

IBM TotalStorage SAN Volume Controller
Storage Software for Cisco MDS 9000



Getting Started

Version 1 Release 1

IBM TotalStorage SAN Volume Controller
Storage Software for Cisco MDS 9000



Getting Started

Version 1 Release 1

Note

Before using this information and the product it supports, read the information in "Notices" on page 71.

First Edition (November 2003)

© Copyright International Business Machines Corporation 2003. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this guide?	v
Who should use this guide?	v
Related publications	v
Related Web sites	vii
How to order IBM publications	viii
How to send your comments	viii

Chapter 1. Virtualization and the IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000.	1
Virtualization	1
The need for virtualization	3
Fabric level virtualization models	4
Symmetric virtualization	4
Introducing the SAN Volume Controller for Cisco MDS 9000	5

Chapter 2. Installation planning	11
Preparing your SAN Volume Controller Storage Software for Cisco MDS 9000 environment	12

Chapter 3. Planning guidelines for using your SAN Volume Controller Storage Software for Cisco MDS 9000 in a SAN environment	13
Storage Area Network	13
SAN fabric	14
Virtual storage area networks (VSANs)	15
Nodes	15
Cluster state	16
Cluster operation and quorum disks	16
Internal battery backup and power domains	17
Disk controllers	17
Migration	18
Image mode virtual disk (VDisk) migration	18
Copy Services	18
FlashCopy	19
FlashCopy mappings	20
FlashCopy consistency groups	22
Remote Copy	23
Synchronous Remote Copy	23
Remote Copy consistency groups	24

Chapter 4. Object descriptions	27
Disk controller systems	28
Managed disks (MDisks)	29
Managed disk (MDisk) groups	31
Virtual disks (VDisks)	33
Virtual disk-to-host mapping	36
Host objects	37

Chapter 5. Planning for your SAN Volume Controller Storage Software for Cisco MDS 9000 configuration	39
Maximum configuration	40
Configuration rules and requirements	42
Configuration rules	43
Configuration requirements	48

Chapter 6. Setting up SAN Volume Controller Storage Software for Cisco MDS 9000	51
Preparing the Cisco MDS switch	51
Setting up the Cisco MDS switch	51
Creating SAN Volume Controller Storage Software for Cisco MDS 9000 interfaces	54
Creating a cluster	54
Creating a cluster on the Console for SAN Volume Controller for Cisco MDS 9000	55
Separating hosts and storage devices	56
Verifying an existing interface	57
Adding nodes to a cluster using the Cisco MDS 9000 CLI	57
Adding nodes to the cluster using the IBM TotalStorage CLI	58
Deleting a node from a cluster using the IBM TotalStorage CLI	60

Chapter 7. SAN Volume Controller Storage Software for Cisco MDS 9000 supported environment	63
Supported host attachments	63
Physical disk storage system support	64
Supported fibre-channel, host bus adapters (HBAs)	64
Supported switches	64

Chapter 8. Getting Started Hints and Tips	65
Checking the battery capacity	65
Hints and tips for troubleshooting the SAN Volume Controller Storage Software for Cisco MDS 9000 environment	66

Appendix. Accessibility	69
--	-----------

Notices	71
Trademarks	72

Glossary	73
---------------------------	-----------

Index	79
------------------------	-----------

About this guide?

This publication introduces the IBM® TotalStorage® SAN Volume Controller for Cisco MDS 9000.

Related topics:

- Chapter 1, "Virtualization and the IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000," on page 1
- Chapter 2, "Installation planning," on page 11
- Chapter 3, "Planning guidelines for using your SAN Volume Controller Storage Software for Cisco MDS 9000 in a SAN environment," on page 13
- Chapter 4, "Object descriptions," on page 27
- Chapter 5, "Planning for your SAN Volume Controller Storage Software for Cisco MDS 9000 configuration," on page 39
- Chapter 6, "Setting up SAN Volume Controller Storage Software for Cisco MDS 9000," on page 51
- Chapter 7, "SAN Volume Controller Storage Software for Cisco MDS 9000 supported environment," on page 63
- Chapter 8, "Getting Started Hints and Tips," on page 65
- "Accessibility," on page 69

Who should use this guide?

This publication is intended for anyone who is planning to set up and configure the IBM TotalStorage SAN Volume Controller for Cisco MDS 9000.

Related publications

The tables in this section list and describe the following publications:

- The publications that make up the library for the IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000
- Other IBM publications that relate to the SAN Volume Controller Storage Software for Cisco MDS 9000
- Cisco publications that relate to the SAN Volume Controller Storage Software for Cisco MDS 9000

SAN Volume Controller Storage Software for Cisco MDS 9000 library:

Table 1 on page vi lists and describes the publications that make up the SAN Volume Controller Storage Software for Cisco MDS 9000 library. Unless otherwise noted, these publications are available in Adobe portable document format (PDF) on a compact disc (CD) that comes with the SAN Volume Controller Storage Software for Cisco MDS 9000. If you need additional copies of this CD, the order number is SK2T-8813. These publications are also available as PDF files from the following Web site:

<http://www.ibm.com/storage/support/2062-2300>

Table 1. IBM Publications in the SAN Volume Controller Storage Software for Cisco MDS 9000 library

Title	Description	Order number
<i>IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000: Command-Line Interface User's Guide</i>	This guide describes the commands that you can use from the SAN Volume Controller Storage Software for Cisco MDS 9000 command-line interface (CLI).	SC26-7555
<i>IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000: Configuration Guide</i>	This guide provides guidelines for configuring your SAN Volume Controller Storage Software for Cisco MDS 9000.	SC26-7554
<i>IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000: Getting Started Guide</i>	This guide introduces the SAN Volume Controller Storage Software for Cisco MDS 9000 and lists the features you can order. It also provides guidelines for planning the installation and configuration of the SAN Volume Controller Storage Software for Cisco MDS 9000.	GA22-1055

Other IBM publications:

Table 2 lists and describes other IBM publications that contain additional information related to the SAN Volume Controller Storage Software for Cisco MDS 9000.

Table 2. Other IBM publications

Title	Description	Order number
<i>IBM TotalStorage Enterprise Storage Server, IBM TotalStorage SAN Volume Controller, IBM TotalStorage SAN Volume Controller for Cisco MDS 9000, Subsystem Device Driver: User's Guide</i>	This guide describes the IBM Subsystem Device Driver Version 1.5 for TotalStorage Products and how to use it with the SAN Volume Controller Storage Software for Cisco MDS 9000. This publication is referred to as the <i>IBM TotalStorage Subsystem Device Driver: User's Guide</i> .	SC26-7608

Cisco Systems publications:

The following list identifies publications from Cisco Systems that contain additional information that is related to the Cisco MDS 9000 Family.

- *Regulatory Compliance and Safety Information for the Cisco MDS 9000 Family*
- *Cisco MDS 9100 Series Quick Start Guide*

- *Cisco MDS 9500 Series and Cisco MDS 9216 Quick Start Guide*
- *Cisco MDS 9100 Series Hardware Installation Guide*
- *Cisco MDS 9216 Switch Hardware Installation Guide*
- *Cisco MDS 9500 Series Hardware Installation Guide*
- *Cisco SAN Volume Controller Configuration Guide*
- *Cisco MDS 9000 Family Command Reference*
- *Cisco MDS 9000 Family Configuration Guide*
- *Cisco MDS 9000 Family Fabric Manager User Guide*
- *Cisco MDS 9000 Family Troubleshooting Guide*
- *Cisco MDS 9000 Family System Messages Guide*
- *Cisco MDS 9000 Family MIB Reference Guide*

You can access the most current technical documentation at the following Web site:
<http://www.cisco.com/>

Click:

- **Products and Services**
- **Storage Networking Products**
- **Cisco MDS 9000 Series Multilayer Switches**

Ordering technical documentation from Cisco Systems:

You can find instructions for ordering documentation at the following Web site:

http://www.cisco.com/univercd/cc/td/doc/es_inpk/pdi.htm

You can order Cisco Systems documentation in the following ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Networking Products MarketPlace at:
<http://www.cisco.com/en/US/partner/ordering/index.shtml>
- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, USA) at 408-526-7208 or, elsewhere in North America, by calling 800-553-NETS (800-553-6387).

Related topics:

- “How to order IBM publications” on page viii
- “How to send your comments” on page viii

Related Web sites

Table 3 lists Web sites that have information about SAN Volume Controller Storage Software for Cisco MDS 9000 or related products or technologies.

Table 3. Web sites

Type of information	Web site
SAN Volume Controller Storage Software for Cisco MDS 9000 support	http://www.ibm.com/storage/support/2062-2300/
Technical support for IBM storage products	http://www.ibm.com/storage/support/

Table 3. Web sites (continued)

Type of information	Web site
Cisco Systems documentation	http://www.cisco.com/univercd/home/home.htm
Cisco Systems product site	http://www.cisco.com/
Cisco Systems international site	http://www.cisco.com/public/countries_languages.shtml

How to order IBM publications

This topic explains how to order copies of IBM publications and how to set up a profile to receive notifications about new or changed publications.

The IBM publications center:

The publications center is a worldwide central repository for IBM product publications and marketing material.

The IBM publications center offers customized search functions to help you find the publications that you need. Some publications are available for you to view or download free of charge. You can also order publications. The publications center displays prices in your local currency. You can access the IBM publications center through the following Web site:

www.ibm.com/shop/publications/order/

Publications notification system:

The IBM publications center Web site offers you a notification system for IBM publications. Register and you can create your own profile of publications that interest you. The publications notification system sends you a daily e-mail that contains information about new or revised publications that are based on your profile.

If you want to subscribe, you can access the publications notification system from the IBM publications center at the following Web site:

www.ibm.com/shop/publications/order/

Related topics:

- “Related publications” on page v

How to send your comments

Your feedback is important to help us provide the highest quality information. If you have any comments about this book or any other SAN Volume Controller Storage Software for Cisco MDS 9000 documentation, you can submit them in one of the following ways:

- e-mail

Submit your comments electronically to the following e-mail address:

starpubs@us.ibm.com

Be sure to include the name and order number of the book and, if applicable, the specific location of the text you are commenting on, such as a page number or table number.

- Mail or fax

Fill out the Readers' Comments form (RCF) at the back of this book. Return it by mail or fax (1-408-256-0488), or give it to an IBM representative. If the RCF has been removed, you can address your comments to:

International Business Machines Corporation
RCF Processing Department
Dept. FTF/Building 050-1
5600 Cottle Road
San Jose, CA 95193-0001
U.S.A.

Related topics:

- "Related publications" on page v

Chapter 1. Virtualization and the IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000

This topic describes why virtualization is needed, what virtualization is, and introduces the IBM TotalStorage SAN Volume Controller for Cisco MDS 9000.

Virtualization

Virtualization is a concept that applies to many areas of the information technology industry. Where data storage is concerned, virtualization includes the creation of a pool of storage that contains several disk subsystems. These subsystems can be from various vendors. The pool can be split into virtual disks that are visible to the host systems that use them. Therefore, virtual disks can use mixed back-end storage and provide a common way to manage storage-area-network (SAN) storage.

Historically, the term *virtual storage* has described the virtual memory techniques that have been used in operating systems. The term *storage virtualization*, however, describes the shift from thinking about physical volumes of data, to thinking about logical volumes of data. This shift can be made on several levels of the components of storage networks. Virtualization separates the representation of storage between the operating system and its users from the actual physical storage components. This technique has been used in mainframe computers for many years through methods such as system-managed storage, and products like the IBM Data Facility Storage Management Subsystem (DFSMS). Virtualization can be applied at four main levels:

- Virtualization at the *server* level is performed by managing volumes on the operating systems servers. An increase in the amount of logical storage over physical storage is more suitable for environments that do not have storage networks.
- Virtualization at the *storage device* level is in common use. Striping, mirroring, and redundant array of independent disks (RAID) arrays are used by almost all disk subsystems. This type of virtualization can range from simple RAID controllers to advanced volume management such as that provided by the IBM TotalStorage Enterprise Storage Server (ESS) or by Log Structured Arrays (LSA). The Virtual Tape Server (VTS) is another example of virtualization at the device level.
- Virtualization at the *fabric* level enables storage pools to be independent of the various types of servers and of the physical components that make up the storage pools. One management interface can be used to manage different storage systems without affecting the servers. The SAN Volume Controller Storage Software for Cisco MDS 9000, can be used to perform virtualization at the fabric level.
- Virtualization at the *file system* level provides the highest level of virtual storage. It can also provide the highest benefit, because it is data that is to be shared, allocated, and protected; not volumes.

Virtualization is a radical departure from traditional storage management. In traditional storage management, storage is attached directly to host systems, and the local host system controls storage management. SANs have introduced the principle of networks of storage, but storage is still primarily created and

maintained at the Redundant Array of Independent Disks (RAID) subsystem level. Multiple RAID controllers of different types require knowledge of, and software that is specific to, the given hardware. Virtualization brings a central point of control for disk creation and maintenance. It brings new ways of handling storage maintenance.

Where storage is concerned, one problematic area that virtualization addresses is that of unused capacity. Rather than individual storage systems remaining islands unto themselves, allowing excess storage capacity to be wasted when jobs do not require it, storage is pooled so that jobs needing the highest storage capacity can use it when they need it. Regulating the amount of storage available becomes easier to orchestrate without computing resource or storage resource having to be turned off and on.

Types of virtualization:

Virtualization can be performed either asymmetrically or symmetrically:

Asymmetric

A virtualization engine is outside the data path and performs a metadata style service.

Symmetric

A virtualization engine sits in the data path, presenting disks to the hosts but hiding the physical storage from the hosts. Advanced functions, such as cache and Copy Services can, therefore, be implemented in the engine itself.

Virtualization at any level provides benefits. When several levels are combined, however, the benefits of those levels can also be combined. An example of how you can gain the highest benefits is if you attach a low-cost RAID controller to a virtualization engine that provides virtual volumes for use by a virtual file system.

Note: The SAN Volume Controller Storage Software for Cisco MDS 9000 implements fabric-level *virtualization*. Within the context of the SAN Volume Controller Storage Software for Cisco MDS 9000 and throughout this document, *virtualization* refers to fabric-level virtualization.




Figure 1. Levels of virtualization

Related topics:

- “Virtual disks (VDisks)” on page 33

The need for virtualization

Storage is a facility that computer users want to access at any time, from any location, with a minimum amount of management. Users expect the storage devices to provide enough capacity and to be reliable. The amount of storage that users require, however, is increasing quickly. Internet users use large amounts of storage daily. Many users are mobile, access patterns cannot be predicted, and the content of the data becomes more and more interactive. Because the amount of data that is handled is large, it can no longer be managed manually. Automatic management is required, as are new levels of bandwidth and load balancing. Also, it is important that all this data can be shared between different types of computer platforms, because the communication networks cannot handle the large replication, download, and copying operations that are required.

Storage area networks (SANs) are high-speed switched networks that let multiple computers share access to many storage devices. SANs allow for the use of advanced software that automatically manages the storage of data. With such advanced software, the computers that are connected to a particular network can, therefore, access storage wherever that storage is available in the network. The user is no longer aware of, and no longer needs to know, which physical devices contain which data. The storage has become virtualized. In a similar way to how virtual memory has solved the problems of the management of a limited resource in application programs, the virtualization of storage has given users a more intuitive use of storage, while software quietly manages the storage network in the background.

Fabric level virtualization models

In traditional storage management, storage devices are connected directly to host systems, and are maintained locally by those host systems. Although SANs have introduced the principle of networks, storage devices are still mainly assigned to individual host systems and storage is still mainly created and maintained at the RAID subsystem level. Therefore, RAID controllers of different types need knowledge of, and software that is specific to, the hardware that is used.

Virtualization provides a complete change from the traditional storage management. It provides a central point of control for disk creation and management, and therefore requires changes to the way in which storage management is done.

Fabric level virtualization is the principle in which a pool of storage is created from more than one disk subsystem. This pool is then used to set up virtual disks that are made visible to the host systems. These virtual disks use whatever storage is available and permit a common way to manage SAN storage.

Fabric level virtualization can be done in either of two ways: asymmetric or symmetric.

With asymmetric virtualization, the virtualization engine is outside the data path and performs a metadata style service. The metadata server contains all the mapping and the locking tables while the storage devices contain only data.

Because the flow of control is separated from the flow of data, input/output (I/O) operations can use the full bandwidth of the SAN. A separate network or SAN link is used for control purposes.

However, there are disadvantages to asymmetric virtualization:

- Data is at risk to increased security exposures, and the control network must be protected with a firewall.
- Metadata can become very complicated when files are distributed across several devices.
- Each host that accesses the SAN must know how to access and interpret the metadata. Specific device driver or agent software must therefore be running on each of these hosts.
- The metadata server cannot run advanced functions, such as caching or Copy Services, because it only knows about the metadata and not about the data itself.

Symmetric virtualization

The SAN Volume Controller Storage Software for Cisco MDS 9000 provides symmetric virtualization. Virtualization splits the physical storage Redundant Array of Independent Disks (RAID) arrays into smaller chunks of storage that is known as extents. These extents are then concatenated together, using various policies, to make virtual disks. With symmetric virtualization, host systems can be isolated from the physical storage. Advanced functions, such as data migration, can run without the need to reconfigure the host. With symmetric virtualization, the virtualization engine is the central configuration point for the SAN.

In symmetric virtual storage networks, see Figure 2 on page 5, data and control both flow over the same path. Because the separation of the control from the data occurs in the data path, the storage can be pooled under the control of the

virtualization engine. The virtualization engine performs the logical to physical mapping.




Figure 2. Symmetrical virtualization

The virtualization engine directly controls access to the storage and to the data that is written to the storage. As a result, locking functions that provide data integrity and advanced functions, such as cache and copy services, can be run in the virtualization engine itself. The virtualization engine is, therefore, a central point of control for device and advanced function management. Symmetric virtualization also allows you to build a kind of firewall in the storage network. Only the virtualization engine can give access through the firewall. Symmetric virtualization does, however, cause some problems.

The main problem that is associated with symmetric virtualization is related to poor performance, because all I/O must flow through the virtualization engine. This problem is one of scalability. You can use an n-way cluster of virtualization engines that has failover capacity to solve this problem. You can scale the additional processor power, cache memory, and adapter bandwidth to get the level of performance that you want. The memory and processing power can be used to run the advanced functions, such as copy services and caching.

The IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000 uses symmetric virtualization. Single virtualization engines, which are known as nodes, are combined to create clusters. A cluster contains 4 nodes.

Related topics:

- “Virtualization” on page 1

Introducing the SAN Volume Controller for Cisco MDS 9000

The IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000 provides symmetric virtualization by creating a pool of managed disks from the attached storage subsystems, which are then mapped to a set of virtual disks for use by various attached host computer systems. The system administrators can view and access a common pool of storage on the SAN which allows them to use storage resources more efficiently, and provides a common base for advanced functions.

The IBM TotalStorage SAN Volume Controller for Cisco MDS 9000 consists of these elements:

- IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000
- Cisco MDS 9000 Caching Services Module (CSM)

The solution offers the following benefits and advantages:

- Reduces complexity
- Lowers the cost of managing SAN-based storage
- Creates a single pool of storage from disparate storage devices to increase capacity utilization
- Implements a cache-based, clustered architecture to provide a highly available solution.
- Provides the scalability and performance required in today's demanding storage environments.

This solution implements a cache-based, clustered architecture and provides a highly available, scalable alternative that is necessary in today's demanding storage environments.

Each Cisco MDS 9000 Series CSM provides two storage engines or nodes. SAN Volume Controller Storage Software for Cisco MDS 9000 runs in a clustered mode. A SAN Volume Controller Storage Software for Cisco MDS 9000 cluster is comprised of four nodes.

