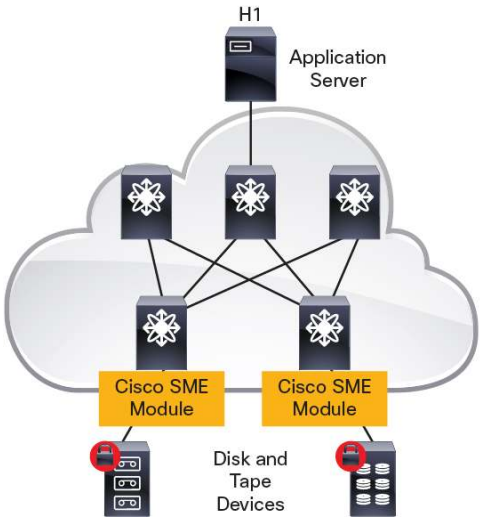

Note: If the Cisco SME line card is on a different switch than the disk or tape library, additional traffic may cross the ISL.

Edge-Core-Edge Topology

In an edge-core-edge topology, the host and target devices are at the two edges of the network connected through core switches.

If the target devices that require Cisco SME services are connected to only one switch on the edge in each fabric (Figure 20), use Cisco SME line cards and provision Cisco SME on that switch in each fabric. The number of Cisco SME line cards depends on the throughput requirements (see “Sizing Guidelines” later in this document).

Figure 20. Edge-Core-Edge Topology: Targets on a Single Edge Switch in Each Fabric

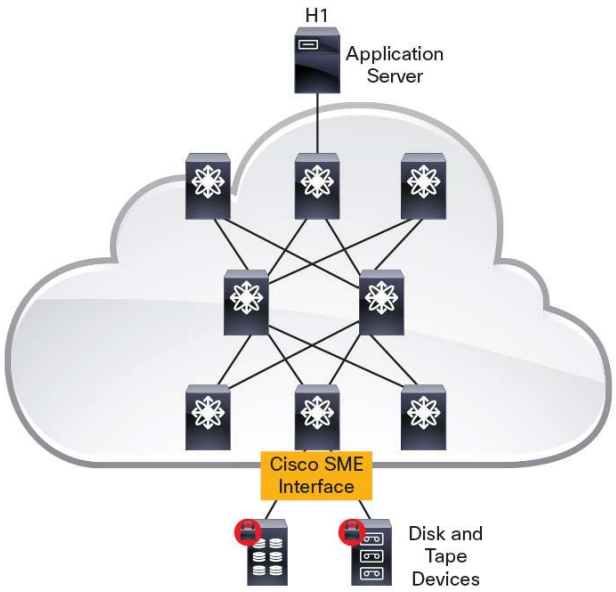

If the target devices that require Cisco SME services are connected to multiple edge switches (Figure 21), connect Cisco SME line cards and provision Cisco SME on the core switches. On the basis of the throughput requirements, derive the total number of Cisco SME line cards and spread them evenly across the core switches. Any additional Cisco SME interfaces (for example, for additional throughput requirements) must go to the core switches. In this scenario, in which Cisco SME interfaces are provisioned on switches not connected to the target devices, do not provision any more Cisco SME interfaces on the switches connected to the target device.

Figure 21. Edge-Core-Edge Topology: Targets on Multiple Edge Switches in Each Fabric

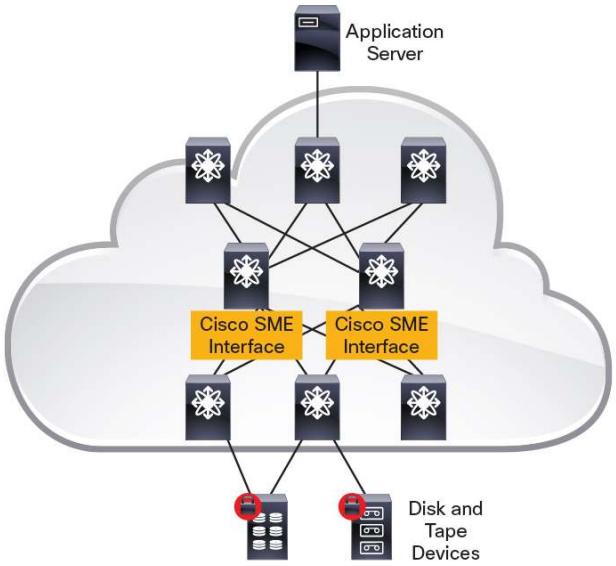

FCIP Topology

The Cisco SME Disk solution can be deployed in an FCIP topology. In a FCIP topology, the host and target devices are located across data centers and connected through an FCIP link.

If the target devices that require Cisco SME services are connected to only the remote fabric (Figure 22), use Cisco SME line cards and provision Cisco SME on that fabric. The number of Cisco SME line cards depends on the throughput requirements (see “Sizing Guidelines” later in this document).

Figure 22. FCIP Topology

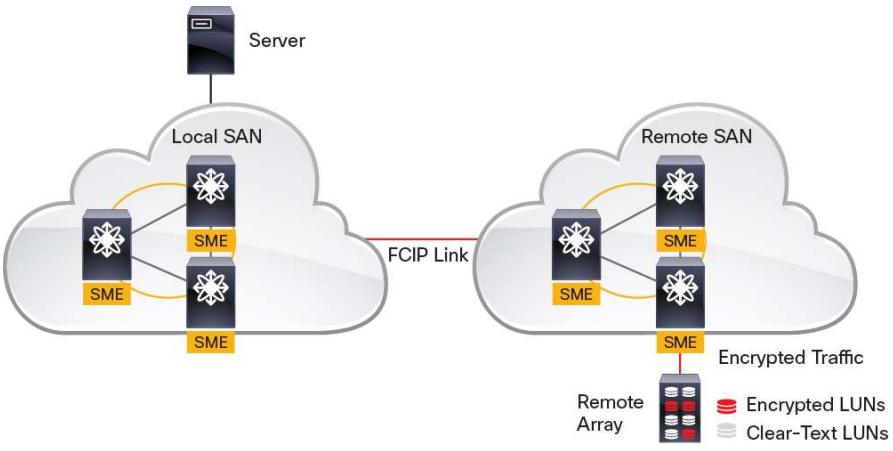

If the target devices that require Cisco SME services are connected to both the local and the remote fabric (Figure 19), connect Cisco SME line cards to provision Cisco SME on both the fabrics. Also deploy separate clusters for the local target devices and the remote FCIP target devices to help prevent unnecessary traffic from traversing the FCIP link.

