IBM TotalStorage SAN Volume Controller

Configuration Guide

Version 2.1.0

SC26-7543-04

IBM TotalStorage SAN Volume Controller

Configuration Guide

Version 2.1.0

SC26-7543-04

Fifth Edition (February 2005)

Before using this information and the product it supports, read the information in ″Notices.″

Contents

Figures

Tables

About this guide

The IBM TotalStorage SAN Volume Controller Configuration guide provides information that helps you configure and use the IBM® TotalStorage®SAN Volume Controller™.

The IBM TotalStorage SAN Volume Controller Configuration guide also describes the configuration tools, both command-line and Web based, that you can use to define, expand, and maintain the storage of the IBM TotalStorage SAN Volume Controller.

Who should use this guide

The IBM TotalStorage SAN Volume Controller Configuration guide is intended for system administrators or others who install and use the SAN Volume Controller.

Before using the IBM TotalStorage SAN Volume Controller, you should have an understanding of storage area networks (SANs), the storage requirements of your enterprise, and the capabilities of your storage units.

Related reference

"About this guide"
The IBM TotalStorage SAN Volume Controller Configuration guide provides information that helps you configure and use the IBM® TotalStorage®SAN Volume Controller™.

Summary of changes

This document contains terminology, maintenance, and editorial changes.

Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change. This summary of changes describes new functions that have been added to this release.

Related reference

"Summary of Changes for SC26-7543-03 SAN Volume Controller Configuration Guide" on page xv
The Summary of Changes provides a list of new, modified, and changed information since the last version of the guide.

Summary of changes for SC26-7543-04 SAN Volume Controller Configuration Guide

This summary of changes provides a list of new, modified, and deleted information since the previous edition of this guide.

New information

This edition includes the following new information:

- The SAN Volume Controller is paired with a new uninterruptible power supply (UPS). The UPS 5115 is the preferred UPS for the SAN Volume Controller. Previous releases of the SAN Volume Controllerwere paired with the UPS 5125, which is still supported.

If text is referring to the “UPS” or the “uninterruptible power supply,” then it is referring to a generic UPS and can be referring to either UPS. When the UPS is referred to as the “UPS 5115” or the “UPS 5125,” then the specific UPS is designated.

- Support for Microsoft Volume Shadow Copy Service has been added.
- Instructions for copying a VDisk between managed disk (MDisk) groups that have different extent sizes using Copy Services have been added.
- Added the following new topics:
  - Extents
  - Input/Output governing
  - Metro Mirror
  - Limiting queue depth in large SANs
  - Queue depth
  - Calculating a queue depth limit
  - Homogeneous queue depth calculation
  - Non-homogeneous queue depth calculation
  - Limiting the queue depth
  - Creating an image mode virtual disk
  - Migration to image mode
  - Deleting an image mode virtual disk
  - Deleting a managed mode virtual disk
  - Creating an image mode VDisk using the CLI
  - Migrating to an image mode virtual disk using the CLI
  - Quorum disk creation and extent allocation
  - Access Logix
  - User interface on Symmetrix
  - User interface on ESS
  - TrueCopy
  - Write protect
  - Configuring settings for Lightning
  - Global settings for Lightning
  - Controller settings for Lightning
  - Port settings for Lightning
  - LU settings for Lightning
  - Zoning guidelines
  - IBM TotalStorage support for Microsoft Volume Shadow Copy service
  - Installation overview
  - System requirements for the IBM TotalStorage support for Microsoft Volume Shadow Copy service
  - Installing the IBM TotalStorage hardware provider
  - Verifying the installation
  - Creating the free and reserved pools of volumes
  - Configuration commands
  - Adding and removing VDisks
  - Error codes
  - Uninstalling the IBM TotalStorage hardware provider

Changed information

This section lists the updates that were made in this edition of the guide.

- Remote Copy is now known as IBM TotalStorage Metro Mirror for SAN Volume Controller.
- The maximum number of Metro Mirror relationships per cluster was increased to 4096.
- Information about image mode virtual disks has been enhanced.
- Information about managed disk (MDisk) extents has been enhanced.
- The SAN Volume Controller now supports connection to SANs that contain switches from different manufacturers.
- Instructions for calculating queue depth have been expanded.
- Complete information about the master console hardware and software have been moved to *IBM TotalStorage Master Console Installation and User's Guide*.

Deleted information

This section lists the information that was removed from this edition of the guide.

- Information about the SAN Volume Controller master console was removed. It now appears in the *IBM TotalStorage Master Console Installation and User's Guide*.

Summary of Changes for SC26-7543-03 SAN Volume Controller Configuration Guide

The Summary of Changes provides a list of new, modified, and changed information since the last version of the guide.

New information

This topic describes the changes to this guide since the previous edition, SC26-7543-02. The following sections summarize the changes that have since been implemented from the previous version.

This version includes the following new information:

- Added the following new topics:
  - Background Copy bandwidth impact on foreground I/O latency
  - Configuration recommendations for large SANs
  - Queue depth
  - Calculating a queue depth limit
  - Homogeneous queue depth calculation
  - Non-Homogeneous queue depth calculation
  - Limiting the queue depth
  - Validating the truststore certificate expiration
  - Using image mode virtual disks
  - Storing preexisting data onto the SAN Volume Controller
  - Exposing LUN's on your existing storage to the cluster via switch zoning
  - Image mode VDisk migration
  - Migration methods
  - Importing data by creating an image mode VDisk

- Reviewing and setting the cluster features using the SAN Volume Controller Console
- Updated the EMC CLARiiON subsystem section with the following new topics:
  - Access Logix
  - User interface on CLARiiON
- Updated the EMC Symmetrix subsystem section with the following new topic:
  - User interface on Symmetrix
- Updated the ESS subsystem section with the following new topics:
  - User interface on ESS
  - Target port groups for the ESS
- Updated the FAStT subsystem section with the following new topic:
  - User interface on FAStT
- Updated the HDS Lightning subsystem section with the following new topics:
  - User interface on Lightning
  - LU configuration for Lightning
  - Logical unit creation and deletion on HDS Lightning
  - Configuring settings for Lightning
  - Global settings for Lightning
  - Controller settings for Lightning
  - Port settings for Lightning
  - LU settings for Lightning
- Updated the HDS Thunder subsystem section with the following new topic:
  - User interface on the HDS Thunder
- Added support for the HPQ Enterprise Virtual Array (EVA) subsystem. This section includes the following new topics:
  - Supported models of the HPQ EVA
  - Supported firmware levels for HPQ EVA
  - User interface on HPQ EVA
  - Sharing the HPQ EVA controller between a host and the SAN Volume Controller
  - Switch zoning limitations for the HPQ EVA subsystem
  - Quorum disks on HPQ EVA
  - Support for HPQ EVA advanced functions
  - Logical unit configuration on the HPQ EVA
  - Logical unit creation and deletion on the HPQ EVA
  - Logical unit presentation
  - Configuration interface for the HPQ EVA
  - Configuring settings for the HPQ EVA
  - Global settings for the HPQ EVA
  - LU settings for the HPQ EVA
  - Host settings for the HPQ EVA

Changed information

This section lists the updates that were made in this document.

- Support for 4-node configurations has been updated to support 8-node.

- Increased the number of supported VDisks per I/O group to 1024.
- Increased the number of supported VDisks per cluster to 4096.
- Updated support information for split controller configurations.
- Updated the number of uninterruptible power supply (UPS) 5115 to support up to 8 nodes.
- Updated the maximum configurations.
- Updated configuration examples to include a 8-node configuration.
- Modified the step-by-step instructions for restoring a cluster configuration.
- Updated the Switch zoning limitations for the EMC CLARiiON subsystem.
- Updated the Global settings for the EMC Symmetrix subsystem.
- Updated the Supported models of the IBM FAStT controller.

Deleted information

No deletions were made in this version of the guide.

Related reference

"Summary of changes" on page xiii
This document contains terminology, maintenance, and editorial changes.

Emphasis

Different typefaces are used in this guide to show emphasis.

The following typefaces are used to show emphasis:

Boldface	Text in **boldface** represents menu items and command names.
Italics	Text in *italics* is used to emphasize a word. In command syntax, it is used for variables for which you supply actual values, such as a default directory or the name of a cluster.
`Monospace`	Text in `monospace` identifies the data or commands that you type, samples of command output, examples of program code or messages from the system, or names of command flags, parameters, arguments, and name-value pairs.

Numbering conventions

A specific numbering convention is used in this guide and product.

The following numbering conventions are used in this guide and in the product:
- 1 kilobyte (KB) is equal to 1024 bytes
- 1 megabyte (MB) is equal to 1 048 576 bytes
- 1 gigabyte (GB) is equal to 1 073 741 824 bytes
- 1 terabyte (TB) is equal to 1 099 511 627 776 bytes
- 1 petabyte (PB) is equal to 1 125 899 906 842 624 bytes

Related reference

"About this guide" on page xiii
The IBM TotalStorage SAN Volume Controller Configuration guide provides information that helps you configure and use the IBM® TotalStorage®SAN Volume Controller™.

"Who should use this guide" on page xiii
The IBM TotalStorage SAN Volume Controller Configuration guide is intended for system administrators or others who install and use the SAN Volume Controller.

SAN Volume Controller library and related publications

A list of other publications that are related to this product are provided to you for your reference.

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
- Other IBM publications that relate to the SAN Volume Controller

SAN Volume Controller library

The following table lists and describes the publications that make up the SAN Volume Controller 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. If you need additional copies of this CD, the order number is SK2T-8811. These publications are also available as PDF files from the following Web site:

http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Title	Description	Order number
IBM TotalStorage SAN Volume Controller: CIM Agent Developer's Reference	This reference guide describes the objects and classes in a Common Information Model (CIM) environment.	SC26-7590
IBM TotalStorage SAN Volume Controller: Command-Line Interface User's Guide	This guide describes the commands that you can use from the SAN Volume Controller command-line interface (CLI).	SC26-7544
IBM TotalStorage SAN Volume Controller: Configuration Guide	This guide provides guidelines for configuring your SAN Volume Controller.	SC26-7543
IBM TotalStorage SAN Volume Controller: Host Attachment Guide	This guide provides guidelines for attaching the SAN Volume Controller to your host system.	SC26-7575
IBM TotalStorage SAN Volume Controller: Installation Guide	This guide includes the instructions the service representative uses to install the SAN Volume Controller.	SC26-7541

Title	Description	Order number
IBM TotalStorage SAN Volume Controller: Planning Guide	This guide introduces the SAN Volume Controller and lists the features you can order. It also provides guidelines for planning the installation and configuration of the SAN Volume Controller.	GA22-1052
IBM TotalStorage SAN Volume Controller: Service Guide	This guide includes the instructions the service representative uses to service the SAN Volume Controller.	SC26-7542
IBM TotalStorage SAN Volume Controller: Translated Safety Notices	This guide contains the danger and caution notices for the SAN Volume Controller. The notices are shown in English and in numerous other languages.	SC26-7577
IBM TotalStorage Master Console Installation and User's Guide	This guide includes the instructions on how to install and use the SAN Volume Controller Console	

Other IBM publications

The following table lists and describes other IBM publications that contain additional information related to the SAN Volume Controller.

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

Related reference

"How to order IBM publications" on page xx
The publications center is a worldwide central repository for IBM product publications and marketing material.

Related information

"How to send your comments" on page xx

Related Web sites

The following Web sites provide information about the SAN Volume Controller or related products or technologies.

Type of information	Web site
SAN Volume Controller support	http://www-1.ibm.com/servers/storage/support/virtual/2145.html
Technical support for IBM storage products	http://www.ibm.com/storage/support/

How to order IBM publications

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

The IBM publications center

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 reference

"SAN Volume Controller library and related publications" on page xviii
A list of other publications that are related to this product are provided to you for your reference.

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 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

  Fill out the Readers' Comments form (RCF) at the back of this book. If the RCF has been removed, you can address your comments to:

International Business Machines Corporation
RCF Processing Department
Department 61C
9032 South Rita Road
Tucson, Arizona 85775-4401
U.S.A.

Related reference

"SAN Volume Controller library and related publications" on page xviii
A list of other publications that are related to this product are provided to you for your reference.

Chapter 1. Overview

This part provides an overview of the SAN Volume Controller.

Related concepts

"SAN Volume Controller"
The SAN Volume Controller is a SAN appliance that attaches open-systems storage devices to supported open-systems hosts.

Related reference

"Object overview" on page 9
The SAN Volume Controller is based on a number of virtualization concepts.

"Copy Services" on page 34
Two Copy Services are supported by the SAN Volume Controller.

"Configuration rules and requirements" on page 51
Ensure that you understand the rules and requirements when configuring the SAN Volume Controller.

SAN Volume Controller

The SAN Volume Controller is a SAN appliance that attaches open-systems storage devices to supported open-systems hosts.

The IBM® TotalStorage® SAN Volume Controller 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 attached host computer systems. System administrators can view and access a common pool of storage on the SAN, which enables them to use storage resources more efficiently and provides a common base for advanced functions.

The SAN Volume Controller is analogous to a logical volume manager (LVM) on a SAN. It performs the following functions for the SAN storage that it is controlling:

- Creates a single pool of storage
- Logical unit virtualization
- Manages logical volumes
- Provides advanced functions for the SAN, such as:
  - Large scalable cache
  - Copy services
    - FlashCopy® (point-in-time copy)
    - Metro Mirror (synchronous copy)
    - Data migration
  - Space management
    - Mapping that is based on desired performance characteristics
    - Quality of service metering

A *node* is a single storage engine. Figure 1 on page 2 provides an illustration of a node. The storage engines are always installed in pairs with one to four pairs of nodes constituting a *cluster*. Each node in a pair is configured to back up the other. Each pair of nodes is known as an *I/O group*. All I/O operations that are managed by the nodes in an I/O group are cached on both nodes for resilience. Each virtual volume is defined to an I/O group. To avoid any single point of failure, the nodes of an I/O group are protected by independent uninterruptible power supply (UPS) 5115 or uninterruptible power supply (UPS) 5125 units.

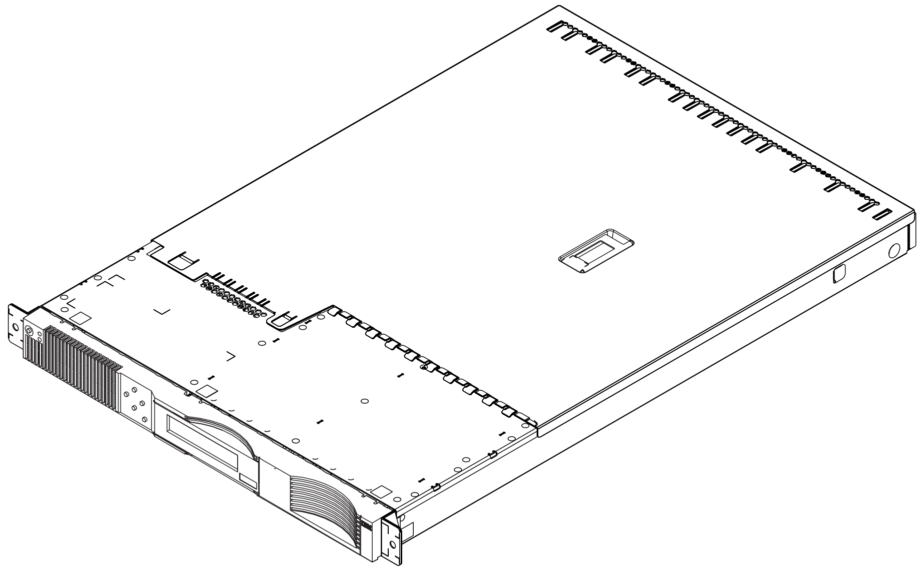

Figure 1. A SAN Volume Controller node

The SAN Volume Controller I/O groups recognize the storage presented to the SAN by the backend controllers as a number of disks known as *managed disks*. The application servers do not recognize 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. Each node must be in only one I/O group and provide access to the virtual disks in the I/O group.

The SAN Volume Controller helps to provide continuous operations and can also optimize the data path to ensure performance levels are maintained. Ensure that you use IBM TotalStorage Multiple Device Manager performance manager to analyze the performance statistics. See *IBM TotalStorage Multiple Device Manager Configuration and Installation Guide* and *IBM TotalStorage Multiple Device Manager CLI Guide* for more information.

The fabric contains two distinct zones: a host zone and a disk zone. In the host zone, the host systems can identify and address the nodes. You can have more than one host zone. Generally, you will create one host zone per operating system type. In the disk zone, the nodes can identify the disk drives. Host systems cannot operate on the disk drives directly; all data transfer occurs through the nodes. Figure 2 on page 3 shows that several host systems can be connected to a SAN fabric. A cluster of SAN Volume Controller nodes is connected to the same fabric and presents virtual disks to the host systems. You create these virtual disks from units of space within a managed disk group. A managed disk group is a collection of managed disks presented by the back-end RAID controllers, providing a storage pool. You choose how each group is made up, and you can combine managed disks from different manufacturers' controllers in the same managed disk group if required.

Note: Some operating systems cannot tolerate other operating systems in the same zone.

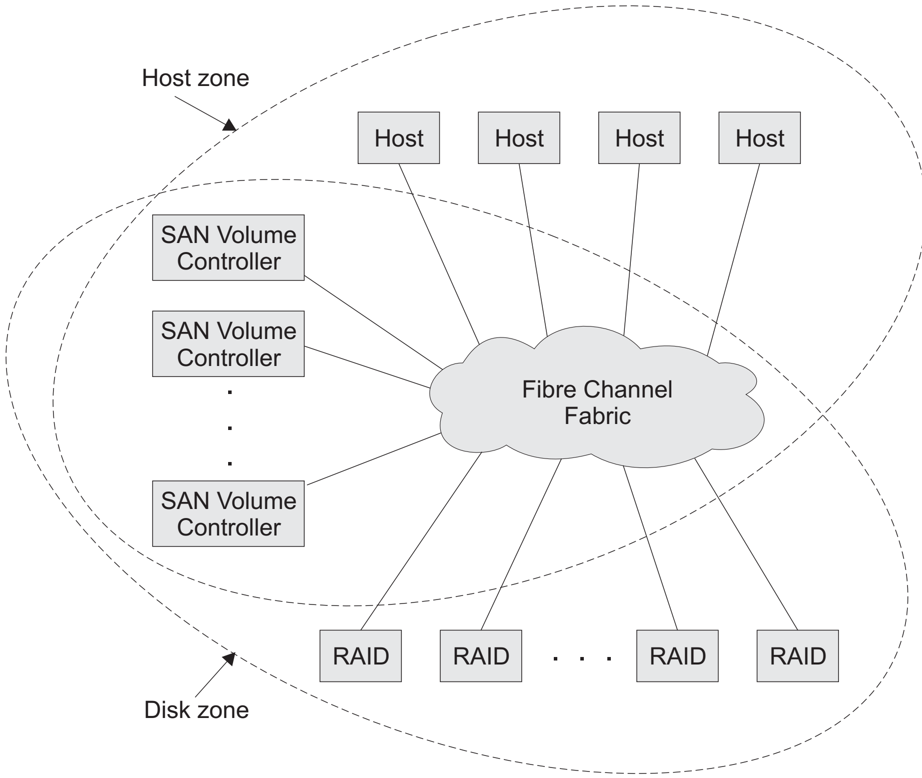

Figure 2. Example of a SAN Volume Controller in a fabric

You can remove one node in each I/O group from a cluster when hardware service or maintenance is required. After you remove the node, you can replace the field replaceable units (FRUs) in the node. All disk drive communication and communication between nodes is performed through the SAN. All SAN Volume Controller configuration and service commands are sent to the cluster through an Ethernet network.

Each node contains its own vital product data (VPD). Each cluster contains VPD that is common to all the nodes on the cluster, and any system connected to the Ethernet network can access this VPD.

Cluster configuration information is stored on every node that is in the cluster to allow concurrent replacement of FRUs. An example of this information might be information that is displayed on the menu screen of the SAN Volume Controller. When a new FRU is installed and when the node is added back into the cluster, configuration information that is required by that node is read from other nodes in the cluster.

SAN Volume Controller operating environment

- Minimum of one pair of SAN Volume Controller nodes
- Minimum of two uninterruptible power supplies
- One master console is required per SAN installation for configuration

Features of a SAN Volume Controller node

- 19-inch rack mounted enclosure
- 4 fibre channel ports
- 2 fibre channel adapters

- 4 GB cache memory

Supported hosts

For a list of supported operating systems, see the IBM TotalStorage SAN Volume Controller Web site at:

http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Multipathing software
- IBM Subsystem Device Driver (SDD)
- Redundant Dual Active Controller (RDAC)

| **Note:** Directly-attached hosts that share a back-end storage controller with a SAN
| Volume Controller can run multipath drivers SDD and RDAC simultaneously.
| Other multipath drivers running with SDD are not supported.

Check the following Web site for the latest support and coexistence information:

http://www-1.ibm.com/servers/storage/support/virtual/2145.html

User interfaces

The SAN Volume Controller provides the following user interfaces:
- IBM TotalStorage SAN Volume Controller Console, a Web-accessible graphical user interface (GUI) that supports flexible and rapid access to storage management information
- A command-line interface (CLI) using Secure Shell (SSH)

Application programming interfaces

The SAN Volume Controller provides the following application programming interface:
- IBM TotalStorage Common Information Model (CIM) Agent for the SAN Volume Controller, which supports the Storage Management Initiative Specification of the Storage Network Industry Association.

Related concepts

"Virtualization"
Virtualization is a concept that applies to many areas of the information technology industry.

"Asymmetric virtualization" on page 7
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.

"Symmetric virtualization" on page 8
The SAN Volume Controller provides symmetric virtualization.

Virtualization

Virtualization is a concept that applies to many areas of the information technology industry.

For data storage, virtualization includes the creation of a pool of storage that contains several disk subsystems. These subsystems can be supplied 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 a storage area network (SAN).

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 managing physical volumes of data to 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 suitable for environments that do not have storage networks.
- Virtualization at the *storage device* level is in common use. Striping, mirroring, and redundant arrays of independent disks (RAIDs) 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 servers and 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 performs virtualization at the fabric level.
- Virtualization at the *file system* level provides the highest benefit because data is shared, allocated, and protected, not volumes.

Virtualization is a radical departure from traditional storage management. In traditional storage management, storage is attached directly to a host system, which controls storage management. SANs introduced the principle of networks of storage, but storage is still primarily created and maintained at the 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 managing storage maintenance.

Where storage is concerned, one problematic area that virtualization addresses is 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 that need the highest storage capacity can use it when they need it. Regulating the amount of available storage 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. Figure 3 on page 6 provides a diagram of the levels of virtualization.

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. For example, you can gain the most benefits 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 implements fabric-level *virtualization*. Within the context of the SAN Volume Controller and throughout this document, *virtualization* refers to symmetric fabric-level virtualization.

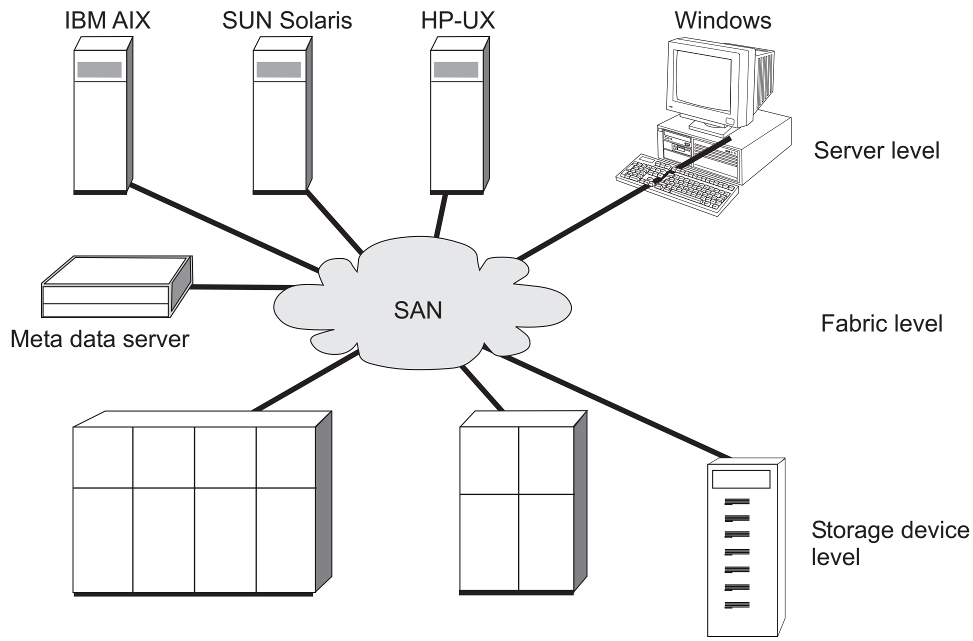

Figure 3. Levels of virtualization

Related concepts

"Asymmetric virtualization" on page 7
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.

"Symmetric virtualization" on page 8
The SAN Volume Controller provides symmetric virtualization.

"Virtual disks" on page 29
A *virtual disk (VDisk)* is a logical disk that the cluster presents to the storage area network (SAN).

Related reference

"Configuring and servicing storage subsystems" on page 233
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

Asymmetric virtualization

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.

In asymmetric virtual storage networks, the data flow, (2) in the figure below, is separated from the control flow, (1). A separate network or SAN link is used for control purposes. The metadata server contains all the mapping and locking tables while the storage devices contain only data. Because the flow of control is separated from the flow of data, I/O operations can use the full bandwidth of the SAN. A separate network or SAN link is used for control purposes. There are disadvantages, however, to asymmetric virtualization.

The disadvantages to asymmetric virtualization include:

- 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 drivers 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. See Figure 4

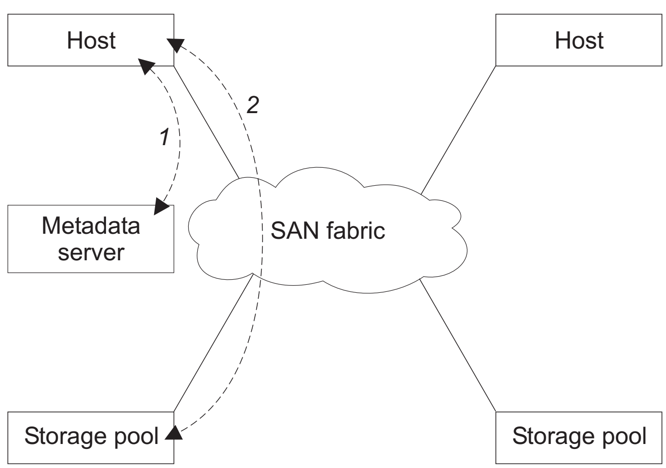

Figure 4. Asymmetrical virtualization

For one, data is at risk to increased security exposures and the control network must be protected with a firewall. In addition, metadata can become very complicated when files are distributed across several devices. Moreover, 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. Finally, 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.

Related concepts

“Symmetric virtualization” on page 8
The SAN Volume Controller provides symmetric virtualization.

Symmetric virtualization

The SAN Volume Controller provides symmetric virtualization.

Virtualization splits the physical storage Redundant Array of Independent Disks (RAID) arrays into smaller chunks of storage that are known as extents. These extents are then concatenated, 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 re-configure the host. With symmetric virtualization, the virtualization engine is the central configuration point for the SAN.

In symmetric virtual storage networks (see Figure 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.

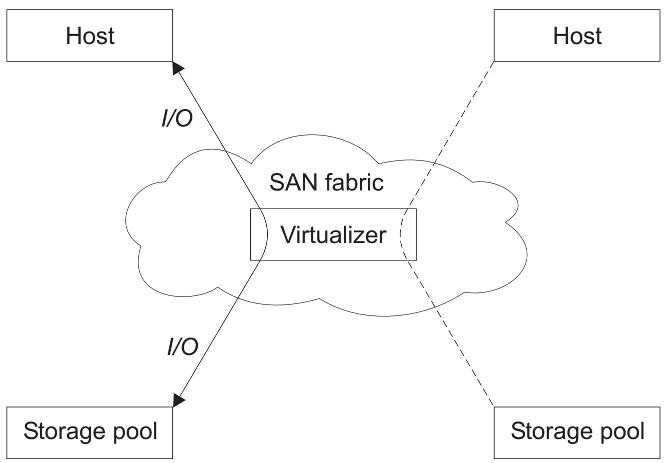

Figure 5. Symmetrical virtualization

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 run the advanced functions, such as Copy Services and caching.

The IBM TotalStorage SAN Volume Controller uses symmetric virtualization. Single virtualization engines, which are known as nodes, are combined to create clusters. Each cluster can contain between two and eight nodes.

Related concepts

“Asymmetric virtualization” on page 7
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.

“Virtualization” on page 4
Virtualization is a concept that applies to many areas of the information technology industry.

Physical links between SAN Volume Controller nodes and a switch

Between the SAN Volume Controller nodes and the switch to which they are connected, SAN Volume Controller will support shortwave SFP transceivers (850 nm with 50 μm or 62.5 μm multimode cables).

The transceivers will run at up to 500 m, limited by the pulse spreading caused by the multimode nature of the transmission. Multimode means that each ray of light can take slightly different paths from one part of the cable to another. The paths are different lengths which means a very short pulse of light put in at one end will somewhat spread out when it reaches the far end. Not all the light associated with that pulse travels exactly the same distance to get out of the of the fibre and the longer the fibre, the more the pulse gets dispersed. For longer distances between the nodes and the local switch, SAN Volume Controller supports long wave SFP transceivers (1300 nm with 9 μm unimode2 cables). These are supported up to 10 km in length.

Support for long links between the local and remote fabric

Ensure you are familiar with the support statement for the ISL between the local and remote fabric.

The following Finisar high power Small Form Factor Pluggable (SFP) transceivers are supported to 10 km:

- FTRJ-1419-7D-2.5

  Note: Finisar supports this up to 35 km
- FTRJ-1519P1BCL

  Note: Finisar supports this up to 80 km

The SAN Volume Controller supports the following DWDM extenders and Dark Fibre solutions. These are *not* supported within the local or remote fabrics.

- CNT Ultranet Edge Storage Router
  - IP (100 Mb/s and 1 Gb), ATM and SONET networks

Object overview

The SAN Volume Controller is based on a number of virtualization concepts.

A SAN Volume Controller consists of **a single node**. Nodes are deployed in pairs to make up a **cluster**. A cluster can have one to four node pairs in it. Each pair of nodes is known as an **I/O group**. Each node must be in only one I/O group.

Virtual disks (VDisks) are logical disks that are presented to the SAN by nodes. Virtual disks are also associated with an 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 provides is only two way.

Each node does not contain any internal battery backup units and therefore must be connected to an **uninterruptible power supply (UPS)** to provide data integrity in the event of a cluster-wide power failure. In such situations, the UPS will maintain power to the nodes while the contents of the distributed cache are dumped to an internal drive.

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

Each managed disk is divided into a number of **extents** which are numbered from zero, sequentially, from the start to the end of the managed disk. The extent size must be specified when creating an MDisk group.

Managed disks are collected into groups, known as **managed disk groups (MDisk groups)**. 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 detects the Fibre Channel ports that are connected to the SAN. These correspond to the Host Bus Adapter (HBA) Fibre Channels worldwide port names (WWPNs) that are present in the application servers. The SAN Volume Controller allows you to create logical host objects that group together WWPNs belonging to a single application server or multiple application servers.

Application servers can only access virtual disks that have been allocated to them. Virtual disks can 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.

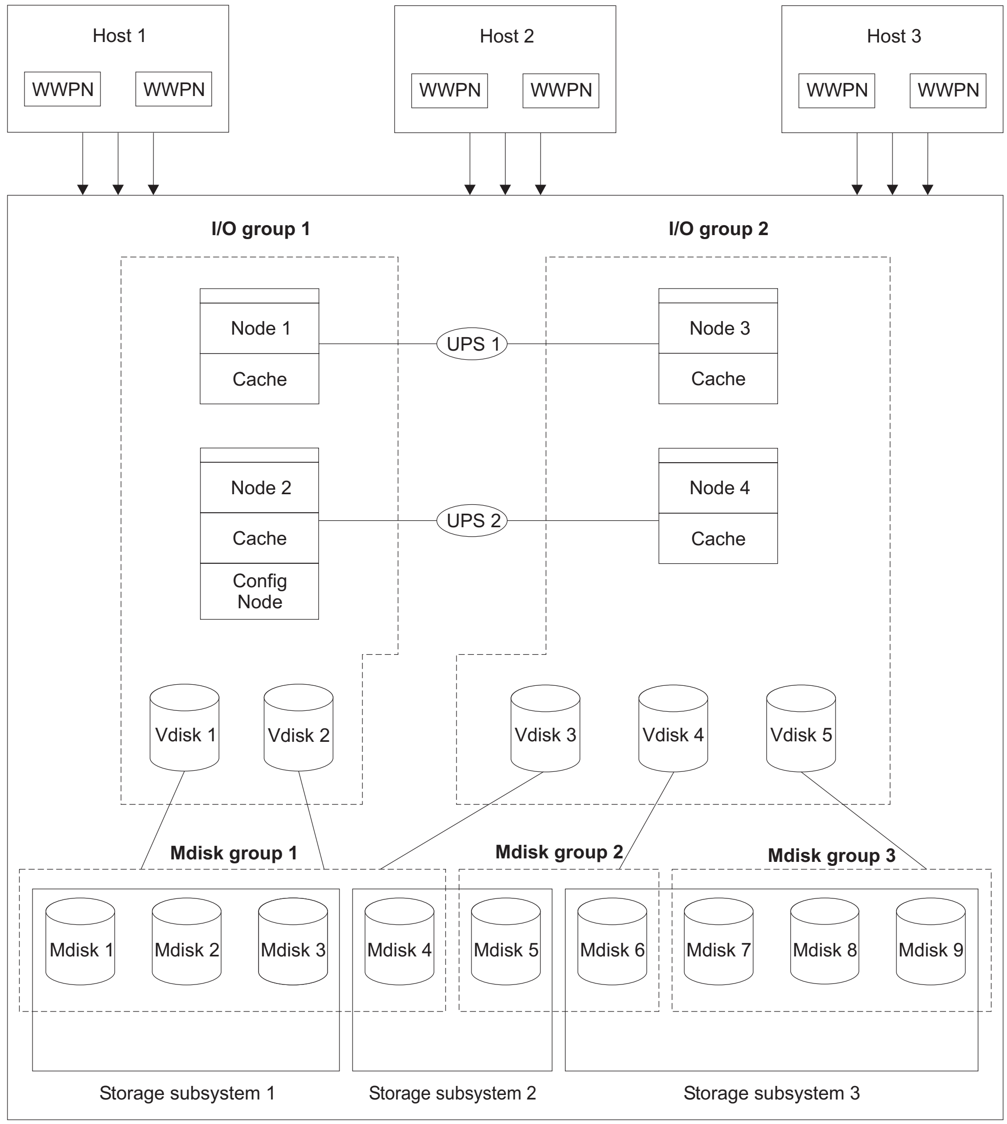

Figure 6. Virtualization

Related concepts

"Nodes and clusters" on page 12
A SAN Volume Controller node is a single processing unit, which provides virtualization, cache, and copy services for the SAN.

"I/O groups and uninterruptible power supply" on page 16
Nodes are deployed in pairs to make up a cluster. Each pair of nodes is known as an **I/O group**. Each node may be in *only* one I/O group.

"Storage subsystems and managed disks" on page 21
The nodes in a cluster see the storage exported by SAN-attached storage subsystems as a number of disks, known as managed disks. The SAN Volume Controller does not attempt to provide recovery from physical disk failures within the storage subsystem. A managed disk is usually, but not necessarily, a RAID array.

"Managed disk groups and virtual disks" on page 26
Managed disks are collected into groups known as managed disk groups. Virtual disks are logical disks that are presented to the SAN by SAN Volume Controller nodes. The maximum number of supported VDisks per I/O group is 1024. The

maximum number of supported VDisks per cluster is 4096. Virtual disks, like nodes, are associated with an I/O group.

"Hosts and virtual (VDisk) mappings" on page 31
Application servers can only access VDisks that have been made accessible to them.

Nodes and clusters

A SAN Volume Controller node is a single processing unit, which provides virtualization, cache, and copy services for the SAN.

Nodes are deployed in pairs called I/O groups. One node in the cluster is designated the configuration node but each node in the cluster holds a copy of the cluster state information.

Related concepts

"Clusters"
All configuration and service is performed at the cluster level.

Related reference

"Nodes" on page 14
A SAN Volume Controller node is a single processing unit within a SAN Volume Controller cluster.

Clusters

All configuration and service is performed at the cluster level.

A cluster can consist of two nodes, with a maximum of eight nodes. Therefore, you can assign up to eight SAN Volume Controller nodes to one cluster.

Some service actions can be performed at node level, but all configuration is replicated across all nodes in the cluster. Because configuration is performed at the cluster level, an IP address is assigned to the cluster instead of each node.

All your configuration and service actions are performed at the cluster level. Therefore, after configuring your cluster, you can take advantage of the virtualization and the advanced features of the SAN Volume Controller.

Cluster state and the configuration node

The cluster state holds all configuration and internal cluster data for the cluster. This cluster state information is held in nonvolatile memory. If the mainline power fails, the two uninterruptible power supplies maintain the internal power long enough for the cluster state information to be stored on the internal disk drive of each node. The read and write cache information is also held in nonvolatile memory. Similarly, if the power fails to a node, configuration and cache data for that node is lost and the partner node attempts to flush the cache. The cluster state is still maintained by the other nodes on the cluster.

Figure 7 on page 13 shows an example cluster containing four nodes. The cluster state shown in the grey box does not actually exist, instead each node holds a copy of the entire cluster state.

The cluster contains a single node that is elected as the configuration node. The configuration node can be thought of as the node that controls the updating of cluster state. For example, a user request is made (item *1*), that results in a change being made to the configuration. The configuration node controls updates to the cluster (item *2*). The configuration node then forwards the change to all nodes (including Node 1), and they all make the state-change at the same point in time (item *3*). Using this state-driven model of clustering ensures that all nodes in the cluster know the exact cluster state at any one time.

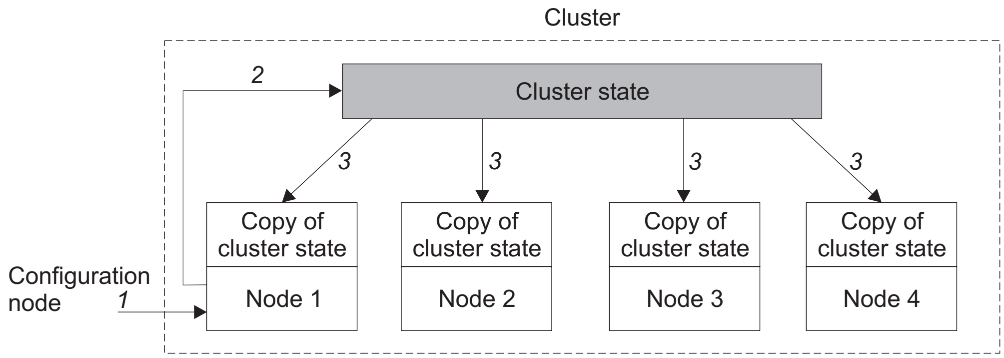

Figure 7. Cluster, nodes, and cluster state.

Cluster configuration backup:

Configuration backup is the process of extracting configuration data from a cluster and writing it to disk.

Backing up the cluster configuration enables you to restore it in the event that configuration data is lost. The data that is backed up is the metadata that describes the cluster configuration, not the data that your enterprise uses to run its business.

The backup configuration files can be saved on the master console or the configuration node.

Objects included in the backup

Configuration data is information about a cluster and the objects that are defined in it. The following objects are copied:

- Storage subsystem
- Hosts
- I/O groups
- Managed disks (MDisks)
- MDisk groups
- Nodes
- Virtual disks (VDisks)
- VDisk-to-host mappings
- SSH key
- FlashCopy mappings
- FlashCopy consistency groups
- Metro Mirror relationships
- Metro Mirror consistency groups

Related concepts

“Clusters” on page 12
All configuration and service is performed at the cluster level.

"Configuration restore"
Configuration restore is the process of using a backup configuration file, or files, on the master console or configuration node to restore a specific cluster configuration.

Configuration restore:

Configuration restore is the process of using a backup configuration file, or files, on the master console or configuration node to restore a specific cluster configuration.

Restoring a cluster configuration involves restoring the metadata that describes the cluster configuration, not the data your enterprise uses to run its business. Restoring the cluster configuration is an important part of a complete backup and disaster recovery solution. However, you must make provision for non-configuration data to be restored as well.

This process consists of two phases:
- Preparing
- Executing

| Before issuing the preparation command or phase, the cluster itself must be reset to a default state. During the preparation phase, the backup data and the new cluster are analyzed for compatibility, and a sequence of commands is prepared.

During the execution phase, the command sequence is run.

Related concepts

"Clusters" on page 12
All configuration and service is performed at the cluster level.

Chapter 5, "Backing up and restoring the cluster configuration," on page 213
You can back up and restore the cluster configuration.

"Cluster configuration backup" on page 13
Configuration backup is the process of extracting configuration data from a cluster and writing it to disk.

Nodes

A SAN Volume Controller node is a single processing unit within a SAN Volume Controller cluster.

For redundancy, nodes are deployed in pairs to make up a cluster. A cluster can have one to four pairs of nodes in it. Each pair of nodes is known as an I/O group. Each node can be in *only* one I/O group. A maximum of four I/O groups each containing two nodes is supported.

At any one time, a single node in the cluster manages 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 takes over its responsibilities.

Table 1 describes the operational states of a node.

Table 1. Node state

State	Description
Adding	The node was added to the cluster but is 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 was assigned to a cluster but is not available on the fibre-channel SAN fabric. Run the Directed Maintenance Procedures to determine the problem.
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, which can take up to 20 minutes. During this time the node will be shown as adding.	

Related concepts

"Configuration node"
At any given time, one node manages configuration activity. This node is the *configuration node*.

Configuration node:

At any given time, one node manages configuration activity. This node is the *configuration node*.

The configuration node is a focal point for configuration commands, and it manages the data that describes the cluster configuration.

If the configuration node fails, the cluster chooses a new configuration node. This action is called configuration node failover. The switch that contains the new node takes over the cluster IP address. Thus you can access the cluster through the same IP address although the original configuration node has failed. During the failover, there is a short period when you cannot use the command line tools or SAN Volume Controller Console.

The figure below shows an example cluster containing four nodes. Node 1 has been designated the configuration node. User requests (*1*) are targeted at Node 1. This may result in requests (*2*) being targeted at the other nodes in the cluster, and data being returned to Node 1. See Figure 8 on page 16 for more information.

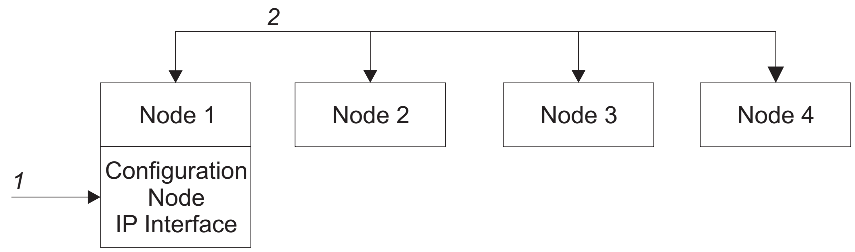

Figure 8. Configuration node

Related reference

"Nodes" on page 14
A SAN Volume Controller node is a single processing unit within a SAN Volume Controller cluster.

I/O groups and uninterruptible power supply

Nodes are deployed in pairs to make up a cluster. Each pair of nodes is known as an **I/O group**. Each node may be in *only* one I/O group.

Virtual disks are logical disks that are presented to the SAN by SAN Volume Controller nodes. Virtual disks are also associated with an I/O group. The SAN Volume Controller does not contain any internal battery backup units and therefore must be connected to an uninterruptible power supply to provide data integrity in the event of a cluster-wide power failure.

Input/Output (I/O) groups

An I/O group is a group that is defined during the cluster configuration process.

It usually contains two SAN Volume Controller nodes for availability purposes. However, depending on the configuration, an I/O group may be empty or just contain a single node. Each node is associated with only one I/O group, and each virtual disk (VDisk) is associated with only one I/O group. The nodes in the I/O group provide access to the VDisks 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. A virtual disk can specify a preferred node. This is specified when the virtual disk is created. This is the node through which a virtual disk should normally be accessed. As each I/O group only has two nodes, the distributed cache in the SAN Volume Controller need only be 2-way. When I/O is performed to a virtual disk, the node that processes the I/O duplicates the data onto the partner node that is in the I/O group.

I/O traffic for a particular virtual disk is, at any one time, handled exclusively by the nodes in a single I/O group. Thus, although a cluster may have many nodes within it, the nodes handle I/O in independent pairs. This means that the I/O capability of the SAN Volume Controller scales well, since additional throughput can be obtained by adding additional I/O groups.

A write operation from a host is shown (item *1*), that is targeted for virtual disk A. This write is targeted at the preferred node, Node 1 (item *2*). The write is cached and a copy of the data is made in the partner node, Node 2's cache (item *3*). The write is now complete, so far as the host is concerned. At some later time, the data is written, or destaged, to storage (item *4*). The figure also shows two uninterruptible power supplies (1 and 2) correctly configured so that each node is in a different power domain.

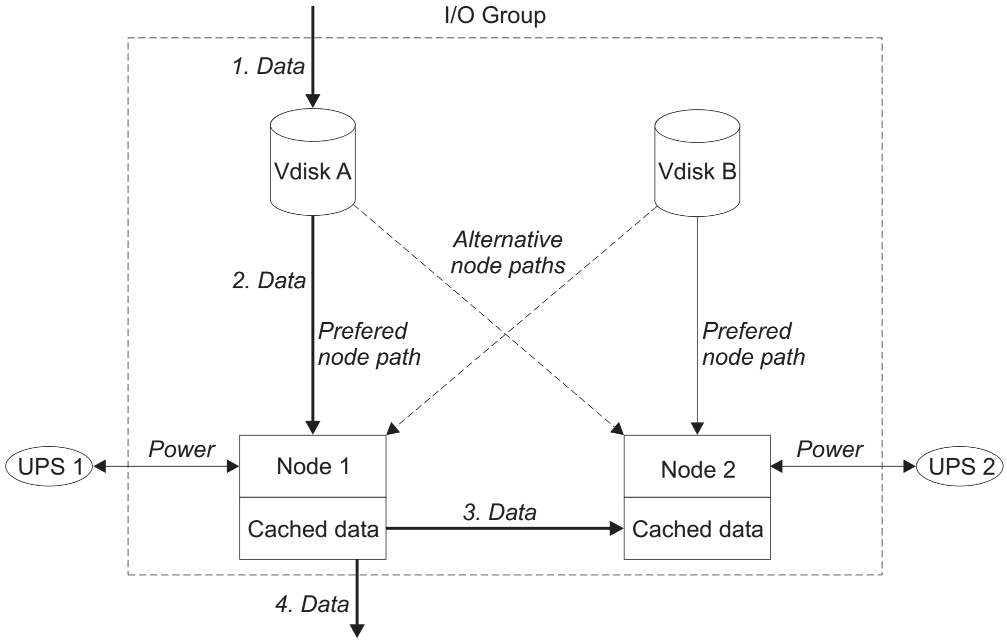

Figure 9. I/O group and uninterruptible power supply

When a node fails within an I/O group, the other node in the I/O group will take over the I/O responsibilities of the failed node. Data loss during a node failure is prevented by mirroring the I/O read/write data cache between the two nodes in an I/O group.

If only one node is assigned to an I/O group, or a node has failed in an I/O group, the cache is flushed to the disk and then goes into write-through mode. Therefore, any writes for the virtual disks that are assigned to this I/O group are not cached; it is sent directly to the storage device. If both nodes in an I/O group go offline, the virtual disks that are assigned to the I/O group cannot be accessed.

When a virtual disk is created, the I/O group that will provide access to the virtual disk must be specified. However, virtual disks can be created and added to I/O groups that contain offline nodes. I/O access will not be possible until at least one of the nodes in the I/O group is online.

The cluster also provides a **recovery I/O group**. This is used when both nodes in the I/O group have suffered multiple failures. This allows you to move the virtual disks to the recovery I/O group and then into a working I/O group. I/O access is not possible when virtual disks are assigned to the recovery I/O group.

Related concepts

"Input/Output (I/O) groups" on page 16
An I/O group is a group that is defined during the cluster configuration process.

"Uninterruptible power supply overview" on page 18
The uninterruptible power supply (UPS) provides the SAN Volume Controller with a secondary power source to be used if you lose power from your primary power source due to power failures, power sags, power surges, or line noise.
| Two types of UPS units can be used with the SAN Volume Controller: the UPS
| 5115 and the UPS 5125.

"I/O groups and uninterruptible power supply" on page 16
Nodes are deployed in pairs to make up a cluster. Each pair of nodes is known as an **I/O group**. Each node may be in *only* one I/O group.

Input/output governing

You can set the maximum amount of I/O activity that a host sends to a virtual disk (VDisk). This amount is known as the *I/O governing rate*. The governing rate can be expressed in I/Os per second or MB per second.

Read, write, and verify commands that access the physical medium are subject to I/O governing.

I/O governing does not effect FlashCopy and data migration I/O rates.

Governing is applied to Metro Mirror primary and secondary VDisks as follows:

- If an I/O governing rate is set on a secondary VDisk, then the same governing rate is applied to the primary.
- If you set an I/O governing on the primary and the secondary VDisk, then the I/O governing rate for the pair is the lowest rate that is set.

Uninterruptible power supply overview

The uninterruptible power supply (UPS) provides the SAN Volume Controller with a secondary power source to be used if you lose power from your primary power source due to power failures, power sags, power surges, or line noise. Two types of UPS units can be used with the SAN Volume Controller: the UPS 5115 and the UPS 5125.

If a power outage occurs, the UPS maintains power long enough to save any configuration and cache data contained in the dynamic random access memory (DRAM). The data is saved to the SAN Volume Controller internal disk. Figure 10 on page 19 and Figure 11 on page 19 provide illustrations of the two types of UPS units.

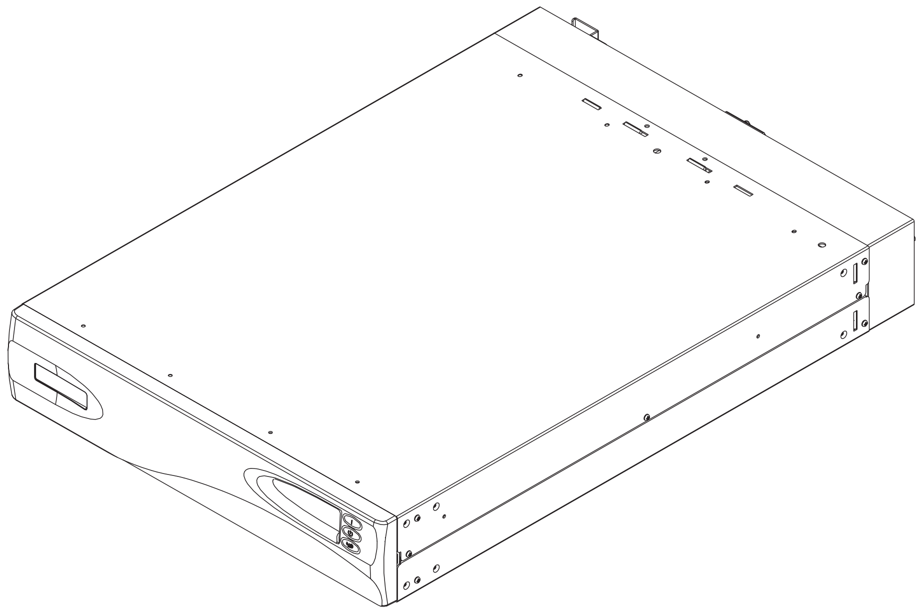

Figure 10. Uninterruptible power supply 5125

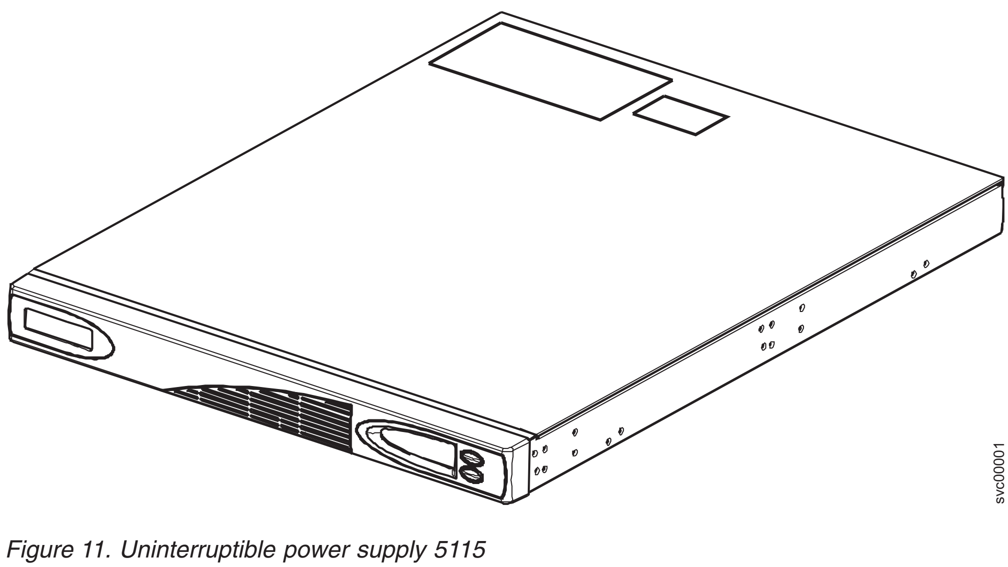

Figure 11. Uninterruptible power supply 5115

Note: The SAN Volume Controller UPS is an integral part of the SAN Volume Controller solution. It maintains continuous SAN Volume Controller specific communications with its attached SAN Volume Controller nodes. The SAN Volume Controller cannot operate without the UPS. The SAN Volume Controller UPS must be used in accordance with documented guidelines and procedures and must not power any equipment other than SAN Volume Controller nodes.