Based on virtualization technology, this solution is designed to support a virtualized pool of storage from the storage subsystems attached to a SAN. It manages the combined storage volumes from a central point, avoids downtime for planned outages, increases capacity utilization, and implements copy services from a single license across multiple storage devices.

This storage pool helps tap unused storage capacity by increasing efficiency. It is designed as an integrated solution supporting high performance and continuous availability in open-system environments. Storage volumes are represented to applications as virtual disks. These virtual disks are created from the pool of managed disks residing behind the storage engines.

SAN Volume Controller Storage Software for Cisco MDS 9000 Features This section explains the features provided by the SAN Volume Controller Storage Software for Cisco MDS 9000 solution.

- A central point for volume management control

Through virtualization, the SAN Volume Controller Storage Software for Cisco MDS 9000 helps create pools of managed disks spanning multiple storage subsystems. These managed disks are mapped to virtual disks used by server applications thus making better use of existing storage. This simple interface incorporates the Storage Management Initiative Specification (SMIS) application programming interface (API), and further demonstrates the IBM focus on open standards.

- Dynamic data migration

The SAN Volume Controller Storage Software for Cisco MDS 9000 includes a dynamic data-migration function that helps administrators migrate storage from one device to another, without taking it offline. This allows administrators to reallocate and scale storage capacity without disrupting applications. SAN Volume Controller Storage Software for Cisco MDS 9000 also leverages the IBM TotalStorage Subsystem Device Driver (SDD) multipathing software.

- Improved resource utilization
This solution enables more efficient use of personnel and technology resources. It helps increase administrator productivity by empowering central management of volumes under disparate storage controllers from a single user interface. It also increases the amount of available storage capacity by pooling storage across multiple devices. Designed to manage up to two petabytes (PB) of total usable storage capacity, SAN Volume Controller Storage Software for Cisco MDS 9000 will support even higher performance by adding storage engine or node pairs. All storage engines within a cluster jointly manage the entire capacity of a storage pool.
- Advanced copy services
With conventional SAN disk arrays, copy operations are limited to in-box or like-box-to-like-box environments. But the SAN Volume Controller Storage Software for Cisco MDS 9000 moves copy services from individual storage controllers to the SAN. Administrators can apply copy services across disparate storage devices within the network. Advanced copy services such as FlashCopy and Peer-to-Peer Remote Copy (PPRC) are supported across the managed storage.
- Cisco MDS 9000 Caching Services Module (CSM)
The CSM integrates two high performance processing nodes which, when part of the SAN Volume Controller Storage Software for Cisco MDS 9000, deliver network hosted virtualization and replication services. Each CSM includes 8 GB of local cache used to hold recently accessed data blocks. There are two disk drives. Dedicated hard drives combined with on board dual redundant batteries provide protection for data in cache in the event of a power loss. To further ensure data availability and integrity, nodes are paired with nodes on other CSMs in high availability clusters.
- Fabric-based virtualization using intelligent networking services
Cisco MDS 9000 fabric-based virtualization provides a level of integration with intelligent SAN services unavailable to host or virtualization appliance-based solutions. The Cisco MDS 9000 platform with integrated CSM delivers the intelligence and advanced features required to make multilayer intelligent storage area networks a reality including hardware-enabled innovations designed to dramatically improve scalability, availability, security, and manageability of storage networks, resulting in increased utility and lower total cost of ownership (TCO).
- Virtual SAN
Virtual SANs (VSANs), a function of the Cisco MDS 9000 switch, allow more efficient SAN utilization by creating hardware-based isolated environments within a single SAN fabric. Each VSAN can be configured with Fibre Channel zones and maintains its own fabric services for added scalability and resilience. VSANs allow the cost of SAN infrastructure to be shared among more users, while assuring absolute segregation and security of traffic and retaining independent control of configuration on a VSAN-by-VSAN basis. VSANs provide a protective barrier between application hosts and physical storage, enhancing data integrity in a virtualized storage environment.
- Comprehensive security
The SAN Volume Controller for Cisco MDS 9000 applies extensive security measures at all possible points of attack. SSH, RADIUS, SNMPv3, and role-based access control are used to block unauthorized access. To guard against compromising control traffic, Fibre Channel Security Protocol (FC-SP) provides confidentiality, data origin authentication, and connectionless integrity across the fabric. Data plane traffic is secured with VSANs, guaranteeing segregation of

traffic across shared fabrics, and with zoning to satisfy traffic segregation requirements within a VSAN. Hardware-based ACLs provide further granularity for advanced security options.

- High availability

Like all other Cisco MDS 9000 Family modules, the CSM is hot-swappable and fully integrates into the Cisco MDS 9000 high availability architecture. The Cisco MDS 9000 software architecture offers an unparalleled level of availability including automatic restart of failed supervisor processes and fabric level availability via Cisco PortChannel capability which allows 16 physical links to be aggregated into one logical interface. The logical interface remains active in the event of a port, ASIC, or module failure, and can sustain the failure of any physical link without causing a reset. Additionally, Fabric Shortest Path First (FSPF) multipathing provides the intelligence to load balance across up to 16 equal cost paths and, in the event of a switch failure, to dynamically reroute traffic. When deployed in clustered pairs and combined with SAN Volume Controller Storage Software for Cisco MDS 9000, availability is extended to the volume level ensuring maximum uptime.

- Management options

The SAN Volume Controller Storage Software for Cisco MDS 9000 provides these principal modes of management of your virtual storage environment:

- The Cisco MDS 9000 command line interface (CLI) used for managing the switch and for setting up the initial cluster, The Cisco CLI can be used to provide virtualization functions as well.
- The IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000 command line interface can be used to provide virtualization functions.
- The IBM-supplied Console for SAN Volume Controller for Cisco MDS 9000 graphical user interface (GUI) can be used to provide virtualization functions.
- The IBM-supplied CIM Agent and Console for SAN Volume Controller for Cisco MDS 9000 graphical user interface (GUI) can be used by other management applications to manage the SAN Volume Controller Storage Software for Cisco MDS 9000.

Basic facts about SAN Volume Controller Storage Software for Cisco MDS 9000:

A *node* is a single storage engine. Each CSM supports two storage engines, or nodes. Four nodes are combined to comprise a cluster. Nodes within the cluster are grouped in pairs known as an I/O group. Nodes within an I/O group back up one another. Data written to the nodes is duplicated across caches in both nodes. Virtual disks are shared between nodes in an I/O group. To eliminate any single point of failure, nodes in an I/O group must be on a separate CSM.

The SAN Volume Controller Storage Software for Cisco MDS 9000 I/O groups see the storage presented to the SAN by the backend controllers as a number of disks, known as *managed disks*. The application services do not see these managed disks. Instead, they see a number of logical disks, known as *virtual disks*, that are presented to the SAN by the SAN Volume Controller Storage Software for Cisco MDS 9000. Each node must only be in one I/O group and provide access to the virtual disks in the I/O group.

The SAN Volume Controller Storage Software for Cisco MDS 9000 helps to provide continuous operations and can also optimize the data path to ensure performance levels are maintained.

Multi-pathing software:

- IBM Subsystem Device Driver (SDD) is an IBM pseudo device driver designed to support the multipath configuration environments in IBM products.

Note: SDD will not coexist on a host with other multi-pathing drivers such as the LSI RDAC driver for FASTT or the QLogic pathing driver.

Interoperability with the IBM TotalStorage SAN Volume Controller:

While the IBM Total Storage SAN Volume Controller for Cisco MDS 9000 provides the same function as the SAN Volume Controller, they do not interoperate at the cluster level. Because the two products do not interoperate, they do not recognize each other's nodes as candidates. Clusters cannot contain nodes managed by the other product. Remote copy relationships cannot be formed between SAN Volume Controller for Cisco MDS 9000 and SAN Volume Controller clusters.

Chapter 2. Installation planning

Before the service representative can start to set up your SAN Volume Controller Storage Software for Cisco MDS 9000, you will want to verify that the prerequisite conditions for the SAN Volume Controller Storage Software for Cisco MDS 9000 are met. The following checklist will help to ensure that all minimum requirements for installation are met.

To prepare for both hardware and software setup, be sure you have installed and set up a functioning Cisco MDS 9000 Family switch fabric. Each Cisco MDS 9000 Family switch needs a switch IP (mgmt0, one for each switch). Refer to the *Cisco MDS 9216 Switch Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide*, and the *Cisco MDS 9000 Family Configuration Guide*.

1. Are the Cisco MDS 9000 switches mounted?
2. Is the Cisco MDS 9000 Caching Services Module (CSM) installed?
3. Are all the fibre-channel cables connected to the switch?
4. Is each Cisco MDS 9000 switch connected to the Ethernet?
5. Has appropriate connectivity been provided?

Hardware requirements:

This program requires at least two Cisco MDS 9000 Caching Services Modules. Both can be in one IBM 2062-D04/T04, or D07/T07; or one each can go in two 2062s - any mix of IBM 2062-D01, D04/T04, and D07/T07.

Installation of the CIM Agent and Console for SAN Volume Controller for Cisco MDS 9000 requires that you provide a Windows 2000-capable host with the following characteristics:

- Intel-based PC running Windows 2000 Server SP-3
- Intel Pentium processor at 1 gigahertz or faster
- Support for a communications adapter
- CD-ROM drive
- Minimum 1 GB memory, recommended

Current support summaries, including specific software, hardware, and firmware levels supported, are maintained at:

<http://www.ibm.com/storage/support/2062-2300>

Software requirements:

This program supports connectivity with Intel-based servers running Windows NT, Windows 2000 Advanced Server, or Red Hat Linux Advanced Server, and with RISC-based servers running AIX, HP-UX, or Sun Solaris operating systems. These servers require the Subsystem Device Driver (SDD) multipathing driver.

Current support summaries, including specific software, hardware and firmware levels supported, are maintained at:

<http://www.ibm.com/storage/support/2062-2300>

Related topics:

- “Preparing your SAN Volume Controller Storage Software for Cisco MDS 9000 environment”

Preparing your SAN Volume Controller Storage Software for Cisco MDS 9000 environment

This topic provides information you need to ensure that your physical site meets the installation requirements for the SAN Volume Controller Storage Software for Cisco MDS 9000.

Refer to these publications from Cisco Systems for information about preparing your environment:

- *Cisco MDS 9216 Switch Hardware Installation Guide*
- *Cisco MDS 9500 Series Hardware Installation Guide*
- *Cisco MDS 9000 Family Command Reference*
- *Cisco MDS 9000 Family Configuration Guide*

Chapter 3. Planning guidelines for using your SAN Volume Controller Storage Software for Cisco MDS 9000 in a SAN environment

This topic provides planning guidelines to help you set up your SAN Volume Controller Storage Software for Cisco MDS 9000 environment. It also discusses major concepts you need to understand to use the SAN Volume Controller Storage Software for Cisco MDS 9000.

Steps:

Consider the following tasks when planning for your SAN Volume Controller Storage Software for Cisco MDS 9000:

1. Plan your configuration.
2. Plan your SAN environment.
3. Plan your fabric set up.
4. Create the RAID resources that you intend to virtualize.
5. Determine if you have a RAID array that contains data that you want to merge into the cluster.
6. Decide if you will migrate data into the cluster or keep them as image-mode VDisk.
7. Do you plan to use copy services?

Related topics:

- Chapter 5, "Planning for your SAN Volume Controller Storage Software for Cisco MDS 9000 configuration," on page 39
- "Storage Area Network"
- "SAN fabric" on page 14
- "Disk controllers" on page 17
- "Image mode virtual disk (VDisk) migration" on page 18
- "Copy Services" on page 18

Storage Area Network

A storage area network (SAN) is a high-speed dedicated network for sharing storage resources. A SAN allows the establishment of direct connections between storage devices and servers. It offers simplified storage management, scalability, flexibility, availability, and improved data access, movement, and backup.

A SAN storage system consists of four nodes arranged in a cluster. These will appear as part of the SAN fabric, along with the host systems, the RAID controllers, and the storage devices, all connected together to create the SAN.

Prerequisites:

To install a SAN Volume Controller Storage Software for Cisco MDS 9000 into an existing SAN that will be in use during installation, you must first ensure that the switch zoning or the VSAN is set to isolate the new SAN Volume Controller

Storage Software for Cisco MDS 9000 connections from the active part of the SAN. See the *Cisco MDS 9000 Family Configuration Guide* for more information about setting up VSANs.

You need to consider the following information before installing the SAN Volume Controller Storage Software for Cisco MDS 9000:

- Consider the design of the SAN according to your requirement for high availability.
- Identify the operating system for each host system that will be connected to the SAN Volume Controller Storage Software for Cisco MDS 9000, ensuring compatibility and suitability.
 - Specify the Host Bus Adapters (HBAs) for each host
 - Define the performance requirements
 - Determine the total storage capacity
 - Determine the storage capacity per host
 - Determine the host LUN sizes
 - Determine if your SAN has enough ports to connect all hosts and backend storage?
 - Determine if your SAN provides enough ports to connect your back-end storage

See the following Web site for specific firmware levels and the latest supported hardware:

<http://www.ibm.com/storage/support/2062-2300>

SAN fabric

The SAN fabric can be arranged with various zones and VSANs to allow the SAN Volume Controller Storage Software for Cisco MDS 9000s to access the RAID controllers, and for the host systems to access the SAN Volume Controller Storage Software for Cisco MDS 9000s. Ports within a VSAN cannot communicate with each other unless they are also in a zone. The hosts should not directly access or operate on the RAID controllers.

- Install and configure your switches to create the zones and VSANs that the SAN Volume Controller Storage Software for Cisco MDS 9000 needs. You must create at least three general zones, each within a VSAN.
 - Host zone or host VSAN
 - Storage zone or storage VSAN
 - Management zone or management VSAN. The management zone or management VSAN provide a zone through which all the nodes in a cluster communicate with one another.
- Each host (or partition of a host) may have between one and four FC ports.

IBM recommends that the subsystem workload to each SAN Volume Controller Storage Software for Cisco MDS 9000 port should be equal. This will typically involve zoning roughly the same number of host FC ports to each SAN Volume Controller Storage Software for Cisco MDS 9000 FC port.

IBM recommends that you manually set the domain IDs prior to building the multiswitch fabric and prior to zoning. There are several reasons for doing this:

- When two switches are joined while active, they will determine if the domain ID is already in use as before; however, if there is a conflict, it cannot be changed in an active switch. This conflict will cause the fabric merging process to fail.
- The domain ID is used to identify switch ports when zoning is implemented using the domain and switch port number. If domain ID's are negotiated at every fabric start up, there is no guarantee that the same switch will have the same ID the next time; therefore, any zoning definitions may become invalid.
- If the domain ID is changed after a SAN is set up, some host systems may have difficulty logging back in with the switch and a host reconfiguration may be required to detect devices on the switch.

Virtual storage area networks (VSANs)

In simple terms, a VSAN is a fabric. A VSAN is a virtual storage area network (SAN). A SAN is a dedicated network that interconnects hosts and storage devices primarily to exchange SCSI traffic. In a SAN, the physical links provide these interconnections. A set of protocols run over the SAN to handle routing, naming, and zoning. You can design multiple SANs with different topologies.

The Cisco MDS series of Fibre Channel switches provide a way for you to divide the SAN into separate sections called VSANs. A VSAN is a completely contained fabric within the SAN. Traffic does not pass from VSAN to VSAN, not even name server updates. A VSAN is completely isolated to itself. VSANs are similar to LPARs. To isolate the traffic in our virtualized solution, you can use VSANs, zoning, or a combination of VSANs and zoning. Ports within a VSAN cannot communicate with each other unless they are also in a zone.

Define two VSANs to set up the logical path for data to flow. The server VSAN contains all of the HBAs from the servers. You can also use zoning in that VSAN to further isolate server types from one another. The storage VSAN has all of the storage systems connections in it. Both of these VSANs have ports from the SAN Volume Controller for Cisco MDS 9000 in them as well. With this configuration, the only path a server will have to a storage system is through the virtualization engine.

Define one VSAN as the communications conduit. The cluster VSAN has only the ports from the virtualization engines. This VSAN is the communication conduit for the cluster to use for housekeeping.

See *Cisco MDS 9000 Family Software Configuration Guide* for more information.

Nodes

A cluster contains 4 nodes. Each pair of nodes is known as an I/O group. Each node may be in *only* one I/O group. A cluster can be composed of nodes that reside on different switches.

At any one time, a single node in the cluster is used to manage configuration activity. This configuration node manages a cache of the configuration information that describes the cluster configuration and provides a focal point for configuration commands. If the configuration node fails, another node in the cluster will take over its responsibilities.

There are five states in which a node can exist, as described by the following table:

Table 4. Node state

State	Description
Adding	The node has been added to the cluster but has not yet synchronized with the cluster state. (See Note.)
Deleting	The node is in the process of being deleted from the cluster.
Online	The node is operational, assigned to a cluster and has access to the Fibre Channel SAN fabric.
Offline	The node is not-operational. The node has been assigned to a cluster but is not available on the Fibre Channel SAN fabric.
Pending	The node is transitioning between states, and in a few seconds will move to one of the other states.
Note: It is possible that a node can stay in the Adding state for a long time. If this is the case, delete the node and then re-add it. However, you should wait for at least 30 minutes before doing this. If the node that has been added is at a lower code level than the rest of the cluster, the node will be upgraded to the cluster code level, this can take up to 20 minutes. During this time the node will be shown as adding.	

Cluster state

The cluster state holds all configuration and internal cluster data for the cluster. This cluster state information is held in volatile memory. If the mainline power fails, the two internal battery backup units maintain the internal power long enough for the cluster state information to be stored on the IDE disk drive of each node. The read and write cache information, which is also held in memory, is stored on the IDE disk drives of the nodes in the I/O group that are using that information.

All nodes in the cluster keep an identical copy of the cluster state. When a change is made to the configuration or internal cluster data the same change is applied to all nodes. For example, a user configuration request is made to the configuration node. This node forwards the request to all nodes in the cluster and they all make the change to the cluster state at the same point in time. This ensures that all nodes are aware of the configuration change.

Cluster operation and quorum disks

The cluster must contain at least half of its nodes to function. That is, when the cluster is formed and becomes stable, only half of the nodes must remain functional for the cluster to continue operating.

The cluster automatically chooses three managed disks to be **quorum disks** and assigns them quorum indices of 0, 1 and 2. One of these disks is used to settle a tie-break condition.

If a tie-break occurs, the first half of the cluster to access the quorum disk after the split has occurred locks the disk and continues to operate. The other side stops. This action prevents both sides from becoming inconsistent with each other.

You can change the assignment of quorum disks at any time by typing the `svctask setquorum` command.

Internal battery backup and power domains

The nodes on a CSM are protected by an internal battery backup unit. If the mainline power fails, the internal battery backup maintains the internal power long enough for the cluster state information to be stored on the IDE disk drive of each node.

Each CSM in the switch is a unique power domain. It is very important that the nodes in the I/O group are not part of the the same power domain. This configuration ensures that the cache and cluster state information is protected against the failure of the internal battery backup or of the mainline power source. Therefore each node of an I/O group must reside on a different CSM. Otherwise, a single failure could result in both nodes in an IO group being unable to write cache and cluster state information to the IDE disk drives causing the loss of cache data.

When nodes are added to the cluster, the I/O group they will join must be specified. The configuration interfaces ensure that the two nodes in the I/O group are not part of the same CSM.

Disk controllers

A disk controller is a device that coordinates and controls the operation of one or more disk drives and synchronizes the operation of the drives with the operation of the system as a whole. These controllers provide the storage that the cluster detects as managed disks (MDisks).

When configuring your disk controllers, ensure that you configure and manage your disk controllers and its devices for optimal performance.

The supported RAID controllers are detected by the cluster and reported by the user interfaces. The cluster can also determine which MDisks each controller has and can provide a view of MDisks filtered by controller. This view enables you to associate the MDisks with the RAID arrays that the controller presents.