In the scenario in which Cisco SME line cards were provisioned on the local fabric and the remote FCIP target device needs to be encrypted, disable the FCIP compression feature on the Cisco MDS 9000 Family switches that provide FCIP connectivity because the traffic passing through the FCIP link is encrypted and hence not highly compressible.

Cisco SME Load-Balancing Considerations

Cisco SME employs target-based load balancing, in which all the hosts for a specific target device use the same Cisco SME interface for encryption. Additionally, to reduce ISL use, if one or more Cisco SME interfaces are available on the switch to which the target device is connected, only these local Cisco SME interfaces are used for encryption for that target device. The placement of Cisco SME interfaces in the SAN must take this behavior into consideration.

- For core-edge topologies, in which the target are in the core, the Cisco SME interfaces should be provisioned in the target switches. For such a scenario, the number of Cisco SME interfaces on each switch should be proportional to the throughput requirement of the disk or tape drives connected to it.
- For edge-core-edge topologies, in which the targets are on a single edge switch, Cisco SME interfaces should be provisioned on that edge switch.
- For edge-core-edge topologies, in which the targets are on multiple edge switches, Cisco SME interfaces should be provisioned in the core switches.
- Cisco SME interfaces should not be provisioned so that some interfaces are on switches that have targets connected and some are on switches that do not have targets connected.

For more information about Cisco SME load balancing, refer to the Cisco SME Configuration Guide.

Cisco KMC Considerations

This section discusses design considerations specific to Cisco KMC.

Cisco KMC High Availability

Cisco KMC high availability can be achieved using primary and secondary Cisco KMC servers. Each Cisco SME cluster can be configured for high availability. The primary and secondary Cisco KMC servers stay in the active-passive state with the solution failing over to the (clusterwide) secondary server in the event of any failure on the primary server. The system automatically fails back to the primary server when it becomes available again. Cisco provides a script that can be run on the primary server with the data synchronized with the secondary server to maintain Cisco KMC high availability (Figure 23).

Figure 23. Cisco KMC 1+1 High Availability and Key Database High Availability for Multisite Deployment

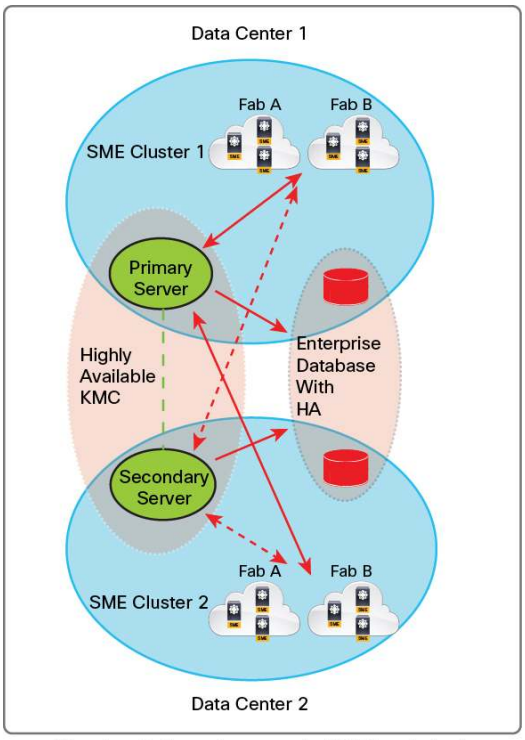

More than 2 data centers supported. KMC is only 1 + 1

Key Database High Availability

By default, Cisco KMC ships with the PostgreSQL database for key management, and the primary and secondary servers share this same database for their operations. However, in the event of a primary server failure, the secondary server will be unable to access the keys from this database and cannot resume normal operation.

As a best practice, you should deploy Oracle Data Guard as the key database instead, in which case the primary server and secondary server will each have its own copy of the key database. The Oracle Data Guard feature will help ensure that all changes to the primary key database are reflected in the secondary database, providing smoother failover when failover becomes necessary.

Cisco DCNM Guidelines

Cisco SME can be managed using Cisco DCNM like any other Cisco MDS 9000 Family feature. Cisco DCNM provides the Cisco SME wizard, which simplifies the provisioning of Cisco SME throughout the fabric, reducing the process to just a few steps, while providing a fabric-level view of the Cisco SME configuration. In addition, the CLI provides comprehensive Cisco SME troubleshooting commands.

In small, dedicated SAN environments of fewer than 5000 ports that use Cisco DCNM for overall SAN management, Cisco DCNM can also be used as the key management server (Figure 24).

Figure 24. Cisco DCNM Deployment for Small SAN Environments

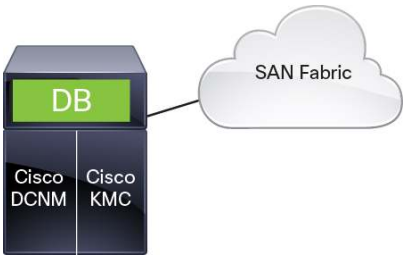

For large SAN deployments of greater than 5000 ports and for environments in which the Cisco DCNM performance manager is being used, a separate server should be used for Cisco KMC management (Figure 25).

Figure 25. Cisco DCNM Deployment for Large SAN Environments

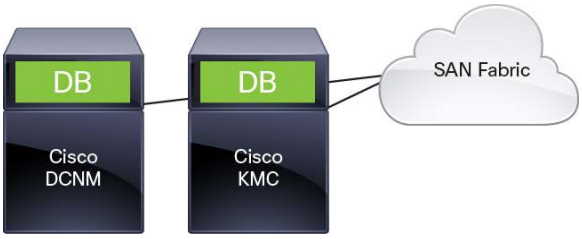

When deploying Cisco SME between two data centers, a best practice is to deploy Cisco KMC with an Oracle database with Data Guard to provide redundant key databases at each data center.

Cisco SME Solution Scalability Limits

Table 3 summarizes Cisco SME Tape scalability limits, and Table 4 summarizes Cisco SME Disk scalability limits.