To provide full redundancy and concurrent maintenance, the SAN Volume Controller must be installed in pairs. Each SAN Volume Controller of a pair must be connected to a different UPS. Each UPS 5125 can support up to two SAN Volume Controller nodes. The UPS 5115 can support one SAN Volume Controller node. It is also recommended that you connect the two UPS units for the pair to different independent electrical power sources. This reduces the chance of an input power failure at both UPS units.

The UPS must be in the same rack as the nodes. Ensure that you are following the UPS guidelines that are provided in Table 2.

Table 2. Uninterruptible power supply support guidelines

Number of nodes	Number of uninterruptible power supply 5125 units	Number of uninterruptible power supply 5115 units
2	2	2
4	2	4
6	4	6
8	4	8

Attention:

1. Do not connect the uninterruptible power supplies to an input power source that does not conform to standards.
2. Each UPS pair must power only one SAN Volume Controller cluster.

Each UPS includes power (line) cords that connect the UPS to either a rack power distribution unit (PDU), if one exists, or to an external power source. Each UPS power input requires the protection of a UL-approved (or equivalent) 250 volt, 15 amp circuit breaker.

The UPS is connected to the SAN Volume Controllers with a power cable and a signal cable. To avoid the possibility of power and signal cables being connected to different UPS units, these cables are wrapped together and supplied as a single field replaceable unit. The signal cables enable the SAN Volume Controllers to read status and identification information from the UPS.

Each SAN Volume Controller monitors the operational state of the UPS to which it is attached. If the UPS reports a loss of input power, the SAN Volume Controller stops all I/O operations and dumps the contents of its DRAM to the internal disk drive. When input power to the UPS is restored, the SAN Volume Controllers restart and restore the original contents of the DRAM from the data saved on the disk drive.

A SAN Volume Controller is not fully operational until the UPS battery charge state indicates that it has sufficient capacity to power the SAN Volume Controller long enough to permit it to save all its memory to the disk drive in the event of a power loss. The UPS has sufficient capacity to save all the data on the SAN Volume Controller at least twice. For a fully-charged UPS, even after battery capacity has been used to power the SAN Volume Controllers while they save DRAM data, sufficient battery capacity remains to let the SAN Volume Controllers become fully operational as soon as input power is restored.

Note: Under normal circumstances, if input power is disconnected from the UPS, the SAN Volume Controller connected to that UPS performs a power down sequence. This operation, which saves the configuration and cache data to an internal disk in the SAN Volume Controller, typically takes about three minutes, at which time power is removed from the output of the UPS. In the event of a delay in the completion of the power down sequence, the UPS output power is removed five minutes after the power was disconnected to the UPS. Because this operation is controlled by the SAN Volume Controller, a UPS that is not connected to an active SAN Volume Controller will not shut off within the five-minute required period. In the case of an emergency, you must manually shut down the UPS by pushing the UPS 5125 power off button, or the UPS 5115 on/off button.

Attention: Data integrity could be compromised by pushing the UPS 5125 power off button or the UPS 5115 on/off button. Never shut down a UPS without first shutting down the SAN Volume Controller nodes that it supports.

It is very important that the two nodes in the I/O group are connected to different uninterruptible power supplies. This configuration ensures that cache and cluster state information is protected in the event of a failure of the UPS or mainline power source.

When nodes are added to the cluster, you must specify the I/O group that they will join. The configuration interfaces also checks the UPS units and ensures that the two nodes in the I/O group are not connected to the same UPS units.

Figure 12 shows a cluster of four nodes, with two I/O groups and two UPS 5125 units.

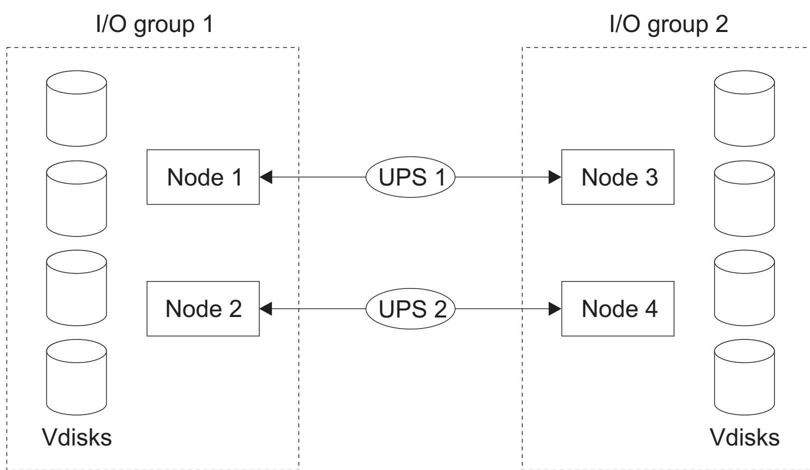

Figure 12. I/O groups and uninterruptible power supply 5125 relationship

Related concepts

“Input/Output (I/O) groups” on page 16
An I/O group is a group that is defined during the cluster configuration process.

“I/O groups and uninterruptible power supply” on page 16
Nodes are deployed in pairs to make up a cluster. Each pair of nodes is known as an **I/O group**. Each node may be in *only* one I/O group.

Storage subsystems and managed disks

The nodes in a cluster see the storage exported by SAN-attached storage subsystems as a number of disks, known as managed disks. The SAN Volume Controller does not attempt to provide recovery from physical disk failures within the storage subsystem. A managed disk is usually, but not necessarily, a RAID array.

Related concepts

“Storage subsystems”
A storage subsystem 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.

“Managed disks” on page 23
A managed disk (MDisk) is a logical disk (typically a RAID array or partition thereof) that a storage subsystem has exported to the SAN fabric to which the nodes in the cluster are attached.

Storage subsystems

A storage subsystem 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.

Storage subsystems that are attached to the SAN fabric provide the physical storage devices that the cluster detects as managed disks. These are usually RAID arrays, because the SAN Volume Controller does not attempt to provide recovery from physical disk failures within the storage subsystem. The nodes in the cluster are connected to one or more fibre-channel SAN fabrics.

The exported storage devices are detected by the cluster and reported by the user interfaces. The cluster can also determine which managed disks each storage subsystem is presenting, and can provide a view of managed disks filtered by the storage subsystem. This allows you to associate the managed disks with the RAID arrays that the subsystem exports.

The storage subsystem can 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, because the namespace is local to the storage subsystem. The storage subsystem will surface these storage devices with a unique ID, the logical unit number (LUN). This ID, along with the storage subsystem serial number or numbers (there may be more than one controller in a storage subsystem), can be used to associate the managed disks in the cluster with the RAID arrays exported by the subsystem.

Storage subsystems export storage to other devices on the SAN. The physical storage associated with a subsystem is normally configured into RAID arrays that provide recovery from physical disk failures. Some subsystems 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 disks.

Many storage subsystems allow the storage provided by a RAID array to be divided up into many SCSI logical units (LUs) that are presented on the SAN. With the SAN Volume Controller, it is recommended that storage subsystems are configured to present each RAID array as a single SCSI LU that will be recognized by the SAN Volume Controller as a single managed disk. The virtualization features of the SAN Volume Controller can then be used to divide up the storage into virtual disks.

Some storage subsystems allow the exported storage to be increased in size. The SAN Volume Controller does not use this extra capacity. Instead of increasing the size of an existing managed disk, add a new managed disk to the managed disk group and the extra capacity will be available for the SAN Volume Controller to use.

Attention: If you delete a RAID that is being used by the SAN Volume Controller, the MDisk group goes offline and the data in that group is lost.

When you are configuring your storage subsystems, ensure that you configure and manage your subsystems and its devices for optimal performance.

The cluster detects and provides a view of the storage subsystems that the SAN Volume Controller supports. The cluster can also determine which MDisks each subsystem has and can provide a view of MDisks filtered by device. This view enables you to associate the MDisks with the RAID arrays that the subsystem presents.

Note: The SAN Volume Controller Console supports storage that is internally configured as a RAID array. However, it is possible to configure a storage subsystem as a non-RAID device. RAID provides redundancy at the disk level. For RAID devices, 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.

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

Related concepts

"Managed disks"
A managed disk (MDisk) is a logical disk (typically a RAID array or partition thereof) that a storage subsystem has exported to the SAN fabric to which the nodes in the cluster are attached.

"Storage subsystems and managed disks" on page 21
The nodes in a cluster see the storage exported by SAN-attached storage subsystems as a number of disks, known as managed disks. The SAN Volume Controller does not attempt to provide recovery from physical disk failures within the storage subsystem. A managed disk is usually, but not necessarily, a RAID array.

Managed disks

A managed disk (MDisk) is a logical disk (typically a RAID array or partition thereof) that a storage subsystem has exported to the SAN fabric to which the nodes in the cluster are attached.

A managed disk might, 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 even if it does not have a one-to-one correspondence with a physical disk.

Each managed disk is divided 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 that the MDisk will be divided into depends on the attribute of the MDisk group to which it has been added.

Access modes

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

Unmanaged
: The MDisk is not used by the cluster.

Managed
The MDisk is assigned to an MDisk group and provides 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 lose the data that it contains. The *image mode* is the only mode that will preserve this data.

Figure 13 shows physical disks and managed disks.

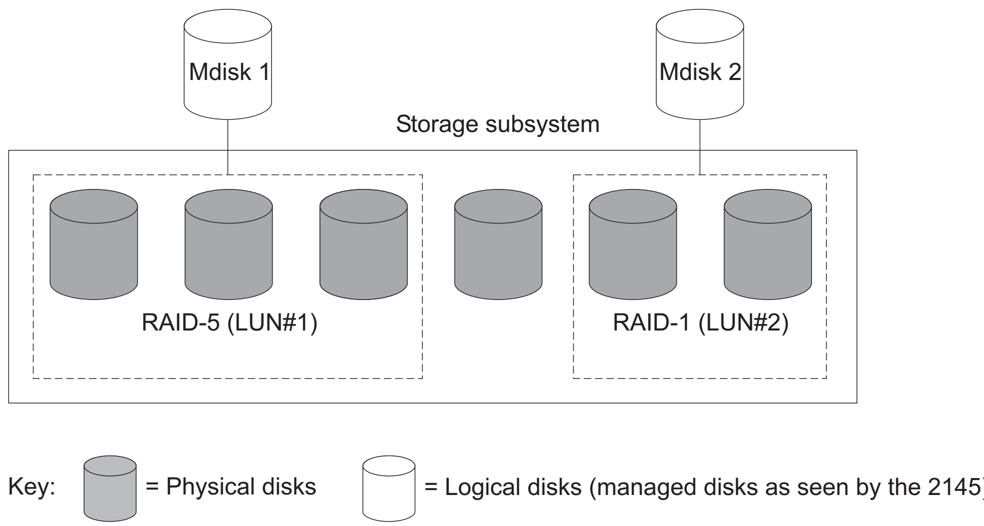

Figure 13. Controllers and MDisks

Table 3 describes the operational states of a managed disk:

Table 3. 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 is online when the following conditions are met: • All timeout error recovery procedures complete and report the disk as online. • LUN inventory of the target ports correctly reported the MDisk. • Discovery of this LUN completed successfully. • All of the managed disk target 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.

Table 3. Managed disk status (continued)

Excluded	The MDisk has been excluded from use by the cluster after repeated access errors. Run the Directed Maintenance Procedures to determine the problem. You can reset an MDisk and include it in the cluster again 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 storage subsystem, or one or more physical disks connected to the storage subsystem. 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 are a unit of mapping the data between MDisks and virtual disks (VDisks).

Attention: If you have observed intermittent breaks in links or if you have been replacing cables or connections in the storage area network fabric, you may have one or more MDisks in degraded status. If an I/O operation is attempted when a link is broken and the I/O operation fails several times, then the system partially excludes the MDisk and it changes the status of the MDisk to degraded. You must 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 SAN Volume Controller Console, or by issuing the following command:

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

Managed disk path Each managed disk has 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 storage 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 might be degraded. That is, one or more nodes might not see the managed disk on the fabric.

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

"Storage subsystems and managed disks" on page 21
The nodes in a cluster see the storage exported by SAN-attached storage subsystems as a number of disks, known as managed disks. The SAN Volume Controller does not attempt to provide recovery from physical disk failures within the storage subsystem. A managed disk is usually, but not necessarily, a RAID array.

Managed disk groups and virtual disks

Managed disks are collected into groups known as managed disk groups. Virtual disks are logical disks that are presented to the SAN by SAN Volume Controller nodes. The maximum number of supported VDisks per I/O group is 1024. The maximum number of supported VDisks per cluster is 4096. Virtual disks, like nodes, are associated with an I/O group.

Virtual disks are created from the extents of managed disks. Only managed disks that are in the same managed disk group can contribute extents to a virtual disk.

Related concepts

"Managed disk groups"
An *MDisk group* is a collection of MDisks that jointly contain all the data for a specified set of virtual disks (VDisks).

"Virtual disks" on page 29
A *virtual disk (VDisk)* is a logical disk that the cluster presents to the storage area network (SAN).

Managed disk groups

An *MDisk group* is a collection of MDisks that jointly contain all the data for a specified set of virtual disks (VDisks).

Figure 14 shows an MDisk group containing four MDisks.

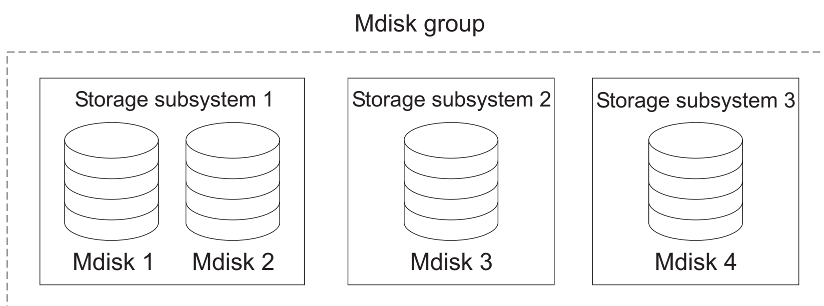

Figure 14. MDisk group

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.

Note: RAID array partitions on HP StorageWorks subsystems controllers are only supported in single-port attach mode. MDisk groups that consist of single-port attached subsystems and other storage subsystems are not supported.

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 the following 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 or the extents that the 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 deletion if VDisks are associated with it.

Table 4 describes the operational states of an MDisk group.

Table 4. Managed disk group status

Status	Description
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 and therefore cannot be seen by any of the online nodes in the cluster, then the MDisk group of which this MDisk is a member goes offline. This causes *all* the VDisks that are being presented by this MDisk group to go offline. Take care when you create MDisk groups to ensure an optimal configuration.

Consider the following guidelines when you create MDisk groups:

- If you are creating image-mode VDisks, do not put all of these VDisks into one MDisk group because a single MDisk failure results in all of these VDisks going offline. Allocate your image-mode VDisks between your MDisk groups.
- Ensure that all MDisks that are allocated to a single MDisk group are the same RAID type. This ensures that a single failure of a physical disk in the storage subsystem does not take the entire group offline. For example, if you have three RAID-5 arrays in one group and add a non-RAID disk to this group, you lose access to all the data striped across the group if the non-RAID disk fails. Similarly, for performance reasons, you must not mix RAID types. The performance of all VDisks will be reduced to the lowest performer in the group.
- If you intend to keep the virtual disk allocation within the storage exported by a storage subsystem, ensure that the MDisk group that corresponds with a single subsystem is presented by that subsystem. This also enables nondisruptive migration of data from one subsystem to another subsystem and simplifies the decommissioning process if you want to decommission a controller at a later time.
- Except when you migrate between groups, you must associate a VDisk with just one MDisk group.
- An MDisk can be associated with just one MDisk group.

| **Extents**

| To track the space that is available on an MDisk, the SAN Volume Controller divides
| each MDisk 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 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 this is not recommended. Different extent sizes place restrictions on the use of data migration. You cannot use the SAN Volume Controller data migration function to move a VDisk between MDisk groups that have different extent sizes.

You can use Copy Services to copy a VDisk between MDisk groups that have different extent sizes. You have the following options:

- Use FlashCopy to copy a VDisk between a source and a destination MDisk group that have different extent sizes.
- Use intracluster Metro Mirror to copy a VDisk between a source and a destination MDisk group that have different extent sizes.

The choice of extent size affects the total amount of storage that will be managed by the cluster. Table 5 shows the maximum amount of storage that can be managed by a cluster for each extent size.

Table 5. 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

A cluster can manage 4 million extents (4 x 1024 x 1024). The choice of extent size determines the amount of storage that will be managed by a cluster. For example, with a 16MB extent size, the cluster can manage up to 16MB x 4M = 64TB of storage.

When choosing an extent size, consider your future needs. For example, if you currently have 40TB of storage, and you specify an extent size of 16MB, the capacity of the MDisk group is limited to 64TB of storage in the future. If you select an extent size of 64MB, the capacity of the MDisk group is 256TB.

Using a larger extent size can waste storage. When a VDisk is created, the storage capacity for the VDisk is rounded to a whole number of extents. If you configure the system to have a large number of small VDisks and you use a large extent size, then this can cause storage to be wasted at the end of each VDisk.

Related concepts

"Virtual disks" on page 29
A *virtual disk (VDisk)* is a logical disk that the cluster presents to the storage area network (SAN).

"Managed disk groups and virtual disks" on page 26
Managed disks are collected into groups known as managed disk groups. Virtual disks are logical disks that are presented to the SAN by SAN Volume Controller nodes. The maximum number of supported VDisks per I/O group is 1024. The

maximum number of supported VDisks per cluster is 4096. Virtual disks, like nodes, are associated with an I/O group.

Virtual disks

A *virtual disk (VDisk)* is a logical disk that the cluster presents to the storage area network (SAN).

Application servers on the SAN access VDisks, not managed disks (MDisks). VDisks are created from a set of extents in an MDisk group. There are three types of VDisks: striped, sequential, and image.

Types

You can create the following types of VDisks:

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.

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.

Attention: Take care when you specify a stripe set if your MDisk group contains MDisks of unequal size. By default, striped VDisks 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 might result in the VDisk not being created.

If you are unsure there is sufficient free space to create a striped VDisk select one of the following options:

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

Figure 15 shows an example of a managed disk group containing three MDisks. This figure also shows a striped virtual disk created from the extents available in the group.

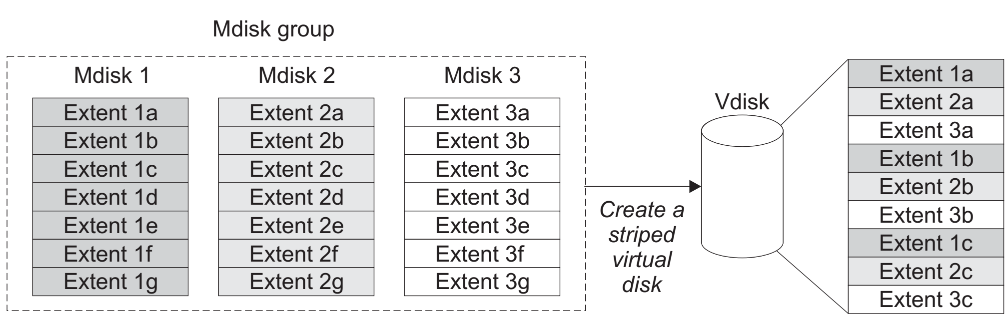

Figure 15. Managed disk groups and VDisks

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 VDisks are special VDisks that have a direct relationship with one managed disk. If you have a managed disk 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. In other words, the logical block address (LBA) x on the managed disk is the same as LBA x on the virtual disk.

When you create an image-mode VDisk, 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 VDisks. 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. After 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 VDisks 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 it contains partitions or data.

A virtual disk can have one of three states. Table 6 describes the different states of a virtual disk:

Table 6. 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 that are 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 VDisks. 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 in which 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 deletion might fail if the virtual disk is still mapped to hosts. The deletion might also fail if the virtual disk is still part of a FlashCopy or a Metro Mirror mapping. If the deletion 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 concepts

"Managed disk groups" on page 26
An *MDisk group* is a collection of MDisks that jointly contain all the data for a specified set of virtual disks (VDisks).

"Managed disk groups and virtual disks" on page 26
Managed disks are collected into groups known as managed disk groups. Virtual disks are logical disks that are presented to the SAN by SAN Volume Controller nodes. The maximum number of supported VDisks per I/O group is 1024. The maximum number of supported VDisks per cluster is 4096. Virtual disks, like nodes, are associated with an I/O group.

Hosts and virtual (VDisk) mappings

Application servers can only access VDisks that have been made accessible to them.

The SAN Volume Controller 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 enables you to create logical hosts that group together WWPNs belonging to a single application server or multiple application servers. 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.

Related concepts

"Host objects"
A host system is an open-systems computer that is connected to the switch through a fibre-channel interface.

"Virtual disk-to-host mapping" on page 32
Virtual disk-to-host mapping is the process of controlling which hosts have access to specific virtual disks (VDisks) within the SAN Volume Controller.

Host objects

A host system is an open-systems computer that is connected to the switch 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) that are 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. 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, 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.

Attention: You must not include a node port in a host object.

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

The number of nodes that can see each port is reported on a per node basis and is known as the node login count. 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.

Related concepts

"Hosts and virtual (VDisk) mappings" on page 31
Application servers can only access VDisks that have been made accessible to them.

"Virtual disk-to-host mapping"
Virtual disk-to-host mapping is the process of controlling which hosts have access to specific virtual disks (VDisks) within the SAN Volume Controller.

Virtual disk-to-host mapping

Virtual disk-to-host mapping is the process of controlling which hosts have access to specific virtual disks (VDisks) within the SAN Volume Controller.

Virtual disk-to-host mapping is similar in concept to logical unit number (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 done at the SAN Volume Controller level.

Application servers can only access VDisks that have been made accessible to them. The SAN Volume Controller 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 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 therefore 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 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 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 Controllers 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, 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 disk 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 16 and Figure 17 on page 34 show two VDisks, and the mappings that exist between the host objects and these VDisks.

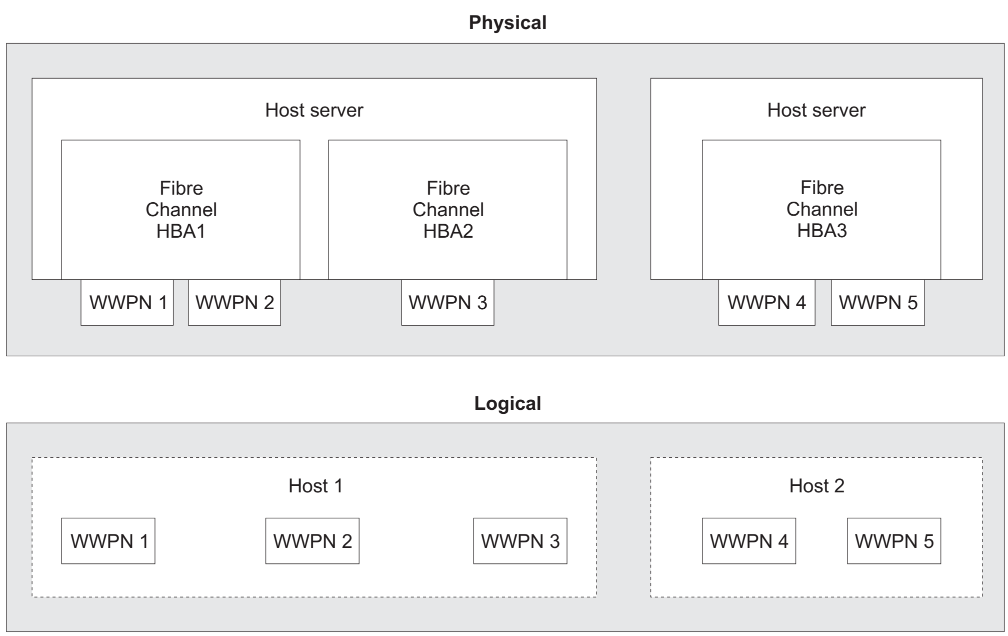

Figure 16. Hosts, WWPNs, and VDisks

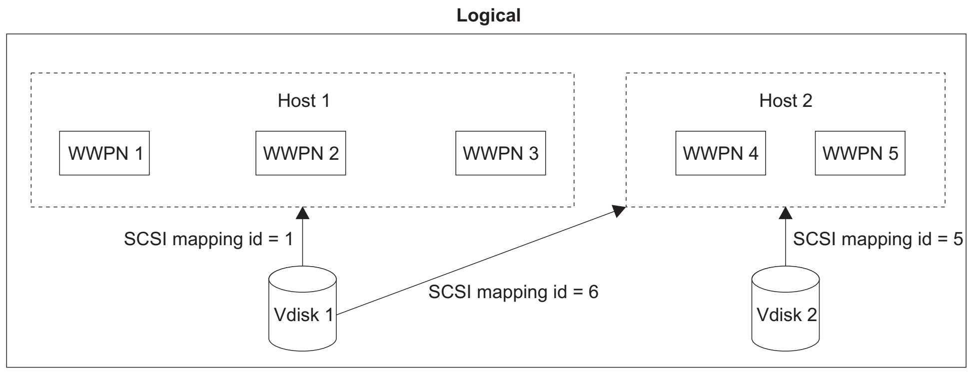

Figure 17. Hosts, WWPNs, VDisks and SCSI mappings

Related concepts

"Host objects" on page 31
A host system is an open-systems computer that is connected to the switch through a fibre-channel interface.

"Hosts and virtual (VDisk) mappings" on page 31
Application servers can only access VDisks that have been made accessible to them.

Copy Services

Two Copy Services are supported by the SAN Volume Controller.

- FlashCopy.
- IBM TotalStorage Metro Mirror for SAN Volume Controller. This service was previously known as synchronous remote copy or peer-to-peer remote copy.

Related concepts

"FlashCopy"
FlashCopy is a copy service that is available with the SAN Volume Controller.

"Metro Mirror" on page 46
Metro Mirror (formerly known as 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 on the other virtual disk.

FlashCopy

FlashCopy is a copy service that is available with the SAN Volume Controller.

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 unless target writes have been performed. 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. Although the FlashCopy operation takes some time, this time is several orders of magnitude less than the time that would be required to copy the data using conventional techniques.

It is difficult to make a consistent copy of a data set that is being constantly updated. Point-in-time copy techniques are used to help solve this problem. If a copy of a data set is created 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, the copy will contain the referenced object at its new location but the reference will point to the old location.

Source VDisks and target VDisks must meet the following requirements:

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

Related concepts

"FlashCopy mappings" on page 37
A FlashCopy mapping defines the relationship between a source virtual disk (VDisk) and a target virtual disk (VDisk).

"FlashCopy consistency groups" on page 41
A consistency group is a container for mappings. You can add many mappings to a consistency group.

Related tasks

"Host considerations for FlashCopy integrity" on page 36
The SAN Volume Controller FlashCopy functionality transfers a point-in-time copy of one virtual disk onto a designated target virtual disk of the same size. Both virtual disks must already be created.

Related reference

"FlashCopy applications"
Uses of FlashCopy in the business environment include taking consistent backups of changing data, application testing, and creating copies for auditing purposes and for data mining.

"FlashCopy indirection layer" on page 43
FlashCopy provides the semantics of a point in time copy by using an indirection layer which intercepts I/Os targeted at both the source and target virtual disks.

"Background copy" on page 45
A FlashCopy mapping has a property background copy rate.

FlashCopy applications

Uses of FlashCopy in the business environment include taking consistent backups of changing data, application testing, and creating copies for auditing purposes and for data mining.

An important use of FlashCopy is to take consistent backups of changing data. In this application, a FlashCopy is created to capture the data at a particular time. The resulting image of the data can be backed up, for example, to a tape device. When the copied data is on tape, the data on the FlashCopy target disks becomes redundant and can now be discarded. Usually in this backup condition, the target data can be handled as read only.

Another use of FlashCopy data is in the application testing. It is often very important to test a new version of an application with real business data before the existing production version of the application is updated or replaced. This testing reduces the risk that the updated application fails because it is not compatible with the actual business data that is in use at the time of the update. Such an application test might require write access to the target data.

Other uses of FlashCopy in the business environment include creating copies for auditing purposes and for data mining.

In the scientific and technical arena one way in which FlashCopy is employed is to create restart points for long running batch jobs. This means that if a batch job fails many days into its run it may be possible to restart the job from a saved copy of its data rather than re-running the entire multi-day job.

Related concepts

"FlashCopy" on page 34
FlashCopy is a copy service that is available with the SAN Volume Controller.

Host considerations for FlashCopy integrity

The SAN Volume Controller FlashCopy functionality transfers a point-in-time copy of one virtual disk onto a designated target virtual disk of the same size. Both virtual disks must already be created.

All the data in the source virtual disk is copied to the destination virtual disk. This includes operating system control information as well as application data and meta-data. Because all the data is copied, some operating systems will not allow a source disk and a target disk to reside on the same host. In order to ensure the integrity of the copy made, it is necessary to completely flush the host cache of any outstanding reads or writes before proceeding with the FlashCopy. Host cache flushing is ensured by unmounting the source virtual disks from the source host before starting the FlashCopy.

Because the target disks will be overwritten with a complete image of the source disks, it is important that any data held in the host operating system (or application) caches for the target disks is discarded before the FlashCopy mappings are started. The easiest way to ensure that no data is held in these caches is to unmount the target disks prior to starting FlashCopy.

Some operating systems and applications provide facilities to stop I/O operations and to ensure that all data is flushed from caches on the host. If these facilities are available then they can be used to prepare and start a FlashCopy in a less disruptive manner. Refer to your host and application documentation for details.

Some operating systems are unable to use a copy of a virtual disk without an additional step, which is called synthesis. Synthesis performs some transformation on the operating system meta-data on the target virtual disk to allow the operating system to use the disk. Refer to your host documentation on how to detect and mount the copied virtual disks.

Perform the following steps to flush data from your host volumes and perform the FlashCopy:

1. If you are using UNIX or Linux operating systems, perform the following steps:
   a. Quiesce all applications to the source volumes you wish to FlashCopy.
   b. Use the **umount** command to unmount the designated drives.
   c. Prepare and start the FlashCopy for those unmounted drives.
   d. Mount back your volumes with the mount command and resume your applications.
2. If you are using Windows operating system using drive letter changes, perform the following steps:
   a. Quiesce all applications to the source volumes you wish to FlashCopy.
   b. Go into your disk management window and remove the drive letter on each drive to be copied (this unmounts the drive).
   c. Prepare and start the FlashCopy for those unmounted drives.

d. Mount back your volumes by restoring the drive letters, and resume your applications.

If you are using the **chkdsk** command, perform the following steps:

a. Quiesce all applications to the source volumes you wish to FlashCopy.
b. Issue the **chkdsk /x** command on each drive to be copied (the /x option will unmount, scan, and remount the volume).
c. Ensure that all applications to the source volumes are still quiesced.
d. Prepare and start the FlashCopy for those unmounted drives.

Note: If you can ensure that no reads and writes will be issued to the source volumes after unmounting, you can immediately remount and then perform the FlashCopy.

Related concepts

"FlashCopy" on page 34
FlashCopy is a copy service that is available with the SAN Volume Controller.

FlashCopy mappings

A FlashCopy mapping defines the relationship between a source virtual disk (VDisk) and a target virtual disk (VDisk).

Because FlashCopy copies one VDisk to another VDisk, the SAN Volume Controller Console must be aware of that relationship. A particular VDisk can take part in only one mapping; that is, a VDisk 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 VDisk at the time that it is started. To create a FlashCopy of a virtual disk, you must first create a mapping between the source VDisk (the disk that is copied) and the target VDisk (the disk that receives the copy). The source and target must be of equal size.

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

A FlashCopy mapping can be created between any two VDisks in a cluster. It is not necessary for the VDisks 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 VDisk. 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 VDisk has been copied. Each bit in the bitmap represents one region of the source VDisk. Such a region is called a *grain*.

After a FlashCopy operation starts, read operations to the source VDisk continue to occur. If new data is written to the source (or target)VDisk, then the existing data on the source is copied to the target VDisk before the new data is written to the source (or target) VDisk. The bitmap is updated to mark that the grain of the source VDisk has been copied so that later write operations to the same grain do not recopy the data.

Similarly, during a read operation to the target VDisk the bitmap is used to determine the grain has been copied. If the grain has been copied, the data is read from the target VDisk. If the grain has not been copied, the data is read from the source VDisk.

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 all the data that is on the source VDisk to the target VDisk.

When a mapping is started and the background copy rate is greater than zero (or a value other than `NOCOPY` is selected in the SAN Volume Controller Console's Creating FlashCopy Mappings panel), 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 VDisk, the whole VDisk 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 (or `NOCOPY`), 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 VDisk 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. Data that is written to the source is not cached in the SAN Volume Controllers; 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 VDisk. You must coordinate the prepare operation with the operating system. Depending on the type of data that is on the source VDisk, 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.

| **Note:** The **svctask startfcmap** command can take some time to process.

For customers who do not need the complexity of consistency groups, the SAN Volume Controller 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 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 reference

"FlashCopy mapping events"
FlashCopy mapping events details the events that modify the state of a FlashCopy mapping.

FlashCopy mapping events:

FlashCopy mapping events details the events that modify the state of a FlashCopy mapping.

Table 7. FlashCopy mapping events

Create	A new FlashCopy mapping is created between the specified source virtual disk and the specified target virtual disk. The various supported parameters are also described there. The operation fails if either the source or target virtual disks are already a member of a FlashCopy mapping. The operation fails if the SAN Volume Controller has insufficient bitmap memory. The operation also fails if the source and target virtual disks are different sizes.
Prepare	The prepare command is directed to either a consistency group for FlashCopy mappings which are members of a normal consistency group or to the mapping name for FlashCopy mappings which are members of the special consistency group 0. The prepare command places the FlashCopy mapping in the preparing state. It is important to note that the act of preparing for start may corrupt any data which previously resided on the target virtual disk since cached writes are discarded. Even if the FlashCopy mapping is never started, the data from the target may have been logically changed by the act of preparing for start.
Flush done	The FlashCopy relationship moves from the preparing state to the prepared state automatically once all cached data for the source has been flushed and all cached data for the target has been invalidated.

Table 7. FlashCopy mapping events (continued)

Start	When all the FlashCopy mappings in a consistency group are in the prepared state, the FlashCopy relationships can be started. Some other FlashCopy products refer to this event as starting the FlashCopy. The start of all of the FlashCopy mappings in the consistency group must be synchronized correctly with respect to I/Os directed at the virtual disks to preserve the cross volume consistency group. This is achieved as follows: During the **start** command: • New reads and writes to all source virtual disks in the consistency group are paused in the cache layer until all ongoing reads and writes below the cache layer have been completed. • Once all FlashCopy mappings in the consistency group have been paused, internal cluster state is set to allow FlashCopy operations. • Once all FlashCopy mappings in the consistency group have had their cluster state set, read and write operations are unpaused on the source virtual disks. • The target virtual disks are brought online. As part of the **start** command, read and write caching is enabled for both the source and target virtual disks.
Modify	A FlashCopy mapping has two properties which can be modified. These are the background copy rate and the consistency group. The background copy rate can be modified in any state but attempting to modify the consistency group in any state other than idling, copied, or stopped will fail.
Stop	There are two mechanisms by which a FlashCopy mapping can be stopped: 1. Either you have issued a command; or 2. An input/output (I/O) error has occurred.
Delete	This command requests that the specified FlashCopy mapping be deleted. If the FlashCopy mapping is in the stopped state, the force flag must be used. Deleting a FlashCopy mapping in the stopped state may allow unflushed write data from the cache to be destaged to what was the target virtual disk. This does not affect the data integrity of the system because following a forced delete, nothing can be certain about the contents of the target virtual disk. The data contained in the target virtual disk could be anything. The destaging of old data to what was the target virtual disk does not affect the future use of the virtual disk because any new data will be written over this old data, in the cache or on disk.
Flush failed	If the flush of data from the cache cannot be completed then FlashCopy mapping will enter the stopped state.
Copy complete	Once every grain of the source and target has been copied, the source and target are independent and the state machine enters the copied state. The FlashCopy mapping is not automatically deleted at this time and can be reactivated by preparing and starting again.
Bitmap Online/Offline	The node has failed.

Related concepts

"FlashCopy mappings" on page 37
A FlashCopy mapping defines the relationship between a source virtual disk (VDisk) and a target virtual disk (VDisk).

FlashCopy consistency groups

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 virtual disks (VDisks). Do not place mappings that you want to control at an individual level instead of at a consistency group level into a consistency group. These individual mappings are known as stand-alone mappings.

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

When you copy data from one 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.

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 reference

"Dependent writes"
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.

"Operations on consistency groups" on page 43
You can create, change, and delete consistency groups.

"FlashCopy limits" on page 43
The SAN Volume Controller supports up to 2048 FlashCopy mappings. The SAN Volume Controller supports up to 512 FlashCopy mappings per consistency group.

Dependent writes:

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.

Think about the following typical sequence of write operations for a data base update transaction.

1. Run a write operation to update the data base log so that it indicates that a data base update is about to take place.
2. Run a second write operation to update the data base.
3. Run a third write operation to update the data base log so that it indicates that the data base update has completed successfully.

The database ensures correct ordering of these writes by waiting for each step to complete before starting the next. If, however, the database log (updates 1 and 3) and the database itself (update 2) are on different virtual disks and a FlashCopy mapping is started during this update, then the possibility that the database itself is copied slightly before the database log resulting in the target virtual disks seeing writes (1) and (3) but not (2) must be excluded. In this case, if the database were restarted from a backup made from the FlashCopy target disks, the data base log would indicate that the transaction had completed successfully when, in fact, that is not the case. The transaction would be lost and the integrity of the data base would be in question.

It may thus be the case that in order to create a consistent image of user data it is necessary to perform a FlashCopy operation on multiple virtual disks as an atomic operation. In order to meet this need, the SAN Volume Controller supports the concept of a consistency group. A consistency group contains a number of FlashCopy mappings. A consistency group can contain an arbitrary number of FlashCopy mappings up to the maximum number of FlashCopy mappings supported by a SAN Volume Controller cluster. The SAN Volume Controller allows the **start** command which causes the point in time copy to occur, to be directed at a consistency group. In this case all of the FlashCopy mappings in the consistency group are started at the same time, resulting in a point in time copy which is consistent across all of the FlashCopy mappings which are contained in the consistency group. The SAN Volume Controller supports 128 consistency groups per cluster.

Related concepts

"FlashCopy consistency groups" on page 41
A consistency group is a container for mappings. You can add many mappings to a consistency group.

Operations on consistency groups:

You can create, change, and delete consistency groups.

You can create, change, and delete consistency groups by using the command line tool that is described in *IBM TotalStorage SAN Volume Controller: Command-Line Interface User's Guide* or you can use the SAN Volume Controller Console.

Related concepts

"FlashCopy consistency groups" on page 41
A consistency group is a container for mappings. You can add many mappings to a consistency group.

FlashCopy limits:

The SAN Volume Controller supports up to 2048 FlashCopy mappings. The SAN Volume Controller supports up to 512 FlashCopy mappings per consistency group.

Related concepts

"FlashCopy consistency groups" on page 41
A consistency group is a container for mappings. You can add many mappings to a consistency group.

FlashCopy indirection layer

FlashCopy provides the semantics of a point in time copy by using an indirection layer which intercepts I/Os targeted at both the source and target virtual disks.

The act of starting a FlashCopy mapping causes this indirection layer to become active in the I/O path. This occurs as an atomic command across all FlashCopy mappings in the consistency group.

The indirection layer makes a decision about each I/O. This decision is based upon the following:

- the virtual disk and LBA to which the I/O is addressed,
- its direction (read or write)
- the state of an internal data structure, the flash copy bitmap.

The indirection layer either allows the I/O through to the underlying storage, redirects the I/O from the target virtual disk to the source virtual disk, or stalls the I/O while it arranges for data to be copied from the source virtual disk to the target virtual disk.

Related reference

"Grains and the FlashCopy bitmap" on page 44
When data is copied from the source virtual disk to the target virtual disk, it is copied in units of address space known as grains.

"Source and target reads" on page 44
The source and target must be of equal size. 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.

"Writes to the source or target" on page 44
Where writes occur to source or target to an area (or grain) which has not yet been copied, these will usually be delayed while a copy operation is performed to copy data from the source to the target, to maintain the illusion that the target contains its own copy.

“FlashCopy limits” on page 45
There are FlashCopy limits as to how many mappings are supported in a single cluster.

Grains and the FlashCopy bitmap:

When data is copied from the source virtual disk to the target virtual disk, it is copied in units of address space known as grains.

In the SAN Volume Controller, the grain size is 256 KB. The FlashCopy bitmap contains one bit for each grain. The bit records whether the associated grain has yet been split by copying the grain from the source to the target.

Related reference

“FlashCopy indirection layer” on page 43
FlashCopy provides the semantics of a point in time copy by using an indirection layer which intercepts I/Os targeted at both the source and target virtual disks.

Source and target reads:

The source and target must be of equal size. 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.

Source reads

Reads of the source are always passed through to the underlying source disk.

Target reads

In order for FlashCopy to process a read from the target disk it must consult its bitmap. If the data being read has already been copied to the target then the read is sent to the target disk. If it has not, then the read is sent to the source disk. Clearly this algorithm requires that, while this read is outstanding, no writes are allowed to execute which would change the data being read from the source. The SAN Volume Controller satisfies this requirement by using a cluster wide locking scheme.

FlashCopy limits the number of concurrent reads to an unsplit target grain to one. If more than one concurrent read to an unsplit target grain is received by the FlashCopy mapping layer, they will be serialized.

Related reference

“FlashCopy indirection layer” on page 43
FlashCopy provides the semantics of a point in time copy by using an indirection layer which intercepts I/Os targeted at both the source and target virtual disks.

Writes to the source or target:

Where writes occur to source or target to an area (or grain) which has not yet been copied, these will usually be delayed while a copy operation is performed to copy data from the source to the target, to maintain the illusion that the target contains its own copy.

A specific optimization is performed where an entire grain is written to the target virtual disk. In this case the new grain contents are written to the target virtual disk and if this succeeds the grain is marked as split in the flash copy bitmap without a copy from the source to the target having been performed. If the write fails, the grain is not marked as split.

Related reference

"FlashCopy indirection layer" on page 43
FlashCopy provides the semantics of a point in time copy by using an indirection layer which intercepts I/Os targeted at both the source and target virtual disks.

FlashCopy limits:

There are FlashCopy limits as to how many mappings are supported in a single cluster.

Up to 512 FlashCopy mappings are supported in a single cluster. A maximum of 16 TB of VDisk space (both source and target) may be participating in FlashCopy mappings in any one I/O group on a single cluster.

Related reference

"FlashCopy indirection layer" on page 43
FlashCopy provides the semantics of a point in time copy by using an indirection layer which intercepts I/Os targeted at both the source and target virtual disks.

Background copy

A FlashCopy mapping has a property background copy rate.

This is a value between 1 and 100. The background copy rate can be changed when the FlashCopy mapping is in any state.

If "NOCOPY" is specified, then background copy is disabled. One use of this is for short-lived FlashCopy mappings which are to be used for backup purposes only. Since the source data set is not expected to change much during the lifetime of the FlashCopy mapping, it is more efficient in terms of managed disk I/Os not to perform a background copy.

Table 8 provides the relationship of the background copy rate value to the attempted number of grains to be split per second. (A grain is the unit of data represented by a single bit, which is 256K.)

Table 8. Background copy

User-specified value	KB/sec	Grains/sec
1 - 10	128	0.5
11 - 20	256	1
21 - 30	512	2
41 - 50	2048	8
91 - 100	64 MB	256

The grains/sec numbers represent goals that the code tries to achieve. The SAN Volume Controller will be unable to achieve these goals if insufficient bandwidth is available from the SAN Volume Controller nodes to the physical disks making up the managed disks after taking into account the requirements of foreground I/O. If this situation arises, then background copy I/O will contend for resources on an equal basis with I/O arriving from hosts. Both will tend to see an increase in latency and consequential reduction in throughput with respect to the situation had the bandwidth not been limited.

Degradation will be graceful. Both background copy and foreground I/O will continue to make forward progress, and will not stop, hang or cause the node to fail.

The background copy is performed by one of the nodes belonging to the I/O group in which the source virtual disk resides. This responsibility is failed over to the other node in the I/O group in the event of the failure of the node performing the background copy.

The background copy is performed backward; that is, it starts with the grain containing the highest logical block numbers (LBAs) and works backwards towards the grain containing LBA 0. This is done to avoid any unwanted interactions with sequential write streams from the using application.

Related concepts

"FlashCopy" on page 34
FlashCopy is a copy service that is available with the SAN Volume Controller.

Metro Mirror

Metro Mirror (formerly known as 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 on the other virtual disk.

Although the application only writes to a single virtual disk (VDisk), the SAN Volume Controller maintains two copies of the data. If the copies are separated by a significant distance, the Metro Mirror copy can be used as a backup for disaster recovery. A prerequisite for the SAN Volume Controller Metro Mirror 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 secondary VDisk. Normally, host applications do not perform input or output operations to the secondary VDisk. When a host writes to the primary VDisk, it will not receive confirmation of I/O completion until the write operation has completed for the copy on the secondary disk, as well as on the primary.

Metro Mirror 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 Metro Mirror relationships with itself and one other cluster.
- Intercluster and intracluster Metro Mirror can be used concurrently within a cluster.
- The intercluster link is bidirectional. This means that it can copy data from clusterA to clusterB for one pair of VDisks while copying data from clusterB to clusterA for a different pair of VDisks.
- The copy direction can be reversed for a consistent relationship by issuing a simple **switch** command. See *IBM TotalStorage SAN Volume Controller: Command-Line Interface User's Guide*.

- Metro Mirror consistency groups are supported for ease of managing a group of relationships that need to be kept in sync for the same application. This also simplifies administration, as a single command issued to the consistency group will be applied to all the relationships in that group.

Related concepts

"IBM TotalStorage Metro Mirror for SAN Volume Controller"
IBM TotalStorage Metro Mirror for SAN Volume Controller (formerly known as synchronous Remote Copy) provides a *consistent* copy, which means that the primary virtual disk (VDisk) is always the exact match of the secondary VDisk.

"Metro Mirror partnerships"
With Metro Mirror, you can copy a VDisk in one cluster to a VDisk in another cluster.

"Metro Mirror relationships" on page 48
A Metro Mirror relationship defines the relationship between two virtual disks: a master VDisk and an auxiliary VDisk.

"Metro Mirror consistency groups" on page 50
Metro Mirror provides the facility to group a number of relationships into a Metro Mirror consistency group so that they can be manipulated in unison.

IBM TotalStorage Metro Mirror for SAN Volume Controller

IBM TotalStorage Metro Mirror for SAN Volume Controller (formerly known as synchronous Remote Copy) provides a *consistent* copy, which means that the primary virtual disk (VDisk) is always the exact match of the secondary VDisk.

The host application writes data to the primary VDisk but does not receive the final status on the write operation until the data is written to the secondary VDisk. For disaster recovery, this mode is the only practical mode of operation because a consistent copy of the data is maintained. However, Metro Mirror is slower than Global Mirror because of the latency time and bandwidth limitations that are imposed by the communication link to the secondary site.

Related concepts

"Metro Mirror" on page 46
Metro Mirror (formerly known as 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 on the other virtual disk.

Metro Mirror partnerships

With Metro Mirror, you can copy a VDisk in one cluster to a VDisk in another cluster.

The SAN Volume Controller needs to know not only about the relationship between the two VDisks but also about the relationship between the two clusters. A Metro Mirror partnership defines the relationship between the two clusters.

To establish a cluster partnership between two clusters it is necessary to issue the **svctask mkpartnership** command from both clusters. For example, to establish a partnership between clusterA and clusterB, you would first issue the **svctask mkpartnership** command from clusterA, specifying clusterB as the remote cluster. At this point the partnership is partially configured, and sometimes described as one-way. Next, you would issue the **svctask mkpartnership** command from clusterB, specifying clusterA as the remote cluster. When this completes, the partnership is fully configured for two-way communication between the clusters. See *IBM TotalStorage SAN Volume Controller: Command-Line Interface User's Guide*.

Background copy management

You can specify the rate at which the initial background copy from the local cluster to the remote cluster is performed. The bandwidth parameter controls this rate.

Related concepts

"Metro Mirror" on page 46
Metro Mirror (formerly known as 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 on the other virtual disk.

Background copy bandwidth impact on foreground I/O latency

The background copy bandwidth determines the rate at which the background copy for the IBM TotalStorage Metro Mirror for SAN Volume Controller will be attempted.

The background copy bandwidth can affect foreground I/O latency in one of three ways:

- The following result can occur if the background copy bandwidth is set too high for the Metro Mirror intercluster link capacity:
  - The background copy I/Os can back up on the Metro Mirror intercluster link.
  - There is a delay in the synchronous secondary writes of foreground I/Os.
  - The foreground I/O latency will increase as perceived by applications.
- If the background copy bandwidth is set too high for the storage at the *primary* site, background copy read I/Os overload the primary storage and delay foreground I/Os.
- If the background copy bandwidth is set too high for the storage at the *secondary* site, background copy writes at the secondary overload the secondary storage and again delay the synchronous secondary writes of foreground I/Os.

In order to set the background copy bandwidth optimally, make sure that you consider all three resources (the primary storage, the intercluster link bandwidth and the secondary storage). Provision the most restrictive of these three resources between the background copy bandwidth and the peak foreground I/O workload.

This provisioning may be done by calculation as above or alternatively by determining experimentally how much background copy can be allowed before the foreground I/O latency becomes unacceptable and then backing off to allow for peaks in workload and some safety margin.

Related concepts

"FlashCopy" on page 34
FlashCopy is a copy service that is available with the SAN Volume Controller.

Metro Mirror relationships

A Metro Mirror relationship defines the relationship between two virtual disks: a master VDisk and an auxiliary VDisk.

In most cases, the master VDisk contains the production copy of the data and is the VDisk that the application normally accesses. The auxiliary VDisk typically contains a backup copy of the data and is used in disaster recovery scenarios.

The master and auxiliary VDisks are defined when the relationship is created, and these attributes never change. However, either VDisk may operate in the primary or secondary role according to the circumstances. The primary VDisk is the one currently receiving updates from the application, analogous to source VDisk . The secondary VDisk receives a copy of any updates to the primary VDisk, because these updates are all transmitted across the Metro Mirror link. Therefore, the secondary VDisk is analogous to a continuously updated target VDisk.

Primary
: Contains a valid copy of the application data, and it is accessible for application write operations.

Secondary
: Might contain a valid copy of the application data, but it is not available for application write operations.

When a relationship is created, the master VDisk is assigned the role of primary VDisk and the auxiliary VDisk is assigned the role of secondary VDisk. Therefore, the initial copying direction is from master to auxiliary. When the relationship is in a consistent state, the copy direction can be reversed by issuing the **svctask switchrcrelationship** command, and specifying the auxiliary disk as the primary.

The two VDisks in a relationship must be the same size. When the two VDisks are in the same cluster, they must be in the same input/output (I/O) group.

A relationship can be added to a Metro Mirror consistency group, for ease of application management (see consistency groups below).

Note: Membership of a consistency group is an attribute of the relationship, not the consistency group. Therefore, the **svctask chrcrelationship** command is used to add or remove a relationship to or from a consistency group. See *IBM TotalStorage SAN Volume Controller: Command-Line Interface User's Guide*.

Metro Mirror states

When a Metro Mirror relationship is created with two virtual disks in different clusters, the distinction between the connected and disconnected states is important. These states apply to both the cluster, relationships, and consistency groups.

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

Inconsistent (Copying)
: The primary VDisk is accessible for read and write I/O operations but the secondary VDisk is 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 VDisk contains a consistent image, but it might be out-of-date with respect to the primary VDisk. 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 VDisk is accessible for read and write I/O operations. The secondary VDisk is accessible for read-only I/O operations.

Idling A master VDisk and a auxiliary VDisk operates in the primary role. Consequently the VDisk is accessible for write I/O operations.

Idling (Disconnected)
The VDisk is 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 VDisk is 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 VDisk is 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

Related concepts

"Metro Mirror" on page 46
Metro Mirror (formerly known as 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 on the other virtual disk.

Metro Mirror consistency groups

Metro Mirror provides the facility to group a number of relationships into a Metro Mirror consistency group so that they can be manipulated in unison.