Note: The SAN Volume Controller Storage Software for Cisco MDS 9000 supports RAID controllers, but it is possible to configure a controller as a non-RAID controller. RAID controllers provide redundancy at the disk level. Therefore, a single physical disk failure does not cause an MDisk failure, an MDisk group failure, or a failure in the virtual disks (VDisks) that were created from the MDisk group.

The controller may have a local name for the RAID arrays or single disks that it is providing. However it is not possible for the nodes in the cluster to determine this name as the namespace is local to the controller. The controller will surface these disks with a unique ID, the controller LUN or LU number. This ID, along with the controller serial number or numbers (there may be more than one controller), can be used to associate the managed disks in the cluster with the RAID arrays presented by the controllers.

To prevent loss of data, virtualize only those RAID arrays that provide some form of redundancy, that is, RAID 1, RAID 10, RAID 0+1 or RAID 5. Do not use RAID 0 because a single physical disk failure might cause the failure of many VDIs.

Unsupported disk controller systems (generic controller):

When a disk controller system is detected on the SAN, the SAN Volume Controller Storage Software for Cisco MDS 9000 attempts to recognize it using its Inquiry data. If the disk controller system is recognized as one of the explicitly supported storage models, then the SAN Volume Controller Storage Software for Cisco MDS 9000 uses error recovery programs that can be tailored to the known needs of the disk controller system. If the storage controller is not recognized, then the SAN Volume Controller Storage Software for Cisco MDS 9000 configures the disk controller system as a generic controller. A generic controller may or may not function correctly when addressed by a SAN Volume Controller Storage Software for Cisco MDS 9000. In any event, the SAN Volume Controller Storage Software for Cisco MDS 9000 does not regard accessing a generic controller as an error condition and, consequently, does not log an error. MDisks presented by generic controllers are not eligible to be used as quorum disks.

Migration

Migration changes how the mapping of the extents for a virtual disk (VDisk) are mapped to the extents for a managed disk (MDisk). Your host can still access the VDisk during this process.

Migration applications:

There are several uses for migration:

- Redistribution of workload within a cluster across your managed disks:
 - Movement of workload onto newly installed storage
 - Moving workload off of old or failing storage, ahead of replacing it
 - Movement of workload to re-balance a changed workload
- Migrating data from legacy disks to MDisks managed by the SAN Volume Controller Storage Software for Cisco MDS 9000.

Image mode virtual disk (VDisk) migration

Image mode VDIsks have the special property that the last extent in the VDisk can be a partial extent. Managed mode disks do not have this property.

Once your data is migrated off a partial extent, you cannot migrate data back onto the partial extent.

Copy Services

The SAN Volume Controller Storage Software for Cisco MDS 9000 provides two Copy Services that enable you to copy virtual disks (VDisks): FlashCopy™ and Remote Copy. These Copy Services are available for all supported hosts that are connected to the SAN Volume Controller Storage Software for Cisco MDS 9000.

FlashCopy

Makes an instant, point-in-time copy from a source VDisk to a target VDisk.

Remote Copy

Provides a consistent copy of a source VDisk on a target VDisk. Data is written to the target VDisk synchronously after it is written to the source VDisk, so the copy is continuously updated.

FlashCopy applications:

You can use FlashCopy to backup data that changes frequently. After creating the point-in-time copy, it can be backed up to tertiary storage such as tape.

Another use of FlashCopy is for application testing. It is often important and useful to test a new version of an application on real business data before you move the application into production. This reduces the risk that the new application will fail because it is not compatible with actual business data.

You can also use FlashCopy to create copies for your auditing and data mining purposes.

In the scientific and technical arena, FlashCopy can create restart points for long-running batch jobs. Therefore, if a batch job fails many days into its run, you might be able to restart the job from a saved copy of its data. This is preferable to rerunning the entire multiday job.

Remote Copy applications:

Disaster recovery is the primary application for Remote Copy. Because an exact copy of your business data can be maintained at a remote location, you can use your remote location as a recovery site in the event of a local disaster.

Related topics:

- “FlashCopy”
- “Remote Copy” on page 23

FlashCopy

FlashCopy is a Copy Service available with the SAN Volume Controller Storage Software for Cisco MDS 9000. It copies the contents of a source virtual disk (VDisk) to a target VDisk. Any data that existed on the target disk is lost and is replaced by the copied data. After the copy operation has been completed, the target virtual disks contain the contents of the source virtual disks as they existed at a single point in time. Although the copy operation takes some time to complete, the resulting data on the target is presented in such a way that the copy appears to have occurred immediately. FlashCopy is sometimes described as an instance of a Time-Zero copy (T 0) or point-in-time copy technology.

Source VDIsks and target VDIsks must meet the following requirements:

- They must be the same size.
- The same cluster must manage them.

To copy a VDisk, it must be part of a FlashCopy mapping or of a consistency group.

It is difficult to make a consistent copy of a data set which is being constantly updated. Point-in-time copy techniques are used to help solve the problem. If a copy of a data set is taken using a technology that does not provide point in time techniques and the data set changes during the copy operation, then the resulting copy may contain data which is not consistent. For example, if a reference to an object is copied earlier than the object itself and the object is moved before it is itself copied then the copy will contain the referenced object at its new location but the reference will point to the old location.

Related topics:

- “FlashCopy consistency groups” on page 22
- “FlashCopy mappings”
- “Virtual disks (VDisks)” on page 33

FlashCopy mappings

Because FlashCopy copies one VDisk to another VDisk, the SAN Volume Controller Storage Software for Cisco MDS 9000 needs to be aware of that relationship. A FlashCopy mapping defines the relationship between a source VDisk and a target VDisk. A particular virtual disk can take part in only one mapping; that is, a virtual disk can be the source or target of only one mapping. You cannot, for example, make the target of one mapping the source of another mapping.

FlashCopy makes an instant copy of a virtual disk at the time it is started. To create a FlashCopy of a virtual disk, you must first create a mapping between the source virtual disk (the disk that is copied) and the target virtual disk (the disk that receives the copy). The source and target must be of equal size.

A FlashCopy mapping can be created between any two virtual disks in a cluster. It is not necessary for the virtual disks to be in the same I/O group or managed disk group. When a FlashCopy operation is started, a checkpoint is made of the source virtual disk. No data is actually copied at the time a start occurs. Instead, the checkpoint creates a bitmap that indicates that no part of the source virtual disk has yet been copied. Each bit in the bitmap represents one region of the source virtual disk. Such a region is called a grain.

After a FlashCopy operation has been started, read operations to the source virtual disk operate as normal; that is, they read data from the source virtual disk. If data is written to the source virtual disk, the existing data is copied to the target virtual disk before the new data is written to the source virtual disk. The bitmap is updated to mark that the grain of the source virtual disk has been copied so that later write operations to the same grain do not recopy the data.

It is also possible to read and write to the target virtual disk. Read operations to the target virtual disk use the bitmap to determine whether the grain has been copied or not. If the grain has been copied, data is read from the target virtual disk, otherwise data is read from the source. A write operation to the target virtual disk is similar to a target virtual disk before the new data is written to the target virtual disk. The bitmap is updated to mark that the grain of the source virtual disk has been copied.

When you create a mapping, you specify the background copy rate. This rate determines the priority that is given to the background copy process. If you want to end with a copy of the whole source at the target (so that the mapping can be deleted, but the copy can still be accessed at the target), you must copy to the target virtual disk all the data that is on the source virtual disk. When a mapping is started and the background copy rate is greater than zero, the unchanged data is copied to the target, and the bitmap is updated to show that the copy has occurred. After a time, the length of which depends on the priority given and the size of the virtual disk, the whole virtual disk is copied to the target. The mapping returns to the idle/copied state. You can restart the mapping at any time to create a new copy at the target; the process copy starts again.

If the background copy rate is zero, only the data that changes on the source is copied to the target. The target never contains a copy of the whole source unless every extent is overwritten at the source. You can use this copy rate when you need only a temporary copy of the source.

You can stop the mapping at any time after it has been started. This action makes the target inconsistent and therefore the target virtual disk is taken offline. You must restart the mapping to correct the target.

FlashCopy mapping states:

At any point in time, a FlashCopy mapping is in one of the following states:

Idle or copied

The source and target VDisks act as independent VDisks even if a FlashCopy mapping exists between the two. Read and write caching is enabled for both the source and the target.

Copying

The copy is in progress.

Prepared

The mapping is ready to start. While in this state, the target VDisk is offline.

Preparing

Any changed write data for the source VDisk is flushed from the cache. Any read or write data for the target VDisk is discarded from the cache.

Stopped

The mapping is stopped because either you issued a command or an input/output (I/O) error occurred. Preparing and starting the mapping again can restart the copy.

Suspended

The mapping started, but it did not complete. The source VDisk might be unavailable, or the copy bitmap might be offline. If the mapping does not return to the copying state, stop the mapping to reset the mapping.

Before you start the mapping, you must prepare it. By preparing the mapping, you ensure that the data in the cache is destaged to disk and that a consistent copy of the source exists on disk. At this time the cache goes into write-through mode. That is, data that is written to the source is not cached in the SAN Volume Controller Storage Software for Cisco MDS 9000s; it passes straight through to the managed disks. The prepare operation for the mapping might take you a few minutes; the actual length of time depends on the size of the source virtual disk. You must coordinate the prepare operation with the operating system. Depending on the type of data that is on the source virtual disk, the operating system or application software might also cache data write operations. You must flush, or synchronize, the file system and application program before you prepare for, and finally start, the mapping.

For customers who do not need the complexity of consistency groups, the SAN Volume Controller Storage Software for Cisco MDS 9000 allows a FlashCopy mapping to be treated as an independent entity. In this case the FlashCopy mapping is known as a stand alone mapping. For FlashCopy mappings which have been configured in this way, the **Prepare** and **Start** commands are directed at the FlashCopy mapping name rather than the consistency group ID.

Veritas Volume Manager:

For FlashCopy target VDisks, the SAN Volume Controller Storage Software for Cisco MDS 9000 sets a bit in the inquiry data for those mapping states where the target VDisk could be an exact image of the source VDisk. Setting this bit enables the Veritas Volume Manager to distinguish between the source and target VDisks and thus provide independent access to both.

Related topics:

- “FlashCopy” on page 19
- “FlashCopy consistency groups”
- “Virtual disks (VDisks)” on page 33

FlashCopy consistency groups

This topic provides an overview of FlashCopy consistency groups.

When you copy data from one virtual disk (VDisk) to another, that data might not include all that you need to enable you to use the copy. Many applications have data that spans multiple VDisks and that include the requirement that data integrity is preserved across VDisks. For example, the logs for a particular database usually reside on a different VDisk than the VDisk that contains the data.

Consistency groups address the problem when applications have related data that spans multiple VDisks. In this situation, FlashCopy must be performed in a way that preserves data integrity across the multiple VDisks. One requirement for preserving the integrity of data being written is to ensure that dependent writes are run in the intended sequence of the application.

A consistency group is a container for mappings. You can add many mappings to a consistency group. The consistency group is specified when the mapping is created. You can also change the consistency group later. When you use a consistency group, you prepare and trigger that group instead of the various mappings. This ensures that a consistent copy is made of all the source VDisks. Mappings that you want to control at an individual level instead of at a consistency group level, should not be put into a consistency group. These mappings are known as stand-alone mappings.

FlashCopy consistency-group states:

At any point in time, a FlashCopy consistency group is in one of the following states:

Idle or copied

The source and target VDisks act independently even if a FlashCopy consistency group exists. Read and write caching is enabled for the source VDisks and target VDisks.

Copying

The copy is in progress.

Prepared

The consistency group is ready to start. While in this state, the target VDisks are offline.

Preparing

Any changed write data for the source VDisks is flushed from the cache. Any read or write data for the target VDisks is discarded from the cache.

Stopped

The consistency group is stopped because either you issued a command or an input/output (I/O) error occurred. Preparing and starting the consistency group again can restart the copy.

Suspended

The consistency group was started, but it did not complete. The source VDisks might be unavailable, or the copy bitmap might be offline. If the consistency group does not return to the copying state, stop the consistency group to reset the consistency group.

Related topics:

- “FlashCopy” on page 19
- “FlashCopy mappings” on page 20
- “Virtual disks (VDisks)” on page 33

Remote Copy

Remote Copy is a copy service available with the SAN Volume Controller Storage Software for Cisco MDS 9000. Remote Copy enables you to set up a relationship between two virtual disks, so that updates that are made by an application to one virtual disk are mirrored to the other virtual disk. Although the application only writes to a single virtual disk, the SAN Volume Controller Storage Software for Cisco MDS 9000 is effectively maintaining two copies of this data. If the copies are separated by a significant distance, then the remote copy may be useful in disaster recovery scenarios. A pre-requisite for the SAN Volume Controller Storage Software for Cisco MDS 9000 Remote Copy operations between two clusters is that the SAN fabric to which they are attached provides adequate bandwidth between the clusters.

One VDisk is designated the primary and the other VDisk is designated the secondary. Host applications write data to the primary VDisk, and updates to the primary VDisk are copied to the target VDisk. Normally, host applications do not perform input or output operations to the secondary VDisk.

Remote Copy supports the following features:

- Intracluster copying of a VDisk, in which both VDisks belong to the same cluster and I/O group within the cluster.
- Intercluster copying of a VDisk, in which one VDisk belongs to a cluster and the other VDisk belongs to a different cluster

Note: A cluster can only participate in active Remote Copy relationships with itself and with a single other cluster.

- Intercluster and intracluster Remote Copy can be used concurrently within a cluster for different relationships

Synchronous Remote Copy

In the synchronous mode, Remote Copy provides a *consistent* copy, which means that the target VDisk is always the exact match of the source VDisk. The host application writes data to the source VDisk but does not receive the final status on the write operation until the data is written to the target VDisk. For disaster recovery, this mode is the only practical mode of operation because a consistent copy of the data is maintained. However, synchronous mode is slower than asynchronous mode because of the latency time and bandwidth limitations imposed by the communication link to the secondary site.

Related topics:

- “Remote Copy” on page 23

Remote Copy consistency groups

Certain uses of Remote Copy require the manipulation of more than one relationship. Remote Copy provides the ability to group relationships so that they are manipulated in unison. To address this requirement, consistency groups have been created.

For some uses it might be that the relationships share some loose association and that the grouping simply provides a convenience for the administrator. But a more significant use arises when the relationships contain VDisks that have a tighter association. One example is when the data for an application is spread across more than one VDisk. A more complex example is when multiple applications run on different host systems. Each application has data on different VDisks, and these applications exchange data with each other. Both these examples are cases in which specific rules exist as to how the relationships must be manipulated, in unison. This ensures that the set of secondary VDisks contains usable data. The key property is that these relationships be consistent. Hence, the groups are called consistency groups.

A relationship can be part of a single consistency group or not be part of a consistency group at all. Relationships that are not part of a consistency group are called stand-alone relationships. A consistency group can contain zero or more relationships. All the relationships in a consistency group must have matching master and auxiliary clusters. All relationships in a consistency group must also have the same copy direction and state.

Remote Copy consistency group states:

Inconsistent (Stopped)

The primary VDisks are accessible for read and write input/output (I/O) operations but the secondary VDisks are not accessible for either. A copy process needs to be started to make the Secondary VDisks consistent.

Inconsistent (Copying)

The primary VDisks are accessible for read and write I/O operations but the secondary VDisk are not accessible for either. This state is entered after a **Start** command is issued to an consistency group in the InconsistentStopped state. This state is also entered when a **Start** command is issued, with the force option, to a consistency group in the Idling or ConsistentStopped state.

Consistent (Stopped)

The secondary VDisks contain a consistent image, but it might be out-of-date with respect to the primary VDisks. This state can happen when a relationship was in the ConsistentSynchronized state and experiences an error which forces a freeze of the consistency group. This state can also happen when a relationship is created with the CreateConsistentFlag set to TRUE.

Consistent (Synchronized)

The primary VDisks are accessible for read and write I/O operations. The secondary VDisks are accessible for read-only I/O operations.

Idling Master VDisks and Auxiliary VDisks are operating in the primary role. Consequently the VDisks are accessible for write I/O operations.

Idling (Disconnected)

The VDIs in this half of the consistency group are all operating in the primary role and can accept read or write I/O operations.

Inconsistent (Disconnected)

The VDIs in this half of the consistency group are all operating in the secondary role and will not accept read or write I/O operations.

Consistent (Disconnected)

The VDIs in this half of the consistency group are all operating in the secondary role and will accept read I/O operations but not write I/O operations

Empty The consistency group contains no relationships.

Related topics:

- “Remote Copy” on page 23
- “Virtual disks (VDIs)” on page 33

Chapter 4. Object descriptions

This topic describes the objects in a SAN Volume Controller Storage Software for Cisco MDS 9000 and the relationships between these objects.

The SAN Volume Controller Storage Software for Cisco MDS 9000 is based on the following virtualization concepts which are discussed more fully later in this chapter.

The smallest processing unit in a SAN Volume Controller Storage Software for Cisco MDS 9000 is a single node. There are 4 nodes in a SAN Volume Controller Storage Software for Cisco MDS 9000 cluster. Each pair of nodes is known as an **I/O group**. Each node may be in only one I/O group.

Virtual disks (Vdisks) are logical disks that are presented to the SAN by the SAN Volume Controller Storage Software for Cisco MDS 9000. Each virtual disk is associated with a particular I/O group. The nodes in the I/O group provide access to the virtual disks in the I/O group. When an application server performs I/O to a virtual disk, it has the choice of accessing the virtual disk via either of the nodes in the I/O group. As each I/O group only has two nodes, the distributed cache the SAN Volume Controller Storage Software for Cisco MDS 9000 provides is only 2-way.

Each Cisco MDS 9000 CSM contains an internal battery backup unit to provide data integrity in the event of a power failure.

The nodes in a cluster see the storage presented by backend **disk controllers** as a number of disks, known as **managed disks (MDisks)**. Because the SAN Volume Controller Storage Software for Cisco MDS 9000 does not attempt to provide recovery from physical disk failures within the backend disk controllers, a managed disk is usually, but not necessarily, a RAID array.

Each managed disk is divided up into a number of **extents** (default size is 16MB) which are numbered, from 0, sequentially from the start to the end of the managed disk.

Managed disks are collected into groups, known as **managed disk groups (MDisk group)**. Virtual disks are created from the extents contained by a managed disk group. The managed disks that constitute a particular virtual disk must all come from the same managed disk group.

At any one time, a single node in the cluster is used to manage configuration activity. This **configuration node** manages a cache of the information that describes the cluster configuration and provides a focal point for configuration.

The SAN Volume Controller Storage Software for Cisco MDS 9000 allows you to create logical host objects that group together WWPNs belonging to a single application server or a set of them.

Application servers can only access virtual disks that have been allocated to them. Virtual disks can then be mapped to a host object. The act of mapping a virtual disk to a host object makes the virtual disk accessible to the WWPNs in that host object, and hence the application server itself.

The SAN Volume Controller Storage Software for Cisco MDS 9000 provides block level aggregation and volume management for disk storage within the SAN. In simpler terms, this means that the SAN Volume Controller Storage Software for Cisco MDS 9000 manages a number of back-end storage controllers and maps the physical storage within those controllers into logical disk images that can be seen by application servers and workstations in the SAN. The SAN is configured in such a way that the application servers cannot see the back-end physical storage; this prevents any possible conflict between the SAN Volume Controller Storage Software for Cisco MDS 9000 and the application servers both trying to manage the back-end storage.

Items such as virtual disks, managed disks, and managed disk groups are known as objects. It is important that you understand what these objects are, how you use them, and their relationships to one another, so that you can effectively configure and use the SAN Volume Controller Storage Software for Cisco MDS 9000.

Disk controller systems

A disk controller system is a device that coordinates and controls the operation of one or more disk drives and synchronizes the operation of the drives with the operation of the system as a whole.