Table 3. Cisco SME Tape Configuration Limits

Configuration	Limit
Clusters per switch	1
Switches in a cluster	4
Switches in the fabric	10
Fabrics in a cluster	2
Modules in a switch	11
Cisco MDS 9000 18/4-Port MSM modules in a cluster	32
Initiator target LUNs (ITLs)	1024
LUNs behind a target	32
Host and target ports in a cluster	128
Hosts per target	128
Tape backup groups per cluster	4
Volume groups in a tape backup group	32
Keys in a tape volume group	8000
Disk groups	128
Cisco SME disks LUNs	2000
Cisco KMC keys	32,000

Configuration	Limit
Targets per switch that can be redirected with FC-Redirect	32
IT connections per Cisco SME interface (soft limit)	256 **Note:** Beyond this limit, a syslog message will be displayed. A best practice is to provision more Cisco SME interfaces in the cluster.[1]
IT connections per Cisco SME interface (hard limit)	512 **Note:** Beyond this limit, new IT connections will not be assigned to that particular Cisco SME interface, and a critical syslog will be displayed.[2]

[1] Applicable to Cisco MDS 9000 NX-OS Release 4.2(1) and later
[2] Applicable to Cisco MDS 9000 NX-OS Release 4.2(1) and later

Table 4. Cisco SME Disk Configuration Limits

Configuration	Per Cluster	Per Switch	Per Cryptographic Node
Clusters	-	2	1
Physical fabrics	2	-	-
Switches	8	-	-
Modules (line cards: Cisco MDS 9000 16-Port SSN or 18/4-Port MSM modules)	-	11	-
Cisco SME interfaces (cryptographic nodes used for encryption)	32	32	-
ITLs	2048	2048	512
LUNs behind a target	512	512	512
Initiator ports	128	-	-
Target ports	128	-	-
Maximum number of initiator and target pairs	128	-	-
Paths per LUN (physical paths per Cisco SME disk)	8	8	8
Disk groups	128	128	128
Number of Cisco SME disks (LUNs)	2048	2048	512
Cisco KMC keys	32,000	32,000	32,000
Maximum number of concurrent data preparations (offline data preparations)	-	-	64
Total number of disk key replication relationships	2048		

Encryption Throughput Guidelines

Cisco SME Tape allows each cryptographic interface to be configured for both compression and encryption of data. When configured in this way, each interface supports up to 500 MBps of throughput.

Cisco SME Disk supports only encryption of data and allows up to 1250 MBps of throughput per cryptographic interface.

The throughput increases linearly with the number of cryptographic interfaces. Therefore, the Cisco MDS 9000 16-Port SSN module (which incorporates four Cisco SME interfaces) can support up to 2000 MBps of throughput for Cisco SME Tape and up to 5000 MBps for Cisco SME Disk.

Disk Encryption Latency Guidelines (Cisco SME Disk Only)

Read Operations

Each Cisco SME interface, when configured for disk encryption, adds:

- 30 microseconds of latency to 512-byte block reads
- 45 microseconds of latency to 2048-byte block reads

Latency increases slightly with increased I/O operations.

Write Operations

Write operations have no effect on latency because Cisco SME proxies the transfer ready from the host while waiting for a response from the target, and when the target responds, the data is already encrypted and ready to be transmitted.

Sizing Guidelines

- You can size the Cisco SME environment in two ways:
  - Provision a sufficient number of Cisco SME line cards to get the best line-rate throughput for each disk array or tape drive.
  - Provision a sufficient number of Cisco SME line cards to get the aggregate throughput needed to for the backup window.
- The peak throughput to each disk or tape drive with Cisco SME is about the same as that without Cisco SME. The addition of Cisco SME has little or no negative effect on the throughput to each disk or tape drive.
- The peak throughput to each tape drive depends on the type of the drive, the server performance, the host bus adapter (HBA) speed (2, 4 or 8 Gbps), and the compressibility of the backup data. For Canterbury Corpus data, the observed compression ratio using Cisco SME is 4.7:1.
  - For example, for an LTO-3 tape drive, the peak throughput is the range 80 to 120 MBps, depending on other factors.
- You should use the actual observed per-drive throughput for sizing calculations.
- The sizing calculations must be performed independently for each fabric in a dual-fabric SAN.

Sizing Using Line-Rate Disk or Tape Drive Throughput

- Estimate the per-drive peak throughput in the existing SAN environment.
- Use the interface throughput per Cisco SME to derive the number of disk or tape drives that can be serviced by each Cisco SME interface.
  - Case 1: With an LTO-3 drive that achieves 80-MBps peak throughput (as observed in the existing SAN), five such drives can be serviced by each Cisco SME interface.
  - Case 2: With a disk array that achieves 3000-MBps peak throughput (as observed in the existing SAN), one such disk array can be serviced by three Cisco SME interfaces.
- Derive the total number of Cisco SME interfaces required to support all the disk arrays or tape drives in the SAN.

Note: Each Cisco MDS 9000 18/4-Port MSM or Cisco MDS 9222i supervisor module contain one Cisco SME interface; each Cisco MDS 9000 16-Port SSN contains four Cisco SME interfaces.

- To account for failures, add 20 percent or one Cisco SME interface per switch (whichever is greater) to the total number calculated above.
- Detailed examples are provided in the appendix.

Sizing Using Backup Window Calculations

- Estimate the aggregate backup throughput requirement to accommodate the backup window. Include projected growth in backup data.
  - For example, assume that a total of 90 TB needs to be backed up in a timeframe of 18 hours; this translates to 5 TB per hour, or about 1456 MBps.
- Derive the total number of Cisco SME interfaces required to meet this aggregate throughput.
  - In the preceding example, a total of three Cisco SME interfaces are required (1456/500).
- To account for failures, add 20 percent or one Cisco SME interface per switch (whichever is greater) to the total number calculated above.
  - In the preceding above, if only one Cisco SME switch is used, a total of four Cisco SME interfaces should be provisioned. If two Cisco SME switches are used, a total of five Cisco SME interfaces should be provisioned.
- Detailed examples are provided in the appendix.

Inserting Cisco SME Tape into Existing Cisco SANs

The Cisco SME Disk solution can be added to existing SAN fabrics in either of two ways:

- Upgrade switches connected to the target devices: Upgrade the Cisco MDS 9000 Family switches connected to the targets to Cisco MDS 9000 NX-OS Software Release 3.2(3) or later and add Cisco SME line cards to these switches. Also consider the configuration and zoning requirements specified in “Cisco SME Requirements” earlier in this document.
- Add new switches to the fabric and move the target devices: Add new Cisco MDS 9000 Family switches with Cisco SME capabilities (using Cisco SME line cards) to the SAN and move the target devices needing Cisco SME to the new switch. This switch must be running Cisco MDS 9000 NX-OS Software Release 3.2(3) or later.

In both these solutions, the host-connected or tape-connected switches should be upgraded to Cisco MDS 9000 NX-OS Software Release 3.2(3) or later.

Cisco SME Disk Deployment in Storage Environments

Replication is divided into two categories: local and remote.