Certain uses of Metro Mirror require the manipulation of more than one relationship. A command issued to the consistency group will be applied to all of the relationships in the group simultaneously.

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 virtual disks (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.

Metro Mirror 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 must 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 a 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 VDisks 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 VDisks 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 VDisks 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 concepts

"Metro Mirror" on page 46
Metro Mirror (formerly known as 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 on the other virtual disk.

Configuration rules and requirements

Ensure that you understand the rules and requirements when configuring the SAN Volume Controller.

The following terms and definitions will guide you in understanding the rules and requirements:

Table 9. Configuration Terms and Definitions

Term	Definition
ISL hop	A hop on an Inter-Switch Link (ISL). With reference to all pairs of N-ports or end-nodes that are in a fabric, the number of ISL hops 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 or 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 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. Create a redundant SAN by splitting 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. The SAN Volume Controller is typically connected to a redundant SAN that is made out of two counterpart SANs.
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 supports Metro Mirror, 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 supports Metro Mirror, 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. There may be significant distances between the components in the local cluster and those in the remote cluster. 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 fibre-channel port fan in	The number of hosts that can see any one port. Some controllers recommend that the number of hosts using 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	In an invalid configuration, an attempted operation will fail and will generate an error code to indicate what caused it to become invalid.

Table 9. Configuration Terms and Definitions (continued)

Term	Definition
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 entry.
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.
Fibre Channel Extender	A device for long distance communication connecting other SAN fabric components. Generally these may involve protocol conversion to ATM, IP or some other long distance communication protocol.
Mesh Configuration	A network that contains a number of small SAN switches configured to create a larger switched network. With this configuration, four or more switches are connected together in a loop with some of the paths short circuiting the loop. An example of this configuration is to have four switches connected together in a loop with ISL's for one of the diagonals. The SAN Volume Controller does not support this configuration.

Related tasks

"Configuration requirements" on page 73
You *must* perform these steps before you configure the SAN Volume Controller.

Related reference

"Configuration rules"
SAN configurations that contain SAN Volume Controller clusters can be set up in various ways.

"Maximum configuration" on page 75
Ensure that you are familiar with the maximum configurations of the SAN Volume Controller.

Configuration rules

SAN configurations that contain SAN Volume Controller clusters can be set up in various ways.

Some configurations 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 Controllers is valid if it observes *all* of the following rules. These rules are discussed in the following section.

Related concepts

"Storage subsystems" on page 54
Follow these rules when you are planning the configuration of storage subsystems in the SAN fabric.

Related reference

"Host bus adapters" on page 58
Follow these configuration rules for host bus adapters (HBAs).

"Nodes" on page 59
Follow these configuration rules for nodes.

"Power requirements" on page 60
Note the power requirements for the SAN Volume Controller.

"Fibre-channel switches" on page 60
Follow these guidelines for configuring the fibre-channel switches that are supported on the SAN.

Storage subsystems

Follow these rules when you are planning the configuration of storage subsystems in the SAN fabric.

All SAN Volume Controller nodes of a cluster must be able to see the same set of storage subsystem ports on each device. Any operation that is in this mode in which two nodes do not see the same set of ports on the same device is degraded, and the system logs errors that request a repair action. This rule can have important effects on a storage subsystem 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 bridges a separate host device and a RAID array is notsupported. Typical compatibility matrixes are shown in a document titled *Supported Hardware List* on the following Web page:

http://www-1.ibm.com/servers/storage/support/virtual/2145.html

The SAN Volume Controller clusters must not share its storage subsystem logical units (LUs) with hosts. A storage subsystem can be shared with a host under certain conditions as described in this topic.

Certain storage controllers can be configured to safely share resources between the SAN Volume Controller and direct attached hosts. This type of configuration is described as a split controller. To determine support for a particular storage controller (and the restrictions on this support) by the SAN Volume Controller, refer to the section of this guide that describes configuring SAN devices and switches. In all cases, it is critical that the controller and SAN is configured so that the SAN Volume Controller cannot access logical units (LUs) that a host or another SAN Volume Controller can also access. This can be arranged by controller logical unit number (LUN) mapping and masking. If this is not guaranteed, then data corruption can occur.

Where SAN Volume Controller supports a controller being split between a SAN Volume Controller and a host as described above, the SAN Volume Controller also supports configurations in which a controller is split between two SAN Volume Controller clusters. In all cases, it is critical that the controller and SAN is configured so that the SAN Volume Controller cannot access LUs that a host or another SAN Volume Controller can also access. This can be arranged by controller LUN mapping and masking. If this is not guaranteed then data corruption can occur. This configuration is not recommended because of the risk of data corruption.

One storage subsystem device should not be configured to present the same LU to more than one SAN Volume Controller cluster. This configuration is not supported and is very likely to cause undetected data loss or corruption.

The SAN Volume Controller must be configured to manage only LUNs that are presented by supported disk controller systems. Operation with other devices is not supported.

Unsupported storage subsystem (generic device)

When a storage subsystem is detected on the SAN, the SAN Volume Controller attempts to recognize it using its Inquiry data. If the device is recognized as one of the explicitly supported storage models, the SAN Volume Controller uses error recovery programs that are potentially tailored to the known needs of the storage subsystem. If the device is not recognized, the SAN Volume Controller configures the device as a generic device. A generic device might not function correctly when it is addressed by a SAN Volume Controller. In any event, the SAN Volume Controller does not regard accessing a generic device as an error condition and, consequently, does not log an error. Managed disks (MDisks) presented by generic devices are not eligible to be used as quorum disks.

Split controller configurations

The SAN Volume Controller is configured to manage LUs exported only by RAID controllers. Operation with other RAID controllers is illegal. While it is possible to use the SAN Volume Controller to manage just a bunch of disks (JBOD) LUs presented by supported RAID controllers, it should be noted that the SAN Volume Controller itself provides no RAID function, so these LUs would be exposed to data loss in the event of a disk failure.

If a single RAID controller presents multiple LUs, either by having multiple RAID arrays configured or by partitioning one or more RAID arrays into multiple LUs, each LU can be owned by either SAN Volume Controller or a directly attached host. Suitable LUN masking must be in place to ensure that LUs are not shared between SAN Volume Controllers and direct attached hosts.

In a split controller configuration, a RAID array presents LUs to both a SAN Volume Controller (which treats the LU as an MDisk) and to another host. The SAN Volume Controller presents virtual disks (VDisks) that are created from the MDisk to another host. There is no requirement for the pathing driver in the two hosts to be the same (although, if the RAID controller were an ESS, both hosts would use SDD). Figure 18 on page 56 shows that the RAID controller is a FAStT, with RDAC used for pathing on the directly attached host, and SDD used on the host that is attached with the SAN Volume Controller. Hosts can simultaneously access LUs, that are provided by the SAN Volume Controller and directly by the device.

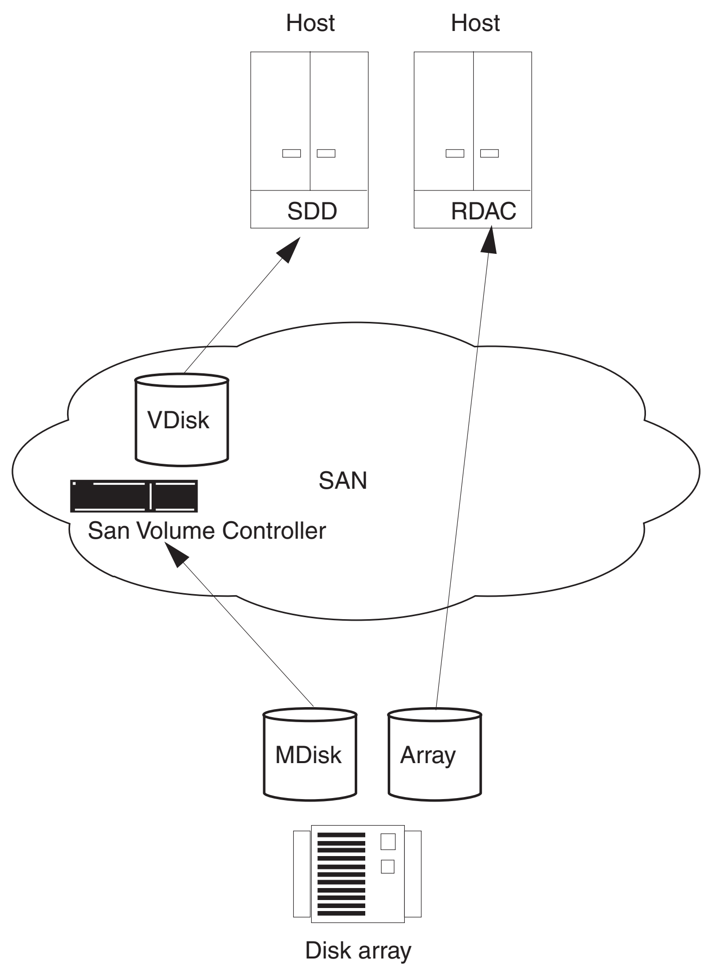

Figure 18. Disk controller system shared between SAN Volume Controller and a host

It is also possible to split a host so that it accesses some of its LUNs through the SAN Volume Controller and some directly. In this case, the multi-pathing software used by the controller must be compatible with SAN Volume Controllers multi-pathing software. Figure 19 on page 57 is a supported configuration because the same pathing driver is used for both direct and VDisks.

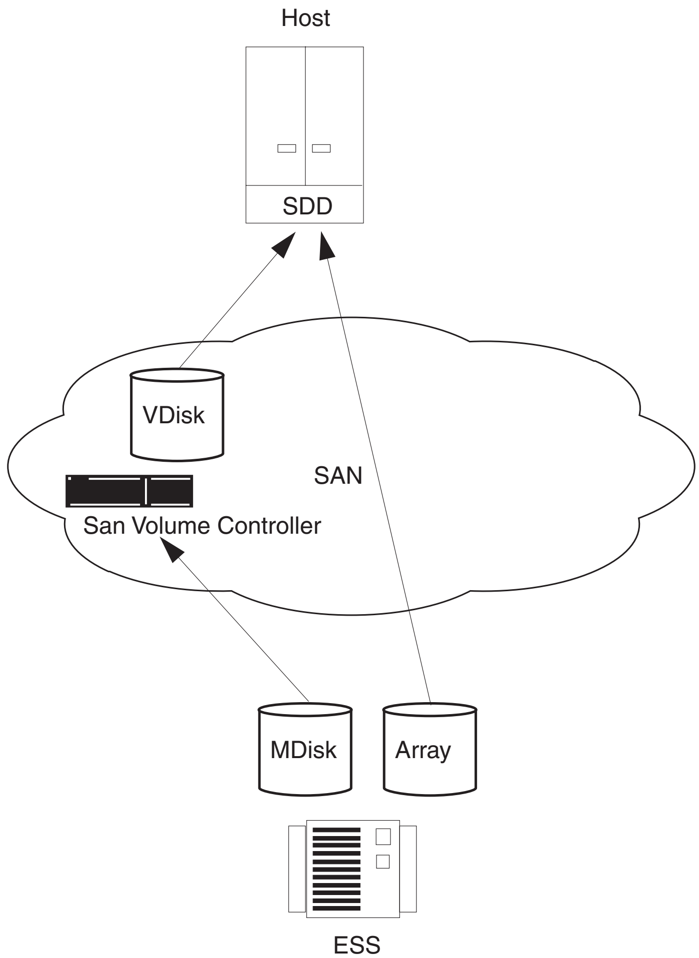

Figure 19. ESS LUs accessed directly with a SAN Volume Controller

In the case where the RAID controller uses multi-pathing software which is compatible with SAN Volume Controllers multi-pathing software, (see Figure 20 on page 58) it is possible to configure a system such that some LUNs are mapped directly to the host and others are accessed through SAN Volume Controller. One example would be ESS which uses the same multi-pathing driver as SAN Volume Controller.

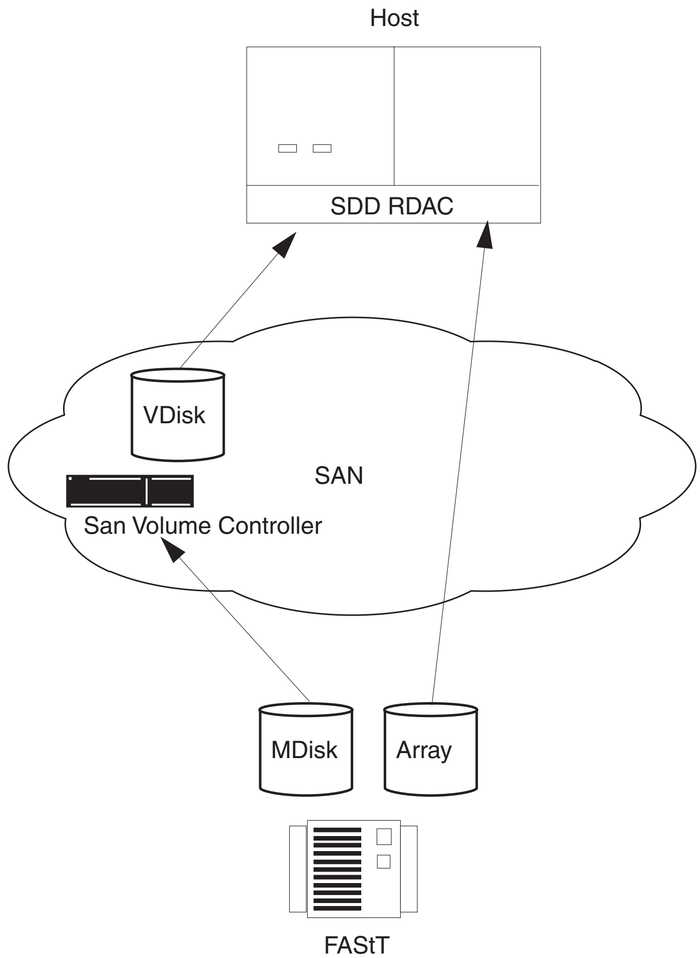

Figure 20. FAStT direct connection with a SAN Volume Controller on one host

Related reference

"Configuration rules" on page 53
SAN configurations that contain SAN Volume Controller clusters can be set up in various ways.

Host bus adapters

Follow these configuration rules for host bus adapters (HBAs).

SAN Volume Controller nodes always contain two HBAs. Each HBA must present two ports. If an HBA fails, the configuration is still valid, and the node operates in degraded mode. If an HBA is physically removed from a SAN Volume Controller node, the configuration is not supported.

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 HP/UX® host and a Windows® 2000 server host, those hosts must be in separate zones. Here, *dissimilar* means that the hosts are running different operating systems or that they are different hardware platforms. Different levels of the same operating system are considered to be similar. This requirement ensures that different SANs can operate with each other. A configuration that breaks this requirement is not supported.

The SAN Volume Controller must be configured to export virtual disks (VDisks) 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-1.ibm.com/servers/storage/support/virtual/2145.html

Operation with other HBAs is not supported.

The number of paths from the SAN Volume Controller nodes to a host must not exceed eight. The maximum number of HBA ports must not exceed four (for example, no more than two 2-port HBAs or four 1-port HBAs). Each SAN Volume Controller node in an I/O group presents four images of a VDisk onto the SAN, and each host SAN attachment has up to four HBA ports. Therefore, with more simplified zoning, the number of paths can equal up to 32: 4 SAN Volume Controller ports x 2 nodes per I/O group x 4 HBA ports. If you want to restrict the number of paths to a host, zone the switches so that each HBA port is zoned with one SAN Volume Controller port for each node in the cluster. If a host has multiple HBA ports, zone each port to a different set of SAN Volume Controller ports to maximize performance and redundancy.

Related reference

"Configuration rules" on page 53
SAN configurations that contain SAN Volume Controller clusters can be set up in various ways.

Nodes

Follow these configuration rules for nodes.

The SAN Volume Controller 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 SAN Volume Controller node always contains two host bus adapters (HBAs), each of which presents two ports. If an HBA fails, this remains a valid configuration, and the node operates in degraded mode. If an HBA is physically removed from a SAN Volume Controller node, then the configuration is unsupported.

Each SAN Volume Controller node presents a virtual disk (VDisk) to the SAN through four ports. Since each VDisk is accessible from the two SAN Volume Controller nodes in an I/O group, this means that a host HBA may see up to eight paths to each logical unit (LU) presented by the SAN Volume Controller. The hosts must run a multi-pathing device driver to resolve this back to a single device.

Support for optical connections is based on the fabric rules that the manufacturers impose for the following connection methods:

- Node to a switch
- Host to a switch
- Backend to a switch
- Switch to an Inter-Switch Link (ISL)

The optical connections supported between node and switch, host and switch, back-end and switch, and switch ISL should be determined by the fabric rules imposed by the vendors of the components used to connect the cluster.

For connections between an SAN Volume Controller node and its switch, the following optical connections are supported:

- Short wave optical fibre
- Long wave optical fibre up to 10 KM

High-power Gigabit Interface Converters (GBICs) and longwave fibre connections beyond 10 KM are not supported.

To ensure cluster failover operations, all nodes in a cluster must be connected to the same IP subnet.

The number of paths through the network from the SAN Volume Controller to a host must not exceed eight. Configurations in which this number is exceeded are unsupported. The SAN Volume Controller has four ports per node with two nodes in an I/O group. Therefore, without any zoning, the number of paths to a VDisks is eight * the number of host ports.

The nodes must be connected to the uninterruptible power supply using the supplied cable which joins together the signal and power cables.

All nodes in the cluster should be connected to the same IP subnet to ensure cluster failover operation.

Related reference

"Configuration rules" on page 53
SAN configurations that contain SAN Volume Controller clusters can be set up in various ways.

Power requirements

Note the power requirements for the SAN Volume Controller.

The uninterruptible power supply must be in the same rack that contains the SAN Volume Controller nodes that it supplies. The combination power and signal cable for connection between SAN Volume Controller and uninterruptible power supply units is two meters long. The SAN Volume Controller and uninterruptible power supply must connect with both the power and the signal cable to function correctly.

Related reference

"Configuration rules" on page 53
SAN configurations that contain SAN Volume Controller clusters can be set up in various ways.

Fibre-channel switches

Follow these guidelines for configuring the fibre-channel switches that are supported on the SAN.

The SAN must contain only supported switches. The SAN Volume Controller supports specific IBM 2109, McData, and InRange switch models and the Cisco MDS 9000 switch and switches supported by the Cisco MDS 9000.

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

http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Operation with other switches is not supported.

Ideally, the SAN would consist of two independent switches (or networks of switches) so that the SAN includes a redundant fabric, and has no single point of failure. If one SAN fabric fails, the configuration is in a degraded mode, but it is still valid. If the SAN contains only one fabric, it is still a valid configuration, but a failure of the fabric might cause a loss of access to data. Such a SAN, therefore, is considered a possible single point of failure.

Within an individual SAN fabric, switches must be from the same manufacturer with the specific exceptions for BladeCenter.

Note: Where a pair of counterpart fabrics (for example, Fabric A and Fabric B) are used to provide a redundant SAN, it is supported for Fabric A to be comprised of one vendors switches and Fabric B to use a different vendors switch so long as each individual fabric complies with the following:

- Different vendor switches may be mixed in a SAN Volume Controller configuration, provided that each fabric contains only switches from a single vendor. Thus, the two counterpart SANs can have different vendor switches.

SANs that are created from a mesh of switches are not supported.

Configurations with more than two SANs are not supported.

The fibre-channel SAN connections between SAN Volume Controller and the switches are optical fibre running at 2 Gb/s, however, SAN Volume Controller is also supported in 1 Gb/s fibre-channel fabrics. The fibre-channel switch will need to be zoned to permit the hosts to see the SAN Volume Controller nodes and the SAN Volume Controller nodes to see the RAID controllers. The SAN Volume Controller nodes within a cluster must be able to see each other.

On the fibre-channel SAN, the SAN Volume Controller nodes must always and only be connected to SAN switches. Each node must be connected to each of the counterpart SANs that are in the redundant fabric. Any operation that uses direct connections between host and node, or controller and node, is not supported.

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 mini-hubs 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.

The connections between the switches and the SAN Volume Controllers can operate at 1 Gbps or at 2 Gbps and are made with optical fibre. All the ports for the SAN Volume Controllers that are in a single cluster, however, must run at one speed. Any operation that runs different speeds on the node-to-switch connections that are in a single cluster is not valid. Mixed speeds are permitted within the fabric. You can use lower speeds to extend distance or to make use of 1 Gb/s legacy components.

Attention: The default transfer rate in the SAN Volume Controller is 2 Gbps. If your environment is set up to use 1 Gbps switches, the switch rate must be set at the transfer rate.

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

Each device has a fibre-channel connection and an Ethernet connection for configuration and error reporting. These connections are aggregated together through an Ethernet hub.

When attaching SAN Volume Controller to a SAN fabric containing core directors and edge switches, it is preferable to connect the SAN Volume Controller ports to the core directors and to connect the host ports to the edge switches. In such a fabric, the next priority for connection to the core directors is the storage controllers, leaving the host ports connected to the edge switches.

The switch configuration of a SAN Volume Controller 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 nodes can see the back-end storage and the front-end HBAs. However, the front-end HBAs and the back-end storage must not be in the same zone. Any operation that runs outside these zoning rules is not supported.

All SAN Volume Controller nodes in a SAN Volume Controller cluster must be able to see the same set of back-end storage ports on each back-end controller. Operation in a mode where two nodes see a different set of ports on the same controller is degraded and the system will log errors requesting a repair action. This could occur if inappropriate zoning was applied to the fabric. If could also occur if inappropriate LUN masking is used. This rule has important implications for back-end storage, such as FAStT which imposes exclusive rules on which HBA WWN's a storage partition can be mapped to.

Because each SAN Volume Controller has four ports, the switches can be zoned so that a particular SAN Volume Controller port is used only for internode communication, for communication to the host, or for communication to back-end storage. Whatever the configuration, each SAN Volume Controller node must remain connected to the full SAN fabric. Zoning must not be used to split the SAN into two parts.

Fibre-channel switches and Inter-Switch Links

The local or remote fabric must not contain more than three Inter-Switch Links (ISLs) hops in each fabric. Any operation that uses more than three hops is not supported. When a local fabric is connected to a remote fabric for Metro Mirror, the hop count between a local node and a remote node must not exceed seven. Therefore, some hops can 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.

Note: You must use trunking when you use multiple ISLs in parallel.

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.

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

With Inter-Switch Links between nodes in the same cluster, the ISLs are considered a single point of failure. This is illustrated in Figure 21.

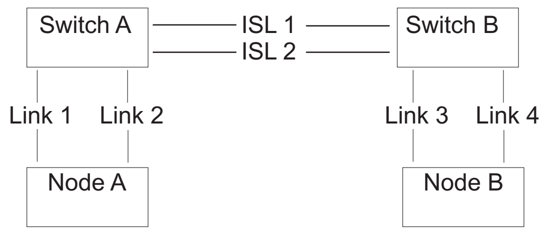

Figure 21. Fabric with Inter-Switch Links between nodes in a cluster

If Link 1 or Link 2 fails, the cluster communication does not fail.

If Link 3 or Link 4 fails, the cluster communication does not fail.

If ISL 1 or ISL 2 fails, the communication between Node A and Node B will fail for a period of time, and the node will not be recognized, even though there is still a connection between the nodes.

To ensure that a fibre-channel link failure does not cause nodes to fail when there are ISLs between nodes, it is necessary to use a redundant configuration. This is illustrated in Figure 22.

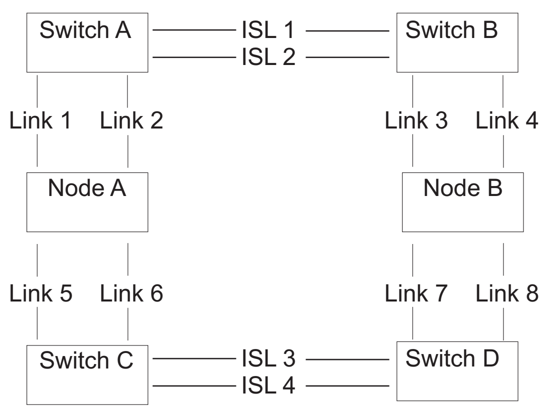

Figure 22. Fabric with Inter-Switch Links in a redundant configuration

With a redundant configuration, if any one of the links fails, then communication on the cluster will not fail.

SAN Volume Controller in a SAN with director class switches

When a large number of RAID controllers and test hosts are to be connected to a SAN Volume Controller cluster, it is possible that director class switches will be employed within the SAN. Director class switches can offer internal redundancy so it may be possible to replace a SAN that employs multiple switches to provide two redundant SAN fabrics with a single director class switch. While this may give adequate SAN network redundancy, this does not protect against physical damage (for example, flood or fire), which would destroy the entire function in the director class switch. It may be more attractive therefore to use tiered networks of smaller switches or a core-edge topology, with multiple switches in the core. The redundancy in the network can then be physically distributed across a wide area protecting against physical damage.

Related concepts

"BladeCenter fabric support"
The following table details the fabrics and switch modules in which the SAN Volume Controller environment supports.

Related reference

"Configuration rules" on page 53
SAN configurations that contain SAN Volume Controller clusters can be set up in various ways.

"Switch zoning for the SAN Volume Controller" on page 65
Consider these constraints when you zone a switch.

BladeCenter fabric support

The following table details the fabrics and switch modules in which the SAN Volume Controller environment supports.

Check the following Web site for the latest support and information:

http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Table 10. BladeCenter fabric support

<table>
<tr><th>Switch module</th><th>Supported fabrics</th><th>Notes</th></tr>
<tr><td rowspan="4">Optical Pass Through Module</td><td>Brocade</td><td rowspan="4">The BladeCenter Optical Pass Through Module is essentially a connection between fibre-channel HBAs in the BladeCenter and the external SAN. As the BladeCenter HBAs are functionally equivalent to a QLogic 231x type adapter, this optical pass through mode of attachment may be used with all SAN Volume Controller SAN configurations that support QLogic HBA attachment to the SAN Volume Controller.</td></tr>
<tr><td>McData</td></tr>
<tr><td>CNT</td></tr>
<tr><td>Cisco</td></tr>
<tr><td rowspan="2">Internal QLogic 2-port fibre-channel Switch Module for BladeCenter</td><td>Brocade</td><td rowspan="3">The BladeCenter HBAs are attached to external devices through either the QLogic or Brocade switch within the BladeCenter. When operating in this configuration the BladeCenter internal switch may only be attached to switches produced by a single vendor.</td></tr>
<tr><td>McData</td></tr>
<tr><td>Internal Brocade SAN Switch Modules for BladeCenter</td><td>Brocade</td></tr>
</table>

Port switches

The SAN Volume Controller can be used in a SAN with either 16 or 32 port switches.

Using a 16 port switch SAN

Use a SAN with 16 port switches if your configuration requires only two SAN Volume Controller nodes. A typical configuration has two SAN Volume Controller nodes and up to four RAID controller pairs. With this configuration, the SAN Volume Controller nodes and RAID controllers use eight of the port switches and leave the remaining eight port switches for host connections. You can also adjust the configuration to include more RAID controllers and less host connections. If you want to include an optional service node, you can attach the service node to one FC port on each of the switches.

Using a 32 port switch SAN

Use a SAN with 32 port switches if your configuration requires more than two SAN Volume Controller nodes. An example configuration for a large SAN Volume Controller based SAN uses a dual switch, redundant SAN fabric and four SAN Volume Controller nodes. The SAN Volume Controller nodes use 16 of the ports and the remaining 16 ports are used for RAID controllers and host connections.

Switch zoning for the SAN Volume Controller

Consider these constraints when you zone a switch.

Overview

The number of virtual paths to each virtual disk is limited. Implementation of the following rules can help you achieve the correct number of virtual paths.

- Each host (or partition of a host) can have between one and four fibre-channel ports.
- Brocade and McData switches can be configured in Vendor Interopability Mode or in Native Mode. Cisco switches are not currently supported in Vendor Interopability Mode.
- The default timeout values for fabric is the following:
  - E_A_TOV=10
  - E_D_TOV=2

Operation with values other then these specified is unsupported. IBM recommends that you manually set the domain IDs prior to building the multiswitch fabric and prior to zoning for the following reasons:

- When two switches are joined while active, they can determine if the domain ID is already in use as before. If there is a conflict the domain ID 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 IDs are negotiated at every fabric start up, there is no guarantee that the same switch will have the same ID the next time. Therefore, zoning definitions can 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 it might be necessary to reconfigure the host in order to detect devices on the switch again.

The maximum number of paths from the SAN Volume Controller nodes to a host is eight. The maximum number of host bus adapter (HBA) ports is four (for example, no more than two two-port HBAs or four one-port HBAs).

Consider the SAN environment in the following example:

- Two SAN Volume Controller nodes, nodes A and B
- Nodes A and B have four ports each
  1. Node A has ports A0, A1, A2, and A3
  2. Node B has ports B0, B1, B2, and B3
- Four hosts called P, Q, R, and S
- Each of the four hosts has four ports, as described in Table 11.

Table 11. Four hosts and their ports

P	**Q**	**R**	**S**
P0	Q0	R0	S0
P1	Q1	R1	S1
P2	Q2	R2	S2
P3	Q3	R3	S3

- Two switches called X and Y
- One storage controller
- The storage controller has four ports on it called I0, I1, I2, and I3

An example configuration would be the following:

```
1. Attach ports 1 (A0, B0, P0, Q0, R0, and S0) and 2 (A1, B1, P1, Q1, R1, and S1)
   of each node and host to switch X.
2. Attach ports 3 (A2, B2, P2, Q2, R2, and S2) and 4 (A3, B3, P3, Q3, R3, and S3)
   of each node and host to switch Y.
3. Attach ports 1 and 2 (I0 and I1) of the storage controller to switch X.
4. Attach ports 3 and 4 (I2 and I3) of the storage controller to switch Y.
```

On switch X we would create the following host zones:

```
5. Create a host zone containing ports 1 (A0, B0, P0, Q0, R0, and S0)
   of each node and host.
6. Create a host zone containing ports 2 (A1, B1, P1, Q1, R1, and S1)
   of each node and host.
```

Similarly, on switch Y we would create the following host zones:

```
7. Create a host zone on switch Y containing ports 3
   (A2, B2, P2, Q2, R2, and S2) of each node and host.
8. Create a host zone on switch Y containing ports 4
   (A3, B3, P3, Q3, R3, and S3) of each node and host.
```

Last, we would create the following storage zone:

```
9. Create a storage zone that is configured on each switch. Each
storage zone contains all the SAN Volume Controller and storage ports on that
switch.
```

In the following example, the SAN environment is similar to the first example, with two additional hosts with two ports each.

- Two SAN Volume Controller nodes called A and B
- Nodes A and B have four ports each
  1. Node A has ports A0, A1, A2, and A3
  2. Node B has ports B0, B1, B2, and B3
- Six hosts called P, Q, R, S, T and U
- Four hosts have four ports each, and two hosts have four ports each as described in Table 12 on page 67.

Table 12. Six hosts and their ports

P	Q	R	S	T	U
P0	Q0	R0	S0	T0	U0
P1	Q1	R1	S1	T1	U1
P2	Q2	R2	S2	—	—
P3	Q3	R3	S3	—	—

- Two switches called X and Y
- One storage controller
- The storage controller has four ports on it called I0, I1, I2, and I3

An example configuration would be the following:

```
1.  Attach ports 1 (A0, B0, P0, Q0, R0, S0 and T0) and 2
    (A1, B1, P1, Q1, R1, S1 and T0) of each node and host to switch X.
2.  Attach ports 3 (A2, B2, P2, Q2, R2, S2 and T1) and 4
    (A3, B3, P3, Q3, R3, S3 and T1) of each node and host to switch Y.
3.  Attach ports 1 and 2 (I0 and I1) of the storage controller to switch X.
4.  Attach ports 3 and 4 (I2 and I3) of the storage controller to switch Y.
```

Attention: Hosts T and U (T0 and U0) and (T1 and U1) are zoned to different SAN Volume Controller ports so that each SAN Volume Controller port is zoned to the same number of host ports.

On switch X we would create the following host zones:

```
5.  Create a host zone containing ports 1 (A0, B0, P0, Q0, R0, S0 and T0) of each
    node and host.
6. Create a host zone containing ports 2 (A1, B1, P1, Q1, R1, S1 and U0) of each
    node and host.
```

Similarly, on switch Y we would create the following host zones:

```
7.  Create a host zone on switch Y containing ports 3
    (A2, B2, P2, Q2, R2, S2 and T1) of each node and host.
8.  Create a host zone on switch Y containing ports 4
    (A3, B3, P3, Q3, R3, S3 and U1) of each node and host.
```

Last, we would create the following storage zone:

```
9.  Create a storage zone configured on each switch. Each
storage zone contains all the SAN Volume Controller and
storage ports on that
switch.
```

Related reference

"Switch zoning limitations for the EMC CLARiiON" on page 260
There are limitations in switch zoning for the SAN Volume Controller and EMC CLARiiON.

"Fibre-channel switches" on page 60
Follow these guidelines for configuring the fibre-channel switches that are supported on the SAN.

Zoning guidelines

Ensure you are familiar with the following zoning guidelines.

Guidelines

- The number of paths through the network from the SAN Volume Controller nodes to a host must not exceed 8. Configurations in which this number is exceeded are unsupported.
  - SAN Volume Controller has 4 ports/nodes with 2 nodes in an I/O group. Therefore, without any zoning the number of paths to a VDisk would be 8 * the number of host ports.
  - This rule exists to limit the number of paths that need to be resolved by SDD.

Controller zones

Switch zones containing controller ports must not contain more than 40 ports. A configuration that breaches this rule is unsupported.

SAN Volume Controller zones

The switch fabric must be zoned so that the SAN Volume Controller nodes can see the back-end storage and the front end host HBA's. Usually, the front-end host HBA's and the back-end storage will not be in the same zone. The exception to this would be where split host and split controller configuration is in use.

It is possible to zone the switches in such a way that a particular SAN Volume Controller port is used solely for inter node communication, or for communication to host, or for communication to back-end storage. This is possible since each SAN Volume Controller contains 4 ports. Each SAN Volume Controller node must still remain connected to the full SAN fabric. Zoning may not be used to separate the SAN into two parts.

With Metro Mirror, 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.

If a SAN Volume Controller node can see another SAN Volume Controller node through multiple paths some of which use ISLs and some of which do not then zoning should be used where possible to ensure that the SAN Volume Controller to SAN Volume Controller communication does not travel over an ISL. Likewise if SAN Volume Controller can see a storage controller through multiple paths, some of which travel over ISLs and some of which do not then consideration should be given to using zoning to restrict communication to those paths which do not travel over ISLs.

Host zones

Switch zoning containing host HBA's must not contain either host HBA's in dissimilar hosts or dissimilar HBA's in the same host which are in separate zones. Dissimilar hosts means that the hosts are running different operating systems or are different hardware platforms, thus different levels of the same operating system are regarded as similar.

Switch zones containing host HBA's must contain no more than 40 initiators in total including the SAN Volume Controller ports which act as initiators. A configuration that breaches this rule is unsupported.

Note: If the switch vendor recommends less ports per zone for a particular SAN then the more strict rules imposed by the fibre-channel vendor takes precedence over the SAN Volume Controller rules.

A valid zone would be 32 host ports plus 8 SAN Volume Controller ports.

It is strongly recommended that switch zoning be used to ensure that each host fibre-channel port is zoned to exactly 1 fibre-channel port of each SAN Volume Controller node in the cluster. For configurations with more than 64 hosts, this is required in order for the configuration to be supported. Zoning in this way also reduces the number of paths to no more than two per host adapter.

To obtain the best performance from a host with multiple fibre-channel ports, the zoning should ensure that each fibre-channel port of a host is zoned with a different group of SAN Volume Controller ports.

To obtain the best overall performance of the subsystem and to prevent overloading, the workload to each SAN Volume Controller port should be equal. This will typically involve zoning approximately the same number of host fibre-channel ports to each SAN Volume Controller fibre-channel port.

Zoning considerations for Metro Mirror

Consider these constraints when you zone a switch to support the Metro Mirror service.

SAN configurations that use the Metro Mirror feature between two clusters require the following additional switch zoning considerations:

- Additional zones for Metro Mirror. For Metro Mirror operations involving two clusters, these clusters must be zoned so that the nodes in each cluster can see the ports of the nodes in the other cluster.
- Use of extended fabric settings in a switched fabric.
- Use of Inter Switch Link (ISL) trunking in a switched fabric.
- Use of redundant fabrics.

Note: These considerations do not apply if the simpler, intracluster mode of Metro Mirror operation is in use, when only a single cluster is needed.

For intracluster Metro Mirror relationships, no additional switch zones are required. For intercluster Metro Mirror relationships, you must perform the following steps:

1. Form a SAN that contains both clusters that are to be used in the Metro Mirror relationships. If cluster A is in SAN A originally, and cluster B is in SAN B originally, there must be at least one fibre-channel connection between SAN A and SAN B. This connection will be one or more inter-switch links. The fibre-channel switch ports that are associated with these inter-switch ports must not appear in any zone.
2. A single SAN can only be formed out of combining SAN A and SAN B if the domain numbers of the switches in each SAN are different, prior to the connection of the two SANs. Ensure that each switch has a different domain ID before you connect the two SANs.
3. Once the switches in SAN A and SAN B are connected, configure them so that they operate as a single group of switches. Each cluster must retain the same set of zones that were required to operate in the original single SAN configuration.

4. Add a new zone that contains all the switch ports that are connected to SAN Volume Controller ports. This zone contains switch ports that were originally in SAN A and in SAN B.
5. You can adjust the switch zoning so that the hosts originally in SAN A can see cluster B. This allows a host to examine data in both the local and remote cluster if required. This view of both clusters is purely optional and in some cases can complicate the way you operate the overall system. Therefore, unless it is specifically needed, this should not be implemented.
6. Verify that the switch zoning is such that cluster A cannot see any of the back-end storage that is owned by cluster B. Two clusters cannot share the same back-end storage devices.

The following zones are needed in a typical intercluster Metro Mirror configuration:
- A zone in the local cluster that contains all the ports in the SAN Volume Controller nodes in that local cluster and the ports on the backend storage that are associated with that local cluster. These zones are required whether or not Metro Mirror is in use.
- A zone in the remote cluster that contains all the ports in the SAN Volume Controller nodes in that remote cluster and the ports on the back-end storage that are associated with that remote cluster. These zones are required whether or not Metro Mirror is in use.
- A zone that contains all the ports in the SAN Volume Controller nodes in both the local and remote cluster. This zone is required for intercluster communication and is specifically required by Metro Mirror operations.
- Additional zones that contain ports in host HBAs and selected ports on the SAN Volume Controller nodes in a particular cluster. These are the zones that allow a host to see VDisks presented by an I/O group in a particular cluster. These zones are required whether or not Metro Mirror were in use.

Note:
1. While it is normal to zone a server connection so that it is only visible to the local or remote cluster, it is also possible to zone the server so that the host HBA can see nodes in both the local and remote cluster at the same time.
2. Intracluster Metro Mirror operation does not require any additional zones, over and above those needed to run the cluster itself.

Switch operations over long distances

Some SAN switch products provide features that allow the users to tune the performance of I/O traffic in the fabric in a way that can affect Metro Mirror performance.

The two most significant features are ISL trunking and extended fabric.

ISL trunking	Trunking enables the switch to use two links in parallel and still maintain frame ordering. It does this by routing all traffic for a given destination over the same route even when there may be more than one route available. Often trunking is limited to certain ports or port groups within a switch. For example, in the IBM 2109-F16 switch, trunking can only be enabled between ports in the same quad (for example, same group of four ports).For more information on trunking with the MDS, refer to ″Configuring Trunking″ on the Cisco Systems Web site. Some switch types can impose limitations on concurrent use of trunking and extended fabric operation. For example, with the IBM 2109-F16 switch, it is not possible to enable extended fabric for two ports in the same quad. Thus, extended fabric and trunking are effectively mutually exclusive. (Although it is possible to enable extended fabric operation one link of a trunked pair, this does not offer any performance advantages and adds complexity to the configuration setup. This mixed mode of operation is therefore not recommended.)
Extended fabric	Extended fabric operation allocates extra buffer credits to a port. This is important over long links usually found in intercluster Metro Mirror operation. Because the time it takes for a frame to traverse the link, it is possible to have more frames in transmission at any instant in time than would be possible over a short link. The additional buffering is required to allow for the extra frames. For example, the default license for the IBM 2109-F16 switch has two extended fabric options: Normal and Extended Normal. • Normal is suitable for short links and Extended Normal is suitable for links up to 10 km long. (With the additional Extended fabric license the user gets two extra options: Medium, up to 10-50 km and Long, 50-100 km.) • The Extended Normal setting gives significantly better performance for the links up to 10 km long. Medium and Long settings are not recommended for use in the intercluster Metro Mirror links that are currently supported.

Limiting queue depth in large SANs

If you are designing a configuration for a large SAN, you must estimate the queue depth for each node in order to avoid application failures.

The queue depth is the number of I/O operations that can be executed in parallel on a device.

If a SAN Volume Controller node reaches the maximum number of queued commands, many operating systems cannot recover if the situation persists for more than 15 seconds. This can result in one or more servers presenting errors to applications and application failures on the servers.

A large SAN is one in which the number of virtual disks (VDisks) mapped to each server and the number of servers is greater than 1 000. For example, 50 servers with each server addressing 20 VDisks.

Queue depth

The queue depth is the number of I/O operations that can be executed in parallel on a device. It is usually possible to set a limit on the queue depth on the subsystem device driver paths (or equivalent) or the host bus adapter.

Ensure that you configure the servers to limit the queue depth on all of the paths to the SAN Volume Controller disks in configurations that contain a large number of servers or VDisks.

You may have a number of servers in the configuration that will be idle or will not initiate the calculated quantity of I/O operations. If so, you may not need to limit the queue depth.

Calculating a queue depth limit

Several factors are considered in the formula for calculating the queue depth limit.

In the queue depth calculation formula, the following factors have been considered:

1. The maximum number of queued commands is per node, and there are two nodes in an I/O group. The system must continue to function when one of the nodes in an I/O group is not available. Thus, an I/O group is considered to have the same number of queued commands as a node. If a node fails, the number of paths to each disk is cut in half.
2. If a VDisk is mapped so that it can be seen by more than one server, then each of the servers may send commands to it.
3. If a device driver times out a command, it immediately reissues the command. The SAN Volume Controller will have both commands in its command queue.

Homogeneous queue depth calculation

For homogeneous queues, use the following calculation.

Homogeneous queues meet one of the following criteria:

- Queued commands are shared among all paths rather than giving some servers additional resources.
- VDisks are distributed evenly among the I/O groups in the cluster.

The queue depth for each VDisk should be set on the servers using the following calculation:

$q = ((n * 7000) / (v * p * c))$

q = Queue depth per device path.

n = Number of nodes in the cluster.

v = Number of VDisks configured in the cluster.

p = Number of paths per VDisk per host. A path is a route from a server fibre-channel port to a SAN Volume Controller fibre-channel port that provides the server access to the VDisk.

c = Number of hosts that can concurrently access each VDisk. Very few applications support concurrent access from multiple hosts to a single VDisk. This number typically will be 1.

Example

Consider the following example.

- An eight-node SAN Volume Controller cluster (n = 8)
- 4096 VDisks (v = 4096)

- One server with access to each VDisk (c = 1)
- Each host has four paths to each VDisk (p = 4)

The calculation is ((8*7 000)/(4096*4*1)) = 4.

The queue depth in the operating systems must be set to four concurrent commands per path.

Related reference

"Maximum configuration" on page 75
Ensure that you are familiar with the maximum configurations of the SAN Volume Controller.

Non-homogeneous queue depth calculation

For non-homogeneous queues, use the following calculation.

Non-homogeneous queues meet one of the following criteria:

- One or more servers are allocated additional resources so that they can queue additional commands.
- VDisks are not distributed evenly among the I/O groups in the cluster.

The queue depth for each VDisk should be set on the servers using the following calculation.

For each VDisk, consider each server to which that VDisk has a mapping. This gives a set of server/VDisk pairs. If the sum of the server and VDisk queue depth for all of the pairs is less than 7 000, then the server should not experience problems due to a full queue.

Limiting the queue depth

Once you have calculated the queue depth limit, you must apply it.

Each operating system has a way to limit the queue depth on a per VDisk basis.

An alternative to setting a limit per VDisk is to set a limit on the host bus adapter. Thus, if the queue depth per path limit is five, the server has access to 40 VDisks through two adapters (four paths). It may be appropriate to place a queue depth limit of (40*(4*5))/2 = 400 on each adapter. This allows sharing the queue depth allocation between VDisks.

Configuration requirements

You *must* perform these steps before you configure the SAN Volume Controller.

1. Your IBM service representative must have installed the SAN Volume Controller.
2. 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 six and eight member disks.

If reasonably small RAID arrays are used, it is easier to extend an MDisk group by adding a new RAID array of the same type. Construct multiple RAID devices of the same type, when possible.

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, 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.

3. Install and configure your switches to create the zones that the SAN Volume Controller requires. One zone must contain all the disk controller systems and the SAN Volume Controller nodes. For hosts with more than one port, use switch zoning to ensure that each host fibre-channel port is zoned to exactly one fibre-channel port of each SAN Volume Controller node in the cluster. Set up a zone on each fibre-channel switch that includes the master console and all of the SAN Volume Controller ports that are connected to that switch.
4. If you want the SAN Volume Controller to export redundant paths to disks, you must install the Subsystem Device Driver (SDD) on all of the hosts that are connected to the SAN Volume Controller. 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

   Version 1.4.x.x or later is required.
5. Install and configure the SAN Volume Controller master console (see the *IBM TotalStorage Master Console Installation and User's Guide*). The communication between the master console and the SAN Volume Controller runs under a client-server network application called Secure Shell (SSH). Each SAN Volume Controller cluster is equipped with SSH Server software and the master console comes to you equipped with the SSH Client software called PuTTY. You will need to configure the SSH client key pair using PuTTY on the master console. Once you have installed your master console, you can configure and administer the SAN Volume Controller using a graphical interface or a command-line interface.
   a. You can configure the SAN Volume Controller using the SAN Volume Controller Console Web-based application that is preinstalled on the master console.

      Note: You can also install the master console on another machine (which you provide) using the CD-ROM provided with the master console.

b. You can configure the SAN Volume Controller using the command-line interface (CLI) commands.
c. 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 master console, ensure that the host has an SSH client installed on it.

Note:
1) AIX comes with an installed SSH client.
2) Linux comes with an installed SSH client.
3) PuTTY is recommended for Windows.

When you and the IBM service representative have completed the initial preparation steps, you must perform the following steps:
1. Add nodes to the cluster and set up the cluster properties.
2. Create managed disk groups from the managed disks to make pools of storage from which you can create virtual disks.
3. Create host objects from the HBA fibre-channel ports to which you can map virtual disks.
4. Create virtual disks from the capacity that is available in your managed disk groups.
5. Map the virtual disks to the host objects to make the disks available to the hosts, as required.
6. Optionally, create Copy Services (FlashCopy and Metro Mirror) objects, as required.

Related reference

"Configuration rules and requirements" on page 51
Ensure that you understand the rules and requirements when configuring the SAN Volume Controller.

Maximum configuration

Ensure that you are familiar with the maximum configurations of the SAN Volume Controller.

Table 13 shows the maximum configuration values to consider when you are planning the SAN Volume Controller installation.

Table 13. SAN Volume Controller maximum configuration values

Objects	Maximum number	Comments
Cluster Properties		
Nodes	8	Arranged as four I/O groups.
I/O groups	4	Each containing two nodes.
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.

Table 13. SAN Volume Controller maximum configuration values (continued)

Objects	Maximum number	Comments
LU size	2 TB	Defined by 32-bit LBA limit.
Concurrent SCSI tasks (commands) per node	10000	– – –
Concurrent commands per node	2500	Assumes a backend latency of 100 ms.
Concurrent commands per FC port	2048	– – –
SDD	512 SAN Volume Controller vpaths per host	One vpath is created for each VDisk mapped to a host. Although the SAN Volume Controller permits 512 VDisks to be mapped to a host, the SDD limit can be exceeded by either: • Creating two (or more) host objects for one physical host and mapping more than 512 VDisks to the host using the multiple host objects. • Creating two (or more) clusters and mapping more than 512 VDisks to the host using the multiple clusters. **Note:** Both of these operations are unsupported.
VDisks per MDisk Group		Cluster limit applies.
Front-end Properties		
SAN ports	512	Maximum size of fabric, including all SAN Volume Controller nodes.
Fabrics	2	Dual fabric configurations.
Host IDs per cluster	256	A host ID is associated with a map table that associates SCSI LUNs with VDisks. It is also associated with one or more host worldwide port names.
Host ports per cluster	512	Up to 512 distinct host worldwide port names are recognized.
Host LUN size	2 TB	Defined by 32-bit LBA limit.
Virtual disks (VDisks)	4096	Includes managed-mode VDisks and image-mode VDisks.
VDisks per I/O group	1024	– – –
VDisks per host ID	512	The limit may be different based on host operating system.
VDisks-to-host mappings	20 000	– – –
Maximum persistent reservation keys	132 000	– – –
Back-end Properties		

Table 13. SAN Volume Controller maximum configuration values (continued)

Objects	Maximum number	Comments
Managed Disks (MDisks)	4096	Represents an average of 64 per worldwide node name.
Back-end Storage WWNNs	64	Maximum number of device fabric worldwide node name.
Back-end Storage WWPNs	256	16 ports per controller
LUs per back-end WWNN	4096	Maximum of 512 LUs presented for each worldwide node name.
WWNNs per subsystem	4	– – –
WWPNs per WWNN	16	The maximum number of ports per worldwide node name.
Preferred ports per subsystem	4	
Copy Services Properties		
Metro Mirror relationships per cluster	4096	– – –
Metro Mirror consistency groups	32	– – –
Metro Mirror VDisk per I/O group	16 TB	– – –
FlashCopy mappings	2048	Supports up to 512 FlashCopy mappings per consistency group.
FlashCopy consistency groups	128	– – –
FlashCopy VDisk per I/O group	16 TB	– – –

Related reference

"Configuration rules and requirements" on page 51
Ensure that you understand the rules and requirements when configuring the SAN Volume Controller.

Supported fibre-channel extenders

The SAN Volume Controller supports the CNT UltraNet Edge Storage Router, the Cisco MDS 9506 and 9509 Director, and Cisco 9516 Fabric Switches to support synchronous copy services.

The maximum one-way latency supported is 10 milliseconds when using Brocade fabric, and 34 milliseconds when using McData fabric. The relationship between latency and distance is dependent on the network and number of hops. The distance is approximately 100-150 kilometers per millisecond.

Note: Performance of copy services degrades as the distance increases.

The supported hardware for the SAN Volume Controller frequently changes. See the following Web site for the latest supported hardware:

http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Related concepts

"Performance of fibre-channel extenders"
When you are planning to use fibre-channel extenders, it is important to be aware that the performance of the link to the remote location decreases as the distance to the remote location increases.

Performance of fibre-channel extenders

When you are planning to use fibre-channel extenders, it is important to be aware that the performance of the link to the remote location decreases as the distance to the remote location increases.

For fibre-channel IP extenders, throughput is limited by latency and bit error rates. Typical I/O latency can be expected to be 10 milliseconds per kilometer. Bit error rates vary depending on the quality of the circuit provided.

You must review the total throughput rates that might be expected for your planned configuration with the vendor of your fibre-channel extender and your network provider.

Related reference

"Supported fibre-channel extenders" on page 77
The SAN Volume Controller supports the CNT UltraNet Edge Storage Router,
| the Cisco MDS 9506 and 9509 Director, and Cisco 9516 Fabric Switches to support synchronous copy services.

Chapter 2. Preparing to configure the SAN Volume Controller

Once you've completed the two phase creation of the cluster, you can begin to configure the SAN Volume Controller.

Fundamentally, the configuration of the SAN Volume Controller begins by completing a two phase creation (initialization) of the cluster. The first phase is performed from the front panel of the cluster. The completion of the creation of the cluster is performed from the SAN Volume Controller Console which is accessible from a Web server running on the master console.

Related tasks

"Create cluster from the front panel"
After you have created a pair of nodes, you can now create a cluster from the front panel.

Create cluster from the front panel

After you have created a pair of nodes, you can now create a cluster from the front panel.

Ensure that the SAN Volume Controller nodes are installed. Prior to configuring the cluster, should you choose to have the customer engineer (CE) initially create the cluster, ensure that you have supplied the following information to your customer engineer:

1. Ensure that you have the correct license. The license will show you whether you are permitted to use flash copy or metro mirror. It will also show how much virtualization you are permitted to use.
2. You must supply the following information to enable the customer engineer to start the configuration procedure:
   - Cluster IP address. This address must be unique, otherwise communication problems can occur.
   - Subnet mask
   - Gateway IP address

   The customer engineer uses the front panel of the SAN Volume Controller to enter the information that you have supplied. The SAN Volume Controller generates a random password on the display panel that the customer engineer will give to you.
3. Make a note of the password and the IP address. You need it when you connect to the web application program to create the cluster.