Disk controller systems attached to the SAN fabric provide the physical storage that the cluster detects as managed disks. These are usually RAID controllers as the SAN Volume Controller Storage Software for Cisco MDS 9000 does not attempt to provide recovery from physical disk failures within the back-end controllers. The nodes in the cluster can be in only one Fibre Channel SAN or in up to 4 virtual SANs (VSANs).

The supported RAID controllers are detected by the cluster and reported by the user interfaces. The cluster can also determine which managed disks each controller is presenting, and can provide a view of managed disks filtered by controller. This allows you to associate the managed disks with the RAID arrays that the controller presents.

The controller may have a local name for the RAID arrays or single disks that it is providing. However it is not possible for the nodes in the cluster to determine this name as the namespace is local to the controller. The controller will surface these disks with a unique ID, the controller LUN number. This ID, along with the controller serial number or numbers (there may be more than one controller), can be used to associate the managed disks in the cluster with the RAID arrays presented by the controller.

Back-end disk controller systems present storage to other devices on the SAN. The physical storage associated with a back-end controller is normally configured into RAID arrays which provide recovery from physical disk failures. Some back-end disk controller systems also allow physical storage to be configured as RAID-0 arrays (striping) or as JBODs; however, this does not provide protection against a physical disk failure and with virtualization can lead to the failure of many virtual disk.

Many back-end disk controller systems allow the storage provided by a RAID array to be divided up into many SCSI LUs which are presented on the SAN. With the SAN Volume Controller Storage Software for Cisco MDS 9000 it is recommended that back-end disk controller systems are configured to present each RAID array as a single SCSI LU which will be recognized by the SAN Volume

Controller Storage Software for Cisco MDS 9000 as a single managed disk. The virtualization features of the SAN Volume Controller Storage Software for Cisco MDS 9000 can then be used to divide up the storage into virtual disks.

Some back-end controller systems allow the storage to be increased in size. The SAN Volume Controller Storage Software for Cisco MDS 9000 will not use this extra capacity. Instead of increasing the size of an existing managed disk, a new managed disk should be added to the managed disk group and the extra capacity will be available for the SAN Volume Controller Storage Software for Cisco MDS 9000 to use.

Attention: If you delete a RAID that is being used by the SAN Volume Controller Storage Software for Cisco MDS 9000, the MDisk group will go offline and the data in that group will be lost.

When configuring your disk controller systems, ensure that you configure and manage your disk controller systems and their devices for optimal performance.

The cluster detects and provides a view of the disk controllers that the SAN Volume Controller Storage Software for Cisco MDS 9000 supports. The cluster can also determine which MDisks each controller has and can provide a view of MDisks filtered by controller. This view enables you to associate the MDisks with the RAID arrays that the controller presents.

Note: The SAN Volume Controller Storage Software for Cisco MDS 9000 supports RAID controllers, but it is possible to configure a controller as a non-RAID controller. RAID controllers provide redundancy at the disk level. Therefore, a single physical disk failure does not cause an MDisk failure, an MDisk group failure, or a failure in the virtual disks (VDisks) that were created from the MDisk group.

Disk controller systems reside on the SAN fabric and are addressable by one or more fibre-channel worldwide ports. Each port has a unique name known as a worldwide port name (WWPN).

Related topics:

- “Managed disks (MDisks)”
- “Managed disk (MDisk) groups” on page 31
- “Virtual disks (VDisks)” on page 33

Managed disks (MDisks)

A managed disk (MDisk) is a logical disk (typically a RAID array or partition thereof) that a storage controller has offered on the SAN fabric that the nodes in the cluster are attached. A managed disk may therefore consist of multiple physical disks that are presented as a single logical disk to the SAN. A managed disk always provides usable blocks of physical storage to the cluster whether it has a one to one correspondence with a physical disk or not.

Each managed disk is divided up into a number of **extents** which are numbered, from 0, sequentially from the start to the end of the managed disk. The extent size is a property of managed disk groups. When an MDisk is added to an MDisk group, the size of the extents the MDisk will be broken into depends on the attribute of the Mdisk group it has been added to.

Access modes:

The access mode determines how the cluster uses the MDisk. The possible modes are:

Unmanaged

The MDisk is not used by the cluster.

Managed

The MDisk is assigned to an MDisk group and is providing extents that virtual disks (VDisks) can use.

Image The MDisk is assigned directly to a VDisk with a one-to-one mapping of extents between the MDisk and the VDisk.

Attention: If you add a managed disk that contains existing data to a managed disk group, you will lose the data that it contains. The only mode that will preserve this data is **image** mode.

The status of a managed disk consists of four settings. The following table describes the different states of a managed disk:

Table 5. Managed Disk status

Status	Description
Online	The MDisk can be accessed by all online nodes. That is, all the nodes that are currently working members of the cluster can access this MDisk. The MDisk will appear online when the following conditions are met: <ul style="list-style-type: none">• All timeout error recovery procedures completed and reported the disk as online.• LUN inventory of the target ports correctly reported the MDisk.• Discovery of this LUN created successfully.• All of the managed disk controller ports report this LUN as available with no fault conditions.
Degraded	The MDisk cannot be accessed by all the online nodes. That is, one or more (but not all) of the nodes that are currently working members of the cluster cannot access this MDisk. The MDisk may be partially excluded, that is some of the paths to the MDisk (but not all) have been excluded.
Excluded	The MDisk has been excluded from use by the cluster after repeated access errors. You can reset an MDisk and re-include it in the cluster by running the svctask includemdisk command.
Offline	The MDisk cannot be accessed by any of the online nodes. That is, all of the nodes that are currently working members of the cluster cannot access this MDisk. This state can be caused by a failure in the SAN, the RAID controller, or one or more physical disks connected to the RAID controller. The MDisk will only be reported as offline if all paths to the disk fail.

Extents:

Each MDisk is divided into chunks of equal size called extents. Extents manage the mapping of data between MDisks and virtual disks (VDisks).

Attention: If your fabric is undergoing transient link breaks or you have been replacing cables or connections in your fabric, you may see one or more MDisk change to the degraded status. If I/O was attempted during the link breaks and the same I/O was failed several times the MDisk will be partially excluded and will change to a status of degraded. You should include the MDisk to resolve the problem. You can include the MDisk by either selecting the Include MDisk task from the **Work with Managed Disks - Managed Disk** panel in the Console for SAN Volume Controller for Cisco MDS 9000, or issue the following command:

```
svctask includemdisk <mdiskname/id>
```

Managed disk path Each managed disk will have an online path count which is the number of nodes that have access to that managed disk; this represents a summary of the I/O path status between the cluster nodes and the particular controller device. The maximum path count is the maximum number of paths that have been detected by the cluster at any point in the past. Thus if the current path count is not equal to the maximum path count then the particular managed disk may be degraded. That is, one or more nodes may not see the managed disk on the fabric.

Related topics:

- “Disk controller systems” on page 28

Managed disk (MDisk) groups

An MDisk group is a collection of MDisks that jointly contain all the data for a specified set of virtual disks (VDisks). All MDisks in a group are split into extents of the same size. VDisks are created from the extents that are available in the group. You can add MDisks to an MDisk group at any time. This way you increase the number of extents that are available for new VDisks or to expand existing VDisks.

You can add MDisks to an MDisk group at any time either to increase the number of extents that are available for new VDisks, or to expand existing VDisks. You can add only MDisks that are in unmanaged mode. When MDisks are added to a group, their mode changes from unmanaged to managed.

You can delete MDisks from a group under these conditions:

- VDisks are not using any of the extents that are on the MDisk.
- Enough free extents are available elsewhere in the group to move any extents that are in use from this MDisk.

Attention: If you delete an MDisk group, you destroy all the VDisks that are made from the extents that are in the group. If the group is deleted, you cannot recover the mapping that existed between extents that are in the group and the extents that VDisks use. The MDisks that were in the group are returned to unmanaged mode, and can be added to other groups. Because the deletion of a group can cause a loss of data, you must force the delete if VDisks are associated with it.

The status of an MDisk group consists of three settings. The following table describes the different states of a MDisk group:

Table 6. Managed Disk Group status

Status	Description
--------	-------------

Table 6. Managed Disk Group status (continued)

Online	The MDisk group is online and available. All the MDisks in the group are available.
Degraded	The MDisk group is available; however, one or more nodes cannot access all the MDisks in the group.
Offline	The MDisk group is offline and unavailable. No nodes in the cluster can access the MDisks. The most likely cause is that one or more MDisks are offline or excluded.

Attention: If a single MDisk in an MDisk group is offline, that is, it cannot be seen by all of the online nodes in the cluster, the MDisk group that this MDisk is a member of goes offline. This causes *all* the VDIs that are being presented by this MDisk group to go offline. Care should be taken when creating MDisk groups to ensure an optimal configuration.

Consider the following guidelines when you create MDisk groups:

1. If you are creating image-mode VDIs, do not put all of these into one MDisk group because a single MDisk failure results in all of these VDIs going offline. Allocate your image-mode VDIs between your MDisk groups.
2. Ensure that all MDisks allocated to a single MDisk group are of the same RAID type. This ensures that a single failure of a physical disk in the controller does not take the entire group offline. For example, if you had three RAID-5 arrays in one group and added a non-RAID disk to this group, if the non-RAID disk fails then you lose access to all the data striped across the group. Similarly, for performance reasons you should not mix RAID types.
3. If you intend to keep the virtual disk allocation within the storage of one disk controller system, you should ensure that the MDisk group that corresponds with a single controller is presented by that controller. This also enables non-disruptive migration of data from one controller to another controller and simplifies the decommissioning process should you wish to decommission a controller at a later time.

Extent:

To track the space that is available, the SAN Volume Controller Storage Software for Cisco MDS 9000 divides each MDisk in an MDisk group into chunks of equal size. These chunks are called extents, and are indexed internally. Extent sizes can be 16, 32, 64, 128, 256, or 512 MB.

You must specify the extent size when you create a new MDisk group. You cannot change the extent size later; it must remain constant throughout the lifetime of the MDisk group. MDisk groups can have different extent sizes, however different extent sizes can place restrictions on the use of data migration. The choice of extent size affects the total amount of storage that can be managed by a SAN Volume Controller Storage Software for Cisco MDS 9000 cluster. Table 7 on page 33 shows the maximum amount of storage that can be managed by a cluster for each extent size. Because the SAN Volume Controller Storage Software for Cisco MDS 9000 allocates a whole number of extents to each virtual disk that is created, using a larger extent size may increase the amount of wasted storage at the end of each virtual disk. Larger extent sizes also reduces the ability of the SAN Volume Controller Storage Software for Cisco MDS 9000 to distribute sequential I/O workloads across many managed disks. Therefore, larger extent sizes might reduce

the performance benefits of virtualization.

Table 7. Capacities of the cluster given extent size

Extent size	Maximum storage capacity of cluster
16 MB	64 TB
32 MB	128 TB
64 MB	256 TB
128 MB	512 TB
256 MB	1 PB
512 MB	2 PB

The figure below shows an MDisk group containing four MDisks.




Figure 3. MDisk group

Related topics:

- "Managed disks (MDisks)" on page 29
- "Virtual disks (VDisks)"

Virtual disks (VDisks)

A VDisk is a logical disk that the cluster presents to the storage area network (SAN). Application servers on the SAN access VDIs, not managed disks (MDisks). VDIs are created from a set of extents in an MDisk group. There are three types of VDIs: striped, sequential, and image.

Types:

You can create the following types of VDIs:

Striped

The striping is at extent level. One extent is allocated, in turn, from each managed disk that is in the group. For example, a managed disk group that has 10 MDisks takes one extent from each managed disk. The 11th extent is taken from the first managed disk and so on. This procedure, known as a round-robin, is similar to RAID-0 striping.

Attention: Care should be taken when specifying a stripe set if your MDisk group contains MDisks of unequal size. By default, striped VDIs are striped across all MDisks in the group. If some of the MDisks are smaller than others, the extents on the smaller MDisks will be used up before the larger MDisks run out of extents. Manually specifying the stripe set in this case, may result in the VDisk not being created.

If you are unsure about whether there is sufficient free space to create a striped VDisk you can:

- Check the free space on each MDisk in the group, using the `svcinfolsfreeextents` command
- Let the system automatically create the VDisk, by not supplying a specific stripe set.

You can also supply a list of MDisks to use as the stripe set. This list can contain two or more MDisks from the managed disk group. The round-robin procedure is used across the specified stripe set.

The figure below shows an example of a managed disk group containing three MDisks. This also shows a striped virtual disk created from the extents available in the group.

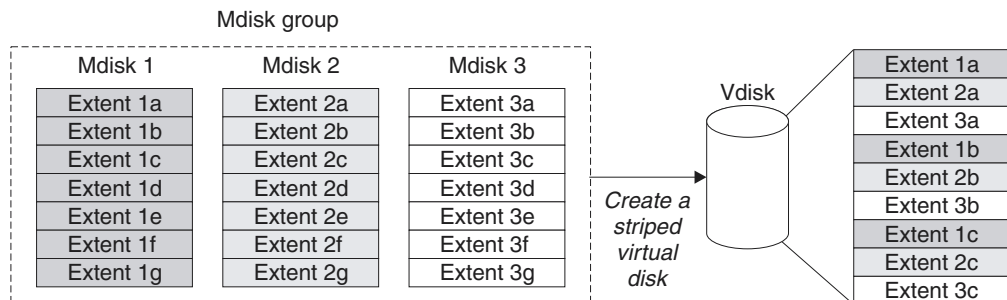


Figure 4. Managed disk groups and VDIs

Sequential

When selected, extents are allocated sequentially on one managed disk to create the virtual disk if enough consecutive free extents are available on the chosen managed disk.

Image Image-mode VDIs are special VDIs that have a direct relationship with one managed disk. If you have a RAID array that contains data that you want to merge into the cluster, you can create an image-mode virtual disk. When you create an image-mode virtual disk, a direct mapping is made between extents that are on the managed disk and extents that are on the virtual disk. The managed disk is not virtualized; meaning, a logical block address (LBA) where x is arbitrary on the managed disk is the same as LBA x on the virtual disk.

When you create an image-mode virtual disk, you must assign it to a managed disk group. An image-mode VDisk must be at least one extent in size. In other words, the minimum size of an image-mode VDisk is the extent size of the MDisk group to which it is assigned.

The extents are managed in the same way as other VDIs. When the extents have been created, you can move the data onto other MDisks that are in the group without losing access to the data. Once you move one or

more extents, the virtual disk becomes a real virtualized disk, and the mode of the managed disk changes from image to managed.

Attention: If you add an MDisk to an MDisk group as a managed disk, any data on the MDisk will be lost. Ensure that you create image mode VDIs from the MDisks that contain data before you start adding any MDisks to groups.

MDisks that contain existing data have an initial mode of unmanaged, and the cluster cannot determine whether they contain partitions or data.

The status of a virtual disk consists of three settings. The following table describes the different states of a virtual disk:

Table 8. Virtual Disk status

Status	Description
Online	The virtual disk is online and available if both nodes in the I/O group can access the virtual disk. A single node will only be able to access a VDisk if it can access all the MDisks in the MDisk group associated with the VDisk.
Offline	The VDisk is offline and unavailable if both nodes in the I/O group are missing or none of the nodes in the I/O group that are present can access the VDisk.
Degraded	The status of the virtual disk is degraded if one node in the I/O group is online and the other node is either missing or cannot access the virtual disk.

You can also use more sophisticated extent allocation policies to create VDIs. When you create a striped virtual disk, you can specify the same managed disk more than once in the list of MDisks that are used as the stripe set. This is useful if you have a managed disk group where not all the MDisks are of the same capacity. For example, if you have a managed disk group that has two 18 GB MDisks and two 36 GB MDisks, you can create a striped virtual disk by specifying each of the 36 GB MDisks twice in the stripe set so that two thirds of the storage is allocated from the 36 GB disks.

If you delete a virtual disk, you destroy access to the data that is on the virtual disk. The extents that were used in the virtual disk are returned to the pool of free extents that is in the managed disk group. The delete might fail if the virtual disk is still mapped to hosts. The delete may also fail if the virtual disk is still part of a FlashCopy or a Remote Copy mapping. If the delete fails, you can specify the force delete flag to delete both the virtual disk and the associated mappings to hosts. Forcing the deletion will also delete the copy services relationship and mappings.

Related topics:

- “Virtual disks (VDIs)” on page 33
- “Virtualization” on page 1

Virtual disk-to-host mapping

Virtual disk-to-host mapping is similar in concept to LUN mapping or masking. LUN mapping is the process of controlling which hosts have access to specific logical units (LUs) within the disk controllers. LUN mapping is typically done at the disk controller level. Virtual disk-to-host mapping is the process of controlling which hosts have access to specific virtual disks (VDisks) within the SAN Volume Controller Storage Software for Cisco MDS 9000. Virtual disk-to-host mapping is done at the SAN Volume Controller Storage Software for Cisco MDS 9000 level.

Application servers can only access VDisks that have been made accessible to them. The SAN Volume Controller Storage Software for Cisco MDS 9000 detects the fibre channel ports that are connected to the SAN. These correspond to the host bus adapter (HBA) worldwide port names (WWPNs) that are present in the application servers. The SAN Volume Controller Storage Software for Cisco MDS 9000 enables you to create logical hosts that group together WWPNs belonging to a single application server. VDisks can then be mapped to a host. The act of mapping a virtual disk to a host makes the virtual disk accessible to the WWPNs in that host, and hence the application server itself.

VDisks and host mappings:

The SAN concept known as LUN masking usually requires device driver software in each host. The device driver software masks the LUNs as instructed by the user. After the masking has been done, only some disks are visible to the operating system. The SAN Volume Controller Storage Software for Cisco MDS 9000 performs a similar function, but, by default, it presents to the host only those VDisks that are mapped to that host. You must therefore map the VDisks to the hosts that are to access those VDisks.

Each host mapping associates a virtual disk with a host object and allows all host HBA ports in the host object to access the virtual disk. You can map a virtual disk to multiple host objects. When a mapping is created, multiple paths might exist across the SAN fabric from the hosts to the SAN Volume Controller Storage Software for Cisco MDS 9000s that are presenting the virtual disk. Most operating systems present each path to a virtual disk as a separate storage device. The SAN Volume Controller Storage Software for Cisco MDS 9000, therefore, needs the IBM Subsystem Device Driver (SDD) software to be running on the host. This software handles the many paths that are available to the virtual disk and presents a single storage device to the operating system.

When you map a virtual disk to a host, you can optionally specify a SCSI ID for the virtual disk. This ID controls the sequence in which the VDisks are presented to the host. Take care when you specify a SCSI ID, because some device drivers stop looking for disks if they find an empty slot. For example, if you present three VDisks to the host, and those VDisks have SCSI IDs of 0, 1, and 3, the virtual disks that has an ID of 3 might not be found because no disk is mapped with an ID of 2. The cluster automatically assigns the next available SCSI ID if none is entered.

Figure 5 on page 37 and Figure 6 on page 37 show two VDisks, and the mappings that exist between the host objects and these VDisks.