- Mirrors or clones: When data for the source disk is being duplicated by the disk array to another disk in the same storage system, the destination disk is called a mirror or clone of the source disk. This process is local replication.
- Remote replication: If data for the source disk is being duplicated by the disk array to another disk in a remote storage system, then the source disk and the remote disk are in a replication relationship. Depending on the distance and bandwidth availability between the local and remote sites, remote replication is of one of two types:
  - Synchronous: The local disk array does not return a response to the write command on the local LUN until the data is also written to the remote LUN.

- Asynchronous: The local disk array does not immediately write the data to the remote LUN. Instead the changes to the local LUN are batched into a delta data set and periodically flushed to the remote LUN.

A snapshot is a point-in-time copy that can be created instantly for a source disk. If any write operations to the source disk occur after a snapshot is created, the previous data will be saved elsewhere before any modification occurs. This behavior allows the disk array to present a specific point-in-time copy of the data of the source disk.

Managing Replication with Cisco SME

Cisco SME supports replication using a feature called disk key replication (DKR). DKR simplifies the key management of the source and destination disks by automating the propagation of the source disk key to the destination disk. Cisco SME Disk clusters have two modes:

- Nonsignature cluster
- Signature cluster

Replication management is the same for both cluster modes and consists of following steps:

1. Extract the replication relationship using array vendor-specific technology. The output of this step identifies the source and destination disk relationship based on the SCSI properties of the vendor, product, and device identifiers.
2. Import the replication relationship information into Cisco SME through DKR using Cisco DCNM.

Caution: All Cisco SME configuration operations on the disks in a DKR relationship must be managed through DCNM only. The CLI must not be used and can cause unpredictable results.

3. Manage key change operations in Cisco DCNM for DKR.

Key change operations involve data preparation or no data preparation:

- No data preparation: Any local key changes will cause DKR to suspend host access to the remote disk. After the local key change is verified for data integrity and the data replication to the remote end is synchronized, the administrator can select the corresponding relationship and perform the synchronization operation for DKR. This operation will synchronize the source and destination keys and resume the host access to the remote disk.
- Data preparation: Before starting data preparation on the source disk, you must perform the following steps:
  - Disable the DKR relationship.
  - Disable replication between the source and destination disks. This operation varies with the disk array vendor.

  After data preparation is completed and verified for data integrity, perform the following two steps:
  - Enable data replication between the source and destination disks using an operation specific to your specific disk array vendor.
  - After data is synchronized between the source and destination disks, enable the DKR relationship. This operation will synchronize the source and destination keys.

Caution: Host access on the destination disk should be halted until the preceding steps are completed. Data corruption can occur otherwise.

Managing Snapshots in Cisco SME

Snapshot management is different for signature and nonsignature clusters.

- Nonsignature clusters: Snapshots are not supported for nonsignature clusters.
- **Signature clusters: DKR must not be used to** manage snapshots in signature clusters. Follow this process to manage **cryptographic** snapshots:
  - Trigger discovery in Cisco SME to configure the snapshot disks.
  - Cisco SME will put the disk in a failed state because it finds the valid Cisco SME metadata on the disk media with no corresponding active key in Cisco KMC.
  - The administrator will have the option of recovering the disk performing recovery from metadata. Select this option.
  - After recovery is performed, the snapshot will come up as a cryptographic disk and can be accessed by the host.

Replication and Mirroring

The Cisco SME Disk solution can be added to LUNs that are locally mirrored or remotely replicated.

Figure 26 shows a data mirroring and replication deployment. In this configuration, Server 1 knows about only the primary storage array and LUN 1 and the I/O path to the primary storage array and LUN 1. When data is sent to the primary storage array and LUN 1, it is locally mirrored to LUN 2 and also remotely replicated to LUN 3 on the remote storage array. When Cisco SME is added to this configuration, data flowing from Server 1 to the primary storage array is encrypted and stored encrypted on the source arrays in LUN 1, mirrored encrypted to LUN 2, and replicated encrypted on LUN 3 on the remote array.

Figure 26. Cisco SME Disk Deployment for Replication and Mirroring

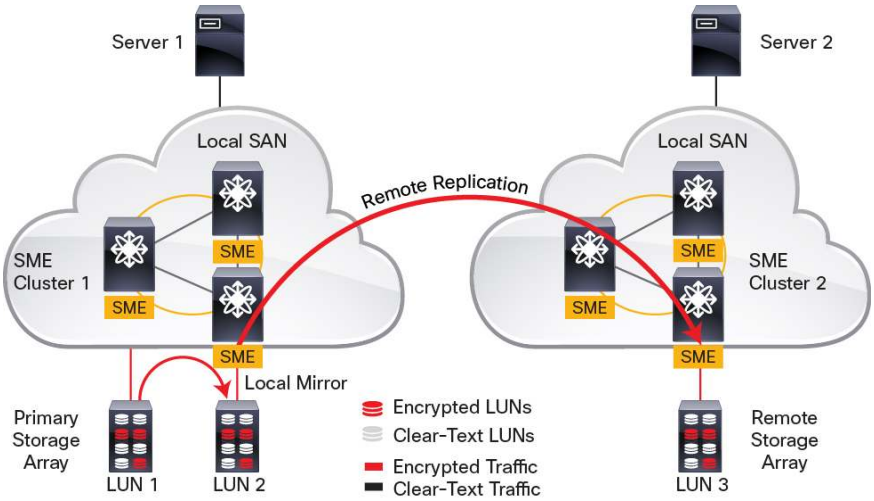

A DKR map file that describes the source, mirrored, and remote LUN relationship must be created and entered in Cisco KMC.

Cisco SME uses DKR information to properly associate local mirrored LUN 2 and remote replicated LUN 3 with the correct LUN key when retrieving data from LUN 2 and LUN 3.

Use the same Cisco KMC, capable of accessing both data centers, to provision Cisco SME in remote replication environments. Use the same Cisco DCNM or DCNM federation, capable of accessing both data centers, to provision Cisco SME in remote replication environments.

Copy-on-Write Snapshots

A snapshot creates a point-in-time copy of the data (Figure 27).