Ensure that you have a new pair of nodes and you want to make a cluster. You also want to gain access to this cluster to start your configuration. The steps are as follows:

1. Choose a node and create a new cluster.
2. Set the IP addresses so you can gain access to the cluster.
3. Configure your cluster.

1. Choose any node that is to become a member of the cluster that you are creating.
2. At the IBM TotalStorage SAN Volume Controller service panel, keep pressing and releasing the up or down navigation button until `Node:is` displayed.

3. Keep pressing and releasing the left or right navigation button until Create Cluster? is displayed.
4. Press the **Select** button.

   If `IP Address:` is displayed on line 1 of the screen, go to step 5.

   If `Delete Cluster?` is displayed in line 1 of the service display screen, this node is already a member of a cluster. Either you have selected the wrong node, or you have already used this node in an existing cluster. The ID of this existing cluster is displayed in line 2 of the service display screen.

   - If you selected the wrong node you can exit this procedure now. The procedure cancels automatically after 60 seconds.

     Attention: When a node is deleted from a cluster, all customer data that is contained in that node is lost.
     If you are sure that the existing cluster is not required:

     a. Press and hold the up button.
     b. Press and release the select button. The node will be restarted. Once the node has been restarted you must then restart this procedure from step 1 on page 79. `IP Address:` is displayed on the service display screen.
     c. Go to step 5.
5. **Changing the fibre-channel port speed** To display the menu that shows the current value of the fibre channel speed setting for the node, press and hold the down button. Then press the select button when the display is showing the status of one of the fibre channel (FC) ports. The setting should be either 1 Gb or 2 Gb. To change the setting, perform the following steps:

   a. Press the up or down buttons to select the speed.
   b. Press the select button when the speed you want is displayed.

   This action changes the speed of all the fibre channel ports on the node.
6. Press the select button.
7. Use the up or down navigation button to change the value of the first field of the IP Address to the value that you have chosen.
8. Use the right navigation button to move to the next field. Use the up or down navigation buttons to change the value of this field.
9. Repeat step 8 for each of the remaining fields of the IP Address.
10. When you have changed the last field of the IP Address, press the select button.
11. Press the right button. `Subnet Mask:` is displayed.
12. Press the select button.
13. Change the fields for Subnet Mask in the same way that you changed the fields for IP Address.
14. When you have changed the last field of Subnet Mask, press the select button.
15. Press the right navigation button. `Gateway:` is displayed.
16. Press the select button.
17. Change the fields for Gateway in the same way that you changed the fields for IP Address.
18. When you have changed the last field of Gateway, press the select button.
19. Keep pressing and releasing the right-hand navigation button until `Create Now?` is displayed.
20. If you are satisfied with your settings, press the select navigation button.

If you want to review your settings before you create the cluster, use the right and left buttons to review those settings. Make any necessary changes, return to `Create Now?`, then press the select button.

If the cluster is created successfully, `Password:` is displayed in line 1 of the service display screen. Line 2 contains a password that you can use to access the cluster. Make a note of this password now. The password is displayed for only 60 seconds, or until the up, down, left or right navigation button is pressed.

Attention: If you do not record the password, you will have to start the cluster configuration procedure again. When the password has been recorded, press the up, down, left, or right navigation button to delete the password from the screen.

21. If the cluster is created successfully:
    a. `Cluster:` is displayed in line 1 of the service display screen,
    b. the cluster IP address is displayed on line 2,
    c. and you have successfully completed the creating a cluster process.

Related information

Chapter 2, “Preparing to configure the SAN Volume Controller,” on page 79
Once you’ve completed the two phase creation of the cluster, you can begin to configure the SAN Volume Controller.

“Scenario: typical usage for the SAN Volume Controller Console” on page 100
This hypothetical example describes a scenario for configuring the SAN Volume Controller using the SAN Volume Controller Console.

Chapter 3. SAN Volume Controller Console

Ensure that you perform the following appropriate procedures on the SAN Volume Controller.

Related reference

"Overview of creating a cluster using the SAN Volume Controller Console" on page 87
The create cluster wizard of the SAN Volume Controller Console enables you to create a cluster.

"Advanced function FlashCopy overview" on page 120
Ensure that you are familiar with the FlashCopy advanced functions.

Related information

"Using the SAN Volume Controller Console"
The SAN Volume Controller is provided with a console that is Web-browser based. It is known as the SAN Volume Controller Console.

"Scenario: typical usage for the SAN Volume Controller Console" on page 100
This hypothetical example describes a scenario for configuring the SAN Volume Controller using the SAN Volume Controller Console.

"Advanced functions overview for the SAN Volume Controller Console" on page 123
You can perform ″advanced″ functions using the SAN Volume Controller Console.

Using the SAN Volume Controller Console

The SAN Volume Controller is provided with a console that is Web-browser based. It is known as the SAN Volume Controller Console.

Overview

The SAN Volume Controller Console can be used to create and maintain the configuration of storage associated with the SAN Volume Controller. It also provides user management and access to multiple clusters.

You can use the SAN Volume Controller Console to perform the following functions:

- Initial setup of the cluster, its nodes, and the I/O groups (or node pairs). This function includes diagnostics and error log analysis of the cluster.
- Setup and maintain managed disks and managed disk groups.
- Setup and maintain SSH keys.
- Setup and maintain virtual disks.
- Setup logical host objects.
- Map virtual disks to hosts.
- Navigate from managed hosts to virtual disk and to managed disk groups, and the reverse direction up the chain.
- Set up and start Copy Services:
  - FlashCopy and FlashCopy Consistency groups.
  - Synchronous Metro Mirror and Metro Mirror Consistency groups.

The SAN Volume Controller Console is Storage Management Initiative Specification (SMI-S) compliant.

Related reference

"Accessing the SAN Volume Controller Console"
You access the SAN Volume Controller Console by pointing a Web browser at a URL.

"SAN Volume Controller Console layout" on page 85
Ensure that you are familiar with the basic frame layout of the SAN Volume Controller Console.

"SAN Volume Controller Console banner area" on page 85
The banner area of the SAN Volume Controller Console is used for product or customer identification.

"SAN Volume Controller Console task bar" on page 86
The task bar of the SAN Volume Controller Console keeps track of all opened primary tasks and allows you to quickly go back to the previous task or move forward to the next task.

"SAN Volume Controller Console portfolio" on page 86
The portfolio area of the SAN Volume Controller Console contains task-based links that open panels in the work area.

"SAN Volume Controller Console work area" on page 86
The work area of the SAN Volume Controller Console is where you work with a cluster and the objects it contains.

Related information

"Upgrading the SAN Volume Controller Console software" on page 86
Before upgrading the software for the SAN Volume Controller Console, you must modify some of the IBM WebSphere Application Server files that are used by the SAN Volume Controller Console and IBM TotalStorage Productivity Center for Fabric.

Accessing the SAN Volume Controller Console

You access the SAN Volume Controller Console by pointing a Web browser at a URL.

The SAN Volume Controller Console is a Web-based application that you can use to manage multiple clusters. Because the application is Web-based, do not set the browser to disable popup windows. This can prevent the windows in the SAN Volume Controller Console from opening.

You access the SAN Volume Controller Console by pointing a Web browser at the following URL on your master console:

http://<svcconsoleip>:9080/ica

where *<svcconsoleip>* is the IP address of your master console.

Log on to the SAN Volume Controller Console using the superuser user name, which is `superuser`, and the superuser password, which is `passw0rd`. (Upon first access, you will be required to change the superuser password.)

Use the SAN Volume Controller Console panels to configure SAN Volume Controller clusters in your environment. Once the cluster has been configured, you can use the View Clusters panel to launch another browser window with specific information for a specific cluster.

Related information

“Using the SAN Volume Controller Console” on page 83
The SAN Volume Controller is provided with a console that is Web-browser based. It is known as the SAN Volume Controller Console.

SAN Volume Controller Console layout

Ensure that you are familiar with the basic frame layout of the SAN Volume Controller Console.

Figure 23 provides, the basic frame layout, which consists of a banner, task bar, portfolio and a work area. An optional frame can be added for embedded task assistance or help.

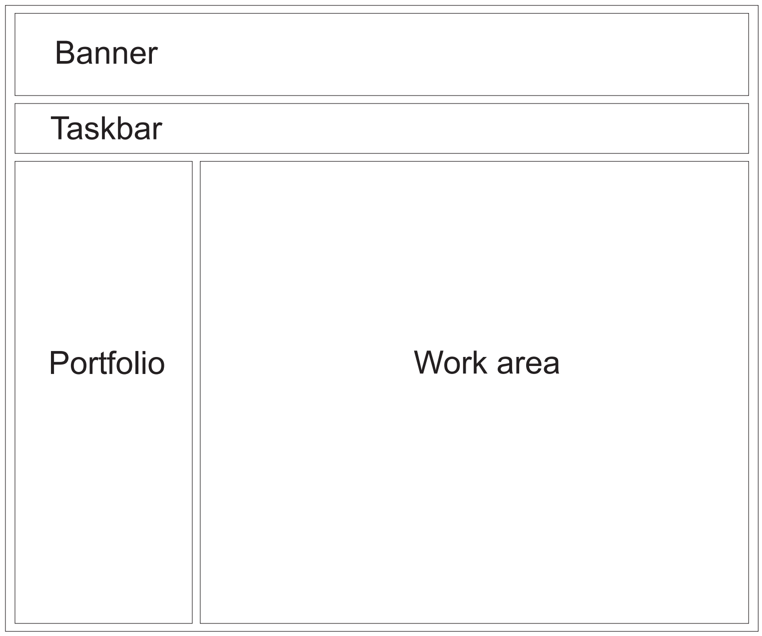

Figure 23. Basic frame layout

Related information

“Using the SAN Volume Controller Console” on page 83
The SAN Volume Controller is provided with a console that is Web-browser based. It is known as the SAN Volume Controller Console.

SAN Volume Controller Console banner area

The banner area of the SAN Volume Controller Console is used for product or customer identification.

Figure 24. Banner area

Related information

“Using the SAN Volume Controller Console” on page 83
The SAN Volume Controller is provided with a console that is Web-browser based. It is known as the SAN Volume Controller Console.

SAN Volume Controller Console task bar

The task bar of the SAN Volume Controller Console keeps track of all opened primary tasks and allows you to quickly go back to the previous task or move forward to the next task.

Clicking the **question mark (?)** icon on the right side brings up the information center in a separate browser window; clicking the (I) icon brings up a help topic for the panel that is currently displayed in the work area.

Figure 25. Task bar

Related information

"Using the SAN Volume Controller Console" on page 83
The SAN Volume Controller is provided with a console that is Web-browser based. It is known as the SAN Volume Controller Console.

SAN Volume Controller Console portfolio

The portfolio area of the SAN Volume Controller Console contains task-based links that open panels in the work area.

Common tasks are grouped under task headings and are expandable and collapsible.

Related information

"Using the SAN Volume Controller Console" on page 83
The SAN Volume Controller is provided with a console that is Web-browser based. It is known as the SAN Volume Controller Console.

SAN Volume Controller Console work area

The work area of the SAN Volume Controller Console is where you work with a cluster and the objects it contains.

The work area is the main area of the application.

Related information

"Using the SAN Volume Controller Console" on page 83
The SAN Volume Controller is provided with a console that is Web-browser based. It is known as the SAN Volume Controller Console.

Upgrading the SAN Volume Controller Console software

Before upgrading the software for the SAN Volume Controller Console, you must modify some of the IBM WebSphere Application Server files that are used by the SAN Volume Controller Console and IBM TotalStorage Productivity Center for Fabric.

For detailed instructions, see the *IBM TotalStorage Master Console Installation and User's Guide*.

Related information

"Using the SAN Volume Controller Console" on page 83
The SAN Volume Controller is provided with a console that is Web-browser based. It is known as the SAN Volume Controller Console.

Overview of creating a cluster using the SAN Volume Controller Console

The create cluster wizard of the SAN Volume Controller Console enables you to create a cluster.

Related concepts

"Browser requirements for the SAN Volume Controller"
Ensure that you are familiar with the Internet browsers and settings when using the SAN Volume Controller.

Related tasks

"Configuring a cluster using the SAN Volume Controller Console" on page 88
Once you've created a pair of nodes, you will then need to create and configure a cluster.

"Launching the SAN Volume Controller Console" on page 96
You can launch the SAN Volume Controller from the Viewing Clusters panel.

"Setting cluster time" on page 98
You can set the cluster time for the SAN Volume Controller from the Setting Cluster Time panel.

"Displaying cluster properties using the SAN Volume Controller Console" on page 99
You can display cluster properties from the View Cluster properties panel.

Browser requirements for the SAN Volume Controller

Ensure that you are familiar with the Internet browsers and settings when using the SAN Volume Controller.

The following Web browser versions are required.

- Windows and UNIX operating systems
  - Netscape version 6.2
    - Netscape is available from the following Web site:

      http://wp.netscape.com/download/archive.html
  - Internet Explorer Version 6 or later
    - You can get version 6 or later from the following Web site:

      http://www.microsoft.com/windows/ie/downloads/ie6/default.asp
- AIX operating system
  - You can get AIX Netscape version 7.0 from the following Web site:

    http://devedge.netscape.com/central/gecko/2002/download/

  You must ensure that the proxy setting is disabled.
  - For Netscape, perform the following steps:
    1. Open your Netscape browser and click **Edit -> Preferences**. The Preferences window is displayed.
    2. From the left side category, click **Advanced** to expand the sub options. The sub option Proxies is displayed.
    3. Click **Proxies**. The Proxies window is displayed.
    4. Select **Direct connection to Internet**.
  - For Internet Explorer, perform the following steps:

1. Click **Tools -> Internet Options -> Connections -> LAN Settings**.
2. Click to clear the **Use a proxy server** box.

Related reference

"Overview of creating a cluster using the SAN Volume Controller Console" on page 87
The create cluster wizard of the SAN Volume Controller Console enables you to create a cluster.

Configuring a cluster using the SAN Volume Controller Console

Once you've created a pair of nodes, you will then need to create and configure a cluster.

If you are creating a cluster using the SAN Volume Controller, you will need to generate a SSH key pair before performing this task. If you are adding an SSH public key to enable your system to use the command-line interface (CLI), you will need to have generated an SSH key pair for this system.

Perform the following steps to create a cluster on the SAN Volume Controller Console:

1. Create a cluster from the front panel of the SAN Volume Controller. A temporary password for the administrator use is generated by the node.
2. Access the SAN Volume Controller Console using a Web browser.
3. Sign on with the superuser name and password. For first-time access, use the superuser name `superuser` and the default password `passw0rd`. You must change this default password the first time you sign on. After you sign on with the superuser name and password, the Welcome panel is displayed.
4. Add a new cluster to the SAN Volume Controller.
5. Complete the Creating a Cluster wizard:
   a. Complete the creation of the cluster
   b. Set up the error notification settings
   c. Set up the featurization attributes
   d. Upload the SSH key
6. Type the IP address of the cluster and select the **Create (Initialize) Cluster** check box. When you click **OK**, the Create a Cluster wizard opens.
7. The Creating a Cluster wizard presents panels to complete the following steps:
   a. Create new cluster information, such as:
      - The new admin password
      - The service password
      - The name of the cluster
      - A service IP address
   b. Set up the error logging attributes.
   c. Set up the featurization attributes.
   d. Upload the SSH key through the wizard.

Once these steps are complete and you have exited out of the wizard, you can now use the Web application for the SAN Volume Controller passwords.

Perform the following steps to create a cluster using the Create a Cluster wizard:

1. Start the SAN Volume Controller Console by clicking on the desktop icon or by pointing your Web browser to http://<svcconsoleip>:9080/ica, where *<svcconsoleip>* is the IP address of the master console.
2. The Enter Network Password window is displayed.Type `superuser` for the user ID and `passw0rd` for the password. The first time you sign on as the superuser, you must change the password for the superuser. After you change the password, the Welcome panel is displayed.
3. If this is the first time you have accessed the SAN Volume Controller Console, go to step 3a. Otherwise, go to step 3b.
   a. The Welcome panel is shown as in Figure 27 on page 90. Click the **Add SAN Volume Controller Cluster** button.

Figure 26. Welcome panel

   b. Select **Clusters** from the portfolio. From the list of tasks, select **Add cluster** and click **Go**.

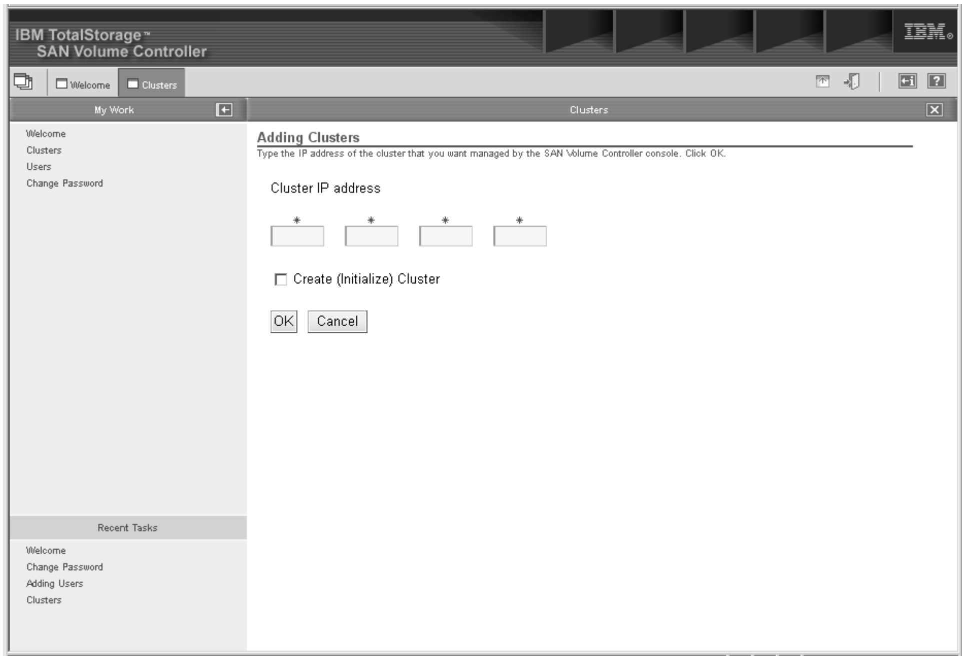

Figure 27. Add Cluster panel

4. Click **Create new cluster**. The SAN Volume Controller creates the new cluster. When the new administrator password is accepted, the cluster displays the password prompt again.
5. Type the user ID`admin` and the new administrator password.
6. Select **Add a Cluster** from the menu, then click **Go** .
7. Enter the IP address of your cluster.

   If the cluster has not been fully created (that is, you have just followed the steps in chapter 5 and created the cluster from the front panel), select the **Create (Initialize) Cluster** check box.

   If the cluster is already in use and you are just adding this cluster to the clusters that this installation of the SAN Volume Controller Console is managing, do not select the Create (Initialize) Cluster check box. Click **OK**.

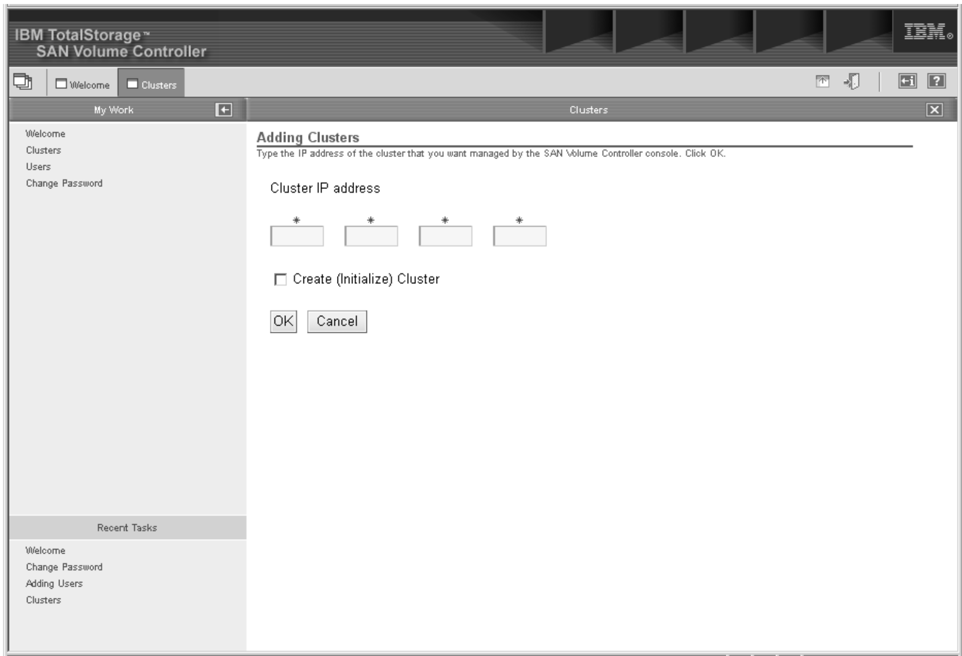

Figure 28. Add Cluster panel

8. You will be prompted to accept the new certificate for the cluster.

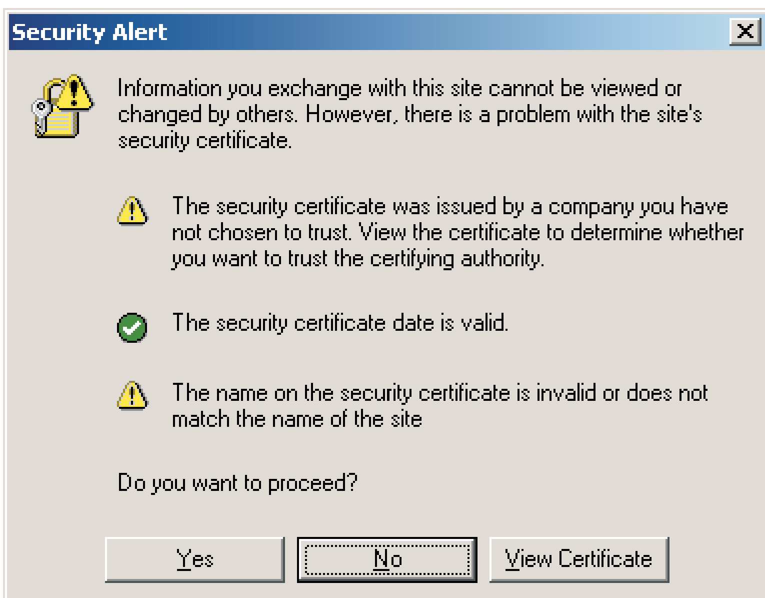

Figure 29. Security alert panel

Click **View Certificate** and on the resulting window click **Install Certificate**.

Figure 30. Certificate Information panel

Click **Next, Next, Install, OK** to complete the install certificate window.

Click **OK** to close the Certificate window as shown in Figure 29 on page 91 and click **Yes** to close the Security Alert window as shown in Figure 30.

9. You will be prompted for the cluster user name and password. The username is `admin` and the password is the one generated by the process described in the topic about Creating a cluster from the front panel. Enter the random password that was generated and click **OK**.
10. The Create a Cluster wizard begins, click **Continue**. If the cluster already existed and you did not check the **Initialize Cluster** checkbox in step 7 on page 90 proceed to step 14 on page 95.
11. Complete the Create New Cluster step by entering a new administrator password and enter a service password. Note these passwords as you will need them in the future to upload new SSH keys via the SAN Volume Controller Console.

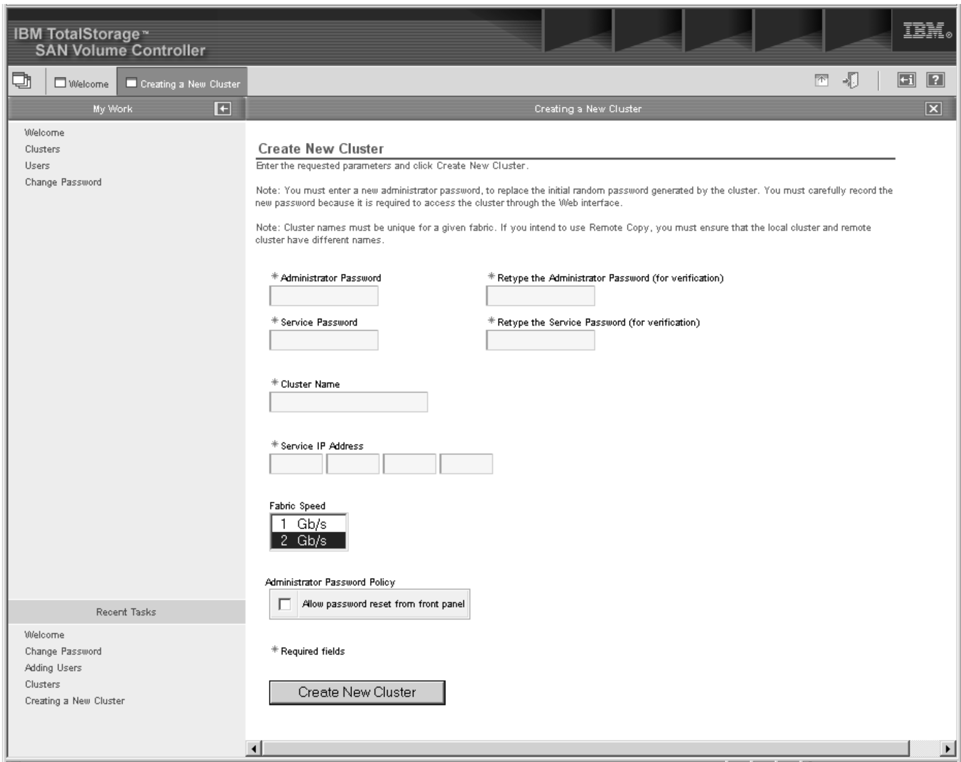

Figure 31. Create New Cluster wizard

a. Enter a name for you cluster. You can modify the cluster name by issuing the following command:

```
svctask chcluster -name <cluster_name>
```

b. Enter the service IP address for the cluster, this is the IP address that will be used if you have to bring a single node up in service mode.

c. Select the speed of your fabric, either 1 or 2 Gb/s

d. If you wish to be able to reset the administrator password from the front panel then check the box.

e. When complete click the Create New Cluster button. The cluster will then be created; this will take a few seconds. When the web page returns, click **Continue**.

12. You will then be notified that the password has been changed. Click **Continue** to proceed to the **Error Notification Settings** panel.

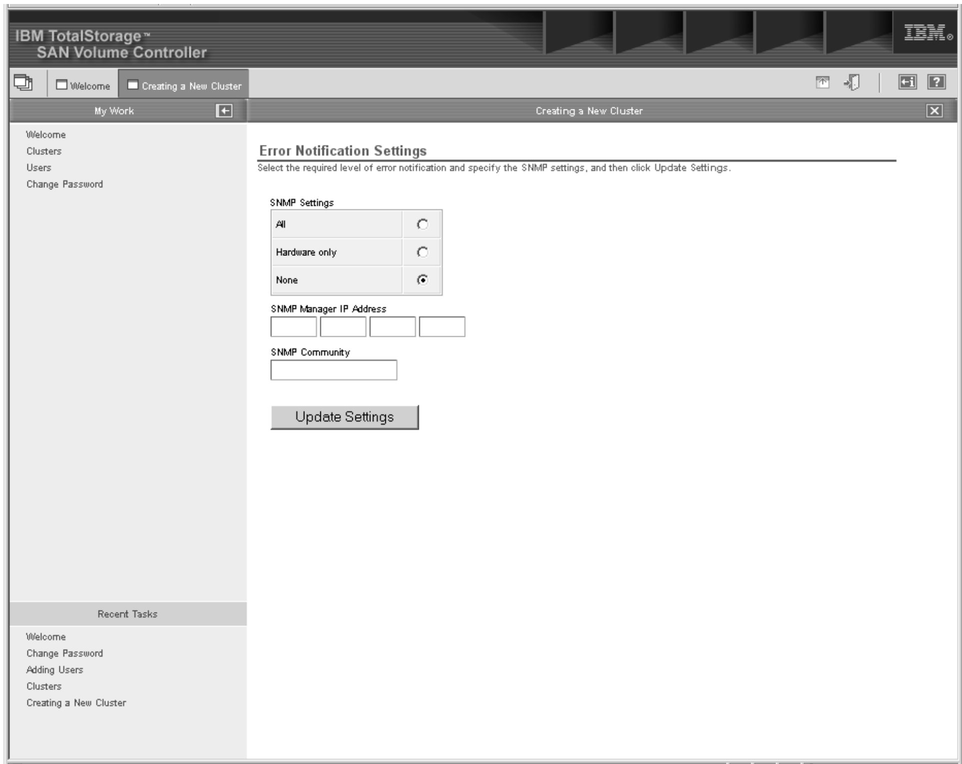

Figure 32. Error Notification Settings panel

a. If you wish errors to be forwarded as SNMP traps, select either **All** or **Hardware only**. `Hardware only` sends SNMP traps for hardware-related errors, `All` sends SNMP traps for all errors, hardware and software.

b. Enter the IP address of the machine that is running your SNMP management software (note if you are using IBM Director on the master console to collect SNMP traps, enter the IP address of the master console here)

c. Enter the SNMP community name.

d. Click **Update Settings** to continue.

13. Click **Continue**. The Featurization window is displayed.

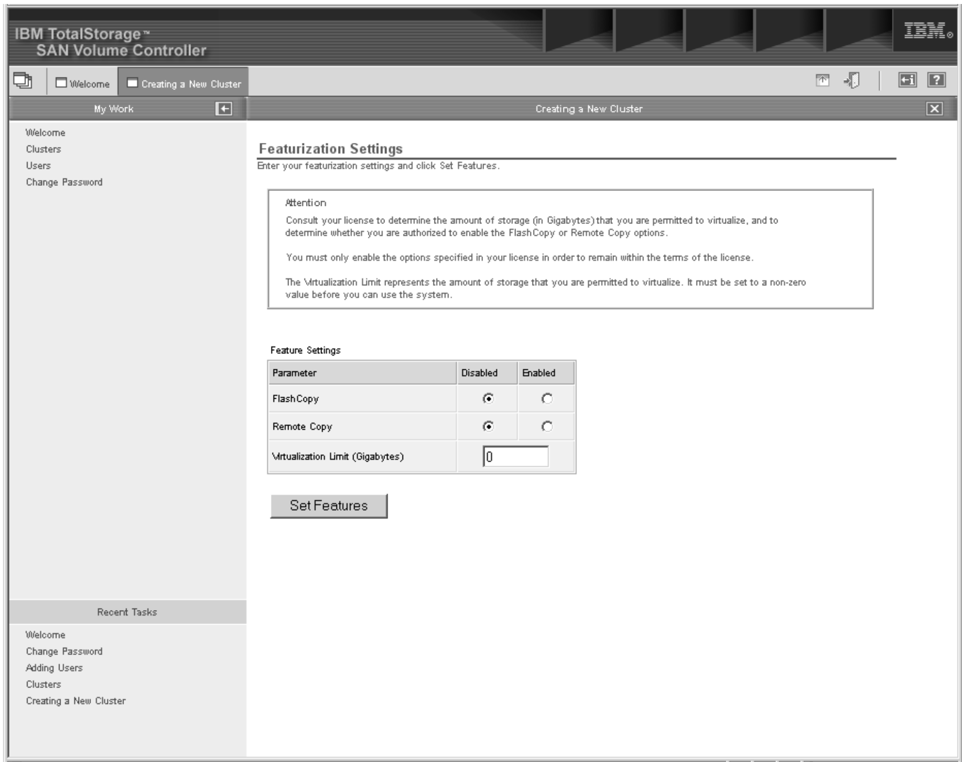

Figure 33. Featurization Settings panel

The allowed setting for each of the parameters is specified in your user's license.

a. Enable the FlashCopy or Metro Mirror options if they are licensed.

b. Enter the virtualization limit as specified in the licence. A zero value is not allowed for this field.

c. Click **Set features**. A featurization screen is displayed.

14. Click **Continue** to display the Add SSH Public Key step.

At this point you may be re-prompted for a username and password. Enter `admin` as the user name and enter the new password you gave during 11 on page 92.

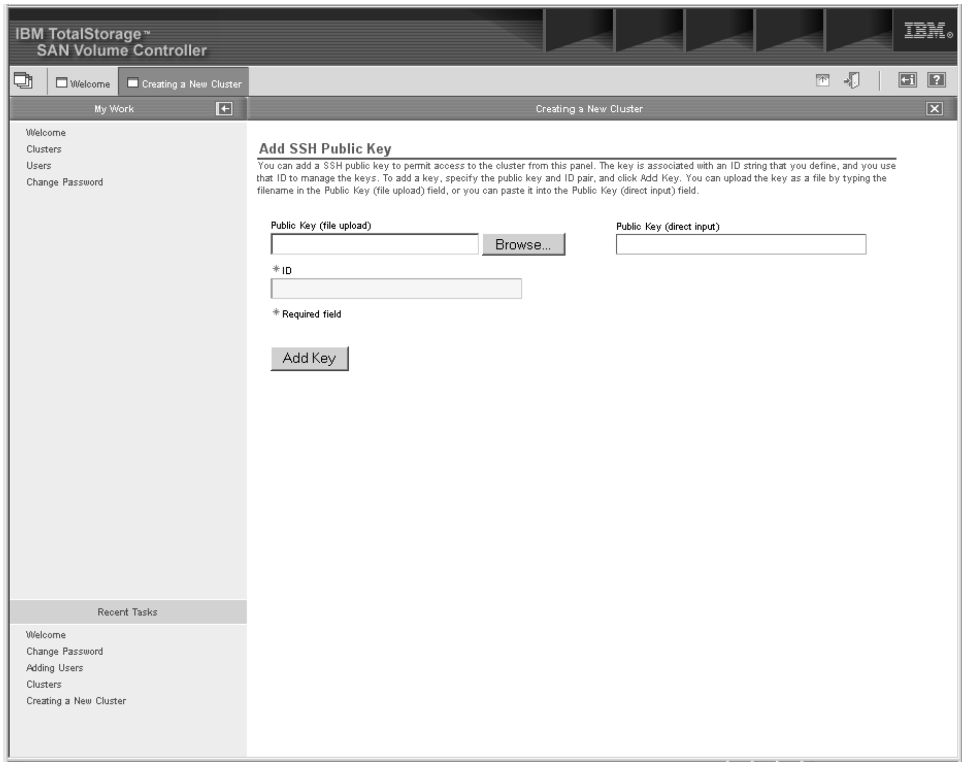

Figure 34. Add SSH public key panel

Click Browse to locate the public key for the master console.

Enter an ID (label) for this key and click **Add Key**.

15. Click on the X in the corner of the window to close the wizard.

You have now successfully connected and configured the cluster.

The cluster is created and it should be listed on the Viewing Clusters panel.

Note: You may have to press **Refresh** on the Viewing Clusters panel to see the new cluster.

Related tasks

"Create cluster from the front panel" on page 79
After you have created a pair of nodes, you can now create a cluster from the front panel.

Related reference

"Overview of creating a cluster using the SAN Volume Controller Console" on page 87
The create cluster wizard of the SAN Volume Controller Console enables you to create a cluster.

Related information

"Scenario: typical usage for the SAN Volume Controller Console" on page 100
This hypothetical example describes a scenario for configuring the SAN Volume Controller using the SAN Volume Controller Console.

Launching the SAN Volume Controller Console

You can launch the SAN Volume Controller from the Viewing Clusters panel.

The SAN Volume Controller is the centralized Web application that is used to manage your clusters. It is preinstalled on the Master console.

This procedure assumes that you are at the Welcome panel for the SAN Volume Controller.

Perform the following steps to launch the SAN Volume Controller:

1. Click **Clusters** from the portfolio. The Viewing Clusters panel is displayed.
2. Select the cluster that you want to manage with the application.
3. Select **Launch the SAN Volume Controller application** from the drop-down list and click **Go**. A secondary browser window opens.

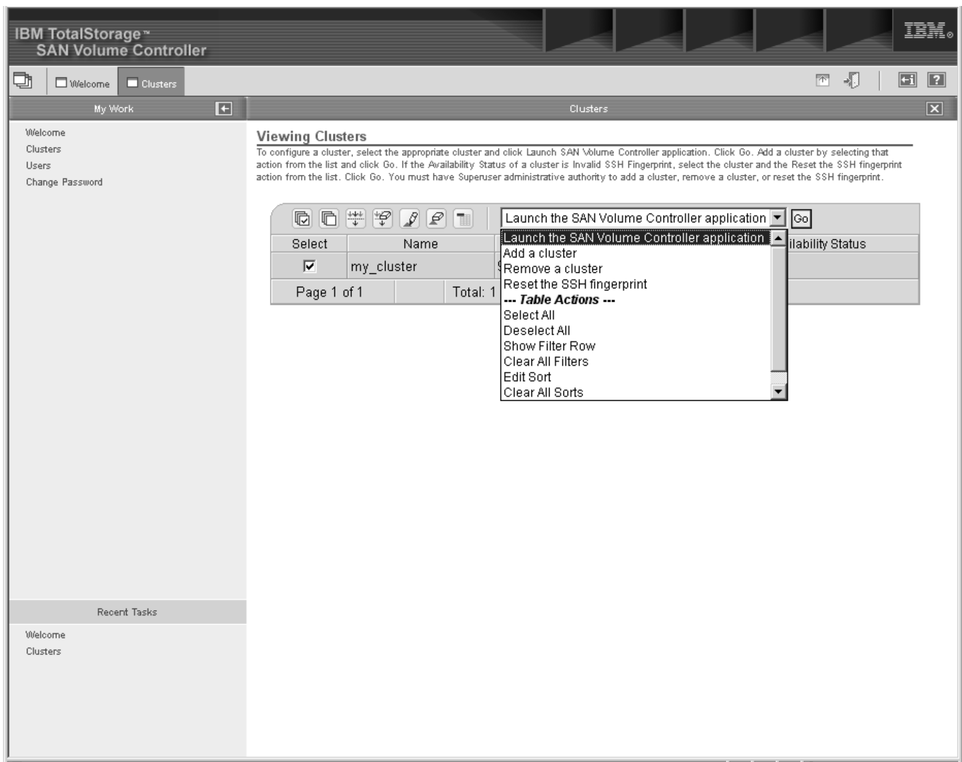

Figure 35. Viewing clusters panels

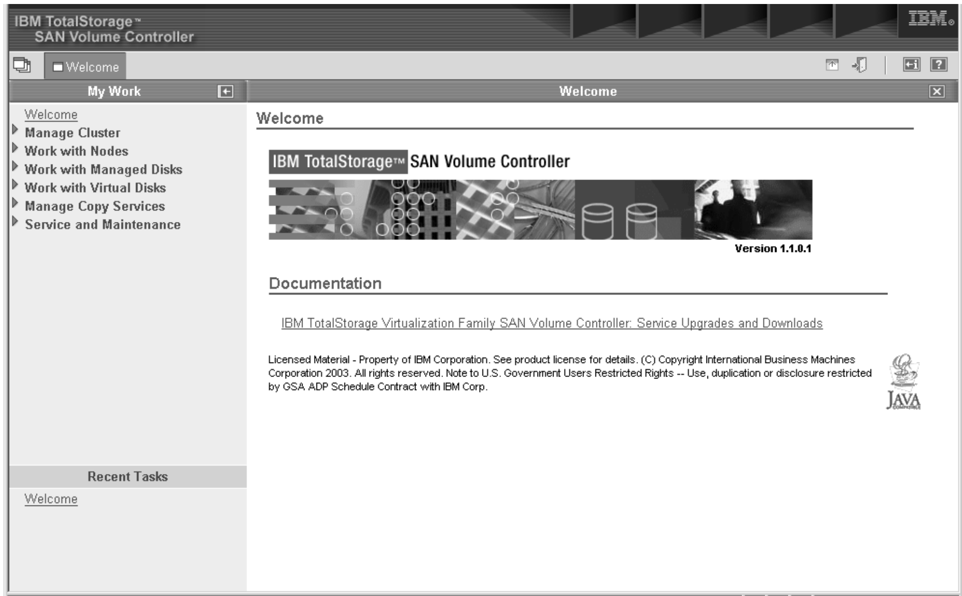

Figure 36. Welcome panel

Related concepts

"Clusters" on page 12
All configuration and service is performed at the cluster level.

Related reference

"Overview of creating a cluster using the SAN Volume Controller Console" on page 87
The create cluster wizard of the SAN Volume Controller Console enables you to create a cluster.

Setting cluster time

You can set the cluster time for the SAN Volume Controller from the Setting Cluster Time panel.

Perform the following steps to set the cluster time:

1. Click **Manage Clusters** from the portfolio.
2. Select **Set Cluster Time** from the list and click **Go**. The Setting Cluster Time panel is displayed.

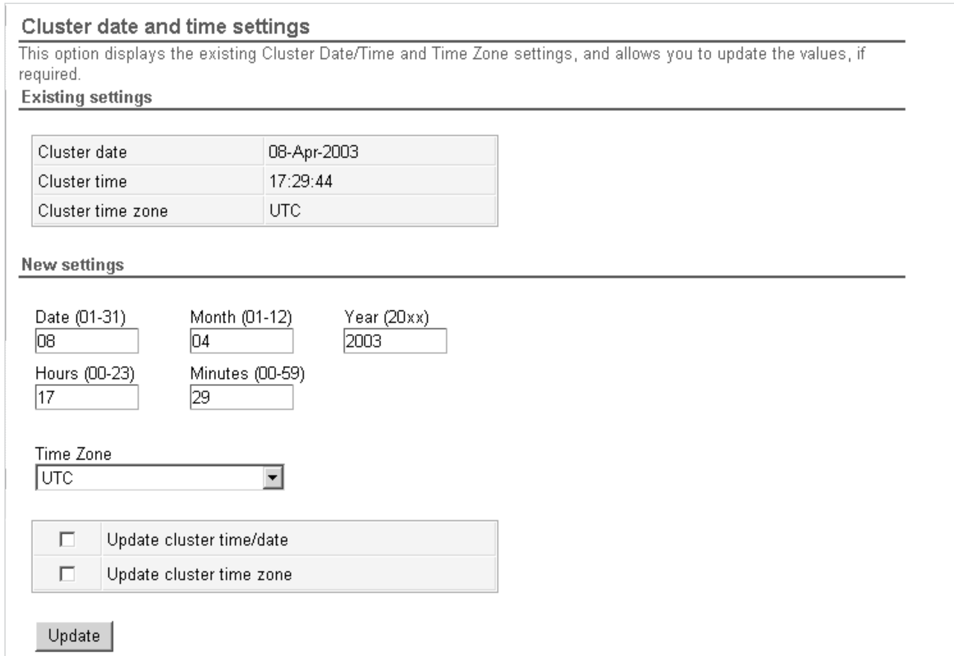

Figure 37. Cluster date and time settings panel

3. Perform the following steps to change the information in the window:
   a. Type the changes to any of the input field parameters or select a new time zone from the list.
   b. When you have made the changes, select the appropriate check-boxes to update the time, the time zone, or both.
   c. Click **Update** to submit the update request to the node.

   Related concepts

   "Clusters" on page 12
   All configuration and service is performed at the cluster level.

   Related reference

   "Overview of creating a cluster using the SAN Volume Controller Console" on page 87
   The create cluster wizard of the SAN Volume Controller Console enables you to create a cluster.

Displaying cluster properties using the SAN Volume Controller Console

You can display cluster properties from the View Cluster properties panel.

Perform the following steps to display the cluster properties:

1. Click**Manage Cluster** from the portfolio.
2. Click **View Cluster properties** to view the properties for the cluster. The Cluster Properties notebook is displayed.

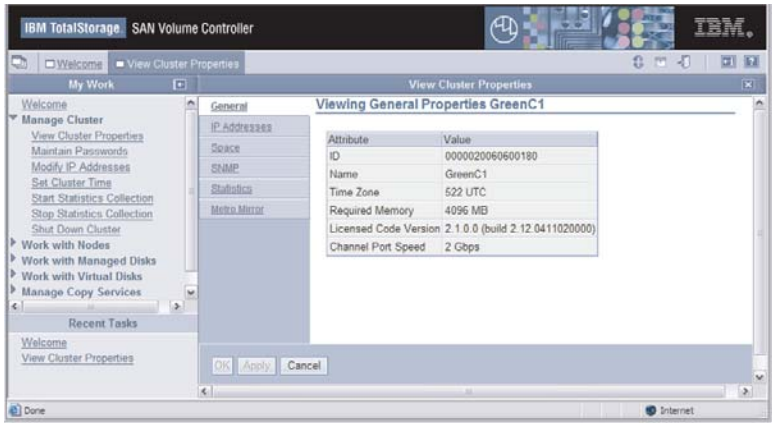

Figure 38. View Cluster properties panel

3. Click:
   a. **General** tab in the notebook to display the general properties.
   b. **IP Addresses** to view, the service IP address, the subnet mask, and the default gateway addresses.
   c. **Space** to view the space and capacity within the virtual disks (VDisks) and managed disk (MDisk) groups.
   d. **SNMP** to view the SNMP details.
   e. **Statistics** to view the cluster statistics details.
   f. **Metro Mirror** to view the Metro Mirror properties of the cluster.

Related reference

"Overview of creating a cluster using the SAN Volume Controller Console" on page 87
The create cluster wizard of the SAN Volume Controller Console enables you to create a cluster.

Scenario: typical usage for the SAN Volume Controller Console

This hypothetical example describes a scenario for configuring the SAN Volume Controller using the SAN Volume Controller Console.

The main focus of the following example is to provide storage to a host system.

For example, you wish to provide a host system with two disks and create a FlashCopy of these two disks. The copy is to be made available to a second host. These two hosts require that the host objects that are created, correspond with the group of WWPNs presented by their fibre-channel HBAs to the SAN. You also need to create four virtual disks, one for each of the disks that are to be presented to the hosts. Once the VDisks are created, you can map two of them to each host. In order to create the VDisks you need to have a managed disk group to be able to create them from. You wish to spread the 8 managed disks across two groups and create the source VDisks from one and the target VDisks from the other. In order to create any of these objects you need to create a cluster and at least one more node to the cluster.

The following steps illustrates how this can be done:

1. Create a cluster.

2. Configure the cluster with an IP address of 9.20.123.456, a fabric speed of 2 Gb/s. Name the cluster examplecluster.
3. Launch the SAN Volume Controller application for the cluster. A secondary browser window opens to the SAN Volume Controller Web application. Now you can work with the specific SAN Volume Controller cluster which you selected.
4. Add nodes
   - knode and lnode to the I/O group called io_group0 in the examplecluster cluster
   - mnode and nnode to the I/O group called io_group1 in the examplecluster cluster
5. Create the managed disk (MDisk) groups maindiskgroup and bkpdiskgroup
6. Create four virtual disks (VDisks)
   - 2 VDisks from maindiskgroup
   - 2 VDisks from bkpdiskgroup
7. Create two host objects
   - A host object called demohost1 with HBAs that have WWPNs of 210100e08b251dd4, and 210100e08b251dd5
   - A host object called demohost2 with HBAs that have WWPNs of 210100e08b251dd6, and 210100e08b251dd7
8. Create the VDisk-to-host mappings
   - Create a VDisk-to-host mapping for demohost1
   - Create a VDisk-to-host mapping for demohost2

   Once this step is complete, you have successfully created storage on your host system.
9. Create a FlashCopy consistency group called maintobkpfcopy and add the two FlashCopy mappings to it.

   Note: You must first create FlashCopy mappings to define the relationships.

Related tasks

"Create cluster from the front panel" on page 79
After you have created a pair of nodes, you can now create a cluster from the front panel.

"Adding nodes to a cluster" on page 102
You can add a node to a cluster from the Adding Nodes to a cluster panel.

"Configuring a cluster using the SAN Volume Controller Console" on page 88
Once you've created a pair of nodes, you will then need to create and configure a cluster.

"Displaying node properties using the SAN Volume Controller Console" on page 107
You can display the node properties using the Viewing general details panel.

"Creating managed disk groups" on page 107
You can create a new managed disk (MDisk) group using the Create a Managed Disk Group wizard.

"Creating virtual disks" on page 114
You can create virtual disks using the Create virtual disks wizard.

"Creating hosts" on page 118
You can create a new host object from the Creating Hosts panel.

"Showing VDisks mapped to a host" on page 117
You can show the VDisks that are mapped to a host by using the Viewing Virtual Disks panel.

"Creating consistency groups" on page 119
You can create a FlashCopy consistency group from the Creating FlashCopy Consistency Groups panel.

"Creating FlashCopy mappings" on page 120
You can create a FlashCopy mapping from the Creating FlashCopy Mappings panel.

Related information

"Using image mode virtual disks" on page 109
Ensure that you are familiar with using image mode virtual disks.

Adding nodes to a cluster

You can add a node to a cluster from the Adding Nodes to a cluster panel.

Attention: Before adding a node to a cluster, make sure that the switch zoning is configured such that the node being added is in the same zone as all other nodes in the cluster. In particular, if you are replacing a node and the switch is zoned by worldwide port name (WWPN) rather than by switch port, make sure the switch is configured such that the node being added is in the same VSAN/zone as you must update the switch configuration.

For availability purposes, you should connect the nodes in an input/output (I/O) group to different Uninterruptible Power Sources (UPSs).

Before adding a node to the cluster check to see if any of the following conditions are true:

- The cluster has more than one I/O group.
- The node being added to the cluster uses physical node hardware or a slot which has previously been used as for a node in the cluster.
- The node being added to the cluster uses physical node hardware or a slot which has previously been used for a node in another cluster and both clusters have visibility to the same hosts and backend storage.

Attention: If any of these conditions are true, then you must perform the following special procedures. Failure to perform the special procedure is likely to result in the corruption of all data managed by the cluster.

Perform the following steps to add a node to a cluster:

1. Click **Work with Nodes** from the portfolio.
2. Click **Nodes** from the portfolio. The Viewing Nodes panel is displayed.
3. Select **Add a Node** from the list and click **Go**. The Adding nodes to a cluster panel is displayed.

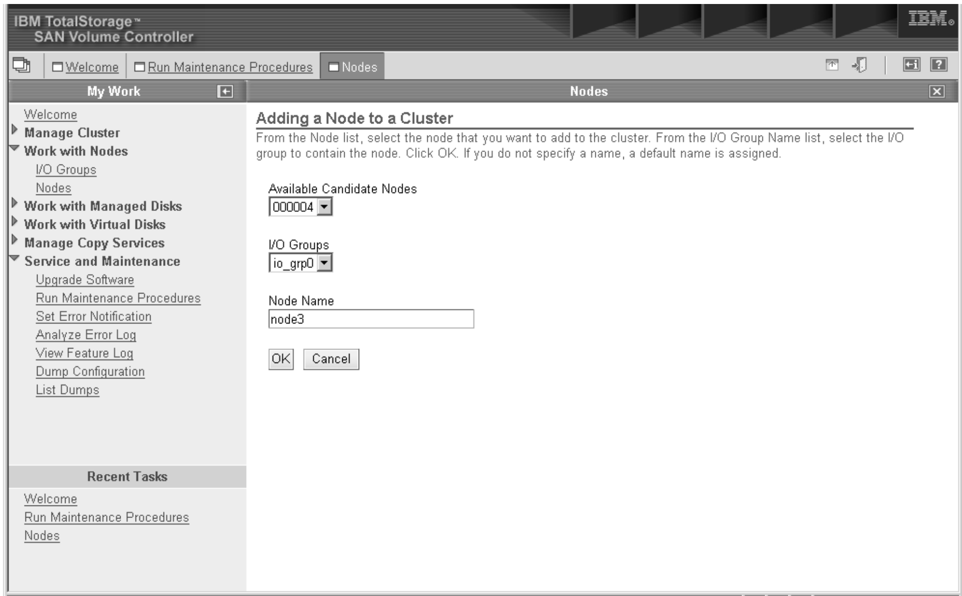

Figure 39. SAN Volume Controller Console Nodes panel

4. From the list of candidate nodes, select the number of the node that you want to add.
5. Select the I/O group for the node.

Special procedures when adding a node to a cluster

If any of the previous conditions are true, then the following special procedures apply. These special procedures apply when you use either the **svctask addnode** command or the SAN Volume Controller Console. When a node is added to a cluster then either:

- The node must be added back to the same I/O group that it was previously in.

  Note: The WWNN of the nodes in the cluster can be determined using the command: `svcinfo lsnode`

  or, if this information is not available, then
- *Before* the node is added back into the cluster all the hosts using the cluster must be shut down. The node must then be added before the hosts are rebooted. Or, if the I/O group information is not available and it is inconvenient to shutdown and reboot all of the hosts using the cluster, then
- On all the hosts connected to the cluster, unconfigure the Fibre Channel adapter device driver, the disk device driver, and the SDD device driver, before you add the node to the cluster.

  Re-configure the Fibre Channel adapter device driver, the disk device driver, and the SDD device driver, after adding the node into the cluster.

  Note: This may not be possible on all operating systems in all circumstances.

Background

Applications on host systems direct I/O operations to filesystems or logical volumes which are mapped by the operating system to vpaths which are pseudo disk objects supported by the SDD driver.

The SDD driver maintains an association between a vpath and a SAN Volume Controller VDisk. This association uses an identifier (UID) which is unique to the VDisk and is never re-used. This allows the SDD driver to unambiguously associate vpaths with VDisks.