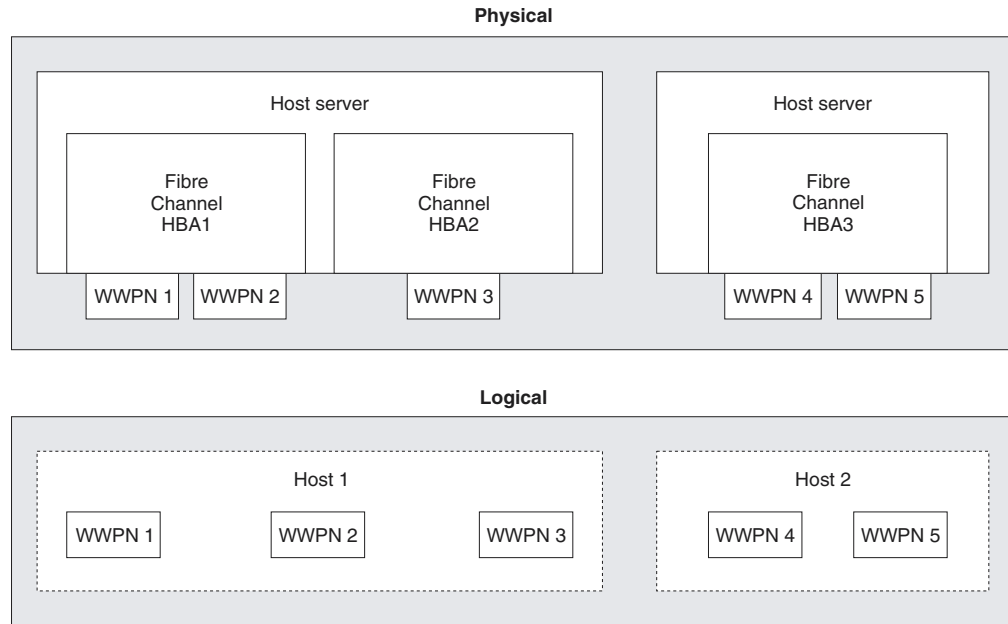


Figure 5. Hosts, WWPNs, and VDisks

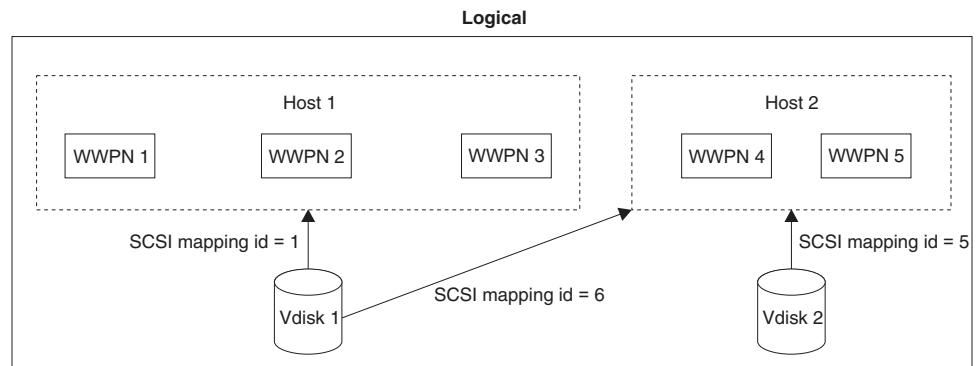


Figure 6. Hosts, WWPNs, VDisks and SCSI mappings

Related topics:

- “Managed disks (MDisks)” on page 29
- “Virtual disks (VDisks)” on page 33

Host objects

A host system is an open-systems computer that is connected to the SAN Volume Controller Storage Software for Cisco MDS 9000 through a fibre-channel interface. Creating a host in a cluster results in the creation of a logical host object. A logical host object has one or more worldwide port names (WWPNs) assigned to it. Generally, a logical host object is associated with a physical host system. However, a single logical host object can have WWPNs from multiple physical host systems that are assigned to it.

A host object is a logical object that groups one or more worldwide port names (WWPNs) of the host bus adapters (HBAs) that the cluster has detected on the SAN. A typical configuration has one host object for each host that is attached to

the SAN. If, however, a cluster of hosts is going to access the same storage, you can add HBA ports from several hosts into the one host object to make a simpler configuration.

The cluster does not automatically present VDisks on the Fibre Channel. You must map each virtual disk to a particular set of ports to enable the virtual disk to be accessed through those ports. The mapping is made between a host object and a virtual disk.

When you create a new host object, by typing the `svctask mkhost` command, the configuration interfaces provide a list of unconfigured WWPNs. These WWPNs represent the fibre channel ports that the cluster has detected.

The cluster can detect only ports that are logged into the fabric. Some HBA device drivers do not let the ports remain logged in if no disks are visible on the fabric. This condition causes a problem when you want to create a host because, at this time, no VDisks are mapped to the host. The configuration interface provides a method by which you can manually enter port names under this condition.

A port can be added to only one host object. When a port has been added to a host object, that port becomes a configured WWPN, and is not included in the list of ports that are available to be added to other hosts.

Node Login Counts:

This is the number of nodes that can see each port and is reported on a per node basis. If the count is less than the number of nodes in the cluster, then there is a fabric problem and not all nodes can see the port.

Chapter 5. Planning for your SAN Volume Controller Storage Software for Cisco MDS 9000 configuration

This topic describes planning tasks you need to perform before you configure the SAN Volume Controller Storage Software for Cisco MDS 9000.

Cluster:

Before you create a SAN Volume Controller Storage Software for Cisco MDS 9000 cluster:

- Define the number of clusters. Each cluster contains 4 nodes. Each pair of nodes (the I/O group) is the container for one or more VDIs.
- Determine the number of hosts you plan to use with the SAN Volume Controller Storage Software for Cisco MDS 9000. Hosts should be grouped by OS and by HBA type.

Host Group:

Hosts have access to specific logical units (LUs) within the disk controllers as a result of LUN masking. Before you create a host group, you must gather the following information:

- List all of the World Wide Port Names (WWPNs) of the fibre channel HBA ports in the hosts.
- Decide on the name you want to assign to the host or host group.
- Decide on the VDIs you want to assign to the host.

Managed disk:

You must define the logical or physical disks (Logical Units) in the backend storage.

Managed disk group:

Before you create your MDisk groups, you must determine the following factors:

- Determine the types of backend controllers you will be using.
- If you want to create VDIs with the sequential policy, either create a separate MDisk group for these VDIs or ensure that you create these VDIs before creating VDIs with the striped policy.
- Create MDisk groups for the backend controllers that provide the same level of performance or reliability, or both. You can group all of the managed disks that are RAID 10 in one MDisk group, all of the MDIs that are RAID 5 in another group.

Virtual disk:

An individual virtual disk is a member of one managed disk group and one I/O group. The managed disk group defines which managed disks provide the backend storage that makes up the virtual disk. The I/O group defines which SAN

Volume Controller Storage Software for Cisco MDS 9000 nodes provide I/O access to the virtual disk. You must determine the following information before creating a virtual disk:

- The name you want to assign to the virtual disk.
- The I/O group you want to use.
- The managed disk group you want to use.
- The capacity you want to define.

Related topics:

- “Configuration rules and requirements” on page 42
- “Configuration rules” on page 43

Maximum configuration

The following table shows the maximum configuration values to be used when planning for your SAN Volume Controller Storage Software for Cisco MDS 9000 installation.

Table 9. SAN Volume Controller Storage Software for Cisco MDS 9000 maximum configuration values

Objects	Maximum number	Comments
Cluster Properties		
Nodes	4	Arranged as pairs
I/O groups	2	
MDisk group	128	
MDisks	4096	Represents an average of 64 per controller
Object MDisks per MDisk group	128	
MDisk size	2 TB	Defined by 32 bit LBA limit
Addressability	2.1 PB	Maximum extent size 512 MB, arbitrary limit of 2 ²² extents in map
LU size	2 TB	Defined by 32 bit LBA limit.
Concurrent commands per node	2500	Assumes a backend latency of 100ms.
Concurrent commands per FC port	2048	
SDD	512 SAN Volume Controller vpaths per host	<p>One vpath is created for each VDisk mapped to a host. Although the SAN Volume Controller Storage Software for Cisco MDS 9000 only permits 512 VDIs to be mapped to a host, the SDD limit can be exceeded by either:</p> <ul style="list-style-type: none"> • Creating two (or more) host objects for one physical host and mapping more than 512 VDIs to the host using the multiple host objects. • Creating two (or more) clusters and mapping more than 512 VDIs to the host using the multiple clusters. <p>Note: Both of these operations are unsupported.</p>

Table 9. SAN Volume Controller Storage Software for Cisco MDS 9000 maximum configuration values (continued)

Objects	Maximum number	Comments
VDisks per MDisk Group		Cluster limit applies
Front-end Properties		
SAN ports	256	Maximum size of fabric, including all SAN Volume Controller Storage Software for Cisco MDS 9000 nodes
Host IDs	64	A host ID is associated with a map table which associates SCSI LUNs with VDIs. It is also associated with one or more host WWPNs
Host ports	128	Up to 128 distinct Host Worldwide Port Names (WWPNs) will be recognized
Host LUN size	2 TB	Defined by 32 bit LBA limit
Virtual disks (VDIs)	1024	Includes managed-mode VDIs and image-mode VDIs
VDIs per Host ID	512	Note: The limit may be different based on host operating system. For HP/UX, the maximum configuration is 8 VDIs per HP/UX host.
VDIs-to-host Mappings	20 000	
Maximum persistent reservation keys	132 000	
Back-end Properties		
Disk controller systems	64	
Disk controller ports	256	16 ports per controller.
LUs per disk controller system	512	
WWNNs per subsystem	4	
WWPNs per WWNN	16	Maximum number of ports per WWNN
MDisks per subsystem		Cluster limit applies
Copy Services Properties		
Data Migration operations	32	
Remote Copy relationships	128	
Remote Copy consistency groups	32	
Remote Copy VDisk per I/O group	16 TB	
FlashCopy mappings	128	
FlashCopy consistency groups	128	

Table 9. SAN Volume Controller Storage Software for Cisco MDS 9000 maximum configuration values (continued)

Objects	Maximum number	Comments
FlashCopy VDisk per I/O group	16 TB	

Configuration rules and requirements

This topic describes the rules and requirements for configuring a SAN Volume Controller Storage Software for Cisco MDS 9000. It also provides a list of defined terms that are referenced in the configuration rules. Before you read the rules, read these definitions, which can help you to understand the rules.

Properties:

ISL hop

A hop on an interswitch link (ISL).

With reference to all pairs of N-ports, or end-nodes, that are in a fabric, an ISL hop is the number of links that are crossed on the shortest route between the node pair whose nodes are farthest apart from each other. The distance is measured only in terms of the ISL links that are in the fabric.

oversubscription

The ratio of the sum of the traffic that is on the initiator N-node connections to the traffic that is on the most heavily-loaded ISLs, where more than one ISL is in parallel between these switches.

This definition assumes a symmetrical network and a specific workload that is applied equally from all initiators and sent equally to all targets. A symmetrical network means that all initiators are connected at the same level and all the controllers are connected at the same level.

The SAN Volume Controller Storage Software for Cisco MDS 9000 makes this calculation difficult, because it puts its back-end traffic onto the same network, and this back-end traffic varies by workload. Therefore, the oversubscription that a 100% read hit gives is different from the oversubscription that 100% write-miss gives.

If you have an oversubscription of 1 or less, the network is nonblocking.

virtual SAN (VSAN)

A VSAN is a virtual storage area network (SAN).

redundant SAN

A SAN configuration in which if any one component fails, connectivity between the devices that are in the SAN is maintained, possibly with degraded performance. The way to make a redundant SAN is to split the SAN into two independent counterpart SANs.

counterpart SAN

A non-redundant portion of a redundant SAN. A counterpart SAN provides all the connectivity of the redundant SAN, but without the redundancy.

local fabric

The fabric that consists of those SAN components (switches and cables) that connect the components (nodes, hosts, and switches) of the local cluster.

Because the SAN Volume Controller Storage Software for Cisco MDS 9000 supports Remote Copy, significant distances might exist between the components of the local cluster and those of the remote cluster.

remote fabric

The fabric that consists of those SAN components (switches and cables) that connect the components (nodes, hosts, and switches) of the remote cluster.

Because the SAN Volume Controller Storage Software for Cisco MDS 9000 supports remote copy, significant distances might exist between the components of the local cluster and those of the remote cluster.

Local/Remote fabric interconnect

The SAN components that connect the local fabrics to the remote fabrics. These components might be single-mode optical fibres that are driven by GigaBit Interface Converters (GBICs), or they might be other, more advanced components, such as channel extenders.

SAN Volume Controller Storage Software for Cisco MDS 9000 Fibre Channel port fan in

The number of hosts that can see any one SAN Volume Controller Storage Software for Cisco MDS 9000.

Some controllers recommend that the number of hosts that use each port be limited to prevent excessive queuing at that port. If the port fails, or the path to that port fails, the host might failover to another port, and the fan-in requirements might be exceeded in this degraded mode.

Invalid configuration

A configuration that refuses to operate and generates an error code to indicate what caused it to become invalid.

Unsupported configuration

A configuration that might operate successfully, but for which IBM does not guarantee to be able to solve problems that might occur.

Usually this type of configuration does not create an error log.

Valid configuration

A configuration that is neither invalid nor unsupported.

Degraded

A valid configuration that has had a failure, but continues to be neither invalid nor unsupported.

Typically, a repair action is required to restore the degraded configuration to a valid configuration.

Configuration rules

SAN configurations that contain SAN Volume Controller Storage Software for Cisco MDS 9000 clusters can be set up in various ways. Some configurations, however, do not work, and are known as *invalid*. You can avoid creating invalid configurations if you follow the rules that are given in this section.

A SAN configuration that contains SAN Volume Controller Storage Software for Cisco MDS 9000s is valid if it observes *all* of the following rules:

Disk controller systems

All SAN Volume Controller Storage Software for Cisco MDS 9000 nodes of a cluster must be able to see the same set of back-end storage ports on each disk controller system. Any operation that is in this mode in which two nodes do not see the same set of ports on the same controller is degraded, and the system logs errors that request a repair action. This rule can have important effects on back-end storage such as FASTT, which has exclusion rules that determine to which host bus adapter (HBA) WWNNs a storage partition can be mapped.

A configuration in which a SAN Volume Controller Storage Software for Cisco MDS 9000 bridges a separate host controller and a RAID controller is supported. Typical compatibility matrixes are shown in a document titled *Supported Hardware List* on the following Web page:

<http://www.ibm.com/storage/support/2062-2300>

The SAN Volume Controller Storage Software for Cisco MDS 9000 clusters must not share its back-end storage devices with hosts. A controller can be shared with a host under certain conditions as described below.

Two SAN Volume Controller Storage Software for Cisco MDS 9000 clusters must not share the same back-end disk controller system. That is, one back-end disk controller system cannot present LUs to two different SAN Volume Controller Storage Software for Cisco MDS 9000 clusters. Although such a configuration is more unsupported than invalid, any operation that runs in this mode causes serious problems because the same managed disk (MDisk) can appear in two different SAN Volume Controller Storage Software for Cisco MDS 9000 clusters and can be concurrently mapped to different virtual disk (VDisks). This condition can cause data corruption.

The SAN Volume Controller Storage Software for Cisco MDS 9000 must be configured to manage only LUNs that are presented by supported disk controller systems. Operation with other disk controller systems is not supported.

Unsupported disk controller systems (generic controller):

When a disk controller system is detected on the SAN, the SAN Volume Controller Storage Software for Cisco MDS 9000 attempts to recognize it using its Inquiry data. If the disk controller system is recognized as one of the explicitly supported storage models, then the SAN Volume Controller Storage Software for Cisco MDS 9000 uses error recovery programs that are potentially tailored to the known needs of the disk controller system. If the storage controller is not recognized, then the SAN Volume Controller Storage Software for Cisco MDS 9000 configures the disk controller system as a generic controller. A generic controller may or may not function correctly when addressed by a SAN Volume Controller Storage Software for Cisco MDS 9000. In any event, the SAN Volume Controller Storage Software for Cisco MDS 9000 does not regard accessing a generic controller as an error condition and, consequently, does not log an error. MDisks presented by generic controllers are not eligible to be used as quorum disks.

Split controller configurations:

In this configuration, a RAID array presents LUs to both a SAN Volume Controller Storage Software for Cisco MDS 9000 (which treats the LU as an MDisk) and to another host. The SAN Volume Controller Storage Software for Cisco MDS 9000 presents VDisks created from the MDisk to another host. There is no requirement for the pathing driver in the two hosts to be the same (though if the RAID controller were an ESS, both hosts would use SDD). In the Figure 7 the RAID controller would be a FAS*T*, with RDAC being used for pathing on the directly attached host, and SDD being used on the host which is attached via SAN Volume Controller Storage Software for Cisco MDS 9000.




Figure 7. Disk controller system shared between SAN Volume Controller Storage Software for Cisco MDS 9000 and a host

In the case where the RAID controller is an ESS, the pathing driver in the host would be IBM Subsystem Device Driver (SDD) for the ESS and SDD for the SAN Volume Controller Storage Software for Cisco MDS 9000 LUs. Figure 8 on page 46 a supported configuration since the same pathing driver is used for both direct and virtual disks.




Figure 8. ESS LUs accessed directly and via SAN Volume Controller Storage Software for Cisco MDS 9000

Host bus adapters (HBAs)

This topic includes information about the configuration rules for HBAs.

HBAs that are in dissimilar hosts, or dissimilar HBAs that are in the same host, must be in separate zones. For example, if you have an AIX[®] host and a Windows NT[®] host, those hosts must be in separate zones. Here, *dissimilar* means either that the hosts are running different operating systems or that they are different hardware platforms. Different levels of the same operating system are, therefore, thought of as similar. This rule helps you to ensure that different SANs can operate with each other. A configuration that breaks this rule is not supported.

The SAN Volume Controller Storage Software for Cisco MDS 9000 must be configured to export virtual disks only to host fibre-channel ports that are on the supported HBAs. See the following Web site for specific firmware levels and the latest supported hardware:

<http://www.ibm.com/storage/support/2062-2300>

Operation with other HBAs is not supported.

The number of paths from the SAN Volume Controller Storage Software for Cisco MDS 9000 nodes to a host must not exceed two. SAN Volume Controller Storage Software for Cisco MDS 9000 nodes present a minimum of three external ports to the SAN. Additional world wide port names are created when a port is put into multiple VSANs. To restrict the number of paths to two, you should place one port in a zone or VSAN with the hosts, one port in a zone or VSAN with the disk controller systems, and one port in a zone or VSAN that just contains the nodes. In this way, the hosts, storage, and inter node traffic are isolated from one another.

Nodes

This topic includes information about the configuration rules for nodes.

The SAN Volume Controller Storage Software for Cisco MDS 9000 nodes must always be deployed in pairs. If a node fails or is removed from the configuration, the remaining node operates in a degraded mode, but the configuration is still valid.

A cluster contains 4 nodes. Each pair of nodes is known as an I/O group. Each node may be in *only* one I/O group. A cluster can be composed of nodes that reside on different switches.

Fibre-channel switches

This topic includes information about the switches that are supported on the SAN.

The SAN must contain only supported switches. The SAN Volume Controller Storage Software for Cisco MDS 9000 supports the Cisco MDS 9000 switch only.

On the fibre-channel SAN, back-end storage must always and only be connected to SAN switches. Multiple connections are permitted from the redundant controllers of the back-end storage, to improve data bandwidth performance. It is not necessary to have a connection between each redundant disk controller system of the back-end storage and each counterpart SAN. For example, in a FAStT configuration in which the FAStT contains two redundant controllers, only two controller minihubs are usually used. Controller A of the FAStT is, therefore, connected to counterpart SAN A, and controller B of the of the FAStT is connected to counterpart SAN B. Any operation that uses a direct connection between the host and the controller is not supported.

Mixed speeds are permitted in the fabric. Lower speeds can be used to extend distances or to make use of 1 Gbps legacy components.

The local or remote fabric must not contain more than three inter-switch links (ISLs) in each fabric. Any operation that uses more than three ISLs is not supported. When a local fabric is connected to a remote fabric for Remote Copy purposes, the ISL count between a local node and a remote node must not exceed seven. Some ISLs can, therefore, be used in a cascaded switch link between local and remote clusters, if the internal ISL count of the local or remote cluster is less than three.