Figure 27. Cisco SME Disk Deployment for Copy-on-Write (COW) Snapshots

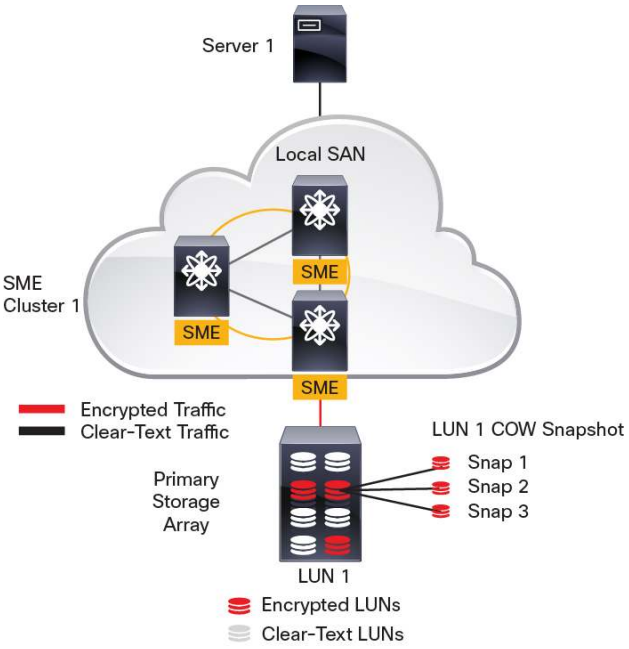

There are two main types of storage snapshot: the copy-on-write snapshot and the split-mirror snapshot.

In a copy-on-write snapshot, only the metadata about where original data is stored is copied. No physical copy of the data is made at the time the snapshot is created. The snapshot copy then tracks the changing blocks on the original volume as write operations to the original volume are performed.

In a split-mirror snapshot, the entire volume or LUN, not only the new or updated data, is copied.

When the Cisco SME Disk solution is inserted into storage topologies in scenarios in which copy-on-write snapshot is deployed, best practices recommend destroying copy-on-write snapshots after a rekey operation completes.

Data preparation and rekey operations change every single block of the LUN, transforming its associated copy-on-write snapshots into split mirrors and possibly consuming all the storage reserved space.

Moreover, for most implementations the device identifier (WWN) of the copy-on-write snapshots never changes, even if a snapshot is associated with a completely different point in time or LUN. This behavior preserves operative systems device handles.

Since the LUN WWN is critical for the Cisco SME Disk solution to identify the cryptographic state of the disk and to associate the appropriate LUN key, copy-on-write snapshots are not supported across data preparation and rekey operations.

EMC Recover Point IO Journal Snapshots

EMC Recover Point IO journal snapshots are bookmarks in the I/O journal, and as with copy-on-write snapshots, the LUN device identifier (WWN) cannot be used to identify the EMC Recover Point I/O journal snapshots (Figure 28). These journal snapshots are not SCSI entities and do not have an associated LUN WWN.

Figure 28. Cisco SME Disk Deployment for EMC Recover Point IO Journal Snapshots

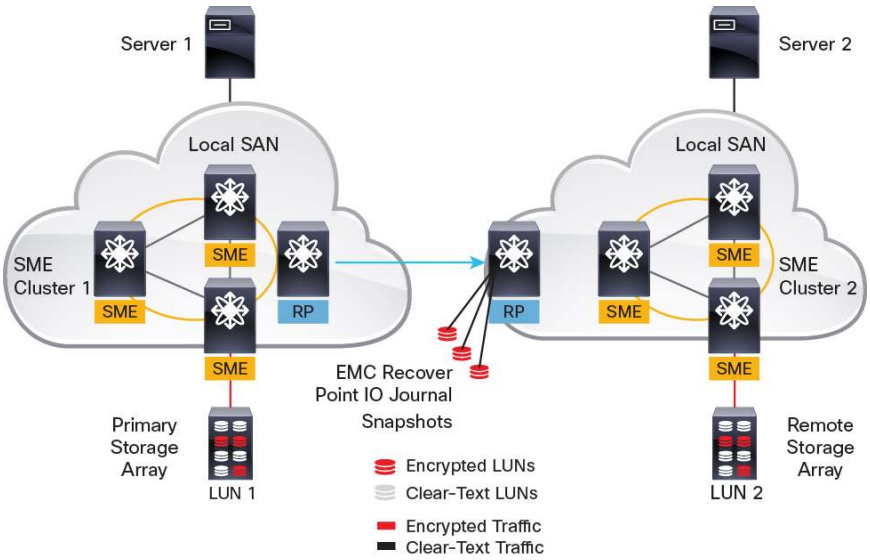

Since the LUN WWN is critical for the Cisco SME Disk solution to identify the cryptographic state of the disk and to associate the appropriate LUN key, EMC Recover Point IO journal snapshots are not supported across data preparation and rekey operations.

Mix of Disk and Tape Devices

The Cisco SME solution supports both disk arrays and tape libraries (Figure 29).

Figure 29. Cisco SME Disk Deployment for a Mix of Disk and Tape Devices

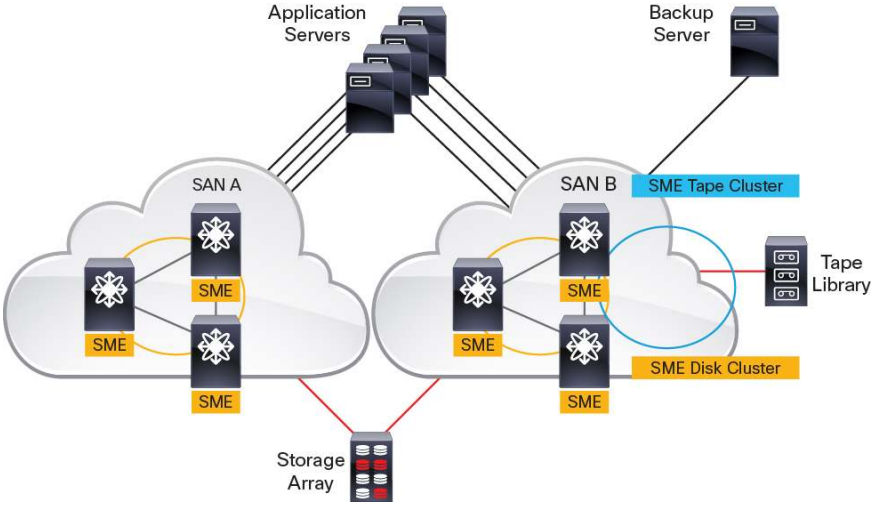

In such a deployment:

- Use separate clusters for tape groups and disk groups.
- Use separate encryption engines, or Cisco SME interfaces, to form tape clusters and disk clusters.

Inserting Cisco SME Disk into Existing Storage Environments

The Cisco SME Disk solution can be added to LUNs containing preexisting data or to LUNs without any data on them. In both cases, the LUN size will not be increased in any way.

LUNs Containing Preexisting Data

If data exists on the LUN undergoing encryption (Figure 30):

- Back up the LUN undergoing encryption.
- Enable encryption and enable the “Data Preparation” option.