The SDD device driver operates within a protocol stack which also contains Disk and Fibre Channel device drivers which allow it to communicate with the SAN Volume Controller using the SCSI protocol over Fibre Channel as defined by the ANSI FCS standard. The addressing scheme provided by these SCSI and Fibre channel device drivers uses a combination of a SCSI Logical unit number (LUN) and the World Wide Name for the Fibre Channel Node and Ports.

In the event of errors occurring, error recovery procedures (ERPs) operate at various tiers in the protocol stack. Some of these ERPs cause I/O to be redriven using the same WWN and LUN numbers which were previously used.

The SDD device driver does not check the association of the VDisk with the VPath on every I/O that it performs.

Perform the following steps to add nodes to the cluster:

1. From the Welcome panel, click **Work with Nodes** in the portfolio.
2. Click **Nodes** in the portfolio. The Nodes panel is displayed.

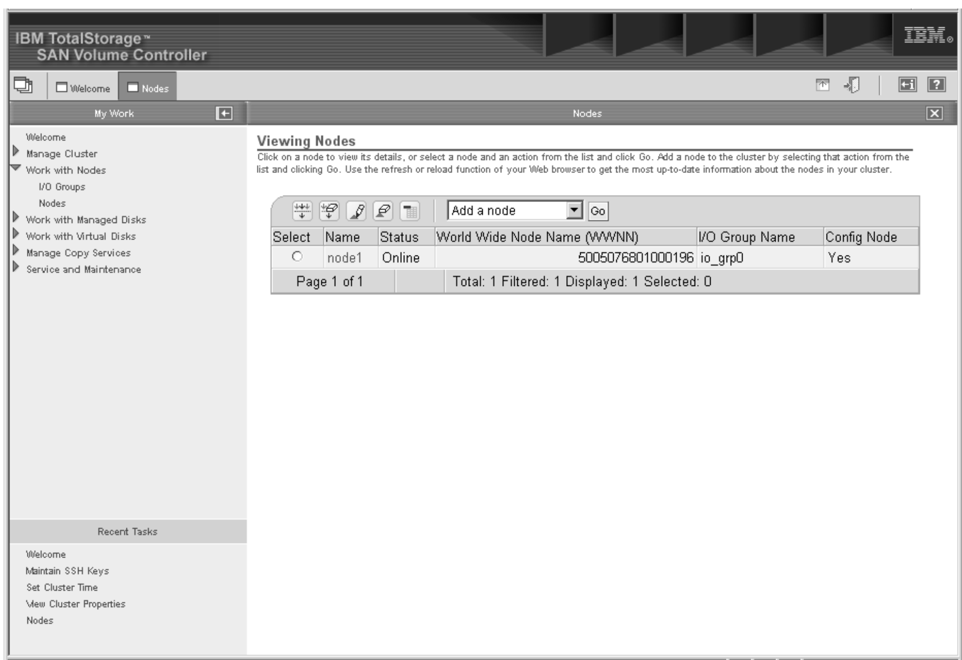

Figure 40. Nodes panel

3. Select **Add Node** from the drop down list and click **Go**.

Figure 41. Add node drop down list

4. 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:
   a. Card **Node** serial number
   a. All WWPNs
   b. 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 SAN Volume Controller Console panel when adding the node.

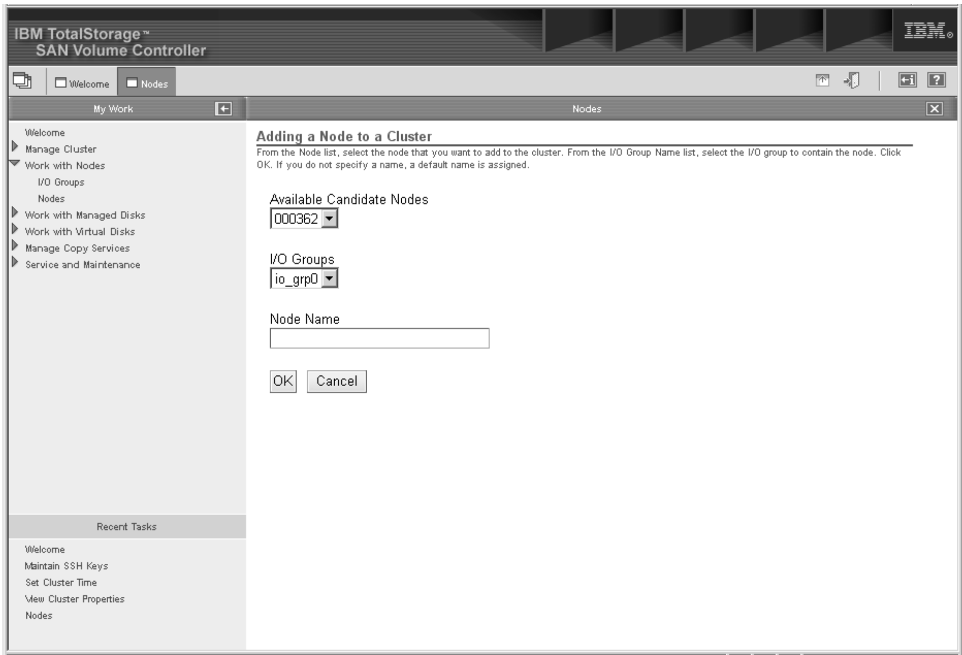

Figure 42. Add Node to Cluster panel

Note: Each node in an I/O group must be connected to a different uninterruptible power supply. If you do not supply a name, the cluster assigns a default name to the object. Whenever possible you should provide a meaningful name for objects to make identifying that object easier in the future.

Hypothetical scenarios where the special procedures may apply.

The following are two hypothetical scenarios where the special procedures may apply:

- Four nodes of an eight-node cluster have been lost because of the failure of a pair of UPS 5115 or four UPS 5125 power supplies. In this case, the four nodes must be added back into the cluster using the **svctask addnode** command or the SAN Volume Controller Console.
- A user decides to delete four nodes from the cluster and add them back into the cluster using the **svctask addnode** command or the SAN Volume Controller Console.

Related concepts

"Clusters" on page 12
All configuration and service is performed at the cluster level.

"Nodes and clusters" on page 12
A SAN Volume Controller node is a single processing unit, which provides virtualization, cache, and copy services for the SAN.

Related information

"Scenario: typical usage for the SAN Volume Controller Console" on page 100
This hypothetical example describes a scenario for configuring the SAN Volume Controller using the SAN Volume Controller Console.

Displaying node properties using the SAN Volume Controller Console

You can display the node properties using the Viewing general details panel.

Perform the following steps to display the node properties:

1. Click **Work with Nodes** from the portfolio.
2. Click **Nodes** from the portfolio. The Nodes panel is displayed.
3. Select the name of the node that you want to view the details for. The Viewing General Details panel is displayed.

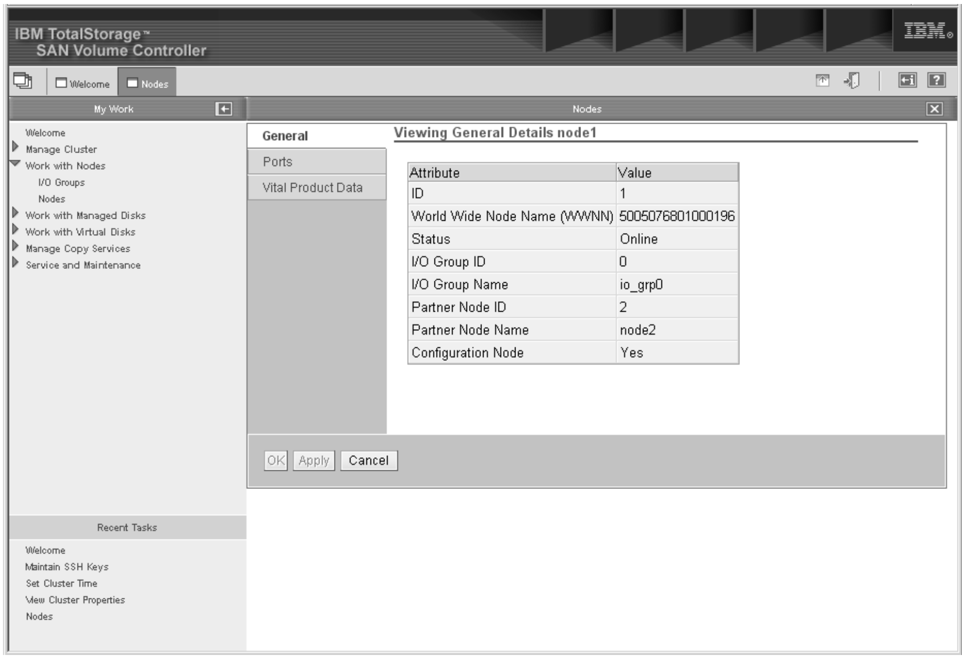

Figure 43. Viewing general details panel

4. Click **Ports** to view the WWPN port details. The Viewing Port Details panel for node1 is displayed.
5. Click **Vital Product Data** to view the node hardware details. The Viewing Vital Product Data panel is displayed.

Related information

"Scenario: typical usage for the SAN Volume Controller Console" on page 100
This hypothetical example describes a scenario for configuring the SAN Volume Controller using the SAN Volume Controller Console.

Creating managed disk groups

You can create a new managed disk (MDisk) group using the Create a Managed Disk Group wizard.

If you intend to keep the virtual disk allocation within one disk controller system, you should ensure that the MDisk group that corresponds with a single disk controller system is presented by that disk controller system. This also enables non-disruptive migration of data from one disk controller system to another disk controller system and simplifies the decommissioning process should you wish to decommission a disk controller system at a later time.

You should also 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 disk controller system 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 will lose access to all the data striped across the group. Similarly, for performance reasons you should not mix RAID types.

Perform the following steps to create a new MDisk group:
1. Click **Work with Managed Disks** from the portfolio.
2. Click **Managed Disk Groups** from the portfolio. The Filtering Managed Disk Groups panel is displayed.

   Note: The filter panels can be used to prefilter the list of objects that are displayed. This reduces the number of objects returned to the . This can be useful when you have a large number of objects (for example 4096 MDisks or 1024 VDisks) and you do not want to display them all. You can bypass the filtering and display all objects by clicking **Bypass Filter**.
3. Specify the filter criteria that you want to use. Click **OK** or click **Bypass Filter** to display all objects of this type. The Viewing Managed Disk Groups panel is displayed.
4. Select **Create MDisk Group** from the list. Click **Go**. The Create a Managed Disk Group wizard is displayed.
5. Type the name of the MDisk group and add the MDisks from the **Managed Disk Candidates** list.

In our hypothetical scenario, type

```
maindiskgroup
```

add MDisks

```
mdsk0, mdsk1, mdsk2, mdsk3
```

from the **Managed Disk Candidates** list.

6. Select the extent size from the list.

In our hypothetical scenario, select

```
32
```

for the extent size used within this MDisk group and click **OK**.

7. Repeat steps 4 through 6 for all of the MDisk groups that you want to create.

In our hypothetical scenario, repeat steps 4 through 6, in which the second MDisk group is named

```
bkpdiskgroup
```

with the following MDisks attached,

```
mdsk4, mdsk5, mdsk6, mdsk7
```

The extent size will be

```
16
```

MB.

Related concepts

"Managed disk groups" on page 26
An *MDisk group* is a collection of MDisks that jointly contain all the data for a specified set of virtual disks (VDisks).

"Managed disks" on page 23
A managed disk (MDisk) is a logical disk (typically a RAID array or partition thereof) that a storage subsystem has exported to the SAN fabric to which the nodes in the cluster are attached.

Related reference

"Configuration guidelines" on page 234
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

"Optimal managed disk group configurations" on page 236
A managed disk group provides the pool of storage from which virtual disks will be created. It is therefore necessary to ensure that the entire pool of storage provides the same performance and reliability characteristics.

Related information

"Scenario: typical usage for the SAN Volume Controller Console" on page 100
This hypothetical example describes a scenario for configuring the SAN Volume Controller using the SAN Volume Controller Console.

Using image mode virtual disks

Ensure that you are familiar with using image mode virtual disks.

An image mode VDisk provides a direct block-for-block translation from the managed disk to the virtual disk with no virtualization. This mode is intended to allow virtualization of managed disks which already contain data which was written directly, not through a SAN Volume Controller. Image mode virtual disks have a minimum size of 1 block (512 bytes) and always occupy at least one extent.

Image mode managed disks are members of a managed disk group, but, they do not contribute to free extents.

In order for a virtual disk to be online it is necessary for ALL managed disks in the managed disk group associated with the virtual disk to be online. This applies to image mode virtual disks as well as managed mode virtual disks. A virtual disk will be offline if any managed disk in the managed disk group is offline even if that managed disk does not contribute any extents to the virtual qisk in question or even if the managed disk has no allocated extents.

An image mode VDisk behaves just as a managed mode VDisk in terms of Metro Mirror and FlashCopy. Image mode VDisks are different from managed mode in two ways:

- Migration. An image mode disk can be migrated to another image mode disk. It becomes managed while the migration is ongoing, but returns to image mode when the migration is complete.
- Quorum disks. Image mode disks cannot be quorum disks. This means that a cluster with only image mode disks does not have a quorum disk.

Related tasks

"Exposing logical units on your existing storage to the cluster via switch zoning"
You can expose logical units on your existing storage to the cluster via switch zoning.

"Virtual disk-to-host mappings" on page 117
You can view the virtual disk-to-host mappings from the Virtual Disk-to-Host Mappings panel.

Related information

"Scenario: typical usage for the SAN Volume Controller Console" on page 100
This hypothetical example describes a scenario for configuring the SAN Volume Controller using the SAN Volume Controller Console.

Exposing logical units on your existing storage to the cluster via switch zoning

You can expose logical units on your existing storage to the cluster via switch zoning.

Ensure that you have already created a cluster.

CAUTION:
Ensure that logical units are not zoned or mapped so that they can be detected by the SAN Volume Controller and a host at the same time. Ensure that all of the logical units can be detected by all of the SAN Volume Controller nodes in the cluster on both storage area networks in a redundant fabric.

Related tasks

"Adding a new storage controller to a running configuration using the SAN Volume Controller Console" on page 247
You can add a new storage controller to your SAN at any time.

Related information

"Using image mode virtual disks" on page 109
Ensure that you are familiar with using image mode virtual disks.

Creating an image mode virtual disk

You can import storage that contains existing data and continue to use this storage but make use of the advanced functions, such as Copy Services, data migration, and the cache. These disks are known as image mode virtual disks (VDisks).

Make sure you are aware of the following before you create image mode VDisks:

1. Unmanaged-mode managed disks (MDisks) that contain existing data cannot be differentiated from unmanaged-mode MDisks that are blank. Therefore, it is vital that you control the introduction of these disks to the cluster. It is recommended that you introduce these disks one at a time. For example, map a single logical unit from your RAID controller to the cluster and refresh the view of managed disks. The newly detected disk is displayed.
2. *Do not* add an unmanaged-mode MDisk that contains existing data to a MDisk group manually. If you do, the data will be lost. When you use the command to convert an image mode virtual disk from an unmanaged-mode disk, you will select the MDisk group where it should be added.

Go to the following Web site for more information:

http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Perform the following steps to create an image mode virtual disk:

1. Stop all I/O operations from the hosts. Un-map the logical disks that contain the data from the hosts.
2. Create one or more managed disk (MDisk) groups. In the SAN Volume Controller Console portfolio, click **Work with Managed Disks –> Managed Disks**.
3. The Viewing Managed Disk Groups panel displays. In the action list, choose **Create an MDisk Group**. The Create Managed Disk Group wizard appears. Use the wizard to great the MDisk group, and ensure that the MDisk groups have enough free capacity to contain all of the migrating data.
4. Map a single RAID array or logical unit from your RAID controller to the cluster. You can do this through a switch zoning or a RAID controller based on your host mappings. The array or logical unit will appear as an unmanaged-mode MDisk to the SAN Volume Controller.
5. Refresh the list of MDisks from the SAN Volume Controller Console. Click **Work with Managed Disks –> Managed Disks**. You can then use the filter to view the unmanaged-mode MDisks.

   If the new unmanaged-mode MDisk is not listed, you can perform a fabric-level discovery. Select **Work with Managed Disks**, and choose **Discover MDisks** in the task list. When this process is complete, refresh the list of MDisks, and the unmanaged-mode MDisk should appear in the list.
6. Convert the unmanaged-mode MDisk to an image mode virtual disk. Click **Work with Managed Disks –> Managed Disks** in the portfolio.
7. The Filtering Managed Disks panel appears, and you can specify filter criteria. Click **OK**. The Viewing Managed Disks panel appears. Select the unmanaged-mode MDisk and choose **Create VDisk in Image Mode** in the task list. Click **Go**.. The create image mode virtual disk wizard appears. Use the wizard to select the MDisk group where the image mode disk should be added and the I/O group that will provide the data path for the virtual disk.
8. Map the new virtual disk to the hosts that were previously using the data that the MDisk now contains. In the SAN Volume Controller Console, select **Work with Virtual Disks --> Virtual Disks**. The Filtering Virtual Disks (VDisks) panel appears, and you can enter filter criteria or click **Bypass filter**. The Viewing Virtual Disks panel appears.
9. Select the VDisk and then choose **Map a VDisk to a host** in the action list. Click **Go**.

Once it is mapped to a host object, the image mode virtual disk is detected as a disk drive with which the host can perform I/O operations.

If you want to virtualize the storage on an image mode VDisk, you can transform it into a striped virtual disk. Migrate the data on the image mode VDisk to managed-mode disks in another MDisk group.

This procedure can only be performed using the command-line interface (CLI). Use the **svctask migratevdisk** command to migrate an entire image mode VDisk from one MDisk group to another MDisk group. For instructions on this command, refer to the IBM TotalStorage SAN Volume Controller: Command-Line Interface User's Guide.

Migration methods

Several methods can be used to migrate image mode VDisks into managed mode VDisks.

In order to perform any type of migration activity on an image mode VDisk, the image mode VDisk must first be converted into a managed mode disk. This happens automatically whenever any kind of migration activity is attempted. Once this image to managed migration operation has occurred, the VDisk becomes a managed mode virtual disk and is treated the same way as any other managed mode VDisk.

If the image mode disk has a partial last extent, then this last extent in the image mode VDisk must be the first to be migrated. This migration is handled as a special case. Once this special migration operation has occurred, the VDisk becomes a managed mode VDisk and is treated in the same way as any other managed mode VDisk. If the image mode disk does not have a partial last extent, no special processing is performed. The image mode VDisk is simply changed into a managed mode VDisk and is treated in the same way as any other managed mode VDisk.

An image mode disk can also be migrated to another image mode disk. It become managed while the migration is ongoing, but returns to image mode when the migration is complete.

The types of migrations that you can perform are the following:
- migrate extents
- migrate a VDisk
- migrate to image mode

The recommended strategy is:
- Dedicate one MDisk group to image mode VDisks.
- Dedicate one MDisk group to managed mode VDisks.
- Use the migrate VDisk function to move the VDisks.

Remember: Ensure that there are enough extents available in the target MDisk group. Once the image mode VDisk has been migrated, you cannot migrate it back.

Related tasks

"Migration to image mode"
You can migrate a virtual disk to a image mode virtual disk. This can be combined with the ability to migrate between managed disk groups. The source for the migrate can be a managed mode or an image mode virtual disk.

Migration to image mode

You can migrate a virtual disk to a image mode virtual disk. This can be combined with the ability to migrate between managed disk groups. The source for the migrate can be a managed mode or an image mode virtual disk.

The following four migration possibilities are now available:
- Migrate image mode to image mode within a managed disk group
- Migrate managed mode to image mode within a managed disk group
- Migrate image mode to image mode between managed disk groups
- Migrate managed mode to image mode between managed disk groups

You can perform these migrations by issuing the **svctask migratetoimage** command.

Note:

- The destination MDisk must be greater than or equal to the size of the VDisk.
- Regardless of the mode that the VDisk starts in, it is reported as managed mode during the migration.
- Both of the MDisks involved are reported as being image mode during the migration.

Related tasks

"Migration methods" on page 111
Several methods can be used to migrate image mode VDisks into managed mode VDisks.

Importing data by creating an image mode VDisk

You can migrate logical units into virtual disks (VDisks).

Ensure that you define which managed disk group you want the image disk to belong too.

CAUTION:
If you import the data onto a SAN Volume Controller in image mode without first creating a quorum disk then you are vulnerable to a loss of access to data should nodes drop out of the cluster.

Before importing disks in image mode, ensure that you create space to use for your quorum disks. Create three MDisks which are NOT to be used as image mode VDisks. Do not have any data on them. The three MDisks should be placed into a MDisk group. The three MDisks should have a size of at least 512 MB each and be allocated from one of the controller types which supports quorum disks.

An image mode VDisk provides a direct block-for-block translation from the managed disk to the virtual disk with no virtualization. This mode is intended to allow virtualization of managed disks which already contain data which was written directly, not through a SAN Volume Controller. Image mode virtual disks have a minimum size of 1 block (512 bytes) and always occupy at least one extent.

Many of the facilities of SAN Volume Controller are available for use on image mode VDisks. You can decide to leave your VDisks as image mode VDisks. In order to expand or reduce the size of the VDisks or if the VDisks need to be migrated, then they will need to be converted into managed mode. This will happen automatically when the VDisks are migrated. Specifically, it can be one in one of the following ways:

- Using the **svctask migratevdisk** command which requires another MDisk group to already exist. This also requires that the other MDisk group has the same extent size and that the MDisk group being migrated to has adequate space, or
- Using the **svctask migrateexts** or the **svctask rmmdisk -force** command which will force a migration of any VDisks occupying space on the MDisk being deleted.

  Note:
  - The **svctask rmmdisk -force** command will only work if there are enough free extents in the group.
  - Image mode MDisks do *not* contribute to the free extent pool because all their extents are used for the corresponding image mode VDisk.
  - Managed MDisks do contribute to free extents, however, it is possible that there won't be enough. You can issue the **svctask addmdisk** command to add more managed MDisks to the group.

– Similarly, if you want to migrate your image mode VDisk to another group, using the **svctask migratevdisk** command, ensure there are enough free extents in the target group.

Exporting data by deleting an image mode virtual disk

The **svctask rmvdisk** will delete an existing image mode virtual disk.

If the command succeeds (without the force flag) for an image mode disk, then the underlying back-end controller logical unit will be consistent with the data that a host could previously have read from the image mode virtual disk. That is, all write data will have been flushed to the underlying LUN. If the force flag is used, then this guarantee does not hold.

If there are any virtual medium errors on the virtual disk then the command will fail. The force flag overrides this behavior, allowing a virtual disk containing medium errors to be deleted, at the expense of creating a potential data integrity issue.

If any IBM TotalStorage Metro Mirror for SAN Volume Controller, FlashCopy or host mappings still exist for this virtual disk then the delete will be failed unless the force flag is specified. When the force flag is specified, any mappings that remain will be deleted and then the virtual disk will be deleted. If there are any undestaged data in the write cache for this virtual disk, then the deletion of the virtual disk will fail. When the force flag is specified, any undestaged data in the write cache will be deleted.

Deleting an image mode virtual disk causes the managed disk associated with the virtual disk to be ejected from the managed disk group. The mode of the managed disk will be returned to unmanaged.

Creating virtual disks

You can create virtual disks using the Create virtual disks wizard.

Perform the following steps to create virtual disks:

1. Click **Work with Virtual Disks** from the portfolio.
2. Click **Virtual Disks** from the portfolio. The Filtering Virtual Disks panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Viewing Virtual Disks panel is displayed.
4. Select **Create virtual disks** from the list. Click **Go**. The Create Virtual Disks wizard is displayed.
5. Perform the following steps to complete the wizard:
   a. Select an I/O group, preferred node, and managed disk group.

      Note: If there are no MDisk groups to select, you will be prompted to create one.
   b. Select the type and quantity of the virtual disks that you want to create.
   c. Type a name for the virtual disks.
   d. Set the attributes, such as MDisk candidates, capacity of the VDisks, and type of VDisks, for the virtual disks.
   e. Verify the attributes.

   Related information

“Scenario: typical usage for the SAN Volume Controller Console” on page 100
This hypothetical example describes a scenario for configuring the SAN Volume Controller using the SAN Volume Controller Console.

Migrating virtual disks

You can migrate a virtual disk (VDisk) from one MDisk group to another from the Migrating VDisks panel.

The SAN Volume Controller provides various data migration features. These can be used to move the placement of data both within MDisk groups and between MDisk groups. These features can be used concurrent with I/O operations. There are two ways in which you can migrate data:

1. Migrating data (extents) from one MDisk to another (within the same MDisk group). This can be used to remove hot or overutilized MDisks. This can only be performed using the CLI.
2. Migrating VDisks from one MDisk group to another. This can be used to remove hot MDisk groups, for example, you can reduce the utilization of a group of MDisks.

You can determine the usage of particular MDisks by gathering I/O statistics about MDisks and VDisks. Once you have gathered this data, you can analyze it to determine which VDisks or MDisks are hot. This procedure then takes you through migrating VDisks from one MDisk group to another.

When a migrate command is issued, a check is made to ensure that the destination of the migrate has enough free extents to satisfy the command. If it does, the command proceeds, but will take some time to complete. During this time, it is possible for the free destination extents to be consumed by another process, for example, by creating a new VDisk in the destination MDisk group or by starting more migrate commands. In this scenario, when all the destination extents have been allocated the migration commands suspend and an error is logged (error id 020005). There are two methods for recovering from this situation:

1. Add additional MDisks to the target MDisk group. This will provide additional extents in the group and will allow the migrations to be restarted (by marking the error as fixed).
2. Migrate one or more VDisks that are already created from the MDisk group to another group. This will free up extents in the group and allow the original migrations to be restarted (again by marking the error as fixed).

Perform the following steps to migrate VDisks between MDisk groups:

1. Isolate any VDisks that are overutilized. You can determine this by requesting an I/O statistics dump and analyzing the output. To start I/O statistics gathering, select **Manage Cluster** from the portfolio, then select the **Start statistics collection** task. Enter `15` minutes for the interval and click **OK**. This will generate a new I/O statistics dump file approximately every 15 minutes. Wait for at least 15 minutes before proceeding to the next step.
2. Select **Service and Maintenance** from the portfolio and then select the **List dumps** task.
3. Click on the **I/O Statistics logs** link in the panel displayed. This will list the I/O statistics files that have been generated. These are prefixed with `m` and `Nm` for MDisk statistics and `v` for VDisk statistics. Click on one of the filenames to view the contents.

4. Analyze the dumps to determine which VDisks are hot. It may be helpful to also determine which MDisks are being heavily utilized as you can spread the data they contain more evenly across all the MDisks in the group.
5. Stop the statistics collection again by selecting **Manage Cluster** from the portfolio and then select **Stop statistics collection** task. Once you have analyzed the I/O statistics data, you can determine which VDisks are hot. You also need to determine which MDisk group you wish to move this VDisk to. Either create a new MDisk group or determine an existing group that is not yet over utilized. You can do this by checking the I/O statistics files generated above and ensuring that the MDisks or VDisks in the target MDisk group are less utilized than the source group.
   a. Click **Work with Virtual Disks** from the portfolio.
   b. Click **Virtual Disks** from the portfolio. The Filtering Virtual Disks panel is displayed.
   c. Specify the filter criteria that you want to use. Click **OK**. The Virtual Disks panel is displayed.
   d. Select the VDisk you want to migrate and select **Migrate a VDisk** from the list. Click **Go**. The Migrating VDisks panel is displayed.

   Related concepts

   "Virtual disks" on page 29
   A *virtual disk (VDisk)* is a logical disk that the cluster presents to the storage area network (SAN).

Shrinking virtual disks

You can shrink a virtual disk (VDisk) from the Shrinking VDisks panel.

VDisks can be reduced in size should it be required. However, if the VDisk contains data that is being used, **under no circumstances should you attempt to shrink a VDisk without first backing up your data**. The SAN Volume Controller arbitrarily reduces the capacity of the VDisk by removing a partial, one or more extents from those allocated to the VDisk. You cannot control which extents are removed and so you cannot guarantee that it is unused space that is removed.

| **Attention:** This feature should *only* be used to make a target or auxiliary VDisk the same size as the source or master VDisk when creating FlashCopy mappings
| or Metro Mirror relationships. You should also ensure that the target VDisk is not mapped to any hosts prior to performing this operation.

Perform the following steps to shrink a VDisk:

1. Validate that the VDisk is not mapped to any host objects. If the VDisk is mapped, data is displayed.
2. You can determine the exact capacity of the source or master VDisk. Issue the following command:

   ```
   svcinfo lsvdisk -bytes <vdiskname>
   ```

   Note: It is not possible to determine the exact size using the SAN Volume Controller Console.
3. Click **Work with Virtual Disks** from the portfolio.
4. Click **Virtual Disks** from the portfolio. The Filtering Virtual Disks panel is displayed.
5. Specify the filter criteria that you want to use. Click **OK**. The Virtual Disks panel is displayed.

6. Select the VDisk you want to shrink and select **Shrink a VDisk** from the list. Click **Go**. The Shrinking VDisks panel is displayed.

Related concepts

"Virtual disks" on page 29
A *virtual disk (VDisk)* is a logical disk that the cluster presents to the storage area network (SAN).

Showing VDisks mapped to a host

You can show the VDisks that are mapped to a host by using the Viewing Virtual Disks panel.

If a number of new VDisks are mapped to a host, and a number of devices are already running I/O operations, then a lot of errors may be logged. At the time the new VDisk is mapped, multiple recoverable errors can be logged in the event log. Decoding of the event log shows the errors to be caused by a check condition. The error states that there has been a change to the device information since the last LUN operation.

Perform the following steps to show the VDisks that are mapped to a host:

1. Click **Work with Virtual Disks** in the portfolio.
2. Click **Hosts** in the portfolio. The Filtering Hosts panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Hosts panel is displayed.
4. Select the host and select **Show the VDisks Mapped to this Host** from the list. Click **Go**.

The virtual disks mapped to this host are displayed in the Viewing Virtual Disks panel

Related concepts

"Host objects" on page 31
A host system is an open-systems computer that is connected to the switch through a fibre-channel interface.

Related information

"Scenario: typical usage for the SAN Volume Controller Console" on page 100
This hypothetical example describes a scenario for configuring the SAN Volume Controller using the SAN Volume Controller Console.

Virtual disk-to-host mappings

You can view the virtual disk-to-host mappings from the Virtual Disk-to-Host Mappings panel.

Perform the following steps to view your virtual disk-to-host mappings:

1. Click **Work With Virtual Disks** in the portfolio.
2. Click **Virtual Disk-to-Host Mappings** in the portfolio. The Filtering Virtual Disk-to-Host Mappings panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Virtual Disk-to-Host Mappings panel is displayed.

Related concepts

"Virtual disk-to-host mapping" on page 32
Virtual disk-to-host mapping is the process of controlling which hosts have access to specific virtual disks (VDisks) within the SAN Volume Controller.

Related information

"Using image mode virtual disks" on page 109
Ensure that you are familiar with using image mode virtual disks.

Deleting a managed mode virtual disk

The **svctask rmvdisk** command will delete an existing managed mode virtual disk.

Any data on the virtual disk will be deleted.

The extents that make up the virtual disk will be returned to the pool of free extents that are available on the managed disk group. If any IBM TotalStorage Metro Mirror for SAN Volume Controller, FlashCopy or host mappings still exist for this virtual disk then the delete will fail unless the force flag is specified. When the force flag is specified, any mappings that remain will be deleted and then the virtual disk will be deleted.

If the virtual disk is currently the subject of a migrate to image mode, then the delete will fail unless the force flag is specified. When the force flag is specified, the migration will be halted and the VDisk deleted. Care should be taken when executing this command to ensure that the virtual disk (and any data that resides on it) is no longer required.

Creating hosts

You can create a new host object from the Creating Hosts panel.

Perform the following steps to create a new host object:
1. Click **Work with Virtual Disks** in the portfolio.
2. Click **Hosts** in the portfolio. The Filtering Hosts panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Hosts panel is displayed.
4. Select **Create Host** from the list. Click **Go**. The Creating Hosts panel is displayed.
5. Type the name of the logical host object. If you do not specify a name, a default name is assigned (for example, host0). Then assign a worldwide port name (WWPN). A WWPN consists of 16 hexadecimal digits (for example, 210100e08b251dd4). You can select a WWPN from the list of candidates, or you can enter a WWPN that is not in the list. You can assign one or more WWPN's to a single logical host object. Click **OK**.

In our hypothetical scenario, because a host name was not specified the default name is:

```
host0
```

The WWPNs assigned are:

```
210100e08b251dd4, 210100e08b251dd5
```

The WWPNs can be found by using the management application for your switch.

6. Repeat steps 4 through 5 for each host object that you want to create.

In our hypothetical scenario, repeat steps 4 on page 118 through 5 on page 118 and name the host:

```
demohost2
```

The WWPNs assigned to the host are:

```
210100e08b251dd6, 210100e08b251dd7
```

Related concepts

"Host objects" on page 31
A host system is an open-systems computer that is connected to the switch through a fibre-channel interface.

Related information

"Scenario: typical usage for the SAN Volume Controller Console" on page 100
This hypothetical example describes a scenario for configuring the SAN Volume Controller using the SAN Volume Controller Console.

Creating consistency groups

You can create a FlashCopy consistency group from the Creating FlashCopy Consistency Groups panel.

Perform the following steps to create a FlashCopy consistency group:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy Consistency Groups** in the portfolio. The Filtering FlashCopy Consistency Groups panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The FlashCopy Consistency Groups panel is displayed.
4. Click **Create a Consistency Group**. The Create FlashCopy Consistency Groups panel is displayed.
5. Type the name of the consistency group in the **FlashCopy Consistency Group Name** field. From the **FlashCopy Mappings** list, select the mappings you want in the consistency group and click **OK**. If you do not specify a name, a default name is assigned.

In our hypothetical scenario, the name of the consistency group is:

```
maintobkpfcopy
```

The mappings that should be added are:

```
main1copy, main2copy
```

Note: You could have created the FlashCopy consistency group before you created the mappings and then added the FlashCopy mappings to the consistency group. To add FlashCopy mappings in this way, you must use the Modifying FlashCopy Mapping panel or the Creating FlashCopy Mappings panel.

Related concepts

"FlashCopy consistency groups" on page 41
A consistency group is a container for mappings. You can add many mappings to a consistency group.

Related information

"Scenario: typical usage for the SAN Volume Controller Console" on page 100
This hypothetical example describes a scenario for configuring the SAN Volume Controller using the SAN Volume Controller Console.

Creating FlashCopy mappings

You can create a FlashCopy mapping from the Creating FlashCopy Mappings panel.

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy mappings** in the portfolio. The Filtering FlashCopy Mappings panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Viewing FlashCopy Mappings panel is displayed.
4. Select **Create Mapping** from the list. Click **Go**. The Creating FlashCopy Mappings panel is displayed.
5. Type the name of the new FlashCopy mapping. In our hypothetical scenario, the name of the FlashCopy mapping is: `maincopy`.
6. Select the source VDisk from the list. In our hypothetical scenario, the name of the source VDisk is: `maindisk1`
7. Select the target VDisk from the list. In our hypothetical scenario, the name of the target VDisk is: `bkpdisk1`
8. Select the priority for the background copy. Click **OK**.

Repeat steps 4 through 8 for each FlashCopy mapping that you want to create.

Related concepts

"FlashCopy mappings" on page 37
A FlashCopy mapping defines the relationship between a source virtual disk (VDisk) and a target virtual disk (VDisk).

Related reference

"Considerations for FlashCopy mappings" on page 237
Ensure that you've considered the type of I/O and frequency of update before creating the virtual disks that you wish to use in FlashCopy mappings.

Related information

"Scenario: typical usage for the SAN Volume Controller Console" on page 100
This hypothetical example describes a scenario for configuring the SAN Volume Controller using the SAN Volume Controller Console.

Advanced function FlashCopy overview

Ensure that you are familiar with the FlashCopy advanced functions.

Overview

The following sections details the advanced FlashCopy functions that you can perform using the SAN Volume Controller Console.

Related tasks

"Starting FlashCopy mappings" on page 121
You can start FlashCopy mappings from the FlashCopy Mappings panel.

"Stopping FlashCopy mappings" on page 121
You can stop FlashCopy mappings from the FlashCopy Mappings panel.

"Deleting FlashCopy mappings" on page 122
You can delete FlashCopy mappings from the FlashCopy mappings panel.

"Starting FlashCopy consistency groups" on page 122
You can start FlashCopy consistency groups from the FlashCopy Consistency Groups panel.

“Stopping FlashCopy consistency groups” on page 122
You can stop FlashCopy consistency groups from the FlashCopy Consistency Groups panel.

“Deleting FlashCopy consistency groups” on page 123
You can delete FlashCopy consistency groups using the FlashCopy consistency groups panel.

Related reference

“Valid combinations of FlashCopy and Metro Mirror functions” on page 366
The following table outlines the combinations of FlashCopy and Metro Mirror functions that are valid for a single virtual disk (VDisk).

Starting FlashCopy mappings

You can start FlashCopy mappings from the FlashCopy Mappings panel.

You can start or trigger a FlashCopy mapping from the Starting FlashCopy Mappings panel.

Perform the following steps to start a FlashCopy mapping:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy Mappings** in the portfolio. The Filtering FlashCopy mappings panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Viewing FlashCopy Mappings panel is displayed.
4. Select the appropriate mapping’s row from the table.
5. Select **Start a Mapping** from the drop down list. Click **Go**. The Starting FlashCopy mappings panel is displayed.

Related concepts

“FlashCopy mappings” on page 37
A FlashCopy mapping defines the relationship between a source virtual disk (VDisk) and a target virtual disk (VDisk).

Stopping FlashCopy mappings

You can stop FlashCopy mappings from the FlashCopy Mappings panel.

You can stop a FlashCopy mapping from the Stopping FlashCopy Mappings panel.

Perform the following steps to stop a FlashCopy mapping:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy Mappings** in the portfolio. The Filtering FlashCopy Mappings panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Viewing Mappings panel is displayed.
4. Select the appropriate mapping’s row from the table.
5. Select **Stop a mapping** from the drop down list. Click **Go**. The Stopping FlashCopy mappings panel is displayed.

Related concepts

“FlashCopy mappings” on page 37
A FlashCopy mapping defines the relationship between a source virtual disk (VDisk) and a target virtual disk (VDisk).

Deleting FlashCopy mappings

You can delete FlashCopy mappings from the FlashCopy mappings panel.

You can delete a FlashCopy mapping from the Deleting FlashCopy Mappings panel.

Perform the following steps to delete a FlashCopy mapping:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy mappings** in the portfolio. The Filtering FlashCopy mappings panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Viewing FlashCopy mappings panel is displayed.
4. Select the appropriate mapping's row from the table.
5. Click **Delete a mapping** and click **Go** . The Deleting FlashCopy mapping panel is displayed.

Related concepts

"FlashCopy mappings" on page 37
A FlashCopy mapping defines the relationship between a source virtual disk (VDisk) and a target virtual disk (VDisk).

Starting FlashCopy consistency groups

You can start FlashCopy consistency groups from the FlashCopy Consistency Groups panel.

You can start or trigger a FlashCopy consistency group from the Starting FlashCopy Consistency Group panel.

Perform the following steps to start or trigger a FlashCopy consistency group:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy Consistency Groups** in the portfolio. The Filtering FlashCopy Consistency Groups panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The FlashCopy Consistency Groups panel is displayed.
4. Select the appropriate mapping's row from the table.
5. Click **Start Consistency Group**. The Starting FlashCopy Consistency Groups panel is displayed.

Related concepts

"FlashCopy consistency groups" on page 41
A consistency group is a container for mappings. You can add many mappings to a consistency group.

Stopping FlashCopy consistency groups

You can stop FlashCopy consistency groups from the FlashCopy Consistency Groups panel.

You can stop a FlashCopy consistency group from the Stopping FlashCopy Consistency Group panel.

Perform the following steps to stop a FlashCopy consistency group:

1. Click **Manage Copy Services** in the portfolio.

2. Click **FlashCopy Consistency Groups** in the portfolio. The Filtering FlashCopy Consistency Groups panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The FlashCopy Consistency Groups panel is displayed.
4. Select the appropriate mapping's row from the table.
5. Click **Stop Consistency Group**. The Stopping Consistency Groups panel is displayed.

Related concepts

"FlashCopy consistency groups" on page 41
A consistency group is a container for mappings. You can add many mappings to a consistency group.

Deleting FlashCopy consistency groups

You can delete FlashCopy consistency groups using the FlashCopy consistency groups panel.

You can delete a FlashCopy consistency group from the Deleting FlashCopy consistency groups panel.

Perform the following steps to delete a FlashCopy consistency groups:

1. Click **Manage Copy Services** in the portfolio.
2. Click **FlashCopy consistency groups** in the portfolio. The Filtering FlashCopy consistency groups panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The FlashCopy consistency groups panel is displayed.
4. Select the appropriate mapping's row from the table.
5. Click **Delete Consistency Groups**. The Delete Consistency Groups panel is displayed.

Related concepts

"FlashCopy consistency groups" on page 41
A consistency group is a container for mappings. You can add many mappings to a consistency group.

"FlashCopy" on page 34
FlashCopy is a copy service that is available with the SAN Volume Controller.

Advanced functions overview for the SAN Volume Controller Console

You can perform ″advanced″ functions using the SAN Volume Controller Console.

Related tasks

"Determining the WWPNs for a node using the SAN Volume Controller Console" on page 124
To determine the WWPNs of a node using the SAN Volume Controller Console, complete these steps.

"Determining the relationship between VDisks and MDisks using the SAN Volume Controller Console" on page 124
Ensure that you understand the relationship between VDisks and MDisks using the SAN Volume Controller Console.

"Determining the relationship between managed disks and RAID arrays or LUNs using the SAN Volume Controller Console" on page 125
Each MDisk corresponds with a single RAID array, or a single partition on a given RAID array. Each RAID controller will define a LUN number for this disk.

The LUN number and controller name or ID are needed to be able to determine the relationship between MDisks and RAID arrays or partitions.

"Increasing the size of your cluster using the SAN Volume Controller Console" on page 125
To increase the size of a cluster, complete these tasks.

"Replacing a faulty node with a spare node using the SAN Volume Controller Console" on page 127
To replace a faulty node in the cluster using the SAN Volume Controller Console, complete these tasks.

"Recovering from offline VDisks after a node or an I/O group failed" on page 131
If you have lost both nodes in an I/O group and have therefore, lost access to all the VDisks that are associated with the I/O group, then you must perform one of the following procedures to regain access to your VDisks. Depending on the failure type, you may have lost data that was cached for these VDisks, therefore, they have gone offline.

Determining the WWPNs for a node using the SAN Volume Controller Console

To determine the WWPNs of a node using the SAN Volume Controller Console, complete these steps.

Perform the following steps to determine a node's WWPN's:

1. List the nodes in the cluster by opening the **Work with Nodes** panel.
2. For the node or nodes in question, select the node name link to view the node details.
3. Select the ports tab and note each WWPN.

Related information

"Advanced functions overview for the SAN Volume Controller Console" on page 123
You can perform ″advanced″ functions using the SAN Volume Controller Console.

Determining the relationship between VDisks and MDisks using the SAN Volume Controller Console

Ensure that you understand the relationship between VDisks and MDisks using the SAN Volume Controller Console.

Perform the following steps to determine the relationship between VDisks and MDisks:

1. Click **Work with VDisks** from the portfolio.
2. Select the VDisk that you want to view the relationship between this VDisk and its MDisks.
3. Select the **Show MDisks** task. The Work with MDisks panel is displayed. This panel lists the MDisks and make up the selected VDisk. Perform the following steps to determine the relationship between MDisks and VDisks:
   a. Click **Work with MDisks** from the portfolio.
   b. Select the VDisk that you want to view the relationship between this VDisk and its MDisks .
   c. Select the **Show VDisks** task. The Work with VDisks panel is displayed. This panel lists the VDisks and make up the selected MDisk.

Related information

“Advanced functions overview for the SAN Volume Controller Console” on page 123
You can perform ″advanced″ functions using the SAN Volume Controller Console.

Determining the relationship between managed disks and RAID arrays or LUNs using the SAN Volume Controller Console

Each MDisk corresponds with a single RAID array, or a single partition on a given RAID array. Each RAID controller will define a LUN number for this disk. The LUN number and controller name or ID are needed to be able to determine the relationship between MDisks and RAID arrays or partitions.

Perform the following steps to determine the relationship between MDisks and RAID arrays:

1. Click **Work with MDisks** from the portfolio.
2. Select the MDisk to view the details. Write down the controller name and controller LUN number.
3. Click the **Work with Disk Controllers** panel.
4. In the filter screen, enter the controller name in the **Name** field. The panel displayed should show just the one controller.
5. Select the name to show the detailed view of the given controller. Write down the vendor ID and the product ID and WWNN and use these to determine the controller that is presented to the MDisk.
6. From the native user interface for the given controller, list the LUNs it is presenting and match the LUN number with that noted in 2. This will tell you the exact RAID array and partition that corresponds with the MDisk.

Related information

“Advanced functions overview for the SAN Volume Controller Console” on page 123
You can perform ″advanced″ functions using the SAN Volume Controller Console.

Increasing the size of your cluster using the SAN Volume Controller Console

To increase the size of a cluster, complete these tasks.

To increase the size of your cluster you need to add nodes in pairs to a new I/O group. Your existing cluster may have become a bottleneck and so you wish to increase throughput by adding more nodes to the cluster.

Perform the following steps to increase the size of your cluster:

1. Add a node to increase the size of your cluster and repeat this procedure for the second node.
2. If you wish to balance the load between the existing I/O groups and the new I/O groups, migrate your VDisks to new I/O groups. Repeat this procedure for all VDisks you want to assign to the new I/O group.

Related tasks

“Adding a node to increase the size of a cluster using the GUI” on page 126
You can add a node to increase the size of a cluster using the command-line interface.

"Migrating a VDisk to a new I/O group"
You can migrate a VDisk to a new I/O group to manually balance the workload across the nodes in the cluster. You may end up with a pair of nodes that are overworked and another pair that are under-worked. Follow this procedure to migrate a single VDisk to a new I/O group. Repeat for other VDisks as required.

Related information

"Advanced functions overview for the SAN Volume Controller Console" on page 123
You can perform ″advanced″ functions using the SAN Volume Controller Console.

Adding a node to increase the size of a cluster using the GUI

You can add a node to increase the size of a cluster using the command-line interface.

CAUTION:

If you are adding a node that was previously removed from a cluster, 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. If you do not know the I/O group name or ID that it was removed from, contact IBM Service to add the node to the cluster without corrupting data.

The stage needs to be set just so.

1. Click **Work with I/O groups** to determine the I/O group where the node will be added. Write down name or ID of the first I/O group that has a node count of zero (0). You will need the ID for the next step.

   Record the following information for future reference:
   - Node serial number.
   - Worldwide node name.
   - All of the worldwide port names.
   - The name or ID of the I/O group that contains the node.
2. Select the node from the list of available candidate nodes and select the I/O group from the list.
3. Optionally enter a node name for this node.
4. Verify that the node is online by refreshing the **Work with Nodes** panel. You may need to close the panel and open it again in order to refresh the panel.

If the disk controller uses mapping to present RAID arrays or partitions to the cluster and the worldwide node names or the worldwide port names have changed, then you must modify the port groups that belong to the cluster.

Migrating a VDisk to a new I/O group

You can migrate a VDisk to a new I/O group to manually balance the workload across the nodes in the cluster. You may end up with a pair of nodes that are overworked and another pair that are under-worked. Follow this procedure to migrate a single VDisk to a new I/O group. Repeat for other VDisks as required.

Attention: This is a disruptive procedure, access to the VDisk will be lost while you follow this procedure. Under no circumstances should VDisks be moved to an offline I/O group. You must ensure the I/O group is online before moving the VDisks to avoid data loss scenarios.

Perform the following steps to migrate a single VDisk:

1. Quiesce all I/O operations for the VDisk. You may need to determine the hosts that are using this VDisk.
2. Before migrating the VDisk, it is essential that for each vpath presented by the VDisk you intend to move, the SDD configuration is updated to remove the vpaths in question. Failure to do this may result in data corruption. See *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* for details about how to dynamically re-configure SDD for the given host operating system.
3. Any FlashCopy mappings or Metro Mirror relationships that use this VDisk should be stopped and deleted. To check if the VDisk is part of a mapping or relationship, perform the following steps:
   a. Click **Work with VDisks**..
   b. Click on the VDisk name to view the details.
   c. Look for the **FlashCopy ID** and **Metro Mirror ID** fields. If these are not blank then the VDisk is part of a mapping or relationship.
4. Migrate the VDisk by selecting the VDisk from the **Work with VDisks** panel and selecting the **Modify** task. Change only the I/O group to the new I/O group name.
5. It is now necessary to follow the SDD procedure to discover the new vpaths and to check that each vpath is now present with the correct number of paths. See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* for details on how to dynamically re-configure SDD for the given host operating system.

Related tasks

"Determining the relationship between VDisks and MDisks using the SAN Volume Controller Console" on page 124
Ensure that you understand the relationship between VDisks and MDisks using the SAN Volume Controller Console.

Related reference

"Advanced function Metro Mirror overview" on page 135
For detailed information about how to perform advanced FlashCopy and Metro Mirror functions, go to the following Web site: www.ibm.com/redbooks

Replacing a faulty node with a spare node using the SAN Volume Controller Console

To replace a faulty node in the cluster using the SAN Volume Controller Console, complete these tasks.

Before you attempt to replace a faulty node with a spare node you must ensure that:

- SAN Volume Controller version 1.1.1 or later is installed on the cluster and on the spare node. To determine the version, you can review the vital product data on the node and the cluster, or you can issue the **svcinfo lsnode** or **svcinfo lscluster** command.
- You know the name of the cluster that contains the faulty node.
- A spare node is installed in the same rack as the cluster that contains the faulty node.
- You make a record of the last five characters of the original worldwide node name (WWNN) of the spare node. You will need this information, if and when, you want to stop using this node as a spare node. In that case, you might prefer to use it as a normal node that can be assigned to any cluster.

Perform the following steps to display and record the WWNN of the spare node:

1. Display the node status on the front panel display of the node. See the *IBM TotalStorage SAN Volume Controller: Service Guide* for more information.
2. With the node status displayed on the front panel, press and hold the **Down** button; press and release the **Select** button; release the **Down** button. `WWNN` is displayed on line-1 of thez display; line-2 of the display contains the last five characters of the WWNN.
3. Record the WWNN in a safe place. It will be needed if you want to stop using the spare node.

If a node fails, the cluster continues to operate with degraded performance, until the faulty node is repaired. If the repair operation is likely to take an unacceptable amount of time, it might be useful to replace the faulty node with a spare node. However, the appropriate procedures must be followed and precautions must be taken, in order not to interrupt I/O operations and to avoid compromising the integrity of your data. The procedures outlined in this topic involve changing the worldwide node name (WWNN) of a SAN Volume Controller. These procedures must be followed with care in order to avoid duplicate WWPNs which can cause data corruption.

Be aware that by performing these procedures the following changes will be made to your configuration:

Front Panel ID
This number will change. It is the number that is printed on the front of the node and used to select the node that is to be added to a cluster.

Node Name
This number might change. If you do not specify a name, the SAN Volume Controller assigns a default name when adding a node to a cluster. The SAN Volume Controller creates a new name each time a node is added to a cluster. If you choose to assign your own names then you need to type in the node name on the Adding a node to a cluster panel. If you are using scripts to perform management tasks on the cluster and those scripts use the node name, then by assigning the original name to a replacement node, you avoid the need to make changes to the scripts.

Node ID
This ID will change. A new node ID is assigned each time a node is added to a cluster; the node name remains the same following service activity on the cluster. You can use the node ID or the node name to perform management tasks on the cluster. However, if you are using scripts to perform those tasks, use the node name rather than the node ID.

Worldwide Node Name
This name will change. The WWNN is used to uniquely identify the node and the fibre-channel ports. The WWNN of the spare node will change to that of the faulty node. The node replacement procedures must be followed exactly to avoid any duplication of WWNNs.

Worldwide Port Names
These names do not change. WWPNs are derived from the WWNN that is written to the spare (replacement) node as part of this procedure. For example, let's say the WWNN for a node is `50050768010000F6`. The four WWPNs for this node would be derived as follows:

```
WWNN                              50050768010000F6
WWNN displayed on front panel     000F6
WWPN Port 1                       50050768014000F6
```

```
WWPN Port 2                                                     50050768013000F6
WWPN Port 3                                                     50050768011000F6
WWPN Port 4                                                     50050768012000F6
```

Perform the following steps to replace a faulty node in the cluster:

1. Verify the name and ID of the node that you wish to replace.

   Perform the following steps to verify the name and ID:

   a. Make sure that the SAN Volume Controller Console application is running on the cluster that contains the faulty node.
   b. Click **Work with Nodes** in the portfolio.
   c. Click **Nodes**.

      If the node is faulty, it will be shown as offline. Ensure the partner node in the I/O group is online.

      1) If the other node in the I/O group is offline, start Directed Maintenance Procedures to determine the fault.
      2) If you have been directed here by the DMPs, and subsequently the partner node in the I/O group has failed, see the procedure for Recovering from offline VDisks.