The local and remote fabric interconnections must be only one ISL hop between a switch that is in the local fabric and a switch that is in a remote fabric. That is, it must be a single-mode fibre up to 10 km (32 810 ft) long. Any operation that uses other local or remote fabric interconnections is not supported.

Where ISLs are used, each ISL oversubscription must not exceed six. Any operation that uses higher values is not supported.

Operation with fibre-channel extenders is not supported.

The switch configuration of a SAN Volume Controller Storage Software for Cisco MDS 9000 SAN must observe the switch manufacturer's configuration rules. These rules might put restrictions on the switch configuration; for example, the switch manufacturer might not permit other manufacturer's switches to be in the SAN. Any operations that run outside the manufacturer's rules is not supported.

The switch must be configured so that the SAN Volume Controller Storage Software for Cisco MDS 9000 nodes can see the back-end storage and the front-end HBAs. The front-end HBAs, however, and the back-end storage must not be in the same zone or VSAN. Any operation that runs outside these zoning rules is not supported.

Each SAN Volume Controller for Cisco MDS 9000 node has 3 ports, one port for communication with hosts (the target port), one port for communication with back-end storage (the initiator port) and one port for communication between nodes (the management port). The switch must be configured so that the proper port can communicate with the hosts, back-end storage and other nodes.

With Remote Copy, additional zones are required that contain only the local nodes and the remote nodes. It is valid for the local hosts to see the remote nodes, or for the remote hosts to see the local nodes. Any zone that contains the local and the remote back-end storage and local nodes or remote nodes, or both, is not valid.

Configuration requirements

This topic describes the rules and steps you *must* perform before you configure the SAN Volume Controller Storage Software for Cisco MDS 9000.

A valid SAN configuration meets the following criteria:

- The two nodes on the same blade cannot be placed in the same I/O group. This means that the minimum SAN Volume Controller Storage Software for Cisco MDS 9000 configuration requires two blades, either in the same switch or different switches.
- The nodes on a blade can belong to different clusters.
- A switch can contain nodes from more than one cluster.
- A cluster can span switches. A cluster spanning more than two switches is not supported.
- The two nodes in an I/O group can be in different switches.
- If a cluster spans switches, the switches must all be in the same subnet. To be in the same subnet, the first three numbers separated by periods of the IP addresses must be the same, and all four numbers separated by periods must be the same for the Subnet Mask.
- If a cluster spans switches, you can synchronize the time across the switches so that the log entry timestamps have the same basis, although this is not a requirement.
- If a cluster spans switches, you can define the same users and passwords on all the switches that the cluster has nodes in or use RADIUS to maintain consistency across all switches.
- A management port can be placed in at most four virtual SANs (VSANs).
- The total number of VSANs in which the three ports of a SAN Volume Controller Storage Software for Cisco MDS 9000 node can be placed must not exceed 64.

(That is, the number of VSANs the target port is in, plus the number of VSANs that the initiator port is in, plus the number of VSANs that the management port is in cannot exceed 64.)

Steps:

Perform the following steps:

1. Ensure that the Cisco MDS 9000 switch is installed.
2. Your IBM service representative must have installed the SAN Volume Controller Storage Software for Cisco MDS 9000.
3. Install and configure your disk controller systems and create the RAID resources that you intend to virtualize. To prevent loss of data, virtualize only those RAID that provide some kind of redundancy, that is, RAID 1, RAID 10, RAID 0+1, or RAID 5. Do *not* use RAID 0 because a single physical disk failure might cause the failure of many virtual disks. RAID 0, like other types of RAID offers cost-effective performance by using available capacity through data striping. However, RAID 0 does not provide a parity disk drive for redundancy (RAID 5) or mirroring (RAID 10).

When creating RAID with parity protection (for example, RAID 5), consider how many component disks to use in each array. The more disks you use, the fewer disks you need to provide availability for the same total capacity (one per array). However, if you use more disks, it will take longer to rebuild a replacement disk after a disk failure. If a second disk failure occurs during the rebuild period, all data on the array is lost. More data is affected by a disk failure for a larger number of member disks resulting in reduced performance while rebuilding onto a hot spare and more data being exposed if a second disk fails before the rebuild has completed. The smaller the number of disks, the more likely it is that write operations span an entire stripe (strip size x number of members minus 1). In this case, write performance is improved because the disk write operations do not have to be preceded by disk reads. The number of disk drives required to provide availability might be unacceptable if the arrays are too small.

When in doubt, create arrays with between 6- and 8-member disks.

When creating RAID with mirroring, the number of component disks in each array does not affect redundancy or performance.

Most back-end disk controller systems enable RAID to be divided up into more than one SCSI logical unit (LU). When configuring new storage for use with the SAN Volume Controller Storage Software for Cisco MDS 9000, you do not need to divide up the array. New storage should be presented as one SCSI LU. This will give a one-to-one relationship between MDisks and RAID.

Attention: Losing an array in an MDisk group can result in the loss of access to *all* MDisks in that group.

4. Configure your Cisco MDS 9000s to create the zones and VSANs that the SAN Volume Controller Storage Software for Cisco MDS 9000 needs. One zone must contain all the disk controller systems and the SAN Volume Controller Storage Software for Cisco MDS 9000 nodes. For hosts, use switch zoning to ensure that each host fibre-channel port is zoned to exactly one fibre-channel port of each SAN Volume Controller Storage Software for Cisco MDS 9000 node in the cluster.
5. If you want the SAN Volume Controller Storage Software for Cisco MDS 9000 to export redundant paths to disks, you must install the Subsystem Device Driver (SDD) on all of your hosts that are connected to the SAN Volume

Controller Storage Software for Cisco MDS 9000. Otherwise, you will not be able to use the redundancy inherent in the configuration. Install the SDD from the following Web site:

<http://www-1.ibm.com/server/storage/support/software/sdd.html>

Be sure to install version 1.5.x.x or higher.

6. Install and configure the SAN Volume Controller Storage Software for Cisco MDS 9000 console workstation. The communication between the console workstation and the SAN Volume Controller Storage Software for Cisco MDS 9000 runs under a client-server network application called Secure Shell (SSH). SSH Server software and the SSH Client software called PuTTY are shipped with SAN Volume Controller Storage Software for Cisco MDS 9000. You will need to configure SSH and PuTTY on the console workstation. Once you have installed your console workstation, you can configure and administer the SAN Volume Controller Storage Software for Cisco MDS 9000 using a graphical interface or a command-line interface.
 - You can configure the SAN Volume Controller Storage Software for Cisco MDS 9000 using the Web application that is preinstalled on the console workstation.
 - You can configure the SAN Volume Controller Storage Software for Cisco MDS 9000 using the command-line interface (CLI) commands.
 - You can install an SSH client if you only want to use the CLI commands. If you want to use the CLI from a host other than the console workstation, ensure that the host has an SSH client installed on it.

Notes:

- a. AIX comes with an installed SSH client.
- b. Linux comes with an installed SSH client.
- c. PuTTY is recommended for Windows.

Result:

When you and the IBM service representative have completed the initial preparation steps, you must:

1. Create the cluster.
2. Add nodes to the cluster and set up the cluster properties.
3. Create managed disk groups from the managed disks to make pools of storage from which you can create virtual disks.
4. Create host objects from the HBA fibre-channel ports to which you can map virtual disks.
5. Create virtual disks from the capacity that is available in your managed disk groups.
6. Map the virtual disks to the host objects to make the disks available to the hosts, as required.
7. Optionally, create Copy Services (FlashCopy and Remote Copy) objects as required.

Related topics:

- “Fibre-channel switches” on page 47

Chapter 6. Setting up SAN Volume Controller Storage Software for Cisco MDS 9000

The topic explains the tasks required to set up and configure your system to run SAN Volume Controller Storage Software for Cisco MDS 9000. You will be performing these tasks using one or more Cisco MDS 9000 Family switches for setup and configuration. Refer to your Cisco and IBM documentation for more detailed procedures.

Related topics:

- “Preparing the Cisco MDS switch”
- “Setting up the Cisco MDS switch”
- “Creating SAN Volume Controller Storage Software for Cisco MDS 9000 interfaces” on page 54
- “Creating a cluster” on page 54
- “Verifying an existing interface” on page 57
- “Adding nodes to a cluster using the Cisco MDS 9000 CLI” on page 57

Preparing the Cisco MDS switch

This topic describes how to prepare the Cisco MDS switch.

To prepare for both hardware and software setup, be sure you have the following in place before getting started:

- Installed and functioning Cisco MDS 9000 Family switch fabric. Each Cisco MDS 9000 Family switch needs a switch IP (mgmt0, one for each switch).
- One subnet for fabric-wide communications: This subnet may be private (only seen by the MDS switches and the IBM TotalStorage server) and should have a range large enough for all supervisor modules in the IP fabric.
- One IP address for each cluster.

Refer to the *Cisco MDS 9216 Switch Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide*, and the *Cisco MDS 9000 Family Configuration Guide*.

Setting up the Cisco MDS switch

This topic describes how to prepare the Cisco MDS switch.

This topic explains the process to set up a new, or update an existing MDS 9000 Family switch in preparation for the SAN Volume Controller Storage Software for Cisco MDS 9000. Before setting up your hardware, ensure you have correctly installed the Cisco MDS chassis and its components as specified in the *Cisco MDS 9216 Switch Hardware Installation Guide, Chapter 2: Installing the Cisco MDS 9000 Family Switch* or the *Cisco MDS 9500 Series Hardware Installation Guide, Chapter 2: Installing the Cisco MDS 9000 Family Switch*.

Be sure to save your work frequently using the Cisco MDS 9000 command:
copy running-config startup-config

Setting up a new Cisco MDS switch:

To set up a new Cisco MDS 9000 Family switch, follow these steps.

1. Follow instructions for preinstallation, installing the chassis in the rack, grounding the chassis, installing modules, installing CompactFlash cards, installing Power Supplies, and installing the Fan Assembly as specified in the *Cisco MDS 9216 Switch Hardware Installation Guide* or the *Cisco MDS 9500 Series Hardware Installation Guide*.
2. Connect to the supervisor module as specified in the *Cisco MDS 9216 Switch Hardware Installation Guide, Chapter 3: Connecting the Cisco MDS 9000 Family Switch* or the *Cisco MDS 9500 Series Hardware Installation Guide, Chapter 3: Connecting the Cisco MDS 9000 Family Switch*.
3. Login to the Cisco MDS 9000 Family switch using the Cisco MDS 9000 Family Command-Line Interface (CLI).
4. Configure the switch as specified in the *Cisco MDS 9000 Family Configuration Guide, Chapter 3: Initial Configuration*.
 - a. Perform the initial setup routine.
 - b. Assign a switch name.
 - c. Access the switch.
 - d. Create at least one user that has the svc-admin role. Users with the svc-admin role have all privileges to administer the cluster and storage. They do not have any access to configure the switch or the SAN.
 - e. Prior to using the Console for SAN Volume Controller for Cisco MDS 9000 with the switch, enable SSH on the switch.

After reviewing the default configuration, you can change it or perform other configuration or management tasks. The initial setup can only be performed at the Cisco MDS 9000 command line interface. You can continue to configure other software features, or access the switch after initial configuration by using either the Cisco MDS 9000 command line interface, the IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000 command line interface, or the Console for SAN Volume Controller Storage Software for Cisco MDS 9000 graphical user interface (GUI). The rest of this procedure uses the Cisco CLI to configure the switch.

5. Verify the module status as specified in the *Cisco MDS 9000 Family Configuration Guide, Chapter 3: Initial Configuration*.
6. Configure the management port as specified in the *Cisco MDS 9000 Family Configuration Guide, Chapter 3: Initial Configuration*.

Updating an existing Cisco MDS switch:

The following steps assume you have a remote FTP, TFTP, SFTP, or SCP server that contains the switch image. Be sure to save your work frequently using the Cisco CLI command:

```
copy running-config startup-config
```

To update a existing Cisco MDS switch, follow these steps.

1. Connect to the supervisor module as specified in the *Cisco MDS 9216 Switch Hardware Installation Guide, Chapter 3: Connecting the Cisco MDS 9000 Family Switch* or the *Cisco MDS 9500 Series Hardware Installation Guide, Chapter 3: Connecting the Cisco MDS 9000 Family Switch*.
2. Login to the Cisco MDS switch using the Cisco MDS 9000 Family CLI.

3. Copy the SAN Volume Controller Storage Software for Cisco MDS 9000 software image file from the remote directory to the bootflash: directory in the active supervisor module.

```
switch# copy ftp://10.1.7.2/m9000-ckg-csm-svc-mz.1.3.1.bin
        bootflash:m9000-ckg-csm-svc-mz.1.3.1.bin
```

The image file can be locally accessible (bootflash/slot0) or on an external machine on the network (scp, sftp, ftp, tftp). This example copies the file from the remote directory using FTP and then installs it from the bootflash directory.

4. Verify that the file was copied in the required directory.

```
switch# dir bootflash: admin 524288 Sep 22 04:27:42 2003
m9000-ckg-csm-svc-mz.1.3.1.bin
```

5. Use the show version image command to verify that you have successfully downloaded the correct version of the SAN Volume Controller Storage Software for Cisco MDS 9000 image.

```
switch# show version image bootflash:m9000-ckg-csm-svc-mz.1.3.1.bin<
image name: m9500-sflek9-svc-mzg.1.0.3.bin
diagnostic: version 1.3(1) [gdb]
compiled: 9/22/2003 12:00:00
```

6. Issue the show install command to verify if the downloaded SAN Volume Controller Storage Software for Cisco MDS 9000 image is compatible with the existing software images (system image, kickstart image).

```
switch# show install bootflash:m9000-ckg-csm-svc-mz.1.3.1.bin
```

Refer to the www.cisco.com web site to verify compatibility issues, or the appropriate Cisco MDS 9000 Family release notes to ensure your system and setup meets the minimum requirements, or images will not install properly.

7. Change to configuration mode.

```
switch# config terminal switch(config)#
```

8. Insert the Cisco MDS 9000 Caching Services Module (CSM) as specified in *Cisco MDS 9216 Switch Hardware Installation Guide, Chapter 2: Installing the Cisco MDS 9000 Family Switch* or the *Cisco MDS 9500 Series Hardware Installation Guide, Chapter 2: Installing the Cisco MDS 9000 Family Switch*. The updated SAN Volume Controller Storage Software for Cisco MDS 9000 software automatically detects the presence of the newly-inserted CSM.

9. Install the new image on each SAN Volume Controller Storage Software for Cisco MDS 9000 node in each module in the fabric. Each module has two nodes.

```
switch# install module 2 node 1 image svc-system
sbootflash:m9000-ckg-csm-svc-mz.1.3.1.bin
SVC reimage going on. Please wait
m9000-ckg-csm-svc-mz.1.3.1.bin 100% |*****|
45408 KB 00:53
svc 2/1 software reimage succeeded

switch# install module 2 node 2 image svc-system
bootflash:m9000-ckg-csm-svc-mz.1.3.1.bin
SVC reimage going on. Please wait
m9000-ckg-csm-svc-mz.1.3.1.bin 100% |*****|
45408 KB 00:55 svc 2/2 software reimage succeeded
```

All prior information on the upgraded node is lost with each upgrade.

10. Verify that the CSM booted successfully, by issuing the command:
show module

The status will progress from testing to ok when it has fully booted.

Mod	Ports	Module-Type	Model	Status
2	0	Caching Services Module	DS-X9560-SMAP	ok
4	0	Caching Services Module		powered-dn
5	0	Supervisor/Fabric-1	DS-X9530-SF1-K9	active *
7	0	Caching Services Module	DS-X9560-SMAP	ok
8	32	1/2 Gbps FC Module	DS-X9032	ok

Mod	Sw	Hw	World-Wide-Name(s) (WNN)
2	1.3(0.58m)	0.702	--
5	1.3(0.58m)	0.602	--
7	1.3(0.58m)	0.602	--
8	1.3(0.58m)	0.3	21:c1:00:05:30:00:00:00 to 21:e0:00:05:30:00:00:00

Mod	MAC-Address(es)	Serial-Num
2	00-05-30-01-37-7a to 00-05-30-01-37-fe	JAB072705JA
5	00-05-30-00-64-be to 00-05-30-00-64-c2	
7	00-05-30-00-b1-2e to 00-05-30-00-b1-b2	JAB071007H9
8	00-05-30-00-2d-e2 to 00-05-30-00-2d-e6	JAB06280ae9

* this terminal session

You might also need to update the Supervisor code to support the Cisco MDS 9000 Caching Services Module (CSM). See the *Cisco MDS 9000 Family Configuration Guide* for more information.

Creating SAN Volume Controller Storage Software for Cisco MDS 9000 interfaces

This topic describes how to create SAN Volume Controller Storage Software for Cisco MDS 9000 interfaces. See the *Cisco MDS 9000 Family Configuration Guide* for more information about VSANs.

Steps:

Each node in a CSM is represented by an interface (SVC interface 1 and SVC interface 2). To create a SAN Volume Controller Storage Software for Cisco MDS 9000 interface and N-port VSANs in a Cisco MDS switch using the Cisco MDS 9000 commands, follow these steps:

1. Enter configuration mode by issuing the command:
switch# config t switch(config)#
2. Create the SAN Volume Controller Storage Software for Cisco MDS 9000 interface for Node in the CSM module that resides in slot 2 by issuing the command:
switch(config)# interface svc 2/1 switch(config-if)#
3. Enable the SAN Volume Controller Storage Software for Cisco MDS 9000 interface by issuing the command:
switch(config-if)# no shutdown

Creating a cluster

This topic describes how to create a cluster.

Prerequisites:

After creating a SAN Volume Controller Storage Software for Cisco MDS 9000 interface, you can create a cluster and assign the IP address. Connecting to the IP address for the cluster creates a connection to the Cisco MDS switch that contains the current configuration node for the cluster. You can use the Cisco MDS 9000 switch commands or the IBM-supplied Console for SAN Volume Controller for Cisco MDS 9000 graphical user interface (GUI) to create a cluster. This topic describes both ways to create a cluster.

Steps:

A cluster name is a string of 1 to 15 characters, composed of the letters A-Z, a-z, 0-9, -, and _. The first letter cannot be numeric. A cluster name cannot be changed once it is set. Once a cluster is created, all other configurations for that cluster, must be performed under the cluster configuration submode identified by the prompt:

```
switch(svc-cluster)#
```

Perform the following steps to create a cluster using the Cisco MDS 9000 commands:

1. Enter configuration mode by issuing:

```
switch# config t switch(config)#
```

2. Enter SAN Volume Controller Storage Software for Cisco MDS 9000 configuration mode for interface 2/1 by issuing:

```
switch(config)# interface svc 2/1  
switch(config-if)#
```

3. Return to EXEC mode by issuing:

```
switch(config-if)# exit  
switch#
```

4. Enter the SAN Volume Controller Storage Software for Cisco MDS 9000 configuration mode by issuing:

```
switch# svc-config  
switch(svc)#
```

5. Create a new cluster (SampleCluster), assign it an IP address, and add the interface 2/1 in this cluster. The cluster IP address must be in the same subnet as the management 0 IP address.

```
switch(svc)# cluster add SampleCluster ip 10.10.0.0  
node svc 2/1
```

6. Enter the cluster configuration submode for the new cluster by issuing:

```
switch(svc)# cluster config SampleCluster  
switch(svc-cluster)#
```

Creating a cluster on the Console for SAN Volume Controller for Cisco MDS 9000

This topic provides information about creating a cluster on the Console for SAN Volume Controller for Cisco MDS 9000.

In order to create a cluster on the Console for SAN Volume Controller for Cisco MDS 9000, you must follow a sequence of steps.