Figure 30. LUNs with Preexisting Data

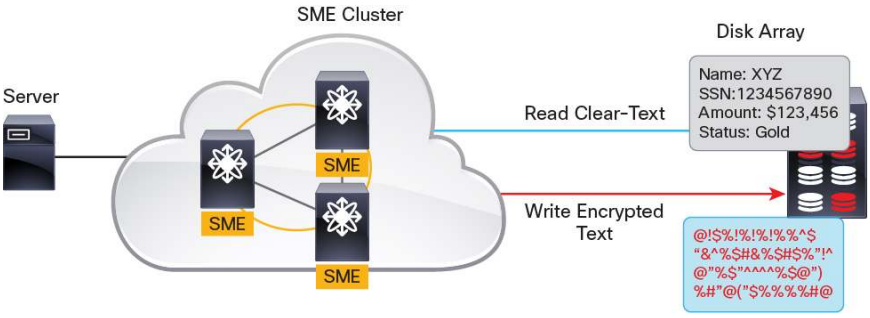

In a data preparation operation, existing clear-text data on a LUN is read, encrypted with the current volume key, and written back to the LUN at the same LBA.

Depending on the size of the LUNs, the data preparation process may take several minutes. During the entire data preparation process, make sure that host I/O operations are halted. Resume host I/O operations after the data preparation process is completed.

LUNs Not Containing Preexisting Data

If no data exists on the LUN undergoing encryption (Figure 31):

- Enable encryption without enabling the Data Preparation option.

The encryption process in this case is instantaneous because there is no data on the LUN to be encrypted.

Figure 31. LUNs Without Any Data

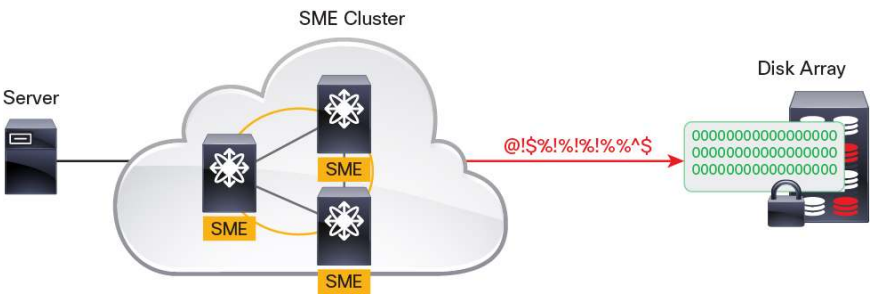

Note: If the data preparation operation fails, Cisco SME Disk will put the disk in a failed state. This disk will not be accessible to hosts. and all paths of the disk will be put in the I/O reject state. The disk can be recovered using Cisco SME. Please refer to the Cisco SME Configuration Guide for details.

Maintaining Data Security

Cisco SME provides two ways to achieve data security: rekey and master-key rekey.

Rekey (Cisco SME Disk Only)

In a rekey operation (Figure 32), previously encrypted data on a LUN is decrypted with the current volume key, reencrypted with a new volume key, and written back to the same LUN at the same LBA.

Figure 32. Rekey

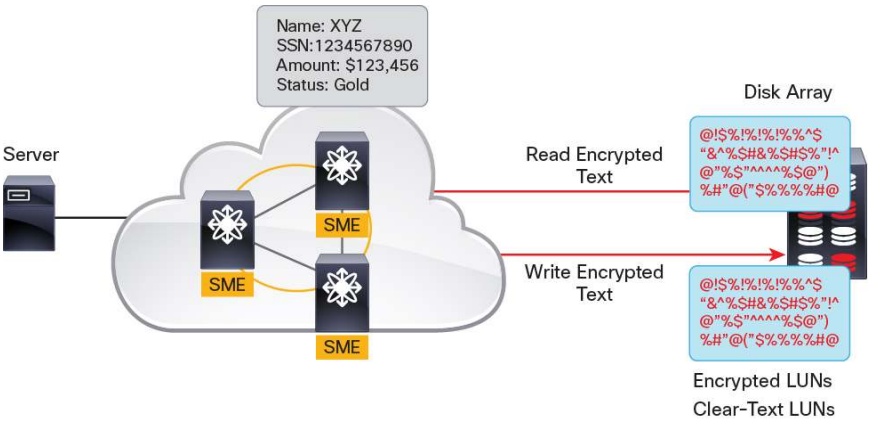

Master-Key Rekey

The master key encrypts all keys for the disk or tape, so it is the most important key that needs to be changed according to the organization's security policies. Cisco SME allows rekeying of the master key itself for additional security.

After a new master key is generated, depending on the security mode chosen, all recovery shares (of the master key, if any) need to be updated. The old master key is archived and can be used to rekey the disk keys. Tape keys, however, cannot be rekeyed because the administrator could not configure the system if they were stored on tape instead of in Cisco KMC (to reduce the size of the Cisco KMC), making it impossible to retrieve every tape key and process it.

Disk Key Replication

Cisco SME Disk supports the use of replication technology such as EMC Symmetrix Remote Data Facility (SRDF). The user can enter the DKR relationships through Cisco DCNM (using the DKR service), and Cisco SME Disk will automatically synchronize the encryption keys across the disks.

To make management easier, Cisco provides two scripts. The first script is device specific and has to be written to meet the requirements of the disk array vendor. This script reads the replication information using commands specific to the disk array vendor and outputs that information to a file that contains the SCSI properties of the source and destination disks.

The second script takes the intermediate file generated by the first script and generates a specially formatted DKR file that can be imported into Cisco KMC using Cisco DCNM. The second script is device independent and can be used as is. After the file is imported, Cisco SME starts synchronizing the encryption keys for these disks and manages them automatically for the life of the relationship.

Be sure that the following conditions are met:

- After a disk is added to a DKR relationship, all Cisco SME operations on that disk must be performed through Cisco DCNM. The CLI must not be used for those disks. Use of the CLI will result in unpredictable results and can put data on the disk at risk.
- Use the same Cisco KMC for source and destination disks that are managed in the disk replication context.
- Disk replication takes care of key replication only and not the data replication, which the storage vendor handles. Follow the proper steps while synchronizing the keys.
- For proper key synchronization, replication switchover requires relationships to be removed and created according to the new data flow of the replication.

Appendix

Note: The following is not a complete checklist, and you are advised to review the design guide thoroughly.