      If you are replacing the node for other reasons, determine the node you wish to replace and again ensure the partner node in the I/O group is online.

      1) If the partner node is offline, you will lose access to the VDisks that belong to this I/O group if you continue. Start the Directed Maintenance Procedures and fix the other node before proceeding.

2. Find and record the following information about the faulty node:

   a. Node name
   b. I/O group name
   c. Last five characters of the WWNN
   d. Front panel ID
   e. Uninterruptible power supply serial number

   a. To find and record the node name and I/O group name, click **Work with Nodes** in the portfolio.
   b. Click **Nodes**.

      The faulty node will be offline.

   c. Record the following information about the faulty node:
      - Node name
      - I/O group name
   d. To find and record the last five characters of the WWNN, click on the name of the offline node.
   e. Click the **General** tab.
   f. Record the last five characters of the WWNN.
   g. To find and record the front panel ID, click the **Vital Product Data** tab.
   h. Find the front-panel-assembly section of the vital product data (VPD).
   i. Record the front panel ID.
   j. To find and record the uninterruptible power supply serial number, click the **Vital Product Data** tab.
   k. Find the uninterruptible power supply section of the VPD.
   l. Record the uninterruptible power supply serial number.

3. Obtain the ID of the faulty node. Disconnect all four fibre-channel cables from the node.

   Important: Do not plug the fibre-channel cables into the spare node until spare node has been configured with the WWNN from the faulty node.
4. Connect the power and signal cables from the spare node to the uninterruptible power supply that has the serial number that you noted in step 5l on page 129.

   Note: The signal cable can be plugged into any vacant position on the top row of serial connectors on the uninterruptible power supply. If no spare serial connectors are available on the uninterruptible power supply, disconnect the cables from the faulty SAN Volume Controller.
5. Power-on the spare node.
6. Display the node status on the service panel. See the *IBM TotalStorage SAN Volume Controller: Service Guide* for more information.
7. Change the WWNN of the spare node.

   Perform the following steps to change the WWNN of the spare node so that it matches the WWNN of the faulty node:
   a. With the node status displayed on the front panel, press and hold the **Down** button; press and release the **Select** button; release the **Down** button. `WWNN` is displayed on line-1 of the display; line-2 of the display contains the last five characters of the WWNN.
   b. With the WWNN displayed on the service panel, press and hold the **Down** button, press and release the **Select** button, release the **Down** button. This switches the display into edit mode. Change the displayed number to match the WWNN recorded in step 5f on page 129.

      Note: To edit the displayed number use the **Up** and **Down** buttons to increase or decrease the numbers displayed. Use the **left** and **right** buttons to move between fields. When the five characters match the number recorded in step 1, press the select button twice to accept the number.
8. Connect the four fibre-channel cables that were disconnected from the faulty node to the spare node.
9. Remove the faulty node from the cluster using the SAN Volume Controller Console.

   Remember: 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 when the node is re-added to the cluster.
10. Add the spare node to the cluster using the SAN Volume Controller Console.
11. Use the Subsystem Device Driver (SDD) management tool on the host systems to verify that all paths are now online. See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* for more information.

    Attention: When the faulty node is repaired do not connect the fibre-channel cables to it. Connecting the cables might cause data corruption.
12. Repair the faulty node.

13. If you want to use the repaired node as a spare node, perform the following steps:
    a. Display the node status on the front panel display of the node. See the *IBM TotalStorage SAN Volume Controller: Service Guide* for more information.
    b. With the node status displayed on the front panel, press and hold the **Down** button; press and release the **Select** button; release the **Down** button. `WWNN` is displayed on line-1 of the display; line-2 of the display contains the last five characters of the WWNN.
    c. With the WWNN displayed on the service panel, press and hold the **Down** button, press and release the **Select** button, release the **Down** button. This switches the display into edit mode. Change the displayed number to `00000`.

       Note: To edit the displayed number use the **Up** and **Down** buttons to increase or decrease the numbers displayed. Use the **left** and **right** buttons to move between fields.

       This SAN Volume Controller can now be used as a spare node.

    Attention: Never connect a SAN Volume Controller with a WWNN of `00000` to the cluster. If this SAN Volume Controller is no longer required as a spare and is to be used for normal attachment to a cluster you must first use the procedure described in the ″Prerequistes″ to change the WWNN to the number you recorded when a spare was created. Using any other number might cause data corruption.

Related tasks

"Deleting a node from a cluster" on page 137
A node might need to be deleted from a cluster if the node has failed and is being replaced with a new node or if the repair that has been performed has caused that node to be unrecognizable by the cluster.

"Launching the SAN Volume Controller Console" on page 96
You can launch the SAN Volume Controller from the Viewing Clusters panel.

"Replacing a faulty node in the cluster using the CLI" on page 186
You can replace a faulty node in the cluster using the command-line interface (CLI).

Related information

"Advanced functions overview for the SAN Volume Controller Console" on page 123
You can perform ″advanced″ functions using the SAN Volume Controller Console.

Recovering from offline VDisks after a node or an I/O group failed

If you have lost both nodes in an I/O group and have therefore, lost access to all the VDisks that are associated with the I/O group, then you must perform one of the following procedures to regain access to your VDisks. Depending on the failure type, you may have lost data that was cached for these VDisks, therefore, they have gone offline.

Data loss scenario 1 One node in an I/O group failed and failover started on the second node. During this time, the second node in the I/O group fails before the data in the write cache has been written to disk. The first node is successfully repaired but its cache data is stale, therefore, it cannot be used. The second node is repaired or replaced and has lost its hard-end data, therefore, the node has no way of recognizing that it is part of the cluster.

Perform the following steps to recover from an offline VDisk:

1. Recover the node and include it back into the cluster.
2. Move all the offline VDisks to the recovery I/O group.
3. Move all the offline VDisks back to their original I/O group.

Data loss scenario 2 Both nodes in the I/O group have failed and have been repaired. The nodes have lost their hard-end data, therefore, the nodes have no way of recognizing that they are part of the cluster.

- Move all the offline VDisks to the recovery I/O group
- Move both recovered nodes back into the cluster
- Move all the offline VDisks back to their original I/O group.

Related tasks

"Recovering a node and including it back into the cluster"
After a node or an I/O group fails, you can use the following procedure to recover a node and include it back into the cluster.

"Moving offline VDisks to the recovery I/O group" on page 133
After a node or an I/O group fails, you can use the following procedure to move offline VDisks to the recovery I/O group.

"Moving offline VDisks to their original I/O group" on page 134
After a node or an I/O group fails, you can use the following procedure to move offline VDisks to their original I/O group.

Related information

"Advanced functions overview for the SAN Volume Controller Console" on page 123
You can perform ″advanced″ functions using the SAN Volume Controller Console.

Recovering a node and including it back into the cluster

After a node or an I/O group fails, you can use the following procedure to recover a node and include it back into the cluster.

Perform the following steps to recover a node and include it back into the cluster:

1. Verify that the node is offline by viewing the **Work with Nodes** panel.
2. Remove the old instance of the offline node from the cluster by selecting the node and selecting the **Delete Node** task.
3. Verify that the node can be seen on the fabric.
4. If the nodes are repaired by replacing the front panel module or a node is repaired by replacing it with another node, then the WWNN for the node will change. In this case, the following additional steps are required:
   a. At the end of the recovery process it will be necessary to follow the SDD procedure to discover the new paths and to check that each vpath is now presenting the correct number of paths. For information about adding paths to existing vpaths. see the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide*.
   b. You may also need to modify the configuration of your disk controller systems. If your disk controller system uses a mapping technique to present its RAID arrays or partitions to the cluster you will need to modify the port groups that belong to the cluster because the WWNN or WWPN's of the node have changed.

Attention: If more than one I/O group is affected, 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. This can avoid a possible data corruption exposure if the node must be removed from and re-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.

Attention: 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

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

5. Add the node back into the cluster by selecting the **Add Node** task from the **Work with Nodes** panel. Select the node from the list of candidate nodes and select the I/O group from the list. Optionally enter a node name for this node.
6. Verify that the node is online by refreshing the **Work with Nodes** panel.

   Note: You may need to close the panel and reopen it for the refresh to take effect.

Related tasks

"Recovering from offline VDisks after a node or an I/O group failed" on page 131
If you have lost both nodes in an I/O group and have therefore, lost access to all the VDisks that are associated with the I/O group, then you must perform one of the following procedures to regain access to your VDisks. Depending on the failure type, you may have lost data that was cached for these VDisks, therefore, they have gone offline.

Moving offline VDisks to the recovery I/O group

After a node or an I/O group fails, you can use the following procedure to move offline VDisks to the recovery I/O group.

Perform the following steps to move offline VDisks to the recovery I/O group:

1. List all VDisks that are offline and that belong to the I/O group in question by selecting **Work with VDisks** from the portfolio. In the filter panel, enter the <iogrpname> in the I/O group filter box and select offline as the status.
2. For each VDisk returned, select the VDisk and select the **Modify** task. In the modify panel, only change the I/O group to **Recovery I/O group**. You may be asked to confirm and force the move, select to force the move.

Related tasks

"Recovering from offline VDisks after a node or an I/O group failed" on page 131
If you have lost both nodes in an I/O group and have therefore, lost access to all the VDisks that are associated with the I/O group, then you must perform one of the following procedures to regain access to your VDisks. Depending on the failure type, you may have lost data that was cached for these VDisks, therefore, they have gone offline.

Moving offline VDisks to their original I/O group

After a node or an I/O group fails, you can use the following procedure to move offline VDisks to their original I/O group.

Attention: Under no circumstances should VDisks be moved to an offline I/O group. Ensure the I/O group is online before moving the VDisks back to avoid any further data loss.

Perform the following steps to move offline VDisks to their original I/O group:

1. For each VDisk, select the VDisk and select the **Modify** task. In the modify panel, only modify the I/O group back to the original <iogrpname>.
2. Verify that the VDisks are now online by closing the Work with VDisks panel and opening it again. This time, in the filter panel, only enter the <iogrpname> in the I/O group filter box. The VDisks should all be online.

Related tasks

"Recovering from offline VDisks after a node or an I/O group failed" on page 131
If you have lost both nodes in an I/O group and have therefore, lost access to all the VDisks that are associated with the I/O group, then you must perform one of the following procedures to regain access to your VDisks. Depending on the failure type, you may have lost data that was cached for these VDisks, therefore, they have gone offline.

Replacing an HBA in a host using the SAN Volume Controller Console

It is sometimes necessary to replace the HBA that connects the host to the SAN, at this time you must notify the SAN Volume Controller of the new WWPN's that this HBA contains.

Ensure your switch is zoned correctly.

This procedure describes how to notify the SAN Volume Controller of a change to a defined host object.

1. Locate the host object that corresponds with the host in which you have replaced the HBA. Click **Work with Hosts** from the portfolio. Select the host object and select the **Add Ports** task.
2. Add the new ports to the existing host object. Select the candidate WWPNs from the list and click **Add**. Complete the task by clicking **OK**.
3. Remove the old ports from the host object. Select the host object and select the **Delete Ports** task. Select the WWPNs you wish to remove (the ones that correspond with the old HBA that was replaced). Click **Add** to add them to the list of WWPNs to be deleted. Complete the task by clicking **OK**.
4. Any mappings that exist between the host object and VDisks will automatically be applied to the new WWPNs. Therefore, the host should see the VDisks as the same SCSI LUNs as before. See *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* for adding paths to existing vpaths.

Related reference

"Configuring and servicing storage subsystems" on page 233
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

Deleting hosts

You can delete a host object using the Deleting hosts panel.

A deletion fails if there are any VDisk-to-host mappings for the host. If you attempt to delete the host and it fails due to the existence of VDisk mappings, then you are presented with the opportunity to perform a forced delete, which will delete VDisk mappings before deleteing the host.

Perform the following steps to delete a host object:
1. Click **Work with Virtual Disks** in the portfolio.
2. Click **Hosts** in the portfolio. The Filtering Hosts panel is displayed.
3. Specify the filter criteria that you want to use. Click **OK**. The Hosts panel is displayed.
4. Select the host that you want to delete and select **Delete Host** from the list.
5. Click **Go**.

When you delete a host object, all active ports are added to the **Available Ports** list.

Related concepts

"Host objects" on page 31
A host system is an open-systems computer that is connected to the switch through a fibre-channel interface.

Advanced function Metro Mirror overview

For detailed information about how to perform advanced FlashCopy and Metro Mirror functions, go to the following Web site: www.ibm.com/redbooks

Related reference

"Valid combinations of FlashCopy and Metro Mirror functions" on page 366
The following table outlines the combinations of FlashCopy and Metro Mirror functions that are valid for a single virtual disk (VDisk).

Advanced function cluster overview

Ensure that you are familiar with the advanced function cluster overview using the SAN Volume Controller Console.

Related tasks

"Analyzing the error log" on page 136
You can analyze the error log from the Analyze Error Log panel.

"Changing the language settings" on page 137
To change the language settings, complete these steps.

"Configuring error notification settings" on page 137
You can configure the error notification settings for the cluster from the Set Error Notification Settings panel.

"Deleting a node from a cluster" on page 137
A node might need to be deleted from a cluster if the node has failed and is being replaced with a new node or if the repair that has been performed has caused that node to be unrecognizable by the cluster.

"Enabling the cluster maintenance procedure using the SAN Volume Controller Console" on page 139
To enable the cluster maintenance procedure using the SAN Volume Controller Console, complete these steps.

"Listing and saving log and dump files" on page 140
You can list the log and dump files on the configuration node on the List Dumps panel.

“Renaming a cluster” on page 141
You can rename a cluster from the Renaming Cluster panel.

“Maintaining cluster passwords using the SAN Volume Controller Console” on page 141
To maintain cluster passwords using the SAN Volume Controller Console, complete these steps.

“Modifying Internet Protocol (IP) addresses” on page 147
You can display and change the IP addresses associated with the cluster from the Modify IP Addresses panel.

“Shutting down a cluster or node” on page 148
You can shut down a cluster from the Shutting Down cluster panel.

“Viewing the feature log” on page 149
You can view the feature log for the cluster from the View Feature Log panel.

“Viewing feature settings and log” on page 150
You can view the feature settings in the Viewing Feature Log panel.

Related information

“Managing SSH keys” on page 142
You can manage SSH keys from the SAN Volume Controller Console.

Analyzing the error log

You can analyze the error log from the Analyze Error Log panel.

Perform the following steps to analyze the error log:

1. Click **Service and Maintenance** in the portfolio.
2. Click **Analyze Error Log** in the portfolio. The Error log analysis panel is displayed. The Error log analysis panel enables you to analyze the cluster error log. You can display the whole log or filter the log so that only errors, events, or unfixed errors are displayed. In addition, you can request the table to be sorted either by error priority or by time. For error priority, the most serious errors are the lowest-numbered errors. They are, therefore, displayed first in the table. For time, either the oldest or the latest entry can be displayed first in the table. You can also select how many error log entries are to be displayed on each page of the table. The default is set to 10, and the maximum number of error logs that can be displayed on each page is 100.
3. After selecting the options, click **Process** to display the filtered error log in the table. The Analyze error log continued panel is displayed. Forward and Backward scroll buttons are displayed, depending on the existing page number and the total number of pages that are in the table. If the table contains more than two pages of entries, a **Go to** input area is displayed in the table footer. This input area enables you to skip to a particular page number.

   If you click on the sequence number of a particular table record, more information about that error log entry is displayed. If the record is an error (instead of an event), you can change the fixed or unfixed status of the record; that is, you can mark an unfixed error as fixed or a fixed error an unfixed.
4. Click **Clear log** to erase the whole cluster error log.

   Note: If you click **Clear log**, this will *not* fix the existing errors.

Related reference

“Advanced function cluster overview” on page 135
Ensure that you are familiar with the advanced function cluster overview using the SAN Volume Controller Console.

Changing the language settings

To change the language settings, complete these steps.

Perform the following steps to change the language settings:

1. Click **View Clusters** and select a cluster that you want to change the language setting for.
2. Click **Launch SAN Volume Controller application**.
3. Click **Manage Cluster**.
4. Click **General Properties**. From this panel you can change the locale setting to the appropriate language.

   Related reference

   "Advanced function cluster overview" on page 135
   Ensure that you are familiar with the advanced function cluster overview using the SAN Volume Controller Console.

Configuring error notification settings

You can configure the error notification settings for the cluster from the Set Error Notification Settings panel.

Perform the following steps to configure the error notification settings:

1. Click **Service and Maintenance** in the portfolio.
2. Click **Error settings** to display the existing error notification settings and to change them. The Modify Error Notification Settings panel is displayed. The Modify Error Notification Settings panel enables you to update your error notification settings. You can select whether the cluster raises an SNMP trap for entries that are added to the cluster error or event log. Three levels of notification are possible:
   a. **None** No error or status changes will be sent.
   b. **Hardware_only** You will be notified of errors, but you will not be notified of status changes.
   c. **All** You will be notified of all errors and status changes.
3. Click **Modify settings** to update the settings.

   Related reference

   "Advanced function cluster overview" on page 135
   Ensure that you are familiar with the advanced function cluster overview using the SAN Volume Controller Console.

Deleting a node from a cluster

A node might need to be deleted from a cluster if the node has failed and is being replaced with a new node or if the repair that has been performed has caused that node to be unrecognizable by the cluster.

Attention: Before deleting or removing 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.

Attention: If you are deleting or removing 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. Proceed to step 3 in the following procedure.

Attention: If you are deleting or removing 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 do not 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.
   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 power off the node that you intend to remove, 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.

   Note:
      a. If the node you are removing is the configuration node, it might take a minute or more before you can perform the delete node request. You must wait for the configuration node failover to occur.
      b. If the node you are removing is the last node in the cluster, the SAN Volume Controller Console might seem to hang for up to 3 minutes because you have removed the last access point to the cluster.

   Attention: Deleting the configuration node or shutting down the configuration node may result in the SSH command hanging. If this happens, you should either wait for the SSH command to timeout or kill the command and ping the cluster IP address until it responds. At this point the failover has completed and you can start issuing commands again.

   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 re-join 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 Multipath Subsystem Device Driver: User's Guide* for details about how to dynamically reconfigure SDD for the given host operating system.
4. Proceed to 1.

For example, if the disk drive or the software on the node has been replaced, that node will no longer be known by the cluster. You can delete or remove a node from the cluster with the Deleting a Node from Cluster panel.

1. Click **Work with Nodes** from the portfolio.

2. Click **Nodes** from the portfolio. The Nodes panel is displayed.

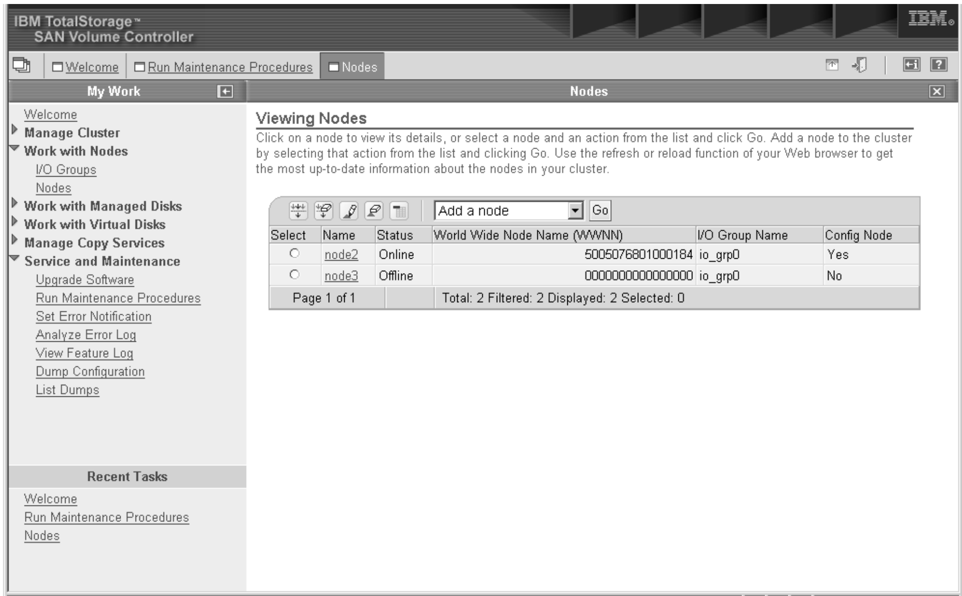

Figure 44. SAN Volume Controller Console nodes panel

3. Select the node you want to delete and Select **Delete a node** from the list. Click **Go**. The Deleting a Node from a Cluster panel is displayed.

Related concepts

"Clusters" on page 12
All configuration and service is performed at the cluster level.

"Nodes and clusters" on page 12
A SAN Volume Controller node is a single processing unit, which provides virtualization, cache, and copy services for the SAN.

Related tasks

"Determining the host that a VDisk is mapped to" on page 183
You can determine the host that a VDisk is mapped to using the command-line interface (CLI).

"Migrating a VDisk to a new I/O group" on page 126
You can migrate a VDisk to a new I/O group to manually balance the workload across the nodes in the cluster. You may end up with a pair of nodes that are overworked and another pair that are under-worked. Follow this procedure to migrate a single VDisk to a new I/O group. Repeat for other VDisks as required.

Related reference

"Advanced function cluster overview" on page 135
Ensure that you are familiar with the advanced function cluster overview using the SAN Volume Controller Console.

Enabling the cluster maintenance procedure using the SAN Volume Controller Console

To enable the cluster maintenance procedure using the SAN Volume Controller Console, complete these steps.

Perform the following steps to enable the maintenance procedure:

1. Click **Service and Maintenance** from the portfolio.

2. Click **Run Maintenance Procedures** to start the online maintenance procedures. The Maintenance procedures panel is displayed. A popup window is also displayed, requesting you to enter the user name and password for the SAN Volume Controller cluster. The Maintenance procedures window enables you to run the maintenance procedure on the cluster.
3. Click **Start Analysis** to analyze the cluster error log. A table of unfixed errors is displayed. The errors are sorted so that the most serious errors (those with the lowest error code) are listed first. The Maintenance panel is displayed. If you click the error code of a particular error log entry, you are guided through a series of actions that help you to estimate the state of the cluster and determine if the error was an isolated event or if a component has failed. If a component has failed, it might be necessary to exchange that component. Where necessary, images of the failing component are displayed. If a repair is performed successfully, the state of an error record in the error log changes from **unfixed error** to **fixed error**.

Related reference

"Advanced function cluster overview" on page 135
Ensure that you are familiar with the advanced function cluster overview using the SAN Volume Controller Console.

Listing and saving log and dump files

You can list the log and dump files on the configuration node on the List Dumps panel.

Dump data can be saved on any node in the cluster. When you use this procedure to display dump data only, the dump files on the configuration node will be displayed. An option is provided on the dumps menu to display data from other nodes. If you choose to display or save data from another node, that data will first be copied to the configuration node.

Perform the following steps to list the various types of log and dump files:

1. Click **Service and Maintenance** from the portfolio.
2. Click **List Dumps** from the portfolio. The List Dumps panel displays. The List dumps (other nodes) continued panel displays the number of log files or dumps of a particular type that are available on the cluster. If there is more than one node in the cluster (as is usual), the **Check other nodes** button is displayed. If you click this button, the log files and dumps for all nodes that are part of the cluster is displayed. Dumps and logs on all nodes in the cluster can be deleted or copied to the node.

   If you click on one of the file types, all the files of that type are listed in a table.

   Note: For error logs and software dumps, the file names include the node name and time and date as part of the file name.

   You can copy the files to a local workstation by right-clicking on the filename and using the **Save target as** (Netscape) or **Save file as** (Internet Explorer) option from the Web browser.

   The file types that the **List dumps** option supports are:
   - Error logs
   - Configuration logs
   - I/O statistic logs
   - I/O trace logs
   - Feature logs
   - Software dumps

Follow the instructions in the right pane to display and save the dumps that you need.

The software dump files contain dumps of the SAN Volume Controller memory. Your service representative might ask for these dumps to debug problems. The software dumps are large files (approximately 300 MB). Consider copying these files to your host using secure copy methods.

Related reference

"Advanced function cluster overview" on page 135
Ensure that you are familiar with the advanced function cluster overview using the SAN Volume Controller Console.

Renaming a cluster

You can rename a cluster from the Renaming Cluster panel.

Perform the following steps to rename a cluster:

1. Click **Clusters** from the portfolio. The Viewing Clusters panel is displayed.
2. Select the cluster that you want to rename and select **Rename a cluster** from the list. Click **Go**. The Renaming Cluster panel is displayed.
3. Complete the Renaming Cluster panel.

The cluster is renamed to the name that you selected.

Related concepts

"Clusters" on page 12
All configuration and service is performed at the cluster level.

Related reference

"Advanced function cluster overview" on page 135
Ensure that you are familiar with the advanced function cluster overview using the SAN Volume Controller Console.

Maintaining cluster passwords using the SAN Volume Controller Console

To maintain cluster passwords using the SAN Volume Controller Console, complete these steps.

Perform the following steps to maintain passwords:

1. Click **Manage Cluster** from the portfolio.
2. Click **Maintain Passwords** to modify the admin or service passwords that control access to the create cluster wizard. The Maintain passwords panel is displayed. The Maintain passwords window enables you to update the passwords that control access to the Web application for admin and service users. Passwords must be typed twice to allow verification. Passwords can consist of A - Z, a - z, 0 - 9, and underscore.
3. Type your admin or service user password and then click **Maintain Passwords** to change the password. If the admin password is changed, a password prompt is displayed and you must re-authenticate the password by entering the new admin password in the password prompt. Make a careful note of the admin password, because without it, you cannot access the cluster through the SAN Volume Controller Console.

Related reference

"Advanced function cluster overview" on page 135
Ensure that you are familiar with the advanced function cluster overview using the SAN Volume Controller Console.

Managing SSH keys

You can manage SSH keys from the SAN Volume Controller Console.

Related tasks

"Adding Secure Shell keys for hosts other than the master console"
Use these step-by-step instructions for adding Secure Shell (SSH) keys on hosts other than the master console.

"Adding subsequent Secure Shell public keys to the SAN Volume Controller"
During the cluster creation wizard, you will have added a Secure Shell (SSH) key to the cluster that allows the master console (where the SAN Volume Controller is running) to access the cluster. If you wish to add more SSH keys, that is, grant SSH access to other servers you need to follow the procedure below.

"Replace the client SSH private key known to the SAN Volume Controller software" on page 144
To replace the client SSH private key known to the SAN Volume Controller software, complete these steps.

"Resetting the SSH fingerprint" on page 146
You can reset the SSH fingerprint for a cluster that is managed by the SAN Volume Controller Console for your configuration by using the Resetting the SSH Fingerprint panel.

Related reference

"Advanced function cluster overview" on page 135
Ensure that you are familiar with the advanced function cluster overview using the SAN Volume Controller Console.

"Replacing the SSH key pair" on page 145
You can replace the SSH key pair using the SAN Volume Controller Console.

Related information

"Resetting a refused SSH key" on page 146
You can reset a refused SSH key relationship between the SAN Volume Controller Console and the SAN Volume Controller cluster.

Adding Secure Shell keys for hosts other than the master console:

Use these step-by-step instructions for adding Secure Shell (SSH) keys on hosts other than the master console.

1. Generate the public private key pair on each host that you want to use the SAN Volume Controller command line interface. See the information that came with your SSH client for specific details about using the key generation program that comes with your SSH client.
2. Copy the public keys from each of these hosts to the master console.
3. Secure copy these public keys from the master console to the cluster.

   Repeat for each public key copied onto the master console in step 2.

Related information

"Managing SSH keys"
You can manage SSH keys from the SAN Volume Controller Console.

Adding subsequent Secure Shell public keys to the SAN Volume Controller:

During the cluster creation wizard, you will have added a Secure Shell (SSH) key to the cluster that allows the master console (where the SAN Volume Controller is running) to access the cluster. If you wish to add more SSH keys, that is, grant SSH access to other servers you need to follow the procedure below.

1. Click **Clusters** in the Portfolio.
2. Click the cluster whose SSH keys you want to maintain.
3. Select Maintain SSH Keys in the drop-down list and click **Go**. The SSH Key Maintenance panel is displayed.

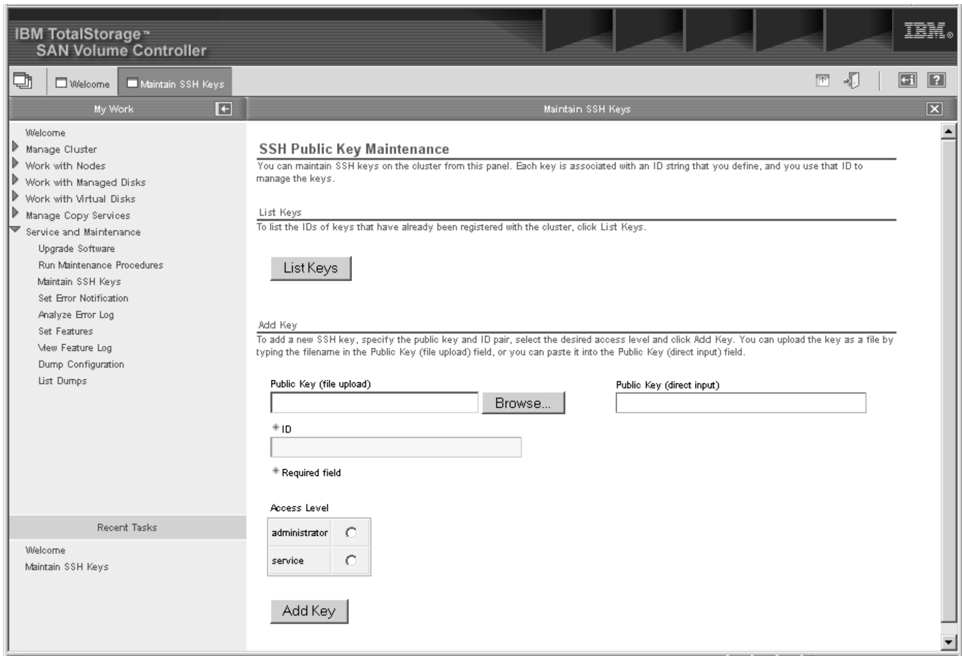

4. Click the **Maintain SSH Keys** option. The window appears to enable you to enter the client SSH public key information to be stored on the cluster. At the SSH key maintenance window, perform the following steps:
   a. If you are adding the SSH client key for the master console, click **Browse** and locate the public key you generated earlier. If you are adding an SSH client key for another system, either click **Browse** and locate the public key or cut and paste the public key into the direct input field.
   b. Click **Administrator**.
   c. Type a name of your choice in the **ID** field that uniquely identifies the key to the cluster.
   d. Click **Add Key**.
   e. Click **Maintain SSH Keys**.

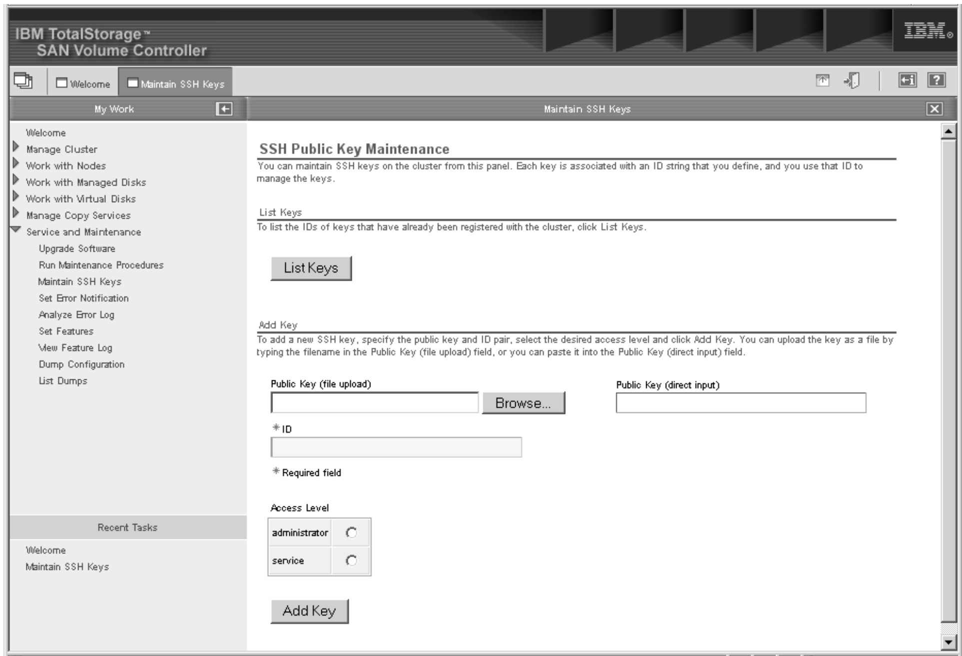

f. Click the **Show IDs** button to see all key IDs loaded on the SAN Volume Controller.

After the initial configuration of the cluster has been performed using the SAN Volume Controller and at least one SSH client key has been added the remainder of the configuration may either be performed using the SAN Volume Controller or the Command Line Interface (CLI).

Related information

"Managing SSH keys" on page 142
You can manage SSH keys from the SAN Volume Controller Console.

Replace the client SSH private key known to the SAN Volume Controller software:

To replace the client SSH private key known to the SAN Volume Controller software, complete these steps.

Attention: If you have successfully contacted other SAN Volume Controller clusters, you will break that connectivity if you replace the client SSH private key known to the SAN Volume Controller software.

Perform the following steps to replace the client SSH private key:

1. Sign off the SAN Volume Controller Console.
2. Using the Windows Services facility, stop the IBM CIM Object Manager. Perform the following:
   a. Click **Start -> Settings -> Control Panel**.
   b. Double-click **Administrative Tools**.
   c. Double-click **Services**.
   d. Select **IBM CIM Object Manager** in the list of services, right click, and select **Stop**.
   e. Leave the Services panel open.

3. Copy the client SSH private key into the appropriate SAN Volume Controller Console directory. Perform the following:
   a. Open a command prompt window by clicking **Start -> Run**.
   b. Type `cmd.exe` in the **Open** field.
   c. Click **OK**.
4. Type the following command:

   ```
   copy <filename> C:\program files\IBM\svcconsole\cimom\icat.ppk
   ```

   where *<filename>* is the path and file name of the client SSH private key.
5. Restart the IBM CIM Object Manager. Select **IBM CIM Object Manager** in the list of services, right click and select **Start**.
6. Log on to the SAN Volume Controller Console.
7. Click **Clusters** in the portfolio.
8. Check the status of the cluster.

   Related information

   "Managing SSH keys" on page 142
   You can manage SSH keys from the SAN Volume Controller Console.

Replacing the SSH key pair:

You can replace the SSH key pair using the SAN Volume Controller Console.

- If you change the SSH keys that will be used by the master console to communicate with the SAN Volume Controller Console, you will have to store the client SSH private key in the SAN Volume Controller Console software and then store the client SSH public key on the SAN Volume Controller cluster.
- If you change the IP address of your SAN Volume Controller cluster after you have added the cluster to SAN Volume Controller Console, the SAN Volume Controller Console will not be aware of the existence of the cluster.

The procedure to correct this is to remove the cluster from the SAN Volume Controller Console and add it back again. To correct these scenarios, perform the following steps:

1. Start the SAN Volume Controller Console by clicking on the desktop icon or by using your Web browser to go to

   http://<IPAddress>:9080/ica

   where *<IPAddress>* is the IP address of the master console. The Sign on window is displayed. This might take a few moments to open.
2. Enter the user ID `superuser` and the password `passw0rd`. The Welcome window is displayed.
3. Click **Clusters** from the portfolio.
4. Check the **Select** box for the cluster for which you wish to replace the key.
5. Click **Remove a cluster** in the selection box.
6. Click **Go**.
7. Click **Clusters** from the portfolio.
8. Select **Add a cluster** from the drop down box.
9. Click **Go**.
10. Input the IP address of the cluster.

11. Do not check the **Create (Initialize Cluster)** box.
12. Click **OK**.
13. Enter the user name and password. When you see the pop-up window, enter the network password and click **OK**.
14. Add the SSH client public key to the SAN Volume Controller cluster:
    a. Click **Browse...** for the key file to upload and locate the public key or input the key in the **Key (direct input)** field.
    b. Type an ID in the **ID** field, which uniquely identifies the key to the cluster.
    c. Select the **administrator** radio button.
    d. Click **Add Key**.
    e. Click **Clusters** from the portfolio to check the status of the cluster. If the cluster status remains **SSH Key Refused**, you do not have a good key pair. You can reset the SAN Volume Controller Console private SSH key. However, if you have successfully contacted other clusters, you will break that connectivity.

Related information

"Managing SSH keys" on page 142
You can manage SSH keys from the SAN Volume Controller Console.

Resetting a refused SSH key:

You can reset a refused SSH key relationship between the SAN Volume Controller Console and the SAN Volume Controller cluster.

Overview

The communication between the SAN Volume Controller Console software and the SAN Volume Controller cluster is through the Secure Shell (SSH) protocol. In this protocol, the SAN Volume Controller Console software acts as the SSH client and the SAN Volume Controller cluster acts as the SSH host server.

As an SSH client, the SAN Volume Controller Console must use an SSH2 RSA key pair composed of a public key and a private key which are coordinated at key generation time. The SSH client public key is stored on each SAN Volume Controller cluster with which the SAN Volume Controller Console communicates. The SSH client private key is known to the SAN Volume Controller Console software by being stored in a specific directory with a specific name. If the SSH protocol detects the key pair is mismatched, the SSH communication fail.

The SAN Volume Controller Console externalizes the status of a mismatched or invalid SAN Volume Controller Console client key pair in the **Availability Status** column of the Cluster panel.

Because the client SSH key pair must be coordinated across two systems, you might have to take one or more actions to reset the pair of keys. Perform one or more of the following steps to reset the refused client SSH key pair:
- Replace the client SSH public key on the SAN Volume Controller cluster
- Replace the client SSH private key known to the SAN Volume Controller software

Related information

"Managing SSH keys" on page 142
You can manage SSH keys from the SAN Volume Controller Console.

Resetting the SSH fingerprint:

You can reset the SSH fingerprint for a cluster that is managed by the SAN Volume Controller Console for your configuration by using the Resetting the SSH Fingerprint panel.

You must have superuser administrator authority to perform the following procedure.

If you have changed the name of the master console, you must also change the master console host name in the IBM WebSphere Application Server files.

The communication between the SAN Volume Controller Console and the cluster is through the Secure Shell (SSH) protocol. In this protocol, the SAN Volume Controller Console acts as the SSH client and the cluster acts as the SSH host server. The SSH protocol requires that credentials are exchanged when communication between the SSH client and server begins. The SSH client places the accepted SSH host server fingerprint in cache. Any change to the SSH server fingerprint in future exchanges results in a challenge to the end user to accept the new fingerprint. When a new code load is performed on the cluster, new SSH server keys can be produced which result in the SSH client flagging the SSH host fingerprint as changed and, therefore, no longer valid.

The SAN Volume Controller Console displays the status of the cluster SSH server key in the **Availability Status** column of the Viewing Clusters panel.

Perform the following steps to reset the SSH fingerprint:

1. Click **Clusters** in the portfolio. The View Clusters panel is displayed.

   Attention: Select a cluster that has an availability status of `Invalid SSH Fingerprint`. In some cases this availability status results from a software upgrade that disrupts normal user operations. In the case of a disruptive software upgrade, follow the procedure for Recovering from a Disruptive Software Upgrade.
2. Select the cluster that you want to reset the SSH fingerprint for and select **Reset SSH Fingerprint** from the list. Click **Go**. The Resetting the SSH Fingerprint panel is displayed.
3. Select **OK** when prompted with the message, CMMVC3201W.

Availability status is changed to `OK`.

Related concepts

"Clusters" on page 12
All configuration and service is performed at the cluster level.

Related information

"Managing SSH keys" on page 142
You can manage SSH keys from the SAN Volume Controller Console.

Modifying Internet Protocol (IP) addresses

You can display and change the IP addresses associated with the cluster from the Modify IP Addresses panel.

Perform the following steps to change the IP addresses:

1. Click **Manage Cluster** in the portfolio.
2. Click **Modify IP Addresses** to check or change the IP address settings for the cluster. The **Modify IP Addresses** panel is displayed. The Modify IP Addresses panel displays the existing value for the following IP Addresses and enables you to change the settings:

a. Cluster IP address
b. Service IP address (used when the node is not part of the cluster)
c. Subnet mask
d. Gateway

Fill in all four fields for the IP address that you want to change. Leave the IP address fields blank if you do not want to change them.

Click **Modify settings** to perform the IP address update. When you specify a new cluster IP address, the existing communication with the cluster is broken. You must use the new cluster IP address to reestablish your browser connection. A new SSL certificate is generated by the cluster (to show the new IP address). This new certificate is displayed when the Web browser first connects to the cluster.

Related concepts

"Clusters" on page 12
All configuration and service is performed at the cluster level.

Related reference

"Advanced function cluster overview" on page 135
Ensure that you are familiar with the advanced function cluster overview using the SAN Volume Controller Console.

Shutting down a cluster or node

You can shut down a cluster from the Shutting Down cluster panel.

If all input power to a SAN Volume Controller cluster is to be removed, for example, if the machine room power is to be shutdown for maintenance, you must shut down the cluster before the power is removed. If you do not shut down the cluster before turning off input power to the uninterruptible power supply, the SAN Volume Controllers will detect the loss of power and continue to run on battery power until all data held in memory is saved to the internal disk drive. This will increase the time required to make the cluster operational when input power is restored and will severely increase the time required to recover from an unexpected loss of power that might occur before the uninterruptible power supply batteries have fully recharged.

When input power is restored to the uninterruptible power supplies they will start to recharge but the SAN Volume Controllers will not permit any I/O activity to be performed to the virtual disks until the uninterruptible power supply is charged enough to enable all the data on the SAN Volume Controller nodes to be saved in the event of an unexpected power loss. This might take as long as three hours. Shutting down the cluster prior to removing input power to the uninterruptible power supply units will prevent the battery power being drained and will make it possible for I/O activity to be resumed as soon as input power is restored.

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

Attention: If you are shutting down the entire cluster, you will lose access to all VDisks being provided by this cluster. Shutting down the cluster will also shut down all nodes. This shutdown will cause the hardened data to be dumped to the internal hard drive.

| **Attention:** Ensure that you have stopped all FlashCopy, Metro Mirror and data migration operations before you attempt a node or cluster shutdown. Also ensure that all asynchronous deletion operations have completed prior to a shutdown operation.

Begin the process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDisks provided by the cluster.

1. If you are unsure which hosts are using the VDisks provided by the cluster, follow the procedure called, Determining the hosts that a VDisk is mapped to.
2. Repeat the previous step for all VDisks.

A cluster or node can be shutdown by stopping I/O activity and either pressing the power buttons on the front of each node or by issuing a shutdown command to the cluster.

Attention: You must press and hold the power button for one second to shutdown the node.

When input power is restored it will be necessary to press the power button on the uninterruptible power supply units before pressing the power buttons on the SAN Volume Controllers.

Perform the following steps to shut down a cluster:

1. Click **Manage Clusters** in the portfolio.
2. Click **Shut down Clusters** in the portfolio. The Shutting Down Clusters panel is displayed.To shut down a node, click **Shut down Nodes**. The Shutting Down Nodes panel is displayed.

Related concepts

"Clusters" on page 12
All configuration and service is performed at the cluster level.

Related tasks

"Determining the host that a VDisk is mapped to" on page 183
You can determine the host that a VDisk is mapped to using the command-line interface (CLI).

Related reference

"Advanced function cluster overview" on page 135
Ensure that you are familiar with the advanced function cluster overview using the SAN Volume Controller Console.

Viewing the feature log

You can view the feature log for the cluster from the View Feature Log panel.

Perform the following steps to view that feature log for the cluster:

1. Click **Service and Maintenance**.
2. Click **View Feature Log**. The View Feature Log panel is displayed.

Related reference

"Advanced function cluster overview" on page 135
Ensure that you are familiar with the advanced function cluster overview using the SAN Volume Controller Console.

Reviewing and setting the cluster features using the SAN Volume Controller Console

You can set up the cluster features using the SAN Volume Controller Console.

Perform the following steps to set up the cluster features:

1. Ensure that you have the correct license. The license will show you whether you are permitted to use FlashCopy or Metro Mirror. It will also show how much virtualization you are permitted to use.
2. From the Creating a Cluster wizard, click **Featurization settings** to set up or modify featurization attributes. The Featurization settings window is displayed.
3. The allowed setting for each of the parameters is specified in your license. The following features can be set:
   a. Enable the FlashCopy or Metro Mirror features if they are licensed.
   b. Enter the virtualization limit as specified in the license. A zero value is not allowed for this field.
   c. Click **Set features**. A featurization screen is displayed.

Viewing feature settings and log

You can view the feature settings in the Viewing Feature Log panel.

Perform the following steps to view the feature settings:

1. Click **Service and Maintenance** from the portfolio.
2. To view the feature settings, click **Set Features** from the portfolio. To view the feature log, click **View Feature Log**.

Related concepts

"Clusters" on page 12
All configuration and service is performed at the cluster level.

Related reference

"Advanced function cluster overview" on page 135
Ensure that you are familiar with the advanced function cluster overview using the SAN Volume Controller Console.

Chapter 4. Command-line interface

This part provides detailed information about using the command-line interface.

Related information

"Using the command-line interface"
The SAN Volume Controller cluster command-line Interface (CLI) is a collection of commands you can use to manage the SAN Volume Controller.

"Scenario: typical usage for the command-line interface" on page 161
The main focus of the following example is to provide storage to your host system using the command-line interface (CLI).

"Advanced functions with the CLI" on page 181
Ensure that you are familiar with the advanced functions that you can perform using the command-line interface (CLI).

Using the command-line interface

The SAN Volume Controller cluster command-line Interface (CLI) is a collection of commands you can use to manage the SAN Volume Controller.

Overview

The vehicle for these commands is the secure shell (SSH) connection between the SSH client software on the host system and the SSH server on the SAN Volume Controller cluster.

Before using the CLI, you must have performed the following initial steps to create and configure a cluster:

- Create a cluster from the front panel
- Complete the creation of a cluster using the SAN Volume Controller Console
- Perform the initial configuration of the cluster using the SAN Volume Controller

In order to use the CLI from a client system you must:

- Install and setup SSH client software on each system which you are going to issue command lines from.
- Generate an SSH key pair on each SSH client.
- Store the SSH public key for each SSH client on to the SAN Volume Controller using the SAN Volume Controller Console.

  Note: After the first SSH public key has been stored further SSH public keys may be added using either the SAN Volume Controller Console or the CLI.

The functions that can be performed with the IBM TotalStorage SAN Volume Controller command-line interface (CLI) are:

- Setup of the cluster, its nodes, and the I/O groups (or node pairs). This function includes diagnostics and error log analysis of the cluster.
- Setup and maintenance of managed disks and managed disk groups.
- Setup and maintenance of client public SSH keys on the cluster.
- Setup and maintenance or virtual disks.
- Setup of logical host objects.
- Mapping of virtual disks to hosts.

- Navigation from managed hosts to virtual disk groups and to managed disks, and the reverse direction up the chain.
- Setup and trigger of Copy Services:
  - FlashCopy and FlashCopy consistency groups
  - Synchronous Metro Mirror and Metro Mirror consistency groups

Related tasks

"Preparing the Secure Shell client system to issue command-line interface commands" on page 153
In order to issue command-line interface (CLI) commands to the cluster from a host, you must prepare the Secure Shell (SSH) client on the host so that the host will be accepted by the SSH server on the cluster and allowed to connect.

"Preparing the SSH client on an AIX host" on page 154
When using AIX hosts, SSH logins are authenticated on the SAN Volume Controller cluster using the RSA based authentication supported in the OpenSSH client available for AIX.

"Issuing CLI commands from a PuTTY SSH Client system" on page 155
You can issue CLI commands from a PuTTY SSH Client system.

"Configuring the cluster using the CLI" on page 158
You can configure the cluster using the command-line interface (CLI).

"Setting the cluster time using the CLI" on page 159
This task provides step-by-step instructions for setting the cluster time using the command-line interface.

"Reviewing and setting the cluster features using the CLI" on page 160
You can set up the cluster features using the command-line interface (CLI).

"Displaying cluster properties using the CLI" on page 160
You can display cluster properties using the command-line interface (CLI).

"Maintaining passwords using the CLI" on page 161
You can maintain passwords using the command-line interface (CLI).

Related reference

"Preparing the Secure Shell client system overview"
Before you can issue CLI commands from the host to the cluster, you must prepare the Secure Shell (SSH) client system.

Related information

Chapter 4, "Command-line interface," on page 151

"Running the PuTTY and plink utilities" on page 156
Ensure that you are familiar with how to run the PuTTY and plink utilities.

Preparing the Secure Shell client system overview

Before you can issue CLI commands from the host to the cluster, you must prepare the Secure Shell (SSH) client system.

Windows operating systems

If you have purchased the master console hardware and software from IBM, then PuTTY for Windows operating systems has been previously installed for you.

If you are installing the master console on your own hardware with a Windows operating system, then you can install PuTTY from the SAN Volume Controller Console CD-ROM.

| Or, you can download PuTTY from the following Web site:
| http://www.chiark.greenend.org.uk/~sgtatham/putty/

The following Web site offers SSH client alternatives for Windows: http://www.openssh.com/windows.html

Cygwin software has an option to install an OpenSSH client. You can download cygwin from the following Web site: http://www.cygwin.com/

AIX operating systems

For AIX® 5L Power 5.1 and 5.2, you can get OpenSSH from the Bonus Packs and you will also need its prerequisite, OpenSSL, from the AIX toolbox for Linux applications for Power Systems. For AIX 4.3.3, you can get the software from the AIX toolbox for Linux applications.

You can also get the AIX installation images from IBM DeveloperWorks at the following Web site: http://oss.software.ibm.com/developerworks/projects/openssh

Linux operating systems

OpenSSH is installed by default on most Linux distributions. If it is not installed on your system, consult your installation media or visit the following Web site: http://www.openssh.org/portable.html

OpenSSH is able to run on a wide variety of additional operating systems. For more information visit the following Web site: http://www.openssh.org/portable.html

Related information

"Using the command-line interface" on page 151
The SAN Volume Controller cluster command-line Interface (CLI) is a collection of commands you can use to manage the SAN Volume Controller.

Preparing the Secure Shell client system to issue command-line interface commands

In order to issue command-line interface (CLI) commands to the cluster from a host, you must prepare the Secure Shell (SSH) client on the host so that the host will be accepted by the SSH server on the cluster and allowed to connect.

If you wish to use a host which requires a different type of SSH client, for example OpenSSH, follow the instructions for that software.

Perform the following steps to enable your host to issue CLI commands:

1. For the master console and Windows hosts:
   a. Generate a SSH key pair using the **PuTTY** key generator.
   b. Store the SSH clients public key on the cluster (using a browser pointing to the SAN Volume Controller Console).
   c. Configure the PuTTY session for the command-line interface
2. For other types of hosts:
   a. Follow the instructions specific to the SSH client to generate an SSH key pair.
   b. Store the SSH clients public key on the cluster (using a browser pointing to the SAN Volume Controller Console or the Command Line Interface from an already established host).

c. Follow the instructions specific to the SSH client to establish an SSH connection to the SAN Volume Controller cluster.
3. For other types of hosts, follow the instructions specific to the SSH client.

Related tasks

"Adding Secure Shell keys for hosts other than the master console" on page 142
Use these step-by-step instructions for adding Secure Shell (SSH) keys on hosts other than the master console.

Related information

"Using the command-line interface" on page 151
The SAN Volume Controller cluster command-line Interface (CLI) is a collection of commands you can use to manage the SAN Volume Controller.

Preparing the SSH client on an AIX host

When using AIX hosts, SSH logins are authenticated on the SAN Volume Controller cluster using the RSA based authentication supported in the OpenSSH client available for AIX.

This scheme is based on public-key cryptography, using a scheme known commonly as RSA. In this system (as in similar OpenSSH systems on other host types) the encryption and decryption is done using separate keys. This means it is not possible to derive the decryption key from the encryption key. Initially, the user creates a public/private key pair for authentication purposes. The server (which is the SAN Volume Controller cluster in this case) knows the public key, and only the user (or the AIX host) knows the private key. Physical possession of the public key allows access to the cluster, so this must be kept in a protected place (typically in the /.ssh directory on the AIX host, with restricted access permissions).

The SAN Volume Controller cluster uses the following mechanism to establish that the AIX host is to be trusted, for example, that the AIX host possesses the private key.

When you log in from the AIX host to the SAN Volume Controller cluster, the SSH program on the SAN Volume Controller cluster tells the AIX host which key pair it would like to use for authentication. The AIX server checks if this key is permitted, and if so, sends the user (actually the SSH program running on behalf of the user) a challenge, a random number, encrypted by the user's public key. The challenge can only be decrypted using the proper private key. The user's client (for example, the AIX host) then decrypts the challenge using the private key, proving that he/she knows the private key but without disclosing it to the server (for example, the SAN Volume Controller cluster) or to anyone who might be intercepting the transmissions between the AIX host and the SAN Volume Controller cluster.