Steps:

Perform the following steps to create a cluster:

1. Access the Console for SAN Volume Controller for Cisco MDS 9000 using a web browser.
2. Sign on with the superuser name and password. For first time access, use the superuser name `superuser` and the default password `passwd`. You must change this default password the first time you sign on. . Sign on with the superuser name and password. The Welcome panel is displayed.
3. Click **Clusters** in the portfolio.
4. Select **Create a cluster** in the list and click **Go**. The Create a Cluster wizard is displayed.
5. Complete the Create a Cluster wizard:
 - a. Identify the node
 - b. Identify the cluster
 - c. Set up the error notification setting
 - d. Enable the product features

After successfully completing the wizard, the cluster is displayed in the list of clusters on the Viewing Clusters panel.

Separating hosts and storage devices

This topic describes how to separate hosts and storage devices.

Context:

Be sure to remove conflicting N ports, if any already exist, before starting this procedure.

Steps:

Each SAN Volume Controller Storage Software for Cisco MDS 9000 interface is assigned a nWWN. A SVC interface consists of three N-ports:

- Target N-ports: used for traffic between SAN Volume Controller Storage Software for Cisco MDS 9000 nodes and hosts.
- Initiator N-ports: used for traffic between SAN Volume Controller Storage Software for Cisco MDS 9000 nodes and disks.
- Management N-ports: used for traffic to and from other SAN Volume Controller Storage Software for Cisco MDS 9000 nodes.

Each N-port can be part of one or more VSANs or zones. Each N-port within an SAN Volume Controller Storage Software for Cisco MDS 9000 interface is assigned a pWWN. The pWWNs and nWWNs are preserved across switch reboots. By default, all N-ports reside in VSAN 1. You must explicitly remove them when necessary. To create a SAN Volume Controller Storage Software for Cisco MDS 9000 interface and N-port VSANs in a Cisco MDS switch, follow these steps:

1. Enter configuration mode by issuing the command:
`switch# config t switch(config)#`
2. Enter the configuration mode for interface 2/1 by issuing the command:
`switch(config)# interface svc 2/1 switch(config-if)#`
3. Configure the initiator VSAN 1 for vdisks by issuing the command:
`switch(config-if)# initiator vsan 1`
4. Configure the target VSAN 2 for hosts by issuing the command:
`switch(config-if)# target vsan 2`

5. Configure the management VSAN 3 for nodes by issuing the command:
switch(config-if)# mgmt vsan 3
6. Enable the interface by issuing the command:
switch(config-if)# no shutdown

Verifying an existing interface

This topic describes verifying an existing interface.

Example:

If interfaces do not exist in the switch, be sure to create the required interfaces using the interface svc command . To verify what interfaces have already been created, use the show nodes local command. This command displays a list of all interfaces in the switch. If a cluster name is not displayed next to an interface, that interface is considered free and can be added to a new or existing cluster.

```
switch(svc)# show nodes local
```

Node	cluster	config node	cluster status	node status	sw version
svc2/1		No	unconfigured	free	1.3(1)
svc2/2		No	unconfigured	free	1.3(1)
svc7/1		No	unconfigured	free	1.3(1)
svc7/2		No	unconfigured	free	1.3(1)

Adding nodes to a cluster using the Cisco MDS 9000 CLI

This topic describes how to configure a node using the Cisco MDS 9000 CLI.

Prerequisites:

Once a cluster is created, you can add nodes to that cluster.

Steps:

You can use the Cisco MDS 9000 switch commands or the IBM-supplied Console for SAN Volume Controller for Cisco MDS 9000 graphical user interface (GUI) to add nodes to a cluster. This topic describes using the Cisco MDS 9000 switch commands to add a node to a cluster. See the *IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000: Configuration Guide* for information about adding a node to a cluster using the IBM-supplied Console for SAN Volume Controller for Cisco MDS 9000 graphical user interface (GUI) and The IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000 command line interface.

1. Find a candidate node.

Use the show cluster cluster-name nodes candidate command to verify the possible nodes that can be added to a cluster. The switch name, the slot number, and the interface number are displayed along with the node WWN for each node that can possibly be added to a cluster.

```
switch(svc)# show cluster SampleCluster nodes candidate
```

NODE	NWWN
switch.2.2	24:0d:00:05:30:00:00:02

2. Add a Node to a Cluster.

When you add a node to an existing cluster, you can also assign one of two characteristics to this node:

- The alias name: Specifies the node using a 16 alphanumeric characters.
- The IO group: Assigns the node to one of four IO groups in the specified cluster. The IO group ID ranges from 1 to 4.

To add a node to a cluster, follow these steps:

- a. Enter configuration mode by issuing:
`switch# config t switch(config)#`
- b. Enter SAN Volume Controller Storage Software for Cisco MDS 9000 configuration mode by issuing:
`switch# svc-config switch(svc)#`
- c. Enter the cluster configuration submode for the new cluster by issuing:
`switch(svc)# cluster config SampleCluster
switch(svc-cluster)#`
- d. Add a node (nWWN 22:10:00:05:30:00:0c:60) to the specified cluster (SampleCluster) and assign it to IO group 1 by issuing:
`switch(svc-cluster)#node nWWN
22:10:00:05:30:00:0c:60 iogroup 1`
- e. Add another node to the specified cluster and specify an alias name (SampleAlias) by issuing:
`switch(svc-cluster)#node nWWN
22:10:00:05:20:00:0c:60 alias SampleAlias`

Related topics:

- “Adding nodes to the cluster using the IBM TotalStorage CLI”
- “Deleting a node from a cluster using the IBM TotalStorage CLI” on page 60

Adding nodes to the cluster using the IBM TotalStorage CLI

This task provides step-by-step instructions you will need to perform to add nodes to the cluster using the IBM TotalStorage CLI.

Context:

Assume that the cluster has been created, that initial configuration has been performed using the Console for SAN Volume Controller for Cisco MDS 9000 and that the necessary setup has been performed to use the Command Line Interface (CLI).

Prerequisites:

Issue the Cisco switch command **svc-ibmcli** to enter a mode where you can issue the IBM TotalStorage CLI commands.

Steps:

Perform the following steps to add nodes to the cluster:

1. Open a command prompt window.
2. Type the **svcinfolnode** command to list the nodes that are currently part of the cluster.

Example:

```
svcinfo lsnode -delim :
```

```
id:name:card_serial_number:WWNN:status:IO_group_id:IO_group_name:
  config_node:card_unique_id
1:node1:JAB0649055d:2541000530003C20:online:
  0:io_grp0:yes:0100064930353564
```

3. Type the **svcinfo lsnodecandidate** command to list nodes that are not assigned to a cluster.

Example:

```
svcinfo lsnodecandidate -delim :
```

```
id:panel_name:card_serial_number:card_unique_id
2543000530003C20:vegas_fvt.7.1:JAB0650081e:0100065030383165
254F000530003C20:vegas_fvt.9.1:JAB071105kr:0100071130356B72
2559000530003C20:vegas_fvt.8.1:JAB06500812:0100065030383132
```

4. **Attention:** If you are re-adding a node to the SAN, ensure that you are adding the node to the same I/O group that it was removed from. Failure to do this can result in data corruption. Use the information that was recorded when the node was originally added to the cluster. If you do not have access to this information, call IBM Service to add the node back into the cluster without corrupting the data.

If you are adding the node into the cluster for the first time, record the following information:

- Node serial number
- WWNN
- All WWPNS
- I/O group that contains the node

This can avoid a possible data corruption exposure if the node must be removed from and re-added to the cluster.

Note: This warning also is displayed on the Console for SAN Volume Controller for Cisco MDS 9000 panel when adding the node.

Type the **svctask addnode** command to add a node to the cluster. Use the output from the previous commands to choose which I/O group to add the node to and to make sure that when adding a second node to a I/O group that it is attached to a different Cisco MDS 9000 Caching Services Module.

Notes:

- a. When adding a node to a cluster you can specify a name for the node. You can also change the name of nodes that are already part of a cluster using the **svctask chnode** command.
- b. When adding a node to a cluster the node can be identified by using the world wide node name of that node.

Example:

Add a second node to the first I/O group.

```
svctask addnode -panelname vegas_fvt.7.1 -iogrp io_grp0 -name group1node2
```

This command will add the node to the cluster. The node will be added to I/O group, **io_grp0**, and called **group1node2**.

Next add two nodes to the second I/O group. Make sure that each node is attached to a different Cisco MDS 9000 Caching Services Module.

```
svctask addnode -wwnodename 2559000530003C20 -iogrp io_grp1 -name group2node1
svctask addnode -wwnodename 254F000530003C20 -iogrp io_grp1 -name group2node2
```

These commands will add the nodes, identified by the WWNN 5005076801000001 and the WWNN 50050768010000F4 to the cluster. The nodes will be added to I/O group, io_grp1 and called group2node1 and group2node2.

Finally change the name of the first node from the default name node1 so that it conforms with your naming convention.

```
svctask chnode -name group1node1 node1
```

5. Verify the final configuration using the **svcinfolnode** command.

Example:

In our hypothetical scenario, the command to list the nodes is:

```
svcinfolnode -delim :
```

```
id:name:card_serial_number:WWNN:status:I/O_group_id:
      I/O_group_name:config_node:card_unique_id
1:group1node1:JAB0649055d:2541000530003C20:online:0:io_grp0:yes:0100064930353564
2:group1node2:JAB0650081e:2543000530003C20:online:0:io_grp0:no:0100065030383165
3:group2node1:JAB06500812:2559000530003C20:online:1:io_grp1:no:0100065030383132
4:group2node2:JAB071105kr:254F000530003C20:online:1:io_grp1:no:0100071130356B72
```

Note: If this command is issued quickly after adding nodes to the cluster the status of the nodes may be adding rather than online indicating that the process of adding the nodes to the cluster is still in progress. You do not however have to wait for all the nodes to become online before continuing with the configuration process.

Remember: Record the following information:

- WWNN
- All WWPNS
- I/O group that contains the node
- Card serial number
- The slot where the CSM resides
- The CSM node (1/2) the node is defined as

This can avoid a possible data corruption exposure if the node must be removed from and re-added to the cluster.

Result:

You have now added four nodes to one cluster. The nodes are split into two I/O groups.

Deleting a node from a cluster using the IBM TotalStorage CLI

This task provides step-by-step instructions about how to delete a node from a cluster using the CLI.

Attention: Before deleting a node from the cluster you should quiesce all I/O operations that are destined for this node. Failure to do so may result in failed I/O operations being reported to your host operating systems.

Prerequisites:

Issue the Cisco switch command `svc-ibmcli` to enter a mode where you can issue the IBM TotalStorage CLI commands.

Attention: If you are deleting a single node, and the other node in the I/O group is online, be aware that the cache on the partner node will go into write-through mode and that you are exposed to a single point of failure should the partner node fail.

Attention: If you are deleting a node, and this is the last node in the I/O group, you will lose access to all VDIsks served by this I/O group. Ensure that all VDIsks are not being accessed or contain data that you wish to continue to access, or ensure that they have been migrated to a different (online) I/O group.

1. Begin by determining the VDIsks that are still assigned to this I/O group:
 - a. Determine the VDIsks in question by requesting a filtered view of VDIsks where the filter attribute is the I/O group in question. This can be done using the following command:

```
svcinfo lsvdisk -filtervalue IO_group_name=<name>
```

where <name> is the name of the I/O group in question.
 - b. Once you have a list of VDIsks, determine the hosts that they are mapped to by following the procedure called, Determining the hosts that a VDisk is mapped to.
 - c. Once you have determined the hosts and are sure that you do not wish to maintain access to these VDIsks proceed to 3.
 - d. If you determine that some or all of the VDIsks assigned to this I/O group do contain data that you wish to continue to access, you should follow the procedure called, Migrating a VDisk to a new I/O group.
2. Before performing the SDD path removal procedure described in 3 you should stop the node that you intend to remove using the `stopcluster node` command, unless this is the last node in the cluster. This ensures that SDD does not re-discover the paths that are manually removed before you issue the delete node request.

Attention:

- If the node being removed is the configuration node, it may take a minute or so before you can perform the delete node request as the configuration node failover has to occur.
- If the node being removed is the last node in the cluster, the Console for SAN Volume Controller for Cisco MDS 9000 may appear to hang for up to 3 minutes because you have removed the last access point to the cluster. Removing the last node in the cluster destroys the cluster. Ensure that this is what you want to do before performing this task.

Note: If you power back on the node that has been removed and it is still connected to the same fabric or zone it will attempt to rejoin the cluster. At this point the cluster will tell the node to remove itself from the cluster and the node will become a candidate for addition to this cluster or another cluster. If you are adding this node back into the cluster, ensure that you add it back to the same I/O group that it was previously a member of. Failure to do so may result in data corruption.

3. Before deleting the node, it is essential that for each vpath presented by the VDIsks you intend to remove, the SDD configuration is updated to remove the vpaths in question. Failure to do this may result in data corruption. See the

IBM TotalStorage Subsystem Device Driver: User's Guide for details about how to dynamically reconfigure SDD for the given host operating system.

4. Proceed to 1.

Steps:

Perform the following steps to delete a node:

1. Open a command prompt window.

Notes:

- a. Before removing a node, be sure this is what you want to do. Any VDisks that are assigned to the I/O group that this node belongs to, will be assigned to the other node in the I/O group, that is, the preferred node will be changed. You cannot change this setting back once this has been done. Also, all VDisks will go into write-through cache mode as there is not a redundant node available to duplicate the cached information.
 - b. If this is the last node in the I/O group or the last node in the cluster, you will be asked to force the delete.
 - c. If this is the last node in the cluster or is currently assigned as the configuration node, all connections to the cluster will be lost. The user interface and any open CLI sessions will halt. This may result in a timeout from the command as the command cannot be completed before the node is deleted.
2. Issue the **svctask rmnode** command to delete a node from the cluster. You can enter this command any time after a cluster has been created.

Chapter 7. SAN Volume Controller Storage Software for Cisco MDS 9000 supported environment

This topic describes the SAN Volume Controller Storage Software for Cisco MDS 9000 supported environment including:

- Host attachments
- Physical disk storage systems support
- HBAs
- Switch support

See the following Web site for specific firmware levels and the latest supported hardware:

<http://www.ibm.com/storage/support/2062-2300>

Related topics:

- “Supported host attachments”
- “Physical disk storage system support” on page 64
- “Supported fibre-channel, host bus adapters (HBAs)” on page 64
- “Supported switches” on page 64

Supported host attachments

The SAN Volume Controller Storage Software for Cisco MDS 9000 supports connectivity with Intel-based servers running Windows NT or Windows 2000 Advanced Server or Red Hat Linux Advanced Server and with RISC-based servers running AIX, HP-UX, or Sun Solaris operating systems.

- AIX 5.1 Required Maintenance Level 4
- AIX 5.2 Required Maintenance Level 1
- Windows[®] 2000 Advanced Server Service Pack 4
- Windows 2000 Server Service Pack 4
- Windows NT[®] 4 Server Service Pack 6a
- Red Hat Linux Advanced Server 2.1 with 2.4.9-e27 enterprise kernel
- SUN Solaris 5.8
- HP-UX 11i version 1.6

Note: Check the SAN Volume Controller Storage Software for Cisco MDS 9000 support Web site for the most current list of supported operating systems.

<http://www.ibm.com/storage/support/2062-2300>

The SAN Volume Controller Storage Software for Cisco MDS 9000 provides heterogeneous host attachments so that you can consolidate storage capacity and workloads for open-systems hosts. The SAN Volume Controller Storage Software for Cisco MDS 9000 supports a maximum of 64 separate hosts and a maximum of 128 host fibre-channel ports, identified by their worldwide port numbers (WWPNs).

Hosts are attached to the SAN Volume Controller Storage Software for Cisco MDS 9000 using a switched, fibre-channel fabric.

Physical disk storage system support

Storage system support is available for the following IBM FAStT models:

- 200
- 500
- 600
- 700
- 900

Storage system support is available for the following IBM TotalStorage[®] Enterprise Storage Server[®] (ESS) models:

- 2105-F20
- 2105-800

Supported fibre-channel, host bus adapters (HBAs)

Ensure that HBAs are at or above the minimum requirements.

See the following Web site for specific firmware levels and the latest supported hardware:

<http://www.ibm.com/storage/support/2062-2300>

Supported switches

Ensure that switches are at or above the minimum requirements.

The SAN must contain only supported switches. The SAN Volume Controller Storage Software for Cisco MDS 9000 supports the Cisco MDS 9000 switch and switches supported by the Cisco MDS 9000.

See the following Web site for the latest models and firmware levels:

<http://www.ibm.com/storage/support/2062-2300>

Chapter 8. Getting Started Hints and Tips

This topic provides hints and tips you might find helpful as you get started with SAN Volume Controller Storage Software for Cisco MDS 9000.

This topic includes the following sections:

- “Checking the battery capacity”
- “Hints and tips for troubleshooting the SAN Volume Controller Storage Software for Cisco MDS 9000 environment” on page 66

Related topics:

- “Related publications” on page v

Checking the battery capacity

This topic describes how to check the battery capacity.

Context:

The purpose of the batteries on the CSM is to provide adequate power to back up data in case of a loss of external power. The CSM batteries are shipped fully charged and should be handled accordingly.

You can check the LEDs that indicate if there is enough charge. If the battery does not have enough charge to save the volatile memory on the nodes, you will see a syslog error message.

The batteries charge with an SMBus compliant Level III smart charger, which enables the batteries to request the appropriate charge voltage and current from the backplane as required. When the CSM is connected to the chassis backplane, the smart charger measures the amount of charge in the battery. If this charge is less than 90% of the battery’s capacity, then the battery charges to full. This is common since the batteries gradually lose charge while in storage.

The batteries are periodically reconditioned, one battery at a time. This extends the life and the reliability of the batteries. During reconditioning all power is drained from the battery and then the battery is completely recharged. The entire reconditioning process takes approximately 6 hours.

Steps:

Perform the following steps to check the battery capacity:

1. To verify the battery status of a CSM at any time, issue the show environment battery module command in EXEC mode.

```
switch# show environment battery module 2
```

Example:

A sample output of the show module command follows:

```
Battery 1:  
-----  
Voltage                : 10.620 V
```

```
Current          : 0.000 A
Temperature      : 38.9 C
Current Capacity : 1839 mAHr
Full Capacity    : 1967 mAHr
Cycle Count      : 11
Last conditioned in : Week 38, 2003
Serial Number    : AMB0722009G
```

Battery 2:

```
Voltage          : 10.660 V
Current          : 0.000 A
Temperature      : 35.4 C
Current Capacity : 1859 mAHr
Full Capacity    : 1971 mAHr
Cycle Count      : 11
Last conditioned in : Week 38, 2003
Serial Number    : AMB0722007C
```

Hints and tips for troubleshooting the SAN Volume Controller Storage Software for Cisco MDS 9000 environment

This topic describes commands that you can use to troubleshoot your SAN Volume Controller Storage Software for Cisco MDS 9000 environment.

Checking boot variables:

Example:

Issue this command to check boot variables.

```
switch0# show boot
```

Sample results:

```
kickstart variable = bootflash:/excal-boot-m17p
system variable = bootflash:/excal-isan-m17p
Module 2
ilce variable = bootflash:/ilc1.bin
lasilc variable = bootflash:/laslcl1.bin
```

Checking the current switch configuration:

Example:

Issue this command to check the current switch configuration.

```
switch0# show running-config
```

Sample results:

```
Building Configuration ...
boot system bootflash:/excal-isan-m17p
boot kickstart bootflash:/excal-boot-m17p
ip default-gateway 9.43.224.1
username admin password 5 Ws/9RT.HQyIwY role network-admin
interface fc1/1
interface fc1/2
interface fc1/3
interface fc1/4
interface fc1/5
interface fc1/6
interface fc1/7
interface fc1/8
interface fc1/9
```

```
interface fc1/10
interface fc1/11
interface fc1/12
interface fc1/13
interface fc1/14
interface fc1/15
interface fc1/16
interface mgmt0
ip address 9.43.224.72 255.255.252.0
```

Rebooting the system:

Example:

Issue this command to reboot the system.

```
switch0# reload
```

Sample results:

None.