Cisco SME Tape Checklist

- Any tape that was used previously and has some data and that has not been encrypted using Cisco SME should not be used for Cisco SME. If the tape must be used, it needs to be formatted or relabeled.
- Any tape that is part of Cisco SME should not be used with other hosts that are not part of Cisco SME. If the user uses the tape anyhow, the data on the tape may be lost.
- The user should make sure that SSH is enabled on all nodes that are part of the Cisco SME cluster.
- FCIP, iSCSI, and Cisco SME interfaces are not supported on the same Cisco SME module.
- When a backup or recovery operation is underway, you should not deactivate or activate zone sets or make any zone set changes.
- The user should never add a disk as a Cisco SME Tape device.
- If a Cisco MDS 9000 Family switch uses a generation-1 supervisor, migrating to generation 2 will result in deletion of the Cisco SME configuration. Before migrating to the new supervisor on the Cisco MDS 9000 Family switch, the user should make a backup copy of all keys and keep the Cisco SME configuration in one place. Delete the Cisco SME cluster and migrate the supervisor; then after the switch is up, the user needs to redo the Cisco SME configuration.
- Upgrading the Cisco MDS 9000 Family software release is disruptive to Cisco SME flows.
- Purging the key is an unrecoverable operation; the key cannot ever be retrieved.
- Manual load balancing is required after a successful In-Service Software Upgrade (ISSU) operation.
- For switches in the same fabric, the fabric membership configured in the CLI should be the same.
- Set the default zone policy to Deny.

Cisco SME Disk Checklist

- By default, all SME Clusters are Tape capable. Disk capability needs to be enabled for SME Disk (for the cluster to be enabled) before the first Cisco SME interface is added.

- Enabling encryption for the first time and disabling decryption without data preparation will lead to the loss of all the existing data on the disk. This process should be used only in a fresh deployment (using a LUN with no existing usable data).
- Use Cisco SME CLI recovery only for recovery of encryption keys and not for data.
- Purging a key is an unrecoverable operation; the key cannot ever be retrieved.
- The path that is undergoing data preparation should be online until data preparation is complete. Any host port or target port flap will result in failure of the data preparation operation.
- All paths (ITLs) from a host to the target LUN should be on the Cisco SME disk to prevent data corruption.
- Manual load balancing is required after a successful ISSU operation.
- Cisco SME Disk does not allow dynamic resizing of LUNs.
- Cisco SME Disk does not support thin provisioning of disks.
- For switches in the same fabric, the fabric membership configured on the CLI should be the same.
- The SSH feature should be enabled.
- Cisco SME Disk clustering is not compatible with IP compression in combination with configuration of IPsec and write acceleration.
- Set the default zone policy to Deny.
- Use FC-Redirect v2 across switches for Cisco SME Disk.
- Cisco SME Disk will not work if the block size is not 512 bytes.

Generic Checklist

- The smart card reader works only in Microsoft Windows XP (32-bit) and Windows Server 2003 (32-bit) platforms.
- The Cisco DCNM password should be the same as the switch password. All the switches in the fabric that need to be managed by Cisco SME and DCNM also should have the same password.
- In the smeserver. properties file located in <fm install path>/dcm/fm/conf/, sme.useIP should be set to True.
- Do not upgrade Cisco DCNM and the switch at the same time if tables are dropped for any reason.
- Do not use the same target LUN (Cisco SME disk) in two different Cisco SME Disk clusters.
- Do not use the same target ports in different clusters.
- Do not add the same Cisco SME interfaces in two different Cisco SME Disk clusters.
- Use different clusters for Cisco SME Disk and Cisco SME Tape.

Sizing and Placement Deployment Examples

This appendix presents deployment examples and derives the requirements for the number and placement of Cisco SME interfaces. The calculations shown here are per fabric.

Example 1: Single-Switch Single-Fabric Tape Backup Environment

- The backup environment consists of 16 application servers and 30 LTO-3 tape drives all connected to a single Cisco MDS 9000 Family switch.
- The switch has Cisco MDS 9000 NX-OS Software Release 5.2(1) or later installed.
- The LTO-3 drives achieve about 80-MBps peak throughput without Cisco SME.

- The backup timeframe is 8 hours, in which about 50 TB of data is backed up every day.

Figure 33 shows the environment for Example 1.

Figure 33. Single-Switch Backup Environment

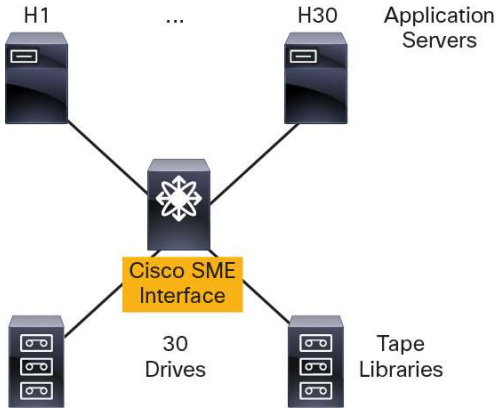

Sizing Using Peak Disk or Tape Drive Throughput

With each LTO-3 drive getting 80 MBps, 5 LTO-3 drives can be supported by one Cisco SME interface (500/80). To support 30 drives, a total of 6 Cisco SME interfaces are required. Adding the capacity to handle failures (20 percent or 1 interface per switch, whichever is greater), a total of about 8 Cisco SME interfaces (and hence 8 Cisco MDS 9000 18/4-Port MSMs) are required, all on the single Cisco MDS 9000 Family switch.

Sizing Using Backup Timeframe Calculations

Fifty TB in 8 hours equals about 1820 MBps (50 x 1024 x 1024)/(8 x 60 x 60). This scenario requires 5 Cisco SME interfaces (1820/500). Adding the capacity for failures (20 percent or 1 interface per switch, whichever is greater), a total of 6 Cisco SME interfaces (and hence 6 Cisco MDS 9000 18/4-Port MSMs) are required, all on the single Cisco MDS 9000 switch.

Zoning Check

The total number of host-target pairs is 16 x 30 = 480, which is within the supported limit of 4096 ITL combinations, so, all 16 hosts can be zoned to all 30 drives.

Example 2: Core-Edge Topology Tape Backup Environment

- The backup environment consists of 16 media servers, 30 LTO-3 tape drives, and Cisco MDS 9000 Family switches.
- The hosts are connected to the edge switches.
- The disk or tape drives are connected to the core switches: 15 drives on each core switch.
- All the switches have Cisco MDS 9000 NX-OS Software Release 5.2(1) or later installed.
- The LTO-3 drives achieve about 80-MBps peak throughput without Cisco SME.
- The backup timeframe is 8 hours, in which about 50 TB of data is backed up every day.

Figure 34 shows the environment for Example 2.