The main steps in setting up a RSA key pair on the AIX host and the SAN Volume Controller cluster are as follows (detailed instructions are given in subsequent paragraphs):
1. Create an RSA key pair by running the ssh-keygen program on the AIX host.
2. Store the private key from this key pair on the AIX host, in the /.ssh directory.
3. Place the public key on the SAN Volume Controller cluster and associate this key with an 'admin' or 'service' type user.

The main steps in setting up a RSA key pair on the AIX host and the SAN Volume Controller cluster are as follows (detailed instructions are given in subsequent paragraphs):

1. Create an RSA key pair by running the ssh-keygen program on the AIX host.
2. Store the private key from this key pair on the AIX host, in the /.ssh directory.
3. Place the public key on the SAN Volume Controller cluster and associate this key with an ’admin’ or ’service’ type user.

At this point, you may then use the ’ssh’ or ’scp’ utilities on the AIX host to establish an SSH session with the SAN Volume Controller cluster or to perform Secure Copy operations between the AIX host and the SAN Volume Controller cluster.

Perform the following steps to set up a RSA key pair on the AIX host and the SAN Volume Controller cluster:

1. Create an RSA key pair by running the ssh-keygen program on the AIX host. This is best done in the $HOME/.ssh directory. This process will generate two user named files. Suppose the user selects the name ’key’, then the files will be named ’key’ and ’key.pub’.
2. Store the private key from this key pair on the AIX host, in the ’$HOME/.ssh directory’, in the ’$HOME.ssh/identity file’. In the simplest case, this means replacing the contents of the ’identity’ file with the contents of the ’key’ file. If multiple keys are to be used however, then all of these keys must appear in the ’identity’ file. This step places the private key on the host.
3. Move the public key, ’key.pub’ to the master console of the SAN Volume Controller cluster of interest. Typically this might be done with ftp however the master console may have ftp disabled for security reasons, in which case an alternative method would be required (for example, a secure copy between the application host and the master console). Then, using the SAN Volume Controller Console, and the SAN Volume Controller Web interface, select the ’Maintain SSH Keys’ panel, and transfer the key.pub to the cluster. Select an access level of ’administrator’ or ’service’ as appropriate. In this example, we will assume the key was associated with an administrative ID and that the cluster IP name is ’mycluster’. This step places the public key on the cluster.
4. You can now access the cluster from the AIX host, using SSH commands similar to the following:

   ```
   ssh admin@mycluster
   ssh admin@mycluster svcinfo lsnode
   ```

Refer to your clients documentation for SSH on your host system for more host specific details of this procedure.

Related tasks

“Adding Secure Shell keys for hosts other than the master console” on page 142
Use these step-by-step instructions for adding Secure Shell (SSH) keys on hosts other than the master console.

Related information

“Using the command-line interface” on page 151
The SAN Volume Controller cluster command-line Interface (CLI) is a collection of commands you can use to manage the SAN Volume Controller.

Issuing CLI commands from a PuTTY SSH Client system

You can issue CLI commands from a PuTTY SSH Client system.

Perform the following steps to issue CLI commands:

1. Open a Command Prompt to open the SSH connection to issue the CLI commands.

2. Make the PuTTY executables available by performing the following:
   a. Change directory into the PuTTY executable directory. For example, on the master console, type the following:

      ```
      C:\Support Utils\putty
      ```

      ```
       C:\Program Files\Putty
      ```
   b. Set the path environment variable to include the PuTTY executable directory. For example, type the following:

      ```
      Set path=c:\Support Utils\putty;%path%
      ```
3. Use the PuTTY plink utility to connect to the SSH server on the cluster.

   Related information

   "Using the command-line interface" on page 151
   The SAN Volume Controller cluster command-line Interface (CLI) is a collection of commands you can use to manage the SAN Volume Controller.

Running the PuTTY and plink utilities

Ensure that you are familiar with how to run the PuTTY and plink utilities.

All CLI commands are run in an SSH session. You can run the commands in either of the following modes:

- An interactive prompt mode
- A single line command mode, which is entered one time to include all parameters

Interactive mode

For interactive mode, you use the PuTTY executable to open the SSH restricted shell. Type the following:

```
C:\support utils\putty>putty admin@<svcconsoleip>
```

If you were to issue the **svcinfo lssshkeys** command, which lists the SSH client public keys that are stored on the SAN Volume Controller cluster, the following output is displayed:

```
IBM_2145:admin>svcinfo lssshkeys -user all -delim :
id:userid:key identifier
1:admin:smith
2:admin:jones
```

Type `exit` and press **Enter** to escape the interactive mode command.

The SSH protocol specifies that the first access to a new host server will result in a *challenge* to the SSH user to accept the SSH server public key. Because this is the first time that you connect to an SSH server, the server is not included in the SSH client list of known hosts. Therefore, there is a fingerprint challenge, which asks you do you accept the responsibility of connecting with this host. If you type `y`, the host fingerprint and IP address is saved by the SSH client. For PuTTY, you answer by typing `y` to accept this host fingerprint. This information is stored in the registry for the user name which is logged onto Windows.

The following is an example of the host fingerprint challenge when running in interactive mode:

```
C:\Program Files\IBM\svcconsole\cimom>plink admin@9.43.225.208
The server’s host key is not cached in the registry. You
have no guarantee that the server is the computer you
think it is.
The server’s key fingerprint is:
ssh-rsa 1024 e4:c9:51:50:61:63:e9:cd:73:2a:60:6b:f0:be:25:bf
If you trust this host, enter ″y″ to add the key to
PuTTY’s cache and carry on connecting.
If you want to carry on connecting just once, without
adding the key to the cache, enter ″n″.
If you do not trust this host, press Return to abandon the
connection.
Store key in cache? (y/n) y
Using username ″admin″.
Authenticating with public key ″imported-openssh-key″
IBM_2145:admin>
```

Single line command

For single line command mode, you can type the following all on one command line:

```
C:\Program Files\IBM\svcconsole\cimom>
 plink admin@9.43.225.208 svcinfo lssshkeys
 -user all -delim :
Authenticating with public key "imported-openssh-key"
/bin/ls: id:userid:key identifier
1:admin:smith
2:admin:jones

C:\Program Files\IBM\svcconsole\cimom>
```

The SSH protocol specifies that the first access to a new host server will result in a *challenge* to the SSH user to accept the SSH server public key. Because this is the first time that you connect to an SSH server, the server is not included in the SSH client list of known hosts. Therefore, there is a fingerprint challenge, which asks you do you accept the responsibility of connecting with this host. If you type `y`, the host fingerprint and IP address is saved by the SSH client. For PuTTY, you answer by typing `y` to accept this host fingerprint. This information is stored in the registry for the user name which is logged onto Windows.

The following is an example of the host fingerprint challenge when running in single line command mode:

```
C:\Program Files\IBM\svcconsole\cimom>
 plink admin@9.43.225.208 svcinfo lssshkeys
 -user all -delim :
The server's host key is not cached in the registry. You
have no guarantee that the server is the computer you
think it is.
The server's key fingerprint is:
ssh-rsa 1024 e4:c9:51:50:61:63:e9:cd:73:2a:60:6b:f0:be:25:bf
If you trust this host, enter "y" to add the key to
PuTTY's cache and carry on connecting.
If you want to carry on connecting just once, without
adding the key to the cache, enter "n".
```

```
If you do not trust this host, press Return to abandon the
connection.
Store key in cache? (y/n) y
Authenticating with public key "imported-openssh-key"
/bin/ls: /proc/20282/exe: Permission denied
dircolors: `/etc/DIR_COLORS': Permission denied
id:userid:key identifier
1:admin:smith
2:admin:jones

C:\Program Files\IBM\svcconsole\cimom>
```

Note: If you are submitting a CLI command with all parameters in single line command mode, you will be challenged upon first appearance of the SSH server host fingerprint. Be careful to ensure that the SSH server host fingerprint is accepted before submitting a batch script file.

The SSH protocol also specifies that once the SSH server public key is accepted, another challenge will be presented if the fingerprint of an SSH server changes from the one previously accepted. In this case, you will need to decide whether to accept this changed host fingerprint. For PuTTY, you answer by typing `y` to accept this host fingerprint. PuTTY stores this information in the registry for the user name which is logged onto Windows.

Note: The SSH server keys on the SAN Volume Controller will be regenerated when a microcode load is performed on the cluster. Due to this behavior, you will see a *challenge* presented because the fingerprint of the SSH server has changed.

Related information

"Using the command-line interface" on page 151
The SAN Volume Controller cluster command-line Interface (CLI) is a collection of commands you can use to manage the SAN Volume Controller.

Configuring the cluster using the CLI

You can configure the cluster using the command-line interface (CLI).

The initial steps in creating and configuring a cluster must be performed using the front panel and the SAN Volume Controller Console. Once the cluster has been createdand a SSH public key has been added all further tasks can be accomplished using the command-line interface (CLI).

Perform the following steps to configure the cluster:

1. Open a command prompt window.
2. To change your time-zones and set your cluster time, you can issue the **svctask settimezone** and **svctask setclustertime** commands.
3. If you wish to use the Command Line Interface (CLI) from additional systems then use the **svctask addsshkey** to add further SSH public keys.
4. If you choose you can review and modify the initial configuration of the cluster that was performed using the front panel and SAN Volume Controller Console:
   a. Issue the command **svcinfo lscluster** to display the cluster properties. To display full details of your cluster's properties, issue the **svcinfo lscluster -delim : <cluster_name>** command.

b. To modify the passwords, fabric speed or cluster IP address issue the command **svctask chcluster**.

   Note: When you issue the **svctask chcluster** command to change the fabric speed, your connection to the command-line interface will be restarted. This occurs because the software on each node is stopped and restarted. When you restart the configuration node, you lose access to the command-line interface for a short time.

c. Issue the command **svctask setpwdreset -show** to view the status of the password reset feature for the front panel, and issue the command **svctask setpwdreset -enable?-disable** to change it.

d. To review and modify your featurization settings you can issue the **svcinfo lslicense** and **svctask chlicense** commands.

e. If you wish to modify the set up for error notifications to help manage your errors from the cluster, you can issue the **svctask setevent** command to set up SNMP traps.

5. Issue the **svctask mkcluster** command.
6. Enable your featurization settings by issuing the **svctask chlicense** command. You will need to specify whether you want the FlashCopy or Metro Mirror enabled or disabled. You can also specify your capacity size for virtualized storage.
7. To change your time-zones and reset your cluster time, you can issue the **svctask settimezone** and **svctask setclustertime** commands.
8. If you choose to set up error notifications to manage your errors from the cluster, you can issue the **svcservicemodetask setevent** command to set up SNMP traps.

Related tasks

"Modifying IP addresses using the CLI" on page 206
You can modify IP addresses using the command-line interface (CLI).

"Maintaining SSH keys using the CLI" on page 206
You can maintain SSH keys using the command-line interface (CLI).

"Modifying passwords using the CLI" on page 207
You can modify the admin and service passwords using the command-line interface (CLI).

"Setting up error notifications using the CLI" on page 206
You can set up error notifications using the command-line interface (CLI).

Related information

"Using the command-line interface" on page 151
The SAN Volume Controller cluster command-line Interface (CLI) is a collection of commands you can use to manage the SAN Volume Controller.

Setting the cluster time using the CLI

This task provides step-by-step instructions for setting the cluster time using the command-line interface.

Perform the following steps to set the cluster time:

1. Open a command prompt window.
2. Issue the **svcinfo showtimezone** command to display the current time-zone settings for the cluster. The cluster ID and its associated time-zone are displayed.

3. Issue the **svcinfo lstimezones** command to list the time-zones available on the cluster. A list of valid time-zones settings are displayed in a list. The specific cluster ID and its assigned time-zone are indicated in the list.
4. Issue the **svctask settimezone** command to set the time zone for the cluster.
5. Issue the **svctask setclustertime** command to set the time for the cluster.

Related information

"Using the command-line interface" on page 151
The SAN Volume Controller cluster command-line Interface (CLI) is a collection of commands you can use to manage the SAN Volume Controller.

Reviewing and setting the cluster features using the CLI

You can set up the cluster features using the command-line interface (CLI).

Perform the following steps to set up the cluster features:

1. Open a command prompt window.
2. Issue the **svctask lslicense** command to return the current license (featurization) settings for the cluster. The output displayed lists the feature functions in a list and displays whether they are enabled or disabled.
3. Issue the **svcinfo lslicense** command to change the licensed settings of the cluster. Because the feature settings are entered when the cluster is first created, you need to update the settings only if you have changed your license. You can change the following values:
   a. FlashCopy: disabled or enabled
   b. Metro Mirror: disabled or enabled
   c. Virtualization limit: number, in gigabytes (1073741824 bytes)

Related information

"Using the command-line interface" on page 151
The SAN Volume Controller cluster command-line Interface (CLI) is a collection of commands you can use to manage the SAN Volume Controller.

Displaying cluster properties using the CLI

You can display cluster properties using the command-line interface (CLI).

Perform the following steps to display cluster properties:

1. Open a command prompt window.
2. Issue the **svcinfo lscluster** command to display a concise view of the cluster.

```
svcinfo lscluster -delim : 10030a007e5
```

   where *10030a007e5* is the name of the cluster. The output from this command will display the following: The output from this command will include the following for each cluster on the fabric:
   a. cluster ID
   b. cluster name
   c. cluster IP address
   d. cluster service mode IP address

Related information

"Using the command-line interface" on page 151
The SAN Volume Controller cluster command-line Interface (CLI) is a collection of commands you can use to manage the SAN Volume Controller.

Maintaining passwords using the CLI

You can maintain passwords using the command-line interface (CLI).

Perform the following steps to maintain passwords:

1. Open a command prompt window.
2. Issue the **svctask setpwdreset** command to view and change the status of the password-reset feature for the display panel. Passwords can consist of A - Z, a - z, 0 - 9, and underscore. Make a careful note of the admin password, because without it, you cannot access the cluster.

   Related information

   "Using the command-line interface" on page 151
   The SAN Volume Controller cluster command-line Interface (CLI) is a collection of commands you can use to manage the SAN Volume Controller.

Scenario: typical usage for the command-line interface

The main focus of the following example is to provide storage to your host system using the command-line interface (CLI).

For example, you wish to provide a host system with two disks and create a FlashCopy of these two disks. The copy is to be made available to a second host. These two hosts require that the host objects that are created, correspond with the group of WWPNs presented by their fibre-channel HBAs to the SAN. You also need to create four virtual disks, one for each of the disks that are to be presented to the hosts. Once the VDisks are created, you can map two of them to each host. In order to create the VDisks you need to have a managed disk group to be able to create them from. You wish to spread the 8 managed disks across two groups and create the source VDisks from one and the target VDisks from the other. In order to create any of these objects you need to create a cluster and at least add one more node to the cluster.

The following steps illustrates how this can be done:

1. Create a cluster.
2. Configure the cluster with an IP address of 9.20.123.456, a fabric speed of 2 GB. Name the cluster examplecluster.
3. Add nodes
   - knode and lnode to the I/O group called io_grp0 in the examplecluster cluster
   - mnode and nnode to the I/O group called io_grp1 in the examplecluster cluster
4. Create the MDisk groups maindiskgroup and bkpdiskgroup
5. Create 4 VDisks
   - 2 VDisks from maindiskgroup
   - 2 VDisks from bkpdiskgroup
6. Create 2 host objects
   - a host object called demohost1 with HBAs that have WWPNs of 10000000C92AD7E5, and 10000000C92F5123
   - a host object called demohost2 with HBAs that have WWPNs of 210000E08B0525D4, and 210100E08B2525D4
7. Create the VDisk-to-host mappings
   - Map the two VDisks from maindiskgroup to demohost1

- Map the two VDisks from bkpdiskgroup to demohost2
8. Create FlashCopy mappings
   - Create a FlashCopy mapping called main1copy that has a background copy rate of 75
   - Create a FlashCopy mapping called main2copy that has a background copy rate of 50
9. Create a FlashCopy consistency group called maintobkpfcopy and add the 2 FlashCopy mappings to it
10. Prepare and trigger (start) the FlashCopy Consistency Group that contains these mappings

    Note: Once this step is complete, you have created and allocated storage to your host systems. You have made two VDisks available to demohost1 and then used FlashCopy to make backup copies on two VDisks which are accessible to demohost2.

Related tasks

“Configuring a cluster using the SAN Volume Controller Console” on page 88
Once you’ve created a pair of nodes, you will then need to create and configure a cluster.

“Create cluster from the front panel” on page 79
After you have created a pair of nodes, you can now create a cluster from the front panel.

“Adding nodes to a cluster using the CLI” on page 163
You can add nodes to a cluster using the CLI.

“Displaying node properties using the CLI” on page 167
You can display node properties using the command-line interface (CLI).

“Discovering MDisks using the CLI” on page 168
You can discover MDisks using the command-line interface (CLI).

“Creating managed disk (MDisk) groups using the CLI” on page 169
You can create an MDisk group using the command-line interface (CLI).

“Adding MDisks to MDisk groups using the CLI” on page 171
You can add MDisks to MDisk groups using the command-line interface (CLI).

“Create virtual disks (VDisks)” on page 172
You can create a VDisk using the command-line interface (CLI).

“Creating host objects using the CLI” on page 175
This task provides step-by-step instructions about how to create host objects.

“Create VDisk-to-host mappings using the CLI” on page 176
You can create VDisk-to-host mappings using the command-line interface (CLI).

“Create FlashCopy mappings using the CLI” on page 176
You can create FlashCopy mappings using the command-line interface (CLI).

“Creating a FlashCopy consistency group and adding mappings using the CLI” on page 177
You can create a FlashCopy Consistency Group and add mappings to it using the command-line interface (CLI).

“Preparing and triggering a FlashCopy mapping using the CLI” on page 178
You can prepare and trigger a FlashCopy to start the FlashCopy process using the command-line interface (CLI).

“Preparing and triggering a FlashCopy Consistency Group using the CLI” on page 179
You can prepare and trigger a FlashCopy Consistency Group to start the flash copy process using the command-line interface (CLI).

Related information

Chapter 4, "Command-line interface," on page 151

Adding nodes to a cluster using the CLI

You can add nodes to a cluster using the CLI.

Before adding a node to a cluster check to see if any of the following conditions are true:

- The cluster has more than one I/O group.
- The node being added to the cluster uses physical node hardware which has previously been used as a node in the cluster.
- The node being added to the cluster uses physical node hardware which has previously been used as a node in *another* cluster and both clusters have visibility to the same hosts.

Attention: If any of these conditions are true, then you must perform the following special procedures. Failure to perform the special procedure is likely to result in the corruption of all data managed by the cluster.

Special procedures when adding a node to a cluster

If any of the previous conditions are true, then the following special procedures apply. These special procedures apply when you use either the **svctask addnode** command or the SAN Volume Controller Console
.

When a node is added to a cluster then either:

- The node must be added back to the same I/O group that it was previously in.

  Note: The WWNN of the nodes in the cluster can be determined using the command:

  ```
  svcinfo lsnode
  ```

  or, if this information is not available, then
- Call IBM Service to ensure that data is not lost during the Adding a node procedure.

  Note:

  - *Before* the node is added back into the cluster all the hosts using the cluster must be shut down.

    The node must then be added before the hosts are rebooted.or, if the I/O group information is not available and it is inconvenient to shutdown and reboot all of the hosts using the cluster, then
  - On all the hosts connected to the cluster, unconfigure the Fibre Channel adapter device driver, the disk device driver, and the SDD device driver, before you add the node to the cluster.

    Re-configure the Fibre Channel adapter device driver, the disk device driver, and the SDD device driver, after adding the node into the cluster.

    Note: This may not be possible on all operating systems in all circumstances.

Hypothetical scenarios where the special procedures may apply.

The following are two hypothetical scenarios where the special procedures may apply:

- Four nodes of an eight-node cluster have been lost because of the failure of a pair of UPS 5125 or four UPS 5115 power supplies. In this case, the four lost nodes must be added back to the cluster using the **svctask addnode** command or the SAN Volume Controller Console.
- A user decides to delete four nodes from the cluster and add them back into the cluster using the **svctask addnode** command or the SAN Volume Controller Console.

Before you begin adding nodes, you must have already created a cluster, performed the initial configuration using the SAN Volume Controller Console and performed the necessary setup to use the Command Line Interface (CLI).

Background

Applications on host systems direct I/O operations to file systems or logical volumes that are mapped by the operating system to vpaths that are pseudo-disk objects supported by the SDD driver. See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide*.

The SDD driver maintains an association between a vpath and a SAN Volume Controller VDisk. This association uses an identifier (UID) that is unique to the VDisk and is never reused. This allows the SDD driver to unambiguously associate vpaths with VDisks.

The SDD device driver operates within a protocol stack that also contains disk and fibre-channel device drivers that allow it to communicate with the SAN Volume Controller using the SCSI protocol over fibre-channel as defined by the ANSI FCS standard. The addressing scheme provided by these SCSI and fibre-channel device drivers uses a combination of a SCSI logical unit number (LUN) and the worldwide name for the fibre-channel node and ports.

In the event of errors occurring, error recovery procedures (ERPs) operate at various tiers in the protocol stack. Some of these ERPs cause I/O to be re-driven using the same WWN and LUN numbers which were previously used.

The SDD device driver does not check the association of the VDisk with the VPath on every I/O that it performs.

Important: Before you begin adding nodes, you must have already created a cluster, performed the initial configuration using the SAN Volume Controller Console and performed the necessary setup to use the Command Line Interface (CLI).

The following examples are all based on our hypothetical scenario of setting up an eight-node cluster. The first node has already been used to create a cluster. Therefore, the remaining seven nodes must be added to the cluster.

Perform the following steps to add nodes to a cluster:

1. Open a command prompt window.

2. Determine the I/O group to which the node should be added. Type the **svcinfo lsnode** command to list the nodes that are currently part of the cluster.

```
svcinfo lsnode -delim :

 id:name:UPS_serial_number:WWNN:status:IO_group_id:
 IO_group_name:config_node:UPS_unique_id
 1:node1:10L3ASH:500507680100002C:online:0:io_grp0:yes:202378101C0D18D8
....
```

The cluster has only just been created so there is only one node in the cluster.

3. Type the **svcinfo lsnodecandidate** command to list nodes that are not assigned to a cluster and to verify that when adding a second node to an I/O group that it is attached to a different UPS.

```
svcinfo lsnodecandidate -delim :

 id:panel_name:UPS_serial_number:UPS_unique_id
 5005076801000001:000341:10L3ASH:202378101C0D18D8
 5005076801000009:000237:10L3ANF:202378101C0D1796
 50050768010000F4:001245:10L3ANF:202378101C0D1796
....
```

There are a total of eight nodes, one of which has been used to create a cluster. Therefore there are seven candidate nodes that can be added to the cluster.

4. Issue the command **svctask addnode** to add a node to the cluster.

   You can specify a name for the node. To change the name of a node that is already part of a cluster, use the **svctask chnode** command.

   If you do not specify the name for the node, the node can later be identified by using the front panel name, which is printed on a label on the front of the SAN Volume Controller, or by using the world wide node name of that node.

   For the new node, 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.

   If you are re-adding a node to the cluster, 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.

Add a second node to the first I/O group by issuing the following command. Note from the output from step 1 that the node that is already in I/O group 0 is attached to a UPS with the serial number 10L3ASH. Each node in an I/O group must be attached to a different UPS and therefore only the nodes with front panel IDs 000237 and 001245 are suitable candidates.

```
svctask addnode -panelname 000237 -iogrp io_grp0 -name group1node2
```

This command adds the node, identified by the front panel name 000237 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. Check the output from step 3 to make sure that each node is attached to a different UPS.

```
svctask addnode -wwnodename 5005076801000001 -iogrp io_grp1 -name group2node1
svctask addnode -wwnodename 50050768010000F4 -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 **svcinfo lsnode** command.

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

```
svcinfo lsnode -delim :

 id:name:UPS_serial_number:WWNN:status:IO_group_id:
 IO_group_name:config_node:UPS_unique_id
 1:group1node1:10L3ASH:500507680100002C:online:0:io_grp0:yes:202378101C0D18D8
 2:group1node2:10L3ANF:5005076801000009:online:0:io_grp0:no:202378101C0D1796
 3:group2node1:10L3ASH:5005076801000001:online:1:io_grp1:no:202378101C0D18D8
 4:group2node2:10L3ANF:50050768010000F4:online:1:io_grp1:no:202378101C0D1796
....
```

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:
- 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.

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

Related tasks

“Registering the SAN Volume Controller ports with the EMC CLARiiON” on page 255
To register SAN Volume Controller ports with an EMC CLARiiON controller that has Access Logix installed, complete these steps.

Displaying node properties using the CLI

You can display node properties using the command-line interface (CLI).

Perform the following steps to display the node properties:

1. Open a command prompt window.
2. Issue the **svcinfo lsnode** command to display a concise list of nodes in the cluster.

Type the following command:

```
svcinfo lsnode -delim :
```

This command displays the following:

```
id:name:UPS_serial_number:WWNN:status:IO_group_id:
 IO_group_name:config_node:UPS_unique_id
1:group1node1:10L3ASH:500507680100002C:online:0:io_grp0:yes:202378101C0D18D8
2:group1node2:10L3ANF:5005076801000009:online:0:io_grp0:no:202378101C0D1796
3:group2node1:10L3ASH:5005076801000001:online:1:io_grp1:no:202378101C0D18D8
4:group2node2:10L3ANF:50050768010000F4:online:1:io_grp1:no:202378101C0D1796
```

3. Issue the **svcinfo lsnode** command again, however, this time, specify the node ID or name of a node to receive the detailed output.

For example, to provide a detailed view of the node named group1node1 type the following:

```
svcinfo lsnode -delim : group1node1
```

This command displays the following:

```
id:1
name:group1node1
UPS_serial_number:10L3ASH
WWNN:500507680100002C
status:online
IO_group_id:0
IO_group_name:io_grp0
partner_node_id:2
partner_node_name:group1node2
config_node:yes
UPS_unique_id:202378101C0D18D8
port_id:500507680110002C
port_status:active
port_id:500507680120002C
port_status:active
port_id:500507680130002C
port_status:active
port_id:500507680140003C
port_status:active
```

The output includes:
- node ID
- node name
- WWNN
- details about the uninterruptible power supply to which the node is attached
- details about the I/O group of which the node is a member
- detailed fibre channel port status information.

Discovering MDisks using the CLI

You can discover MDisks using the command-line interface (CLI).

When backend controllers are added to the fibre-channel SAN and are included in the same switch zone as a SAN Volume Controller Cluster the cluster will automatically discover the backend controller and will integrate the controller to determine what storage it is presented to the SAN Volume Controller. The SCSI LUs presented by the backend controller will be displayed as unmanaged MDisks. If however the configuration of the backend controller is modified after this has occurred then the SAN Volume Controller may be unaware of these configuration changes. This task allows a user to request the SAN Volume Controller to re-scan the fibre-channel SAN to update the list of unmanaged MDisks.

Note: The automatic discovery performed by SAN Volume Controller does not write anything to a unmanaged MDisk. It is only when a the user instructs the

SAN Volume Controller to add a MDisk to a managed disk group or use a MDisk to create an image mode virtual disk that the storage will actually be used.

Perform the following steps to display MDisks:

1. Open a command prompt window.
2. Check to see which MDisks are available by issuing the **svctask detectmdisk** command to manually scan the fibre-channel network for any MDisks.
3. Issue the **svcinfo lsmdiskcandidate** command to show the unmanaged MDisks. These MDisks have not been assigned to an MDisk group. Alternatively, you can issue the **svcinfo lsmdisk** command to view all of the MDisks.

In our hypothetical scenario we have a single backend controller that is presenting eight SCSI LUs to the SAN Volume Controller. Issue the following command:

```
svctask detectmdisk
```

```
svcinfo lsmdiskcandidate
```

This command displays the following:

```
id
0
1
2
3
4
5
6
7
```

Issue the following command:

```
svcinfo lsmdisk -delim : -filtervalue mode=unmanaged
```

This command displays the following:

```
id:name:status:mode:mdisk_grp_id:mdisk_grp_name:
 capacity:ctrl_LUN_#:controller_name
0:mdisk0:online:unmanaged:::273.3GB:0000000000000000:controller0
1:mdisk1:online:unmanaged:::273.3GB:0000000000000001:controller0
2:mdisk2:online:unmanaged:::273.3GB:0000000000000002:controller0
3:mdisk3:online:unmanaged:::273.3GB:0000000000000003:controller0
4:mdisk4:online:unmanaged:::136.7GB:0000000000000004:controller0
5:mdisk5:online:unmanaged:::136.7GB:0000000000000005:controller0
6:mdisk6:online:unmanaged:::136.7GB:0000000000000006:controller0
7:mdisk7:online:unmanaged:::136.7GB:0000000000000007:controller0
```

You have now shown that the backend controllers and switches have been setup correctly and that the SAN Volume Controller can see the storage being presented by the backend controller.

Creating managed disk (MDisk) groups using the CLI

You can create an MDisk group using the command-line interface (CLI).

Attention: If you add a MDisk to a MDisk group as a managed disk, any data on the MDisk will be lost. If you want to keep the data on a MDisk (for example because you want to want to import storage that was previously not managed by a SAN Volume Controller) then you should create image mode VDisks instead.

Assume that the cluster has been set up and that a backend controller has been configured to present some new storage to the SAN Volume Controller.

Before creating managed disk groups consider how you are going to use your storage. The SAN Volume Controller allows you to create up to 128 managed disks groups and to add up to 128 MDisks to an MDisk group. Consider the following factors when deciding how many managed disk groups to create:

- A virtual disk can only be created using the storage from one managed disk group. Therefore if you create small managed disk groups then you may lose the benefits provided by virtualization, namely more efficient management of free space and a more evenly distributed workload to provide better performance.
- If any managed disk in a managed disk group goes offline then all the virtual disks in the managed disk group will go offline. Therefore you might want to consider using different managed disk groups for different backend controllers or for different applications.
- If you anticipate regularly adding and removing backend controllers or storage then this task will be made simpler by grouping all the managed disks presented by a backend controller into one managed disk group.
- All the managed disks in a managed disk group should have similar levels of performance or reliability, or both. If a managed disk group contains managed disks with different levels of performance then the performance of the virtual disks in this group will be limited by the performance of the slowest managed disk. If a managed disk group contains managed disks with different levels of reliability then the reliability of the virtual disks in this group will be that of the least reliably managed disk in the group.

Even with the best planning, circumstances may change and you may wish to re-configure your managed disk groups after they have been created. The data migration facilities provided by the SAN Volume Controller enable you to move data without disrupting I/O.

Choosing a managed disk group extent size: 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 may have different extent sizes however this will 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 Cluster. Table 14 shows the maximum amount of storage that can be managed by a cluster for each extent size. Because the SAN Volume Controller 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 to distribute sequential I/O workloads across many managed disks and hence may reduce the performance benefits of virtualization.

Table 14. Extent size

Extent Size	Maximum storage capacity of cluster
16MB	64TB
32MB	128TB

Table 14. Extent size (continued)

64MB	256TB
128MB	512TB
256MB	1PB
512MB	2PB

Attention: You can specify different extent sizes for different managed disk groups, however you will not be able to migrate virtual disks between managed disk groups with different extent sizes. Therefore if possible create all your managed disk groups with the same extent size.

Perform the following steps to create a MDisk group:

1. Open a command prompt window.
2. Type the **svctask mkmdiskgrp** command to create a MDisk group.

   Example

   In our hypothetical scenario, the command to create a MDisk group is:`svctask mkmdiskgrp -name maindiskgroup -ext 32 -mdisk mdsk0:mdsk1:mdsk2:mdsk3` This command will create a MDisk group called *maindiskgroup*. The extent size used within this group will be *32* MB, and there are four MDisks *mdsk0, mdsk1, mdsk2, mdsk3* added to the group.

   Example

   In our hypothetical scenario, the command to create a second MDisk group is:`svctask mkmdiskgrp -name bkpmdiskgroup -ext 32`

   Note: In this example we will create a second MDisk group first and add MDisks later.

   This command will create a MDisk group called *bkpmdiskgroup*. The extent size used within this group will be *32* MB.

   Example

   To add MDisks to the MDisk group, issue the **svctask addmdisk** command. In our hypothetical scenario, the command to add MDisks to the MDisk group is:`svctask addmdisk -mdisk mdsk4:mdsk5:mdsk6:mdsk7 bkpdiskgroup` This command will add four MDisks *mdsk4, mdsk5, mdsk6, mdsk7* to the MDisk group called *bkpdiskgroup*.

Adding MDisks to MDisk groups using the CLI

You can add MDisks to MDisk groups using the command-line interface (CLI).

The managed disks must be in unmanaged mode. Disks that already belong to a group cannot be added to another group until they have been deleted from their current group. You can delete a managed disk from a group under the following circumstances:

- If the managed disk does not contain any extents in use by a virtual disk
- If you can first migrate the extents in use onto other free extents within the group.

Perform the following steps to add MDisks to MDisk groups:

Note: Do not add the MDisk using this procedure if you want to make an image mode VDisk with it.

1. Open a command prompt window.
2. Type the **svcinfo lsmdiskgrp** command to list the existing MDisks groups.

In our hypothetical scenario, we have two MDisk groups, one with four managed disks and one with no managed disks. Type the following command:

```
svcinfo lsmdiskgrp -delim :
```

This command displays the following:

```
id:name:status:mdisk_count:vdisk_count:
capacity:extent_size:free_capacity
0:mainmdiskgroup:online:4:0:1093.2GB:32:1093.2GB
1:bkpmdiskgroup:online:0:0:0:32:0
```

3. To add MDisks to the MDisk group, issue the **svctask addmdisk** command.

In our hypothetical scenario, the command to add MDisks to the MDisk group is:

```
svctask addmdisk -mdisk mdisk4:mdisk5:mdisk6:mdisk7 bkpmdiskgroup
```

This command will add four MDisks mdisk4, mdisk5, mdisk6 and mdisk7 to the MDisk group called bkpmdiskgroup.

Related tasks

"Importing data by creating an image mode VDisk" on page 113
You can migrate logical units into virtual disks (VDisks).

Related reference

"Optimal managed disk group configurations" on page 236
A managed disk group provides the pool of storage from which virtual disks will be created. It is therefore necessary to ensure that the entire pool of storage provides the same performance and reliability characteristics.

Create virtual disks (VDisks)

You can create a VDisk using the command-line interface (CLI).

Note: If you want to keep the data on a MDisk (for example because you want to want to import storage that was previously not managed by a SAN Volume Controller) then you should create image mode VDisks instead. This task only deals with creating VDisks with a striped virtualization policy. For details of other virtualization policies refer to the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide*.

Assume that the cluster has been setup and that you have created managed disk groups. You must establish an empty managed disk group to hold the MDisks used for image mode VDisks.

Perform the following steps to create VDisks:

1. Open a command prompt window.

2. Decide which managed disk group will provide the storage for the VDisk. Use the **svcinfo lsmdiskgrp** command to list the available managed disk groups and the amount of free storage in each group.

In our hypothetical scenario, issue the following:

```
svcinfo lsmdiskgrp -delim :
```

This command displays the following:

```
id:name:status:mdisk_count:vdisk_count:
capacity:extent_size:free_capacity
0:mainmdiskgroup:online:4:0:1093.2GB:32:1093.2GB
1:bkpmdiskgroup:online:4:0:546.8GB:32:546.8GB
```

3. Decide which I/O group the VDisk should be assigned to. This determines which SAN Volume Controller nodes in the cluster process the I/O requests from the host systems. If you have more than one I/O group then make sure you distribute the VDisks between the I/O groups so that the I/O workload is shared evenly between all SAN Volume Controller nodes. Use the **svcinfo lsiogrp** command to show the I/O groups and the number of virtual disks assigned to each I/O group.

Note: It is normal for clusters with more than one I/O group to have MDisk groups that have VDisks in different I/O groups. FlashCopy can be used to make copies of VDisks regardless of whether the source and destination VDisk are in the same I/O group. If however you plan to use intra-cluster Metro Mirror then make sure that both the master and auxiliary VDisk are in the same I/O group.

In our hypothetical scenario, there are two I/O groups each with two nodes. Neither I/O group has any virtual disks yet. Issue the following command:

```
svcinfo lsiogrp -delim :
```

This command displays the following:

```
id:name:node_count:vdisk_count
0:io_grp0:2:0
1:io_grp1:2:0
2:io_grp2:0:0
3:io_grp3:0:0
4:recovery_io_grp:0:0
```

4. Type the **svctask mkvdisk** command to create a virtual disk (VDisk).

In our hypothetical scenario, the command to create a VDisk is:

```
svctask mkvdisk -name mainvdisk1 -iogrp 0
-mdiskgrp 0 -vtype striped -size 256 -unit gb
```

This command will create a VDisk called mainvdisk1, the VDisk will use I/O group 0 and MDisk group 0 (the ID of maindiskgroup as shown in the output from step 2). The VDisk capacity is 256GB and will be made up of extents from the MDisks in the MDisk group.

In our hypothetical scenario, the command to create a second VDisk is:

Note: This command is the same as the above example, however, here we are specifying the names of the objects instead of the IDs.

```
svctask mkvdisk -name mainvdisk2 -iogrp io_grp0
-mdiskgrp maindiskgroup -vtype striped -size 256 -unit gb
```

This command will create a VDisk called mainvdisk2, the VDisk will use the I/O group named io_grp0 and the MDisk group named maindiskgroup. The VDisk capacity is 256GB and will be made up of extents from the MDisks in the MDisk group.

In our hypothetical scenario, the commands to create a third VDisk is:

Note: This virtual disk is created with an ordered list of MDisks within the MDisk group to allocate extents from.

The following command lists the managed disks in the MDisk group with ID 1 (named bkpmdiskgroup):

```
svcinfo lsmdisk -delim : -filtervalue mdisk_grp_id=1
```

This command displays the following:

```
id:name:status:mode:mdisk_grp_id:
mdisk_grp_name:capacity:ctrl_LUN_#:
controller_name
4:mdisk4:online:managed:1:bkpmdiskgroup:
136.7GB:0000000000000004:controller0
5:mdisk5:online:managed:1:bkpmdiskgroup:
136.7GB:0000000000000005:controller0
6:mdisk6:online:managed:1:bkpmdiskgroup:
136.7GB:0000000000000006:controller0
7:mdisk7:online:managed:1:bkpmdiskgroup:
136.7GB:0000000000000007:controller0
```

Issue the following command:

```
svctask mkvdisk -name bkpvdisk1 -iogrp io_grp1
-mdiskgrp bkpmdiskgrp -vtype striped -size 256
-unit gb -mdisk 4:5
```

This command will create a VDisk called bkpvdisk1, the Vdisk will use the I/O group named io_grp1 and the MDisk group named bkpmdiskgrp. The VDisk capacity is 256GB and will be made up of extents allocated from the MDisks with IDs 4 and 5.

In our hypothetical scenario, the command to create a fourth VDisk is:

```
svctask mkvdisk -name bkpvdisk2 -iogrp io_grp1
-mdiskgrp bkpmdiskgrp -vtype striped -size 256 -unit
gb -mdisk mdisk6:mdisk7
```

This command will create a VDisk called bkpvdisk2, the VDisk will use the I/O group named io_grp1 and the MDisk group named bkpmdiskgrp. The VDisk capacity is 256GB and will be made up of extents allocated from the MDisks with names mdisk6 and mdisk7.

5. To list all the virtual disks that have been created use the **svcinfo lsvdisk** command.

In our hypothetical scenario we have created four VDisks. Issue the following command:

```
svcinfo lsvdisk -delim :
```

This command displays the following:

```
id:name:IO_group_id:IO_group_name:status:
mdisk_grp_id:mdisk_grp_name:capacity:type:FC_id:
FC_name:RC_id:RC_name
0:mainvdisk1:0:io_grp0:online:0:mainmdiskgroup:
512.0GB:striped::::
1:mainvdisk2:0:io_grp0:online:0:mainmdiskgroup:
512.0GB:striped::::
2:bkpvdisk1:1:io_grp1:online:1:bkpmdiskgroup:
512.0GB:striped::::
3:bkpvdisk2:1:io_grp1:online:1:bkpmdiskgroup:
512.0GB:striped::::
```

Creating host objects using the CLI

This task provides step-by-step instructions about how to create host objects.

1. Open a command prompt window.
2. Type the **svctask mkhost** command to create a logical host object. Assign your WWPN for the HBAs in the hosts.

In our hypothetical scenario, the command to create a host is:

```
svctask mkhost -name demohost1 -hbawwpn  210100e08b251dd4
```

This command will create a host called *demohost1* with the HBA WWPN of *210100e08b251dd4*.

3. Type the **svctask addhostport** command to add ports to the host.

In our hypothetical scenario, the command to add a port to the host is:

```
svctask mkhost -name demohost2 -hbawwpn  210100e08b251dd5
```

This command will add another HBA WWPN called *210100e08b251dd5* to the host that we created in step 2.

In our hypothetical scenario, the command to create a second host is:

```
svctask mkhost -hbawwpn 210100e08b251dd6:210100e08b251dd7 -name demohost2
```

This command will create a second host called *demohost2* with the HBA WWPN of *210100e08b251dd6, 210100e08b251dd7*.

Note: If you were to add a host with a faulty WWPN, or the WWPN had been assigned to the wrong host, you will need to issue the **svctask addhostport** command to add that same host with the correct WWPN, then issue the **svctask rmhostport** command to delete the host with the wrong or faulty WWPN. For example, if you had a host called *demohost1* and its WWPN stopped working, you would need to issue the following:

```
svctask addhostport -hbawwpn 210100e08b251dd4 demohost1
```

This would add the host called *demohost1* with the WWPN, *210100e08b251dd4*. You would then need to issue the **svctask rmhostport** command to delete the host with the WWPN that had stopped working. For example, you would issue the following:

```
svctask rmhostport -hbawwpn 210100e08b251dd5 demohost1
```

From these two commands, you have deleted the host with the WWPN *210100e08b251dd5*, and have added the same host with the WWPN *210100e08b251dd4*.

Create VDisk-to-host mappings using the CLI

You can create VDisk-to-host mappings using the command-line interface (CLI).

We are going to map the VDisks named, mainvdisk1 and mainvdisk2, to the host named demohost1. We are also going to map the VDisks named, bkpvdisk1 and bkpvdisk2, to the host named demohost2. The VDisks, mainvdisk1 and mainvdisk2, are contained in the managed disk (MDisk) group, mainmdiskgroup; while the VDisks, bkpvdisk1 and bkpvdisk2, are contained in the MDisk group, bkpmdiskgroup.

Perform the following steps to create VDisk-to-host mappings:

1. Open a command prompt window.
2. Type the **svctask mkvdiskhostmap** to create VDisk-to-host mappings.

In our hypothetical scenario, the commands to create VDisk-to-host mappings are:

```
svctask mkvdiskhostmap -host demohost1 mainvdisk1
svctask mkvdiskhostmap -host demohost1 mainvdisk2
svctask mkvdiskhostmap -host demohost2 bkpvdisk1
svctask mkvdiskhostmap -host demohost2 bkpvdisk2
```

The above set of commands map each VDisk to a host.

Create FlashCopy mappings using the CLI

You can create FlashCopy mappings using the command-line interface (CLI).

We are going to create mappings that enable us to copy the VDisk, mainvdisk1, to bkpvdisk1 and the VDisk, mainvdisk2 to bkpvdisk2.

The mapping specifies the source and destination virtual disks. The destination must be identical in size to the source, or the mapping will fail. Issue the **svcinfo lsvdisk -bytes** command to find the exact size of the source VDisk that you want to create a target disk of the same size. The source and destination cannot be in an existing mapping. That is, a virtual disk can be either a source or a destination disk in **only one** mapping. A mapping is triggered at the point in time when the copy is required.

Perform the following steps to create FlashCopy mappings:

1. Open a command prompt window.
2. Type the **svctask mkfcmap** command to create a FlashCopy mapping.

In our hypothetical scenario, the commands to create FlashCopy mappings are:

```
svctask mkfcmap -source mainvdisk1 -target bkpvdisk1
 -name main1copy -copyrate 75
svctask mkfcmap -source mainvdisk2 -target bkpvdisk2
 -name main2copy
```

The above commands create two FlashCopy mappings. For main1copy the background copy rate is 75; for main2copy, because the rate is not specified in the **mkfcmap** command, the priority is the default, 50.

3. To check the attributes of the mappings that have been created, issue the following **svcinfo lsfcmap** command:

   ```
   svcinfo lsfcmap -delim :
   ```

   ```
   id:name:source vdisk id:source vdisk name:target
     vdisk id:target vdisk name:group id:group
    name:status:progress:copy rate
   0:main1copy:0:mainvdisk1:1:bkpvdisk1:::idle_copied::75
   1:main2copy:2:mainvdisk2:3:bkpvdisk2:::idle_copied::50
   ```

Creating a FlashCopy consistency group and adding mappings using the CLI

You can create a FlashCopy Consistency Group and add mappings to it using the command-line interface (CLI).

If you have created several FlashCopy mappings for a group of VDisks that contain elements of data for the same application, you may find it convenient to assign these mappings to a single FlashCopy Consistency Group. Then you can issue a single prepare or trigger command for the whole group, so that, for example, all the files for a particular database are copied at the same time.

Perform the following steps to create a FlashCopy mappings:

1. Open a command prompt window.
2. Issue the **svctask mkfcconsistgrp** command to create a FlashCopy Consistency Group.

In our hypothetical scenario, the command to create a FlashCopy Consistency group called *maintobkpfcopy* is:

```
svctask mkfcconsistgrp -name maintobkpfcopy
```

Use the **svcinfo lsfcconsistgrp** command to display the attributes of the group you have created.

```
svcinfo lsfcconsistgrp -delim :
```

This command displays the following:

```
id:name:status
1:maintobkpfcopy:idle_copied
```

3. Use the **svctask chfcmap** command to add the two FlashCopy mappings created in the previous section to the new consistency group.

In our hypothetical scenario, the commands to add the mappings called main1copy and main2copy to the consistency group called maintobkpfcopy are:

```
svctask chfcmap -consistgrp maintobkpfcopy main1copy
svctask chfcmap -consistgrp maintobkpfcopy main2copy
```

Use the **svcinfo lsfcmap** command to display the new attributes of the mappings.

```
svcinfo lsfcmap -delim :
id:name:source_vdisk_id:source_vdisk_name:target_vdisk_id:
target_vdisk_name:group_id:group_name:state:progress:copy_rate
0:main1copy:28:maindisk1:29:bkpdisk1:1:maintobkpfcopy:idle_copied::75
1:main2copy:30:maindisk2:31:bkpdisk2:1:maintobkpfcopy:idle_copied::50
```

Notice that the group_name field displays maintobkpfcopy for both mappings.

Use the **svcinfo lsfcconsistgrp** command with the name of the consistency group to display the detailed attributes of the group. This now includes a list of the IDs and names of the mappings that are in the group.