Checking communications between the clusters and the Console for SAN Volume Controller for Cisco MDS 9000:

Example:

Here are some tasks you can perform to identify the reason that the clusters and the Console for SAN Volume Controller for Cisco MDS 9000 are not communicating.

- Ping the cluster IP address.
- Ping the switch IP address.
- Can you login to the cluster IP using the plink.exe from the workstation where the Console for SAN Volume Controller for Cisco MDS 9000 is installed? Check the username, password, and if the fingerprint has changed if you have problems logging in.
- Can you login to the switch IP using the plink.exe from the workstation where the Console for SAN Volume Controller for Cisco MDS 9000 is installed? Check the username, password, and if the fingerprint has changed if you have problems logging in.
- Check that the cluster is up and running by issuing the Cisco MDS 9000 switch commands

```
svc-config->show nodes local)(svc-ibmcli cluster-name
```
- Check the command

```
svinfo -x list -o cluster
```
- Remove the cluster entry from the Console for SAN Volume Controller for Cisco MDS 9000 and add the cluster again.

Related topics:

- “Related publications” on page v

Appendix. Accessibility

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

Features:

These are the major accessibility features in the SAN Volume Controller Storage Software for Cisco MDS 9000 console workstation:

- You can operate all features using the keyboard instead of the mouse.

Navigating by keyboard:

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done through mouse actions. You can navigate the Console for SAN Volume Controller for Cisco MDS 9000 and help system from the keyboard by using the following key combinations:

- To traverse to the next link, button, or topic, press Tab inside a frame (page).
- To expand or collapse a tree node, press → or ←, respectively.
- To move to the next topic node, press V or Tab.
- To move to the previous topic node, press ^ or Shift+Tab.
- To scroll all the way up or down, press Home or End, respectively.
- To go back, press Alt+←.
- To go forward, press Alt+→.
- To go to the next frame, press Ctrl+Tab.
- To move to the previous frame, press Shift+Ctrl+Tab.
- To print the current page or active frame, press Ctrl+P.
- To select, press Enter.

Accessing the publications:

You can view the publications for the SAN Volume Controller Storage Software for Cisco MDS 9000 in Adobe Portable Document Format (PDF) using the Adobe Acrobat Reader. The PDFs are provided on a CD that is packaged with the product or you can access them at the following Web site:

<http://www.ibm.com/storage/support/2062-2300>

Related topics:

- “Related publications” on page v

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.*

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been

estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Related topics:

- "Trademarks"

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

- AIX
- e (logo)
- Enterprise Storage Server
- FlashCopy
- IBM
- TotalStorage

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks of others.

Glossary

This glossary includes terms for the IBM TotalStorage Virtualization Family SAN Volume Controller. This glossary includes selected terms and definitions from:

A Dictionary of Storage Networking Terminology (<http://www.snia.org/education/dictionary>), copyrighted 2001 by the Storage Networking Industry Association, 2570 West El Camino Real, Suite 304, Mountain View, California 94040-1313. Definitions derived from this book have the symbol (S) after the definition.

The following cross-references are used in this glossary:

- See** Refers the reader to one of two kinds of related information:
- A term that is the expanded form of an abbreviation or acronym. This expanded form of the term contains the full definition.
 - A synonym or more preferred term.

See also
Refers the reader to one or more related terms.

Contrast with
Refers the reader to a term that has an opposite or substantively different meaning.

A

auxiliary virtual disk. The virtual disk that contains a backup copy of the data and that is used in disaster recovery scenarios. See also *master virtual disk*.

B

blade. One component in a system that is designed to accept some number of components (blades). Blades could be individual servers that plug into a multiprocessing system or individual port cards that add connectivity to a switch. A blade is typically a hot-swappable hardware device.

block. A unit of data storage on a disk drive.

block virtualization. The act of applying virtualization to one or more block-based (storage) services for the purpose of providing a new aggregated, higher-level,

richer, simpler, or secure block service to clients. Block virtualization functions can be nested. A disk drive, RAID system, or volume manager all perform some form of block-address to (different) block-address mapping or aggregation. See also *virtualization*.

C

cache. A high-speed memory or storage device used to reduce the effective time required to read data from or write data to lower-speed memory or a device. Read cache holds data in anticipation that it will be requested by a client. Write cache holds data written by a client until it can be safely stored on more permanent storage media such as disk or tape.

cluster. In SAN Volume Controller Storage Software for Cisco MDS 9000, a pair of nodes that provides a single configuration and service interface.

CIM. See *Common Information Model*.

Cisco command line interface. An interface that is used to perform functions that are provided on the service panel.

CLI. See *command line interface*.

command line interface (CLI). A type of computer interface in which the input command is a string of text characters.

Common Information Model (CIM). A set of standards developed by the Distributed Management Task Force (DMTF). CIM provides a conceptual framework for storage management and an open approach to the design and implementation of storage systems, applications, databases, networks, and devices.

connected. In a Remote Copy relationship, pertaining to the status condition that occurs when two clusters can communicate.

consistency group. A group of copy relationships between virtual disks that are managed as a single entity.

consistent copy. In a Remote Copy relationship, a copy of a secondary virtual disk (VDisk) that is identical to the primary VDisk from the viewpoint of a host system, even if a power failure occurred while I/O activity was in progress.

copied. In a FlashCopy[®] relationship, a state that indicates that a copy has been started after the copy

relationship was created. The copy process is complete and the target disk has no further dependence on the source disk.

Copy Services. In the SAN Volume Controller, the two services that enable you to copy virtual disks (VDisks): FlashCopy and Remote Copy.

copying. A status condition that describes the state of a pair of virtual disks (VDisks) that have a copy relationship. The copy process has been started but the two virtual disks are not yet synchronized.

counterpart SAN. A nonredundant portion of a redundant storage area network (SAN). A counterpart SAN provides all the connectivity of the redundant SAN but without the redundancy. Each counterpart SAN provides an alternate path for each SAN-attached device. See also *redundant SAN*.

cross-volume consistency. In SAN Volume Controller Storage Software for Cisco MDS 9000, a consistency group property that guarantees consistency between virtual disks when an application issues dependent write operations that span multiple virtual disks.

D

data migration. The movement of data from one physical location to another without disrupting I/O operations.

destage. A write command initiated by the cache to flush data to disk storage.

disk controller. A device that coordinates and controls the operation of one or more disk drives and synchronizes the operation of the drives with the operation of the system as a whole. Disk controllers provide the storage that the cluster detects as managed disks (MDisks).

E

error code. A value that identifies an error condition.

excluded. In SAN Volume Controller Storage Software for Cisco MDS 9000, the status of a managed disk that the cluster has removed from use after repeated access errors.

extent. A unit of data that manages the mapping of data between managed disks and virtual disks.

F

fabric. In fibre-channel technology, a routing structure, such as a switch, that receives addressed information and routes it to the appropriate destination. A fabric can consist of more than one switch. When multiple

fibre-channel switches are interconnected, they are described as cascading. See also *cascading*.

failover. In SAN Volume Controller Storage Software for Cisco MDS 9000, the function that occurs when one redundant part of the system takes over the workload of another part of the system that has failed.

fibre channel. A technology for transmitting data between computer devices at a data rate of up to 4 Gbps. It is especially suited for attaching computer servers to shared storage devices and for interconnecting storage controllers and drives.

FlashCopy mapping. A relationship between two virtual disks.

FlashCopy relationship. See *FlashCopy mapping*.

FlashCopy service. In SAN Volume Controller Storage Software for Cisco MDS 9000, a copy service that duplicates the contents of a source virtual disk (VDisk) to a target VDisk. In the process, the original contents of the target VDisk are lost. See also *point-in-time copy*.

H

HBA. See *host bus adapter*.

host. An open-systems computer that is connected to the SAN Volume Controller Storage Software for Cisco MDS 9000 through a fibre-channel interface.

host bus adapter (HBA). In SAN Volume Controller Storage Software for Cisco MDS 9000, an interface card that connects a host bus, such as a peripheral component interconnect (PCI) bus, to the storage area network.

host ID. In SAN Volume Controller Storage Software for Cisco MDS 9000, a numeric identifier assigned to a group of host fibre-channel ports for the purpose of logical unit number (LUN) mapping. For each host ID, there is a separate mapping of Small Computer System Interface (SCSI) IDs to virtual disks (VDisks).

I

IBM® Subsystem Device Driver (SDD). An IBM pseudo device driver designed to support the multipath configuration environments in IBM products.

idling. The status of a pair of virtual disks (VDisks) that have a defined copy relationship for which no copy activity has yet been started.

image mode. An access mode that establishes a one-to-one mapping of extents in the managed disk (MDisk) with the extents in the virtual disk (VDisk). See also *managed space mode* and *unconfigured mode*.

I/O group. A collection of virtual disks (VDisks) and node relationships that present a common interface to host systems.

inconsistent. In a Remote Copy relationship, pertaining to a secondary virtual disk (VDisk) that is being synchronized with the primary VDisk.

input/output (I/O). Pertaining to a functional unit or communication path involved in an input process, an output process, or both, concurrently or not, and to the data involved in such a process.

Internet Protocol (IP). In the Internet suite of protocols, a connectionless protocol that routes data through a network or interconnected networks and acts as an intermediary between the higher protocol layers and the physical network.

interoperability. The capability to communicate, run programs, or transfer data among various functional units in a way that requires the user to have little or no knowledge of the unique characteristics of those units.

Inter-Switch Link (ISL). A protocol for interconnecting multiple routers and switches in a storage area network.

ISL. See *Inter-Switch Link*.

ISL hop. Considering all pairs of node ports (N-ports) in a fabric and measuring distance only in terms of inter-switch links (ISLs) in the fabric, the number of ISLs traversed on the shortest route between the pair of nodes that are farthest apart in the fabric.

IP. See *Internet Protocol*.

I/O. See *input/output*.

L

line card. See *blade*.

local fabric. In SAN Volume Controller Storage Software for Cisco MDS 9000, those storage area network (SAN) components (such as switches and cables) that connect the components (nodes, hosts, switches) of the local cluster together.

logical unit (LU). An entity to which Small Computer System Interface (SCSI) commands are addressed, such as a virtual disk (VDisk) or managed disk (MDisk).

logical unit number (LUN). The SCSI identifier of a logical unit within a target. (S)

LU. See *logical unit*.

LUN. See *logical unit number*.

M

managed disk (MDisk). A Small Computer System Interface (SCSI) logical unit that a redundant array of independent disks (RAID) controller provides and a cluster manages. The MDisk is not visible to host systems on the storage area network (SAN).

managed disk group. A collection of managed disks (MDisks) that, as a unit, contain all the data for a specified set of virtual disks (VDisks).

managed space mode. An access mode that enables virtualization functions to be performed. See also *image mode* and *unconfigured mode*.

Management Information Base (MIB). Simple Network Management Protocol (SNMP) units of managed information that specifically describe an aspect of a system, such as the system name, hardware number, or communications configuration. A collection of related MIB objects is defined as a MIB.

mapping. See *FlashCopy mapping*.

master virtual disk. The virtual disk (VDisk) that contains a production copy of the data and that an application accesses. See also *auxiliary virtual disk*.

MDisk. See *managed disk*.

migration. See *data migration*.

N

node. An addressable entity connected to an I/O bus or network. Used primarily to refer to computers, storage devices, and storage subsystems. The component of a node that connects to the bus or network is a port. (S)

node rescue. In SAN Volume Controller Storage Software for Cisco MDS 9000, the process by which a SAN Volume Controller Storage Software for Cisco MDS 9000 node that has no valid software installed on its hard disk drive can copy the software from another SAN Volume Controller Storage Software for Cisco MDS 9000 node connected to the same fibre-channel fabric.

non-RAID. Disks that are not in a redundant array of independent disks (RAID).

O

offline. Pertaining to the operation of a functional unit or device that is not under the continual control of the system or of a host.

online. Pertaining to the operation of a functional unit or device that is under the continual control of the system or of a host.

P

PLUN. See *managed disk*.

port ID. An identifier associated with a port.

primary virtual disk. In a Remote Copy relationship, the target of write operations issued by the host application.

PuTTY. A free implementation of Telnet and SSH for Windows® 32-bit platforms

Q

quorum disk. A managed disk (MDisk) that contains quorum data and that a cluster uses to break a tie and achieve a quorum.

R

rack. A free-standing framework that holds the devices and card enclosure.

RAID. See *redundant array of independent disks*.

redundant array of independent disks. A collection of two or more disk drives that present the image of a single disk drive to the system. In the event of a single device failure, the data can be read or regenerated from the other disk drives in the array.

redundant SAN. A storage area network (SAN) configuration in which any one single component might fail, but connectivity between the devices within the SAN is maintained, possibly with degraded performance. This configuration is normally achieved by splitting the SAN into two, independent, counterpart SANs. See also *counterpart SAN*.

relationship. In Remote Copy, the association between a master virtual disk and an auxiliary virtual disk (VDisk) and an auxiliary VDisk. These VDIsks also have the attributes of a primary or secondary VDisk. See also *auxiliary virtual disk*, *master virtual disk*, *primary virtual disk*, and *secondary virtual disk*.

Remote Copy. In SAN Volume Controller Storage Software for Cisco MDS 9000, a copy service that enables host data on a particular source virtual disk (VDisk) to be copied to the target VDisk designated in the relationship.

remote fabric. In Remote Copy, the storage area network (SAN) components (switches and cables) that connect the components (nodes, hosts, and switches) of the remote cluster.

roles. Authorization is based on roles that map to the administrator and service roles in an installation. The switch translates these roles into SAN Volume Controller Storage Software for Cisco MDS 9000 administrator and service user IDs when a connection is made to the node for the SAN Volume Controller Storage Software for Cisco MDS 9000.

S

SAN. See *storage area network*.

SDD. See *IBM Subsystem Device Driver*.

secondary virtual disk. In Remote Copy, the virtual disk (VDisk) in a relationship that contains a copy of data written by the host application to the primary VDisk.

Simple Network Management Protocol (SNMP). In the Internet suite of protocols, a network management protocol that is used to monitor routers and attached networks. SNMP is an application-layer protocol. Information on devices managed is defined and stored in the application's Management Information Base (MIB).

SNMP. See *Simple Network Management Protocol*.

storage area network (SAN). A network whose primary purpose is the transfer of data between computer systems and storage elements and among storage elements. A SAN consists of a communication infrastructure, which provides physical connections, and a management layer, which organizes the connections, storage elements, and computer systems so that data transfer is secure and robust. (S)

IBM Subsystem Device Driver (SDD). An IBM pseudo device driver designed to support the multipath configuration environments in IBM products.

switch. A network infrastructure component to which multiple nodes attach. Unlike hubs, switches typically have internal bandwidth that is a multiple of link bandwidth, and the ability to rapidly switch node connections from one to another. A typical switch can accommodate several simultaneous full link bandwidth transmissions between different pairs of nodes. (S) Contrast with *hub*.

symmetric virtualization. A virtualization technique in which the physical storage in the form of Redundant Array of Independent Disks (RAID) is split into smaller chunks of storage known as *extents*. These extents are then concatenated, using various policies, to make virtual disks (VDisks). See also *asymmetric virtualization*.

synchronized. In Remote Copy, the status condition that exists when both virtual disks (VDisks) of a pair that has a copy relationship contain the same data.

U

uninterruptible power supply. A device connected between a computer and its power source that protects the computer against blackouts, brownouts, and power surges. The uninterruptible power supply contains a power sensor to monitor the supply and a battery to provide power until an orderly shutdown of the system can be performed.

V

valid configuration. A configuration that is supported.

VDisk. See *virtual disk*.

virtual disk (VDisk). A set of disk blocks presented to an operating environment as a range of consecutively numbered logical blocks with disk-like storage and I/O semantics. The virtual disk is the disk array object that most closely resembles a physical disk from the operating environment's viewpoint. (S)

virtual storage area network (VSAN). A fabric within the SAN.

virtualization. In the storage industry, a concept in which a pool of storage is created that contains several disk subsystems. The subsystems can be from various vendors. The pool can be split into virtual disks that are visible to the host systems that use them.

virtualized storage. Physical storage that has virtualization techniques applied to it by a virtualization engine.

VLUN. See *virtual disk*.

VSAN. See *virtual storage area network*.

W

worldwide node name (WWNN). An identifier for an object that is globally unique. WWNNs are used by Fibre Channel and other standards.

WWNN. See *worldwide node name*.

WWPN. See *worldwide port name*.

worldwide port name (WWPN). A unique 64-bit identifier associated with a fibre-channel adapter port. The WWPN is assigned in an implementation- and protocol-independent manner.

Index

A

- accessibility 69
 - keyboard 69
 - shortcut keys 69

C

- cluster operation 16
- cluster state 16
- configuration
 - maximum 40
- configuration rules 42
 - HBAs 46
 - nodes 47
 - switches 47
- consistency group, Remote Copy 24
- consistency groups, FlashCopy 22
- copy services
 - overview 18
- creating
 - clusters
 - from the Console for SAN Volume Controller for Cisco MDS 9000 55

D

- deleting
 - nodes 60
- disability 69
- disk controllers
 - overview 28

F

- fibre channel host bus adapters (HBAs)
 - supported 64
- FlashCopy
 - consistency groups 22
 - mappings 20
 - overview 19

H

- host attachments
 - supported 63
- hosts
 - overview 37

I

- image mode 18
- information
 - center v
- installation
 - planning 11

K

- keyboard 69
 - shortcut keys 69

M

- MDisk (managed disk) groups
 - overview 31
- MDisks (managed disks)
 - overview 29
- migration 18

N

- nodes
 - adding 58
 - status 15
- notices 71

O

- object descriptions 27
- ordering publications viii
- overview
 - creating a cluster 55
 - disk controllers 17

P

- physical disk storage system
 - supported 64
- planning
 - configuration 39
 - installation 11
 - planning guidelines
 - SAN Volume Controller 13
- power domains
 - uninterruptible power supply 17
- product characteristics
 - SAN Volume Controller 12
- publications
 - ordering viii

R

- related information v
- Remote Copy
 - overview 23, 24

S

- SAN fabric 14
- SAN Volume Controller
 - introducing 5
- SAN Volume Controller Storage Software for Cisco MDS 9000 Console
 - using to create a cluster 55

- SAN. See storage area network (SAN) 13
- shortcut keys 69
- specifications
 - SAN Volume Controller 12
- storage area network (SAN) 13
- support
 - Web sites vii
- supported fibre channel host bus adapters (HBAs) 64
- supported host attachments 63
- supported physical disk storage system 64
- supported switches 64
- switches
 - supported 64
- synchronous copy
 - overview 23

T

- trademarks 72

V

- VDisks (virtual disks)
 - overview 33
- virtual disk-to-host mapping
 - description 36
- virtual disks (VDisks)
 - overview 33
- virtualization
 - asymmetric 3
 - overview 1
 - symmetric 4

W

- Web sites vii

Readers' Comments — We'd Like to Hear from You

IBM TotalStorage SAN Volume Controller
Storage Software for Cisco MDS 9000
Getting Started
Version 1 Release 1

Publication No. GA22-1055-00

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? Yes No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name

Address

Company or Organization

Phone No.



Fold and Tape

Please do not staple

Fold and Tape




NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
RCF Processing Department
Dept. FTFA/Building 050-1
5600 Cottle Road
San Jose, CA
U.S.A. 95193-0001



Fold and Tape

Please do not staple

Fold and Tape



Printed in USA

GA22-1055-00



Spine information:

IBM TotalStorage SAN Volume
Controller
Storage Software for Cisco MDS
9000



SAN Volume Controller Storage Software for Cisco
MDS 9000 Getting Started

Version 1 Release 1