Figure 34. Core-Edge Topology

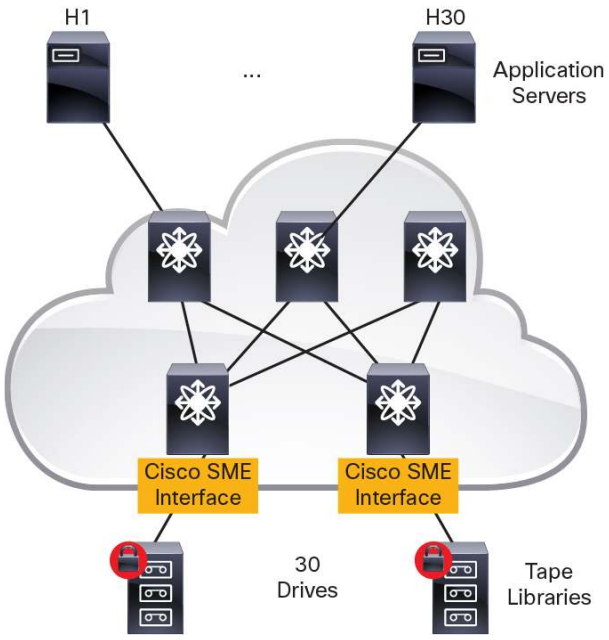

In this topology, the Cisco SME interfaces are provisioned on the core switches to which the disk or tape drives are connected. For such a scenario, the number of Cisco SME interfaces on each switch should be proportional to the throughput requirements of the disk or tape drives connected to it.

Sizing Using Peak Disk or Tape Drive Throughput

With each LTO-3 drive getting 80 MBps, 5 LTO-3 drives can be supported by one Cisco SME interface (500/80). To support 15 drives on each switch, 3 Cisco SME interfaces are required. Similarly, 3 Cisco SME interfaces are required on the other switch. Adding the capacity needed to handle failures (20 percent or 1 interface per switch, whichever is greater), a total of 4 + 4 = 8 Cisco SME interfaces (and hence 8 Cisco MDS 9000 18/4-Port MSMs) are required: 4 on each switch.

Sizing Using Backup Timeframe Calculations

Fifty TB in 8 hours equals about 1820 MBps (50 x 1024 x 1024)/(8 x 60 x 60). This scenario requires 4 Cisco SME interfaces (1820/500): 2 on each switch. Adding the capacity for failures (20 percent or 1 interface per switch, whichever is greater), a total of 6 Cisco SME interfaces (and hence 6 Cisco MDS 9000 18/4-Port MSMs) are required: 3 on each switch.

Zoning Check

The total number of host-target pairs is 16 x 30 = 480, which is within the supported limit of 4096 ITL combinations, so all 16 hosts can be zoned to all 30 drives.

ISL Considerations

All the ISLs between the core and the edge switches should be on generation-2 (or later) Fibre Channel modules.

Target Connectivity

All the disk or tape drives should be connected to generation-2 (or later) Fibre Channel modules.

Example 3: Edge-Core-Edge Topology Tape Backup Environment with Restricted Zoning

- The backup environment consists of 80 media servers, 72 LTO-3 disk or tape drives, and Cisco MDS 9000 Family switches.
- The hosts are connected to the server-edge switches.
- The tape drives are connected to the target-edge switches.
- The core has 2 switches.
- All switches have Cisco MDS 9000 NX-OS Software Release 5.2(1) or later installed.
- The LTO-3 drives achieve about 80-MBps peak throughput without Cisco SME.
- The backup timeframe is 16 hours, in which about 250 TB of data is backed up every day.

Figure 35 shows the environment for Example 3.

Figure 35. Edge-Core-Edge Topology with Restricted Zoning

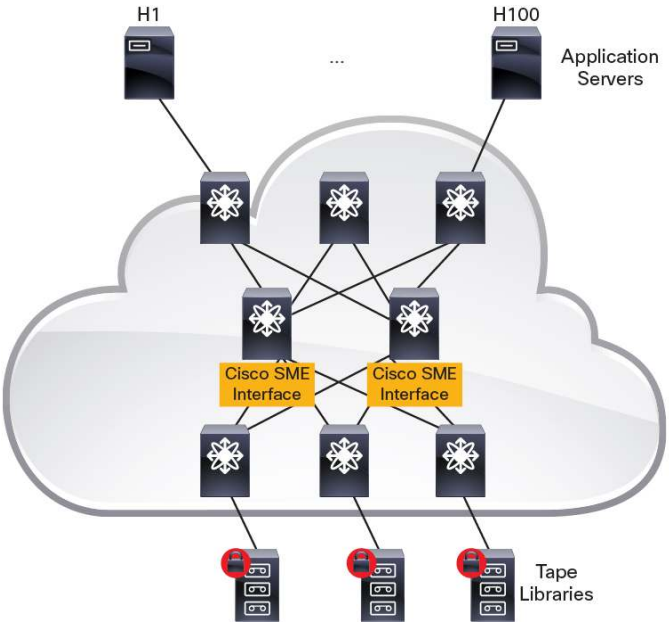

In this topology, the Cisco SME interfaces are provisioned on the core switches, equally distributed.

Sizing Using Peak Disk or Tape Drive Throughput

With each LTO-3 drive achieving 80 MBps, 5 LTO-3 drives can be supported by one Cisco SME interface (500/80). To support 72 drives, 15 Cisco SME interfaces are required. Adding the capacity to handle failures (20 percent or 1 interface per switch, whichever is greater), a total of 18 Cisco SME interfaces (and hence 18 Cisco MDS 9000 18/4-Port MSMs) are required: 9 on each switch.

Sizing Using Backup Timeframe Calculations

Two-hundred and fifty TB in 16 hours equals about 4551 MBps (250 x 1024 x 1024)/(16 x 60 x 60). This scenario requires 11 Cisco SME interfaces (4551/500). Adding the capacity for failures (20 percent or 1 interface per switch, whichever is greater), a total of 14 Cisco SME interfaces (and hence 14 Cisco MDS 9000 18/4-Port MSMs) are required: 7 on each switch.

Zoning Check

The total number of host-target pairs is 80 x 72 = 5760, which is beyond the supported limit of 4096 ITL combinations. In this case, zoning restrictions need to be placed so that the total number of ITLs is below 4096. For example, the 80 hosts can be divided into 2 groups of 40. Similarly, the 72 disk or tape drives can be divided into 2 groups of 36. Each of the groups of 40 hosts and 36 targets can be zoned with each other, resulting in 1440 ITLs for each group, or a total of 1440 x 2 = 2880 ITLs, which is within the limit of 4096.

ISL Considerations

All the ISLs between the core and the edge switches should be on generation-2 (or later) Fibre Channel modules.

Target Connectivity

All the disk or tape drives should be connected to generation-2 (or later) Fibre Channel modules.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Printed in USA C07-464433-01 06/12