```
 svcinfo lsfcconsistgrp  -delim : maintobkpfcopy
id:1
name:maintobkpfcopy
status:idle_copied
FC_mapping_id:0
FC_mapping_name:main1copy
FC_mapping_id:1
FC_mapping_name:main2copy
```

Related reference

"Considerations for FlashCopy mappings" on page 237
Ensure that you've considered the type of I/O and frequency of update before creating the virtual disks that you wish to use in FlashCopy mappings.

Preparing and triggering a FlashCopy mapping using the CLI

You can prepare and trigger a FlashCopy to start the FlashCopy process using the command-line interface (CLI).

This will create a point-in-time copy of the data on the source VDisk and write it to the target VDisk for the mapping.

Perform the following steps to prepare and trigger a FlashCopy mapping:

1. Open a command prompt window.
2. Issue the **svctask prestartfcmap** command to prepare the FlashCopy mapping before the copy process can be started (triggered).

In our hypothetical scenario, the command to prepare a FlashCopy mapping called main1copy is:

```
svctask prestartfcmap main1copy
```

The mapping will enter the preparing state, and then move to the prepared state when it is ready. Issue the **svcinfo lsfccmap** command to check:

```
svcinfo lsfcmap -delim :
id:name:source_vdisk_id:source_vdisk_name:target_vdisk_id:
 target_vdisk_name:group_id:group_name:status:progress:copy_rate
0:main1copy:0:mainvdisk1:1:bkpvdisk1:::prepared:0:50
```

3. Issue the **svctask startfcmap** command to start (trigger) the FlashCopy mapping to make the copy.

In our hypothetical scenario, the command to trigger a FlashCopy mapping called main1copy is:

```
svctask startfcmap main1copy
```

4. Use the **svcinfo lsfcmapprogress** command to check the progress of the mapping.

   ```
   svcinfo lsfcmapprogress -delim : main1copy
   id:progress
   0:47
   ```

   When the copy is complete the output to the **svcinfo lsfcmap** command will show the progress at 100 and the status as idle_or_copied.

   ```
   svcinfo lsfcmap -delim :
   id:name:source_vdisk_id:source_vdisk_name:target_vdisk_id:
    target_vdisk_name:group_id:group_name:status:progress:copy_rate
   0:main1copy:0:mainvdisk1:1:bkpvdisk1:::idle_or_copied:100:50
   ```

   The output to **svcinfo lsfcmapprogress** will show the progress at 100 for example,

   ```
   svcinfo lsfcmapprogress main1copy
   id              progress
   0               100
   ```

You have now made a point-in-time copy of the data on mainvdisk1 which has been written to bkpvdisk1. The data on bkpvdisk1 will be visible to demohost2 because these VDisks are only mapped to demohost2.

Preparing and triggering a FlashCopy Consistency Group using the CLI

You can prepare and trigger a FlashCopy Consistency Group to start the flash copy process using the command-line interface (CLI).

Starting the FlashCopy process will create a point-in-time copy of the data on the source VDisk and write it to the target VDisk for each mapping in the group.

Perform the following steps to prepare and trigger a FlashCopy consistency group:

1. Open a command prompt window.
2. Issue the **svctask prestartfcconsistgrp** command to prepare the FlashCopy Consistency Group before the copy process can be started (triggered). When you have assigned several mappings to a FlashCopy Consistency Group, you only have to issue a single prepare command for the whole group, to prepare all the mappings at once.

In our hypothetical scenario, the command to prepare a FlashCopy Consistency group called *maintobkpfcopy* is:

```
svctask prestartfcconsistgrp maintobkpfcopy
```

The group will enter the preparing state, and then move to the prepared state when it is ready. Issue the **svcinfo lsfcconsistgrp** command to check:

```
svcinfo lsfcconsistgrp -delim :
id:name:status
1:maintobkpfcopy:prepared
```

3. Issue the **svctask startfcconsistgrp** command to start (trigger) the FlashCopy Consistency Group to make the copy. You only have to issue a single start command for the whole group, to trigger all the mappings at once.

In our hypothetical scenario, the command to trigger a FlashCopy Consistency group called *maintobkpfcopy* is:

```
svctask startfcconsistgrp maintobkpfcopy
```

The group will enter the copying state, and then return to the idle_copied state when complete. You can issue the **svcinfo lsfcconsistgrp** command to check the state of the group:

```
svcinfo lsfcconsistgrp -delim :
id:name:state
1:maintobkpfcopy:copying
```

Use the **svcinfo lsfcmapprogress** command to check the progress of each mapping, main1copy and main2copy:

```
svcinfo lsfcmapprogress  -delim : main1copy
id:progress
0:100

svcinfo lsfcmapprogress  -delim : main2copy
id:progress
1:23
```

Finally issue the **svcinfo lsfcconsistgrp** command to display the detailed view of the group maintobkpfcopy, which returns to idle_copied state when both mappings have reached 100% progress:

```
svcinfo lsfcconsistgrp -delim :  maintobkpfcopy
id:1
name:maintobkpfcopy
state:idle_copied
FC_mapping_id:0
FC_mapping_name:main1copy
FC_mapping_id:1
FC_mapping_name:main2copy
```

You have now made a point-in-time copy of the data on mainvdisk1 which has been written to bkpvdisk1, and a copy of the data on mainvdisk2 which has been written to bkpvdisk2. The data on bkpvdisk1 and bkpvdisk2 will be visible to demohost2 because these VDisks are only mapped to demohost2.

Advanced functions with the CLI

Ensure that you are familiar with the advanced functions that you can perform using the command-line interface (CLI).

Related information

Chapter 4, “Command-line interface,” on page 151

Determining the WWPNs of a node using the CLI

You can determine the WWPNs of a node using the command-line interface (CLI).

Perform the following steps to determine the WWPNs of a node:

1. List the nodes in the cluster by issuing the following command:

```
svcinfo lsnode
```

Note: Remember the node name or ID as you will need it in the next step.

2. For the node or nodes in question, issue the following command:

```
svcinfo lsnode <nodename/id>
```

Note: Remember the four port ID's (WWPNs).

Determining the VDisk name from the vpath number on the host

You can determine the VDisk name from the vpath number on the host using the command-line interface (CLI).

Each VDisk exported by the SAN Volume Controller is assigned a unique vpath number. This number uniquely identifies the VDisk and can be used to determine which VDisk corresponds to the volume that the hosts sees. This procedure can only be performed using the command line interface.

Perform the following steps to determine the VDisk name from the vpath number:

1. For the volume in question, find the vpath serial number by issuing the following command:

```
datapath query device
```

2. Find the host object defined to the SAN Volume Controller that corresponds with the host you are working with.
   a. The WWPNs are an attribute of the HBA. You can find these by looking at the device definitions stored by your operating system. For example, on AIX they will be in the ODM, in Windows they will be in the Device Manager details for the given HBA.
   b. Verify which host object defined to the SAN Volume Controller that these ports belong to. The ports are stored as part of the detailed view, so you will need to list each host in turn by issuing the following:

   ```
   svcinfo lshost <name/id>
   ```

   where *<name/id>* is the name or ID of the host. Check for matching WWPNs.

   Note: You should name your hosts accordingly, for example, if the actual host is called *orange* you should also name the host object defined to the SAN Volume Controller as *orange*.

3. Now that you have the <host name> as defined to the SAN Volume Controller and the <vpath serial number>, issue the following command:

```
svcinfo lshostvdiskmap <hostname>
```

where *<hostname>* is the name of the host. A list is displayed.

4. Look for the VDisk UID that matches the <vpath serial number> and remember the VDisk name or ID.

Determining the host that a VDisk is mapped to

You can determine the host that a VDisk is mapped to using the command-line interface (CLI).

Perform the following steps to determine the host that the VDisk is mapped to:

1. Find the VDisk name or ID that you wish to check.
2. List the hosts that this VDisk is mapped, by issuing the following command:

   ```
   svcinfo lsvdiskhostmap <vdiskname/id>
   ```

   where *<vdiskname/id>* is the name or ID of the VDisk. A list is displayed.
3. Look for the host name or ID to determine which host this VDisk is mapped to. If no data is returned, the VDisk is not mapped to any hosts.

Determining the relationship between VDisks and MDisks using the CLI

You can determine the relationship between VDisks and MDisks using the command-line interface (CLI).

Every VDisk is constructed from one or more MDisks. At times you may need to determine the relationship between the two objects. The following procedure allows you to determine the relationships.

1. Perform the following steps to determine the relationship between VDisks and MDisks:
   a. For a given VDisk <vdiskname/id>, issue the following command: `svcinfo lsvdiskmember <vdiskname/id>` where *<vdiskname/id>* is the name or ID of the VDisk. This will return a list of IDs that correspond to the MDisks that make up the VDisk.
2. Perform the following steps to determine the relationship between VDisks and MDisks and the number of extents provided by each MDisk: If you wish more details, you can also determine the number of extents that make are being provided by each MDisk. This procedure can only be performed using the command line interface.
   a. For a given VDisk <vdiskname/id>, issue the following command:`svcinfo lsvdiskextent <vdiskname/id>` where *<vdiskname/id>* is the name or ID of the VDisk. This will return a table of MDisk IDs and the corresponding number of extents each MDisk is providing as storage for the given VDisk.
3. Perform the following steps to determine the relationship between MDisks and VDisks:
   a. For a given MDisk <mdiskname/id>, issue the following command: `svcinfo lsmdiskmember <mdiskname/id>` where *<mdiskname/id>* is the name or ID of the MDisk. This will return a list of IDs that correspond to the VDisks that are using this MDisk.
4. Perform the following steps to determine the relationship between MDisks and VDisks and the number of extents used by each VDisk: If you wish more details, you can also determine the number of extents that this MDisk is providing for each VDisk. This procedure can only be performed using the command line interface.
   a. For a given MDisk <mdiskname/id>, issue the following command: `svcinfo lsmdiskextent <mdiskname/id>` where *<mdiskname/id>* is the name or ID of the MDisk. This returns a table of VDisk IDs and the corresponding number of extents being used by each VDisk.

Determining the relationship between MDisks and RAID arrays or LUNs using the CLI

You can determine the relationship between MDisks and RAID arrays or LUNs using the command-line interface (CLI).

Each MDisk corresponds with a single RAID array, or a single partition on a given RAID array. Each RAID controller will define a LUN number for this disk. The LUN number and controller name or ID are needed to be able to determine the relationship between MDisks and RAID arrays or partitions.

Perform the following steps to determine the relationship between MDisks and RAID arrays:

1. Show the detailed view of the given MDisk <mdiskname>, by issuing the following command:

   ```
   svcinfo lsmdisk <mdiskname>
   ```

   Note: Remember the controller name or controller ID and controller LUN number.
2. Show the detailed view of the controller determined in by issuing the following command:

   ```
   svcinfo lscontroller <controllername>
   ```

   where *<controllername>* is the name of the controller.

   Note: Remember the vendor ID, product ID, and WWNN. Use these to determine what is being presented to the MDisk.
3. From the native user interface for the given controller, list the LUNs it is presenting and match the LUN number with that noted in 1. This will tell you the exact RAID array or partition that corresponds with the MDisk.

Increasing the size of your cluster using the CLI

You can increase the size of a cluster using the command-line interface (CLI).

To increase the size of your cluster you need to add nodes in pairs to a new I/O group. Your existing cluster may have become a bottleneck and so you wish to increase throughput by adding more nodes to the cluster.

Perform the following steps to increase the size of your cluster:

1. Perform the steps in the section and repeat this procedure for the second node.
2. If you wish to balance the load between the existing I/O groups and the new I/O groups, follow the procedure. Repeat this procedure for all VDisks you want to assign to the new I/O group.

   Related tasks

   "Adding a node to increase the size of a cluster using the CLI" on page 185
   You can add a node to increase the size of a cluster using the command-line interface.

   "Migrating a VDisk to a new I/O group" on page 185
   You can migrate a VDisk to a new I/O group to increase the size of your cluster using the command-line interface (CLI).

Adding a node to increase the size of a cluster using the CLI

You can add a node to increase the size of a cluster using the command-line interface.

CAUTION:

If you are adding a node that was previously removed from a cluster, 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. If you do not know the I/O group name or ID that it was removed from, contact IBM Service to add the node to the cluster without corrupting data.

Perform the following steps to add a node and increase the size of a cluster:

1. Issue the following command to verify that the node is detected on the fabric and to obtain the worldwide node names of the nodes on the cluster: `svcinfo lsnodecandidate`. Write down the worldwide node names for the next step.
2. Issue the following command to determine the I/O group where the node should be added: `svcinfo lsiogrp`. Write down name or ID of the first I/O group that has a node count of zero (0). You will need the ID for the next step.

   Record the following information for future reference:
   - Node serial number.
   - Worldwide node name.
   - All of the worldwide port names.
   - The name or ID of the I/O group that contains the node.
3. Issue the following command to add the node to the cluster: `svctask addnode -wwnodename <WWNN> -iogrp <newiogrpname/id> [-name <newnodename>]`, where <newnodename> is the name you want to assign to the node.
4. Issue the following command to verify that the node is online: `svcinfo lsnode`.

If the disk controller uses mapping to present RAID arrays or partitions to the cluster and the worldwide node names or the worldwide port names have changed, then you must modify the port groups that belong to the cluster.

Migrating a VDisk to a new I/O group

You can migrate a VDisk to a new I/O group to increase the size of your cluster using the command-line interface (CLI).

You can migrate a VDisk to a new I/O group to manually balance the workload across the nodes in the cluster. You may end up with a pair of nodes that are overworked and another pair that are under-worked. Follow this procedure to migrate a single VDisk to a new I/O group. Repeat for other VDisks as required.

Attention: This is a disruptive procedure, access to the VDisk will be lost while you follow this procedure. Under no circumstances should VDisks be moved to an offline I/O group. You must ensure the I/O group is online before moving the VDisks to avoid data loss scenarios.

Perform the following steps to migrate a single VDisk:

1. Quiesce all I/O operations for the VDisk. You may need to determine the hosts that are using this VDisk.
2. Before migrating the VDisk, it is essential that for each vpath presented by the VDisk you intend to move, the SDD configuration is updated to remove the vpaths in question. Failure to do this may result in data corruption. See *IBM*

TotalStorage Multipath Subsystem Device Driver: User's Guide for details about how to dynamically re-configure SDD for the given host operating system.

3. Any FlashCopy mappings or Metro Mirror relationships that use this VDisk should be stopped or deleted. Issue the following command, to check if the VDisk is part of a relationship or mapping:

   svcinfo lsvdisk <vdiskname/id> where <vdiskname/id> is the name or ID of the VDisk.
4. Look for the **FC_id** and **RC_id** fields. If these are not blank then the VDisk is part of a mapping or relationship. *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* for details on how to stop or delete the mapping or relationship.
5. Issue the following command to migrate the VDisk:

   ```
   svctask chvdisk -iogrp <newiogrpname/id> <vdiskname/id>
   ```
6. Follow the procedure to discover the new vpaths and to check that each vpath is now presenting the correct number of paths. See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* for details on how to dynamically re-configure SDD for the given host operating system.

Related tasks

"Determining the host that a VDisk is mapped to" on page 183
You can determine the host that a VDisk is mapped to using the command-line interface (CLI).

"Determining the relationship between VDisks and MDisks using the CLI" on page 183
You can determine the relationship between VDisks and MDisks using the command-line interface (CLI).

Replacing a faulty node in the cluster using the CLI

You can replace a faulty node in the cluster using the command-line interface (CLI).

Before you attempt to replace a faulty node with a spare node you must ensure that:

- SAN Volume Controller version 1.1.1 or later is installed on the cluster and on the spare node. To determine the version, you can review the vital product data on the node and the cluster or you can issue the **svcinfo lsnode** or **svcinfo lscluster** command.
- You know the name of the cluster that contains the faulty node.
- A spare node is installed in the same rack as the cluster that contains the faulty node.
- Make a record of the last five characters of the original worldwide node name (WWNN) of the spare node. You will need this information, if and when, you want to stop using this node as a spare node. In that case, you might prefer to use it as a normal node that can be assigned to any cluster.

  Perform the following steps to display and record the WWNN of the spare node:

  1. Display the node status on the front panel display of the node. See the topic, "SAN Volume Controller menu options" in the *IBM TotalStorage SAN Volume Controller: Service Guide* for more information.
  2. With the node status displayed on the front panel, press and hold the **Down** button. At the same time press and release the **Select** button. Then release the **Down** button. WWNN is displayed on line-1 of the display; line-2 of the display contains the last five characters of the WWNN.

3. Record the WWNN in a safe place. It will be needed if you want to stop using the spare node.

If a node fails, the cluster continues to operate with degraded performance until the faulty node is repaired. If the repair operation is likely to take an unacceptable amount of time, it might be useful to replace the faulty node with a spare node. However, the appropriate procedures must be followed and precautions must be taken in order not to interrupt I/O operations and to avoid compromising the integrity of your data. The procedures outlined in this topic involve changing the worldwide node name (WWNN) of a SAN Volume Controller. These procedures must be followed with care in order to avoid duplicate WWPNs which can cause data corruption.

Be aware that by performing these procedures the following changes will be made to your configuration:

Front Panel ID
: This number will change. It is the number that is printed on the front of the node and used to select the node that is to be added to a cluster.

Node Name
: This number might change. If you do not specify a name, the SAN Volume Controller assigns a default name when adding a node to a cluster. The SAN Volume Controller creates a new name each time a node is added to a cluster. If you choose to assign your own names then you need to type in the node name on the Adding a node to a cluster panel. If you are using scripts to perform management tasks on the cluster and those scripts use the node name, then by assigning the original name to a replacement node you avoid the need to make changes to the scripts.

Node ID
: This ID will change. A new node ID is assigned each time a node is added to a cluster; the node name remains the same following service activity on the cluster. You can use the node ID or the node name to perform management tasks on the cluster. However, if you are using scripts to perform those tasks use the node name rather than the node ID.

Worldwide Node Name
: This name will change. The WWNN is used to uniquely identify the node and the fibre-channel ports. The WWNN of the spare node will change to that of the faulty node. The node replacement procedures must be followed exactly to avoid any duplication of WWNNs.

Worldwide Port Names
: These names do not change. WWPNs are derived from the WWNN that is written to the spare (replacement) node as part of this procedure. For example, let’s say the WWNN for a node is `50050768010000F6`. The four WWPNs for this node would be derived as follows:

```
WWNN                                  50050768010000F6
WWNN displayed on front panel         000F6
WWPN Port 1                           50050768014000F6
WWPN Port 2                           50050768013000F6
WWPN Port 3                           50050768011000F6
WWPN Port 4                           50050768012000F6
```

Perform the following steps to replace a faulty node in the cluster:

1. Verify the name and ID of the node that you wish to replace.

   Perform the following steps to verify the name and ID:

a. Open a DOS window.
b. Type the **svcinfo lsnode** command and press **Enter**.

   If the node was faulty it will be shown as offline. Ensure the partner node in the I/O group is online.

   1) If the other node in the I/O group is offline, start the Directed Maintenance Procedures to determine the fault.
   2) If you have been directed here by the DMPs, and subsequently the partner node in the I/O group has failed.

   If you are replacing the node for other reasons, determine the node you wish to replace and again ensure the partner node in the I/O group is online.

   1) If the partner node is offline, you will lose access to the VDisks that belong to this I/O group if you continue. Start the Directed Maintenance Procedures and fix the other node before proceeding.

2. Find and record the following information about the faulty node:
   a. Node name
   b. I/O group name
   c. Last five characters of the WWNN
   d. Front panel ID
   e. UPS serial number
   a. To find and record the node name and I/O group name, type the **svcinfo lsnode** command and hit **Enter**.

      The faulty node will be offline.
   b. Record the following information about the faulty node:
      - Node name
      - I/O group name
   c. To find and record the last five characters of the WWNN, type the **svcinfo lsnodevpd <nodename>** command and hit **Enter**. <nodename> is the name that you recorded in step 1 on page 187.
   d. Find the **WWNN** field in the output. Record the last five characters of the WWNN.
   e. To find and record the front panel ID, type the **svcinfo lsnodevpd <nodename>** command and hit **Enter**. <nodename> is the name that you recorded in step 1 on page 187.
   f. Find the **front_panel_id** field in the output. Record the front panel ID.
   g. To find and record the UPS serial number, type the **svcinfo lsnodevpd <nodename>** command and hit **Enter**. <nodename> is the name that you recorded in step 1 on page 187.
   h. Find the **UPS_serial_number** field in the output. Record the UPS serial number.
3. Obtain the ID of the faulty node. Disconnect all four fibre-channel cables from the node.

   Important: Do not plug the fibre-channel cables into the spare node until spare node has been configured with the WWNN from the faulty node.
4. Connect the power and signal cables from the spare node to the UPS that has the serial number that you noted in step 5h.

Note: The signal cable can be plugged into any vacant position on the top row of serial connectors on the UPS. If no spare serial connectors are available on the UPS, disconnect the cables from the faulty SAN Volume Controller.

5. Power-on the spare node.
6. Display the node status on the service panel. See the topic ″SAN Volume Controller menu options″ in the *IBM TotalStorage SAN Volume Controller: Service Guide* for more information.
7. Change the WWNN of the spare node.

   Perform the following steps to change the WWNN of the spare node so that it matches the WWNN of the faulty node:

   a. With the node status displayed on the front panel, press and hold the **Down** button; press and release the **Select** button; release the **Down** button. `WWNN` is displayed on line-1 of the display; line-2 of the display contains the last five characters of the WWNN.
   b. With the WWNN displayed on the service panel, press and hold the **Down** button, press and release the **Select** button, release the **Down** button. This switches the display into edit mode. Change the displayed number to match the WWNN recorded in 5d on page 188.

      Note: To edit the displayed number use the **Up** and **Down** buttons to increase or decrease the numbers displayed. Use the **Left** and **Right** buttons to move between fields. When the five characters match the number recorded in step 1 on page 187, press the **Select** button twice to accept the number.
8. Connect the four fibre-channel cables that were disconnected from the faulty node to the spare node.
9. Having noted the <nodename> in step 1 on page 187, remove the node from the cluster by issuing the following **svctask rmnode <nodename/id>** command:

   Remember: 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 when the node is re-added to the cluster.
10. Issue the following command to add the node back into the cluster: **svctask addnode -wwnodename <WWNN> -iogrp <IOGRPNAME/ID> -name <NODENAME>**.
11. Use the Subsystem Device Driver (SDD) management tool on the host systems to verify that all paths are now online. See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* for more information.

    Attention: When the faulty node is repaired do not connect the fibre-channel cables to it. Connecting the cables might cause data corruption.
12. Repair the faulty node.
13. If you want to use the repaired node as a spare node, perform the following steps:

    a. Display the node status on the front panel display of the node. See the topic, ″SAN Volume Controller menu options″ in the *IBM TotalStorage SAN Volume Controller: Service Guide* for more information.

b. With the node status displayed on the front panel, press and hold the **Down** button; press and release the **Select** button; release the **Down** button. `WWNN` is displayed on line-1 of the display; line-2 of the display contains the last five characters of the WWNN.
c. With the WWNN displayed on the service panel, press and hold the **Down** button, press and release the **Select** button, release the **Down** button. This switches the display into edit mode. Change the displayed number to `00000`.

Note: To edit the displayed number use the **Up** and **Down** buttons to increase or decrease the numbers displayed. Use the **Left** and **Right** buttons to move between fields.

This SAN Volume Controller can now be used as a spare node.

Attention: Never connect a SAN Volume Controller with a WWNN of `00000` to the cluster. If this SAN Volume Controller is no longer required as a spare and is to be used for normal attachment to a cluster you must first use the procedure described in the prerequisites to change the WWNN to the number you recorded when a spare was created. Using any other number might cause data corruption.

14. Issue the following command to add the node back into the cluster: **svctask addnode -wwnodename <WWNN> -iogrp <IOGRPNAME/ID> -name <NODENAME>**.
15. Issue the **svcinfo lsnode** command to verify that the node is online.

Related tasks

"Deleting a node from a cluster using the CLI" on page 204
You can delete a node from a cluster using the command line interface (CLI).

"Adding nodes to a cluster using the CLI" on page 163
You can add nodes to a cluster using the CLI.

Related reference

"Configuring and servicing storage subsystems" on page 233
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

Recovering from offline VDisks after a node or an I/O group failed using the CLI

You can recover from an offline VDisk after a node or an I/O group has failed using the command-line interface (CLI).

If you have lost both nodes in an I/O group and have therefore lost access to all the VDisks that are associated with the I/O group, then you must perform one of the following procedures to regain access to the VDisks. Depending on the failure type, you may have lost data that was cached for these VDisks, therefore, they have gone offline.

Data loss scenario 1 One node in an I/O group failed and failover started on the second node. During this time, the second node in the I/O group fails before the data in the write cache has been written to disk. The first node is successfully repaired but its cache data is stale, therefore, it cannot be used. The second node is repaired or replaced and has lost its hard-end data, therefore, the node has no way of recognizing that it is part of the cluster.

Perform the following steps to recover from an offline VDisk:

1. Recover the node and include it back into the cluster.
2. Move all the offline VDisks to the recovery I/O group.
3. Move all the offline VDisks back to their original I/O group.

Data loss scenario 2 Both nodes in the I/O group have failed and have been repaired. The nodes have lost their hard-end data, therefore, the nodes have no way of recognizing that they are part of the cluster.

1. Move all the offline VDisks to the recovery I/O group
2. Move both recovered nodes back into the cluster
3. Move all the offline VDisks back to their original I/O group.

Related tasks

"Recovering a node and including it back into the cluster"
You can recover a node and including it back into the cluster using the command-line interface (CLI).

"Moving offline VDisks to the recovery I/O group" on page 192
You can move offline VDisks to the recovery I/O group using the command-line interface (CLI).

"Moving offline VDisks to their original I/O group using the CLI" on page 193
You can move offline VDisks to their original I/O group using the command-line interface (CLI).

Recovering a node and including it back into the cluster

You can recover a node and including it back into the cluster using the command-line interface (CLI).

After a node or an I/O group fails, you can use the following procedure to recover a node and include it back into the cluster.

Perform the following steps to recover a node and include it back into the cluster:

1. Verify that the node is offline. Issue the following command:

   ```
   svcinfo lsnode
   ```
2. Remove the old instance of the offline node from the cluster. Issue the following command:

   ```
   svctask rmnode <nodename/id>
   ```

   where *<NODENAME>* is the name of the node.
3. Verify that the node can be seen on the fabric. Issue the following command:

   `svcinfo lsnodecandidate` You should see the nodes listed as a candidate.

   Note: Remember the WWNNs for each node, you will need it in the following step.
4. If the nodes are repaired by replacing the front panel module or a node is repaired by replacing it with another node, then the WWNN for the node will change. In this case, the following additional steps are required:
   a. At the end of the recovery process it will be necessary to follow the SDD procedure to discover the new paths and to check that each vpath is now presenting the correct number of paths. See the *IBM TotalStorage Multipath*

Subsystem Device Driver: User's Guide sections on dynamic reconfiguration, specifically adding paths to existing vpaths.

b. You may also need to modify the configuration of your disk controllers. If your controller uses a mapping technique to present its RAID arrays or partitions to the cluster you will need to modify the port groups that belong to the cluster because the WWNN or WWPN's of the node have changed.

Attention: If more than one I/O group is affected, 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. This can avoid a possible data corruption exposure if the node must be removed from and re-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.

Attention: 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

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

5. Add the node back into the cluster. Issue the following command:

```
svctask addnode -wwnodename <WWNN> -iogrp
<IOGRPNAME/ID> [-name <NODENAME>]
```

where *<WWNN>* is the worldwide node name *<IOGRPNAME/ID>* is the I/O group name or ID.

6. Verify that the node is online. Issue the following command:

```
svcinfo lsnode
```

Related tasks

"Recovering from offline VDisks after a node or an I/O group failed using the CLI" on page 190
You can recover from an offline VDisk after a node or an I/O group has failed using the command-line interface (CLI).

Moving offline VDisks to the recovery I/O group

You can move offline VDisks to the recovery I/O group using the command-line interface (CLI).

Perform the following steps to move offline VDisks to the recovery I/O group:

1. List all VDisks that are offline and belong to the I/O group in question. Issue the following command:

```
svcinfo lsvdisk -filtervalue IO_group_name=
<IOGRPNAME/ID>:status=offline
```

Note: <IOGRPNAME> = the name of the I/O group that failed.

2. For each VDisk returned, move the VDisk to the recovery I/O group. Issue the following command:

```
svctask chvdisk -iogrp recovery_io_grp -force
<vdiskname/ID>
```

Note: <vdiskname/ID> = the name of one of the VDisks that are offline.

Related tasks

"Recovering from offline VDisks after a node or an I/O group failed using the CLI" on page 190
You can recover from an offline VDisk after a node or an I/O group has failed using the command-line interface (CLI).

Moving offline VDisks to their original I/O group using the CLI

You can move offline VDisks to their original I/O group using the command-line interface (CLI).

After a node or an I/O group fails, you can use the following procedure to move offline VDisks to their original I/O group.

Attention: Under no circumstances should VDisks be moved to an offline I/O group. Ensure the I/O group is online before moving the VDisks back to avoid any further data loss.

Perform the following steps to move offline VDisks to their original I/O group:

1. For each VDisk, move the VDisk back into the original I/O group. Issue the following command:

```
svctask chvdisk -iogrp <IOGRPNAME/ID> -force
<vdiskname/ID>
```

Note: <IOGRPNAME> = the name of the I/O group that failed.

2. Verify that the VDisks are now online. Issue the following command:

```
svcinfo lsvdisk -filtervalue IO_group_name=
<IOGRPNAME/ID>
```

Note: <vdiskname/ID> = the name of one of the VDisks that are offline.

Related tasks

"Recovering from offline VDisks after a node or an I/O group failed using the CLI" on page 190
You can recover from an offline VDisk after a node or an I/O group has failed using the command-line interface (CLI).

Replacing an HBA in a host using the CLI

You can replace an HBA in a host using the command-line interface (CLI).

This procedure describes how to notify the SAN Volume Controller of a change to a defined host object. It is sometimes necessary to replace the HBA that connects the host to the SAN, at this time you must notify the SAN Volume Controller of the new WWPN's that this HBA contains.

Ensure that your switch is zoned correctly.

Perform the following steps to replace an HBA in a host using the CLI:

1. Issue the following command to list the candidate HBA ports:

   ```
   svcinfo lshbaportcandidate
   ```

   You should see a list of the HBA ports that are available to be added to host objects. One or more of these should correspond with the one or more WWPNs that belong to the new HBA.
2. Locate the host object that corresponds with the host in which you have replaced the HBA. The following command lists all the defined host objects:

   ```
   svcinfo lshost
   ```

   To list the WWPNs currently assigned to the host, issue the following: `svcinfo lshost <hostobjectname>`where *<hostobjectname>* is the name of the host object.
3. Add the new ports to the existing host object by issuing the following command:

   ```
   svctask addhostport -hbawwpn <one or more existing WWPNs
   separated by :> <hostobjectname/ID>
   ```
4. Remove the old ports from the host object by issuing the following command:

   ```
   svctask rmhostport -hbawwpn <one or more existing WWPNs
   separated by :> <hostobjectname/ID>
   ```

   where *<one or more existing WWPNs separated by :>* correspond with those listed in step 2 that belong to the old HBA that has been replaced.
5. Any mappings that exist between the host object and VDisks will automatically be applied to the new WWPNs. Therefore, the host should see the VDisks as the same SCSI LUNs as before.
6. See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* for additional information about dynamic reconfiguration.

Related reference

"Configuring and servicing storage subsystems" on page 233
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

Expanding VDisks

You can expand a VDisk from the Expanding VDisks panel or by using the command-line interface (CLI).

This topic lists the supported platforms and requirements if this feature is to be used

A VDisk that is not yet mapped to any hosts and hence does not contain any customer data can be expanded at any time. However, if the VDisk contains data that is being used, only AIX and Windows 2000 hosts can cope with a VDisk being expanded.

The following matrix shows the supported platforms and requirements if this feature is to be used:

Table 15. Supported platforms and requirements

Platform	Supported	Requirement
AIX	Yes	AIX 5.2 onwards only
HP-UX	No	
Linux	No	
SUN Solaris	No	
Windows NT	No	
Windows 2000, 2003	Yes	

Related tasks

"Expanding a virtual disk that is mapped to an AIX host"
You can expand a virtual disk (VDisk) that is mapped to an AIX host.

"Expanding a virtual disk that is mapped to a Windows 2000 host" on page 196
You can expand a virtual disk (VDisk) that is mapped to a Windows 2000 host.

Expanding a virtual disk that is mapped to an AIX host

You can expand a virtual disk (VDisk) that is mapped to an AIX host.

VDisks that are mapped for FlashCopy or that are in Metro Mirror relationships cannot be expanded.

Determine the exact size of the source or master VDisk by issuing the following command-line interface (CLI) command:

```
svcinfo lsvdisk -bytes <vdiskname>
```

This feature can be used in two ways:

- To increase the capacity available on a particular VDisk that is already mapped to a host.
- To increase the size of a VDisk so that it matches the size of the source or master VDisk and can be used in a FlashCopy mapping or Metro Mirror relationship.

Perform the following steps to expand a VDisk that is mapped to an AIX host:

1. Determine the VDisk you wish to expand and remember its <vdiskname>.
2. Verify that this VDisk is mapped to an AIX host.
3. Determine the volume group that contains the VDisk (you must know the VDisk to hdisk relationship)
4. Quiesce all I/O operations to **all** volumes that belong to the volume group and sync the file-systems mounted on this volume group.
5. Check the current type of the VDisk by viewing the VDisk details in the Work with VDisks panel.

   Note:

   a. If the VDisk has a type of image, it cannot be expanded.
   b. If the VDisk has a type of sequential, it becomes a striped VDisk when you expand it.
6. Deactivate the volume group that contains this VDisk. Issue the following command from the command prompt:

   ```
   varyoffvg <volume_group>
   ```
7. Expand the VDisk using either of the following methods:

   a. From the Work with VDisks panel, select the VDisk and select the Expand task. Enter the capacity by which you wish to extend this VDisk and the select the appropriate units. Select one, more, or all of the MDisks from the list. These will be the MDisks that provide the extra capacity. Optionally, select the format check-box if you want this extra capacity to be formatted before use.
   b. From the command prompt issue the following command:

      ```
      svctask expandvdisksize
      ```

8. Reactivate the volume group so that the change in size is detected by the HBA device driver. Issue the following command from the command prompt:

   ```
   varyonvg <volume_group>
   ```

9. Run the **change volume group** command to notify the LVM that the size has changed. Issue the following command from the command prompt:

   ```
   chvg -g <volume_group>
   ```

10. Expand all the file-systems that are mounted on this VDisk (or use the new capacity as required)

Restart I/O operations to the volume group.

Expanding a virtual disk that is mapped to a Windows 2000 host

You can expand a virtual disk (VDisk) that is mapped to a Windows 2000 host.

VDisks that are mapped for FlashCopy or that are in Metro Mirror relationships cannot be expanded.

Ensure that you have run Windows Update and have applied all recommended updates to your system prior to attempting to expand a VDisk that is mapped to a Windows 2000 host

Determine the exact size of the source or master VDisk by issuing the following command-line interface (CLI) command:

```
svcinfo lsvdisk -bytes <vdiskname>
```

This feature can be used in two ways:
- To increase the capacity available on a particular VDisk that is already mapped to a host.
- To increase the size of a VDisk so that it matches the size of the source or master VDisk and can be used in a FlashCopy mapping or Metro Mirror relationship.

VDisks can be expanded under Windows 2000 concurrently with I/O operations.

Perform the following steps to expand a VDisk that is mapped to a Windows 2000 host:
1. Expand the VDisk using either of the following methods:
   a. From the Work with VDisks panel, select the VDisk and select the Expand task. Enter the capacity by which you wish to extend this VDisk and the select the appropriate units. Select one, more, or all of the MDisks from the list. These will be the MDisks that provide the extra capacity. Optionally, select the format check-box if you want this extra capacity to be formatted before use.
   b. From the command prompt issue the following command:

      ```
      svctask expandvdisksize
      ```

2. On the Windows Host, start the Computer Management application and open the Disk Management window under the Storage branch.

You will see the VDisk that you expanded now has some unallocated space at the end of the disk.

Dynamic disks can be expanded without stopping I/O operations in most cases. However, in some applications the operating system may report I/O errors. When this problem occurs, either of the following entries may be recorded in the System event log:

```
Event Type: Information
Event Source: dmio
Event Category: None
Event ID: 31
Description: dmio:
Harddisk0 write error at block ######## due to
disk removal

Event Type: Information
Event Source: dmio
Event Category: None
Event ID: 34
Description: dmio:
Harddisk0 is re-online by PnP
```

Attention: This is a known problem with Windows 2000 and is documented at the Microsoft knowledge base as article Q327020. If either of these errors are seen, run Windows Update and apply the recommended fixes to resolve the problem.

Restart the Computer Management application if it was opened prior to expanding the VDisk.

If the disk is a Windows basic disk you can create a new primary or extended partition from the unallocated space.

If the disk is a Windows dynamic disk you can use the unallocated space to create a new volume (simple, striped, mirrored) or add it to an existing volume.

Related concepts

"Virtual disks" on page 29
A *virtual disk (VDisk)* is a logical disk that the cluster presents to the storage area network (SAN).

Shrinking a VDisk using the CLI

You can shrink a VDisk using the command-line interface (CLI).

VDisks can be reduced in size should it be required. However, if the VDisk contains data that is being used, **under no circumstances should you attempt to shrink a VDisk without first backing up your data**. The SAN Volume Controller arbitrarily reduces the capacity of the VDisk by removing a partial, one or more extents from those allocated to the VDisk. You cannot control which extents are removed and so you cannot guarantee that it is unused space that is removed.

Attention: This feature should *only* be used to make a target or auxiliary VDisk the same size as the source or master VDisk when creating FlashCopy mappings or Metro Mirror relationships. You should also ensure that the target VDisk is not mapped to any hosts prior to performing this operation.

Perform the following steps to shrink a VDisk:

1. Validate that the VDisk is not mapped to any host objects. If the VDisk is mapped, data is displayed.
2. You can determine the exact capacity of the source or master VDisk. Issue the following command:

```
svcinfo lsvdisk -bytes <vdiskname>
```

3. Shrink the VDisk by the required amount. Issue the following command:

```
svctask shrinkvdisksize -size <capacitytoshrinkby> -unit
<unitsforreduction> <vdiskname/ID>
```

Migrating extents using the CLI

To improve performance, you can migrate extents using the command-line interface (CLI).

The SAN Volume Controller provides various data migration features. These can be used to move the placement of data both *within* MDisk groups and *between* MDisk groups. These features can be used concurrent with I/O operations. There are two ways in which you can migrate data:

1. Migrating data (extents) from one MDisk to another (within the same MDisk group). This can be used to remove hot or overutilized MDisks.
2. Migrating VDisks from one MDisk group to another. This can be used to remove hot MDisk groups, for example, reduce the utilization of a group of MDisks.
3. The source MDisk must not currently be the source MDisk for any other migrate extents operation.
4. The destination MDisk must not be the destination MDisk for any other migrate extents operation.

You can determine the usage of particular MDisks by gathering I/O statistics about MDisks and VDisks. Once you have gathered this data, you can analyze it to determine which MDisks are hot. The procedure then takes you through querying and migrating extents to elsewhere in the same MDisk group. This procedure can only be performed using the command line tools.

To migrate extents to remove possible problems, perform the following:

1. Isolate any MDisks that are overutilized. You can determine this by requesting an I/O statistics dump and analyzing the output. To start I/O statistics gathering, issue the following:

```
svctask startstats -interval 15
```

2. This will generate a new I/O statistics dump file approximately every 15 minutes. Wait for at least 15 minutes after issuing the **svctask startstats** command and then issue the following:

```
svcinfo lsiostatsdumps
```

   This will list the I/O statistics files that have been generated. These are prefixed with `m` and `Nm` for MDisk statistics and `v` for VDisk statistics.

3. Use secure copy (scp) to retrieve the dumps files to analyze. For example, issue the following:

```
<AIX HOST PROMPT#>scp <clusterip>:/dumps/iostats/m_*
```

This will copy all the MDisk statistics files to the AIX host in the current directory.

4. Analyze the dumps to determine which MDisks are hot. It may be helpful to also determine which VDisks are being heavily utilized as you can spread the data they contain more evenly across all the MDisks in the group using the procedure below.
5. Stop the statistics collection again by issuing the following command:

```
svctask stopstats
```

Once you have determined which MDisks are hot, you can migrate some of the data onto some less hot MDisks within the same MDisk group.

1. Determine the number of extents that are in use by each VDisk for the given MDisk. Issue the following command:

   ```
   svcinfo lsmdiskextent <mdiskname>
   ```

   This will return the number of extents that each VDisk is using on the given MDisk. You should pick some of these to migrate elsewhere in the group.
2. Determine the other MDisks that reside in the same MDisk group.
   a. To determine the MDisk group that the MDisk belongs to, issue the following command:

      ```
      svcinfo lsmdisk <mdiskname/ID>
      ```

      Look for the mdisk_grp_name attribute.
   b. List the MDisks in the group by issuing the following command:

      ```
      svcinfo lsmdisk -filtervalue mdisk_grp_name=<mdiskgrpname>
      ```
3. Select one of these MDisks as the target MDisk for the extents. You can determine how many free extents exist on an mdisk by issuing the following command:

   svcinfo lsfreeextents <mdiskname>

   You can issue the **svcinfo lsmdiskextent <newmdiskname>** command for each of the target MDisks to ensure that you are not just moving the over-utilization to another MDisk. Check that the VDisk that owns the set of extents to be moved, (see step 1), does not already own a large set of extents on the target MDisk.
4. For each set of extents, issue the following command to move them to another MDisk:

   ```
   svctask migrateextents -source <mdiskname/ID> -exts
   <num_extents_from_step1> -target <newmdiskname/ID>
    -threads 4 <vdiskid_returned_from_step1>
   ```

   where *<num_extents_from_step1>* is the number of extents on the *<vdiskid_returned_from_step1>*, that is, the data that is returned from the command issued in step 1. *<newmdiskname/ID>* is the name or ID of the MDisk to which you want to migrate this set of extents.
5. Repeat steps 2 to 4 for all the sets of extents you wish to move.

6. You can check the progress of the migration(s) by issuing the following command:

   ```
   svcinfo lsmigrate
   ```

   Related tasks

   "Migration methods" on page 111
   Several methods can be used to migrate image mode VDisks into managed mode VDisks.

Migrating VDisks between MDisk groups using the CLI

You can migrate VDisks between MDisk groups using the command-line interface (CLI).

You can determine the usage of particular MDisks by gathering I/O statistics about MDisks and VDisks. Once you have gathered this data, you can analyze it to determine which VDisks or MDisks are hot. This procedure then takes you through migrating VDisks from one MDisk group to another.

When a migrate command is issued, a check is made to ensure that the destination of the migrate has enough free extents to satisfy the command. If it does, the command proceeds, but will take some time to complete. During this time, it is possible for the free destination extents to be consumed by another process, for example, by creating a new VDisk in the destination MDisk group or by starting more migrate commands. In this scenario, when all the destination extents have been allocated the migration commands suspend and an error is logged (error id 020005). There are two methods for recovering from this situation:

1. Add additional MDisks to the target MDisk group. This provides additional extents in the group and allows the migrations to be restarted. You will need to mark the error as fixed in order to reattempt the migration.
2. Migrate one or more VDisks that are already created from the MDisk group to another group. This will free up extents in the group and allow the original migrations to be restarted (again by marking the error as fixed).

Perform the following steps to migrate VDisks between MDisk groups:

1. Isolate any VDisks that are overutilized. You can determine this by requesting an I/O statistics dump and analyzing the output. To start I/O statistics gathering, issue the following command:

   ```
   svctask startstats -interval 15
   ```

2. This will generate a new I/O statistics dump file approximately every 15 minutes. Wait for at least 15 minutes after issuing the **svctask startstats** command and then issue the following command:

   ```
   svcinfo lsiostatsdumps
   ```

   This will list the I/O statistics files that have been generated. These are prefixed with `m` and `Nm` for MDisk statistics and `v` for VDisk statistics.
3. Use secure copy (scp) to retrieve the dumps files for analyzing. For example, issue the following:

   ```
   <AIX HOST PROMPT#>scp <clusterip>:/dumps/iostats/v_*
   ```

   This will copy all the VDisk statistics files to the AIX host in the current directory.

4. Analyze the dumps to determine which VDisks are hot. It may be helpful to also determine which MDisks are being heavily utilized as you can spread the data they contain more evenly across all the MDisks in the group by migrating the extents.
5. Stop the statistics collection again. Issue the following command:

   ```
   svctask stopstats
   ```

   Once you have analyzed the I/O statistics data, you can determine which VDisks are hot. You also need to determine which MDisk group you wish to move this VDisk to. Either create a new MDisk group or determine an existing group that is not yet over utilized. You can do this by checking the I/O statistics files generated above and ensuring that the MDisks or VDisks in the target MDisk group are less utilized than the source group.
6. After having determined which VDisk you wish to migrate, and the new MDisk group you wish to migrate it to, issue the following command:

   ```
   svctask migratevdisk -vdisk <vdiskname/ID> -mdiskgrp
    <newmdiskgrname/ID> -threads 4
   ```
7. You can check the progress of the migration by issuing the following command:

   ```
   svcinfo lsmigrate
   ```

Migrating a VDisk between I/O groups using the CLI

Ensure that you are familiar with migrating a VDisk between I/O groups.

Attention: These migration tasks are disruptive, in that the cached data held within the cluster must first be written to disk, then the allocation of the VDisk can be changed.

Modifying the I/O group that services the virtual disk cannot be done concurrently with I/O operations. It also requires a re-scan at the host level to ensure that SDD gets notified that the allocation of the preferred node has changed and the ports by which the virtual disk is accessed has changed. This should only be done in the situation where one pair of nodes has become over utilized.

Perform the following steps to migrate a VDisk between I/O groups:
1. Sync all file-systems that are mounted on the given virtual disk.
2. Stop all I/O operations to the virtual disk.
3. Type the following:

   ```
   svctask chvdisk -iogrp <new_io_grp_name_or_id>
   <vdisk>
   ```
4. Issue the SDD command to resync the VDisk to host mapping. See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* for more information.
5. Restart the I/O operations to the virtual disk.

Creating an image mode VDisk using the CLI

You can import storage that contains existing data and continue to use this storage but make use of the advanced functions, such as Copy Services, data migration, and the cache. These disks are known as image mode virtual disks (VDisks).

Make sure you are aware of the following before you create image mode VDisks:

1. Unmanaged-mode managed disks (MDisks) that contain existing data cannot be differentiated from unmanaged-mode MDisks that are blank. Therefore, it is vital that you control the introduction of these disks to the cluster. It is recommended that you introduce these disks one at a time. For example, map a single LUN from your RAID controller to the cluster and refresh the view of managed disks. The newly detected disk is displayed.
2. *Do not* add an unmanaged-mode MDisk that contains existing data to an MDisk group manually. If you do, the data will be lost. When you use the command to convert an image mode virtual disk from an unmanaged-mode disk, you will select the MDisk group where it should be added.

Go to the following Web site for more information:

http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Perform the following steps to create an image mode virtual disk. For complete instructions on the CLI commands, refer to the IBM TotalStorage SAN Volume Controller: Command-Line Interface User's Guide.

1. Stop all I/O operations from the hosts. Un-map the logical disks that contain the data from the hosts.
2. Create one or more managed disk (MDisk) groups. Use the **svcinfo mkmdiskgrp** command. Ensure that the MDisk groups have enough free capacity to contain all of the migrating data.
3. Map a single RAID array or logical unit from your RAID controller to the cluster. You can do this through a switch zoning or a RAID controller based on your host mappings. The array or logical unit will appear as an unmanaged-mode MDisk to the SAN Volume Controller.
4. Use the **svcinfo lsmdisk** command to list the unmanaged-mode MDisks.

   If the new unmanaged-mode MDisk is not listed, you can perform a fabric-level discovery. Use the **svctask detectmdisk** command to scan the fibre-channel network for the unmanaged-mode MDisks.
5. Convert the unmanaged-mode MDisk to an image mode virtual disk. Use the **svctask mkvdisk** command to create an image mode virtual disk object.
6. Map the new virtual disk to the hosts that were previously using the data that the MDisk now contains. Use the **svctask mkvdiskhostmap** command to create a new mapping between a virtual disk and a host. This makes the image mode virtual disk accessible for I/O operations to the host.

Once it is mapped to a host object, the image mode virtual disk is detected as a disk drive with which the host can perform I/O operations.

If you want to virtualize the storage on an image mode VDisk, you can transform it into a striped virtual disk. Migrate the data on the image mode VDisk to managed-mode disks in another MDisk group. Use the **svctask migratevdisk** command to migrate an entire image mode VDisk from one MDisk group to another MDisk group. For instructions on this command, refer to the IBM TotalStorage SAN Volume Controller: Command-Line Interface User's Guide.

Migrating to an image mode virtual disk using the CLI

You can migrate to an image mode virtual disk using the command-line interface (CLI).

The **svctask migratetoimage** command allows you to migrate the data from an existing virtual disk onto a different managed disk.

When the **svctask migratetoimage** command is issued, it will migrate the data of the user specified source virtual disk onto the managed disk specified as the target. At completion of the command, the virtual disk will be classified as an image mode disk.

The managed disk specified as the target must be in an unmanaged state at the time the command is run. Execution of this command will result in the inclusion of the MDisk into the user specified MDisk group.

Issue the

```
svctask migratetoimage -vdisk <vdiskname/ID>
-mdisk <newmdiskname/ID> -mdiskgrp <newmdiskgrpname/ID>
 -threads 4
```

where *<vdiskname/ID>* is the name or ID of the VDisk, *<newmdiskname/ID>* is the name or ID of the new MDisk, and *<newmdiskgrpname/ID>* is the name or ID of the new MDisk group name.

Advanced function cluster overview using the CLI

The command-line interface (CLI) offers advanced functions for your cluster.

Overview

The following sections detail the advanced cluster functions that you can perform using the CLI.

Related tasks

"Deleting a node from a cluster using the CLI" on page 204
You can delete a node from a cluster using the command line interface (CLI).

"Performing the cluster maintenance procedure using the CLI" on page 205
You can perform the cluster maintenance procedure using the command-line interface (CLI).

"Modifying IP addresses using the CLI" on page 206
You can modify IP addresses using the command-line interface (CLI).

"Maintaining SSH keys using the CLI" on page 206
You can maintain SSH keys using the command-line interface (CLI).

"Setting up error notifications using the CLI" on page 206
You can set up error notifications using the command-line interface (CLI).

"Modifying passwords using the CLI" on page 207
You can modify the admin and service passwords using the command-line interface (CLI).

"Listing log or dump files using the CLI" on page 207
You can list log or dump files using the command-line interface (CLI).

"Changing the language setting using the CLI" on page 208
You can change the language settings using the command-line interface (CLI).

"Viewing the feature log using the CLI" on page 209
You can view the feature log using the command-line interface (CLI).

"Analyzing the error log using the CLI" on page 209
You can analyze the error log using the command-line interface (CLI).

"Shutting down a cluster or single node using the CLI" on page 210
You can shut down a cluster using the command-line interface (CLI).

Deleting a node from a cluster using the CLI

You can delete a node from a cluster using the command line interface (CLI).

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

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.

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 on page 205.
   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 on page 205 you should power off the node that you intend to remove, unless this is the last node in the cluster. This ensures that SDD does not rediscover 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 SAN Volume Controller Console 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 Multipath Subsystem Device Driver: User's Guide* for details about how to dynamically re-configure SDD for the given host operating system.

Perform the following steps to delete a node:
1. Open a command prompt window.

   Note:
   a. Before removing a node, be sure that this is what you want to do. All of the VDisks in the node will be assigned to the remaining node in the I/O group. That is, the preferred node will be changed. You cannot change this setting once this is done. Also, all VDisks will go into write-through cache mode because 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 prompted to force the deletion.
   c. If this is the last node in the cluster or it is the configuration node, all connections to the cluster will be lost. The user interface and any open CLI sessions will become unavailable. This can result in a timeout from the command, because 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.

   Related reference

   "Advanced function cluster overview using the CLI" on page 203
   The command-line interface (CLI) offers advanced functions for your cluster.

Performing the cluster maintenance procedure using the CLI

You can perform the cluster maintenance procedure using the command-line interface (CLI).

Perform the following steps for cluster maintenance:
1. Open a command prompt window.
2. Issue the **svctask finderr** command to analyze the error log for the highest severity of unfixed errors. This command scans the error log for any unfixed errors. Given a priority ordering defined within the code, the highest priority of unfixed errors is returned.
3. Issue the **svctask dumperrlog** command to dump the contents of the error log to a text file.
4. Locate the error and fix.
5. Issue the **svctask clearerrlog** command to clear all entries from the error log including status events and any unfixed errors.

   Note: Clearing the error log will not fix the errors.

   Attention: You should only use this command when you have either rebuilt the cluster, or have fixed a major problem that has caused many entries in the error log that you do not want to fix individually.
6. Issue the **svctask cherrstate** command to change the state of an error. The state can be changed from `unfixed` to `fixed`, or `fixed` to `unfixed`.

   Related reference

“Advanced function cluster overview using the CLI” on page 203
The command-line interface (CLI) offers advanced functions for your cluster.

Modifying IP addresses using the CLI

You can modify IP addresses using the command-line interface (CLI).

Perform the following steps to modify IP addresses:

1. Open a command prompt window.
2. Issue the **svcinfo lscluster** command to list the IP address of the cluster.
3. Issue the **svctask chcluster** command to modify the IP address. This command enables you to change the settings for the following:
   a. Cluster IP address
   b. Subnet mask
   c. Gateway

   If you specify a new cluster IP address, the existing communication with the cluster is broken.

Related reference

“Advanced function cluster overview using the CLI” on page 203
The command-line interface (CLI) offers advanced functions for your cluster.

Maintaining SSH keys using the CLI

You can maintain SSH keys using the command-line interface (CLI).

Attention: After you add a cluster, close the Maintaining SSH Keys panel.

Perform the following steps to maintain SSH keys:

1. Open a command prompt window.
2. Issue the **svcinfo lssshkeys** command to list the SSH keys that are available on the cluster.
3. Issue the **svctask addsshkey** command to install a new SSH key on the cluster. The key file must first be copied onto the cluster. Each key is associated with an ID string that you define that can consist of up to 30 characters. Up to 100 keys can be stored on a cluster. You can add keys to provide either administrator access or service access. For example, type the following:

   ```
   svctask addsshkey -user service -file /tmp/id_rsa.pub -label testkey
   ```

   where */tmp/id_rsa.pub* is the name of the file that the SSH key will be saved in and *testkey* is the label to associate with this key.
4. You can issue the **svctask rmsshkey** command to remove an SSH key from the cluster.
5. You can issue the **svctask rmallsshkeys** command to remove all of the SSH keys from the cluster.

Related reference

“Advanced function cluster overview using the CLI” on page 203
The command-line interface (CLI) offers advanced functions for your cluster.

Setting up error notifications using the CLI

You can set up error notifications using the command-line interface (CLI).

Perform the following steps to set up error notifications:

1. Open a command prompt window.

2. Issue the **svctask setevent** command to specify what you like to happen when an error or event is logged to the error log. You can select whether the cluster raises an SNMP trap, issues an e-mail notification for entries that are added to the cluster error or event log, or both. Three levels of notification are possible:
   a. **None** No error or status changes will be sent.
   b. **Hardware_only** You will be notified or errors, but you will not be notified of status changes.
   c. **All** You will be notified of all errors and status changes.

   If you have an SNMP manager installed or if you want to be notified by e-mail of errors or events, you should enable error notification. The notification levels for SNMP and e-mail alerts can be set independently. If you choose **All** or **Hardware_only** notification, you must select a destination for the notification.

   Related reference

   "Advanced function cluster overview using the CLI" on page 203
   The command-line interface (CLI) offers advanced functions for your cluster.

Modifying passwords using the CLI

You can modify the admin and service passwords using the command-line interface (CLI).

Passwords only affect access to the cluster via the SAN Volume Controller Console. To restrict access to the command line interface (CLI) you must control the list of SSH client keys installed on the cluster.

Perform the following steps to modify the passwords:

1. Open a command prompt window.
2. Issue the following command to change the administrator users password:

   ```
   svtask chcluster -admpwd <admin_password>
   ```
3. Issue the following command to change the service users password:

   ```
   svtask chcluster -servicepwd <service_password>
   ```

   Related tasks

   "Maintaining passwords using the CLI" on page 161
   You can maintain passwords using the command-line interface (CLI).

   Related reference

   "Advanced function cluster overview using the CLI" on page 203
   The command-line interface (CLI) offers advanced functions for your cluster.

Listing log or dump files using the CLI

You can list log or dump files using the command-line interface (CLI).

Perform the following steps to list log or dump files:

1. Open a command prompt window.
2. You can issue any of the following commands to list error log files:
   a. **svcinfo lserrlogbymdisk**
   b. **svcinfo lserrlogbymdiskgroup**
   c. **svcinfo lserrlogbyvdisk**
   d. **svcinfo lserrlogbyhost**
   e. **svcinfo lserrlogbynode**
   f. **svcinfo lserrlogbyiogrp**

g. **svcinfo lserrlogbyfcconsistgrp**
h. **svcinfo lserrlogbyfcmap**
i. **svcinfo lserrlogbyrcconsistgrp**
j. **svcinfo lserrlogbyrcrelationship**

These commands will list the error log by type. These commands will return a list of dumps in the appropriate directory. For example, issue the **svcinfo lserrlogbymdisk** command, displays the error log by MDisks.

You can display the whole log or filter the log so that only errors, events, or unfixed errors are displayed. In addition, you can request the output to be sorted either by error priority or by time. For error priority, the most serious errors are the lowest-numbered errors. They are, therefore, displayed first in the table. For time, either the older or the latest entry can be displayed first in the output.

3. You can issue any of the following command to list dump files.
   a. **svcinfo lsconfigdumps**
   b. **svcinfo lserrlogdumps**
   c. **svcinfo lsfeaturedumps**
   d. **svcinfo lsiostatsdumps**
   e. **svcinfo lsiotracedumps**
   f. **svcinfo lssoftwaredumps**
   g. **svcinfo ls2145dumps**

   These commands will list the dump file by type. These commands will return a list of dumps in the appropriate directory. For example, issue the **svcinfo lsconfigdumps** command, a list of dumps for configurations will be stored in the /dumps/configs destination directory.

   The software dump files contain dumps of the SAN Volume Controller memory. Your service representative might ask for these dumps to debug problems. The software dumps are large files (approximately 300 MB). Consider copying these files to your host using secure copy (scp) methods.

Related reference

"Advanced function cluster overview using the CLI" on page 203
The command-line interface (CLI) offers advanced functions for your cluster.

Changing the language setting using the CLI

You can change the language settings using the command-line interface (CLI).

Perform the following steps to change the language settings:

1. Open a command prompt window.
2. Issue the **svcservicetask setlocale** command to change the locale setting for the cluster. It changes all interfaces output to the chosen language. For example, if you wanted to change the English default language to Japanese, type the following:

```
svcservicetask setlocale -locale 3
```

   where *3* is the argument that stands for Japanese. The arguments are:
   a. **0** US English (default)
   b. **1** Chinese (simplified)
   c. **2** Chinese (traditional)
   d. **3** Japanese
   e. **4** Korean

f. **5** French
g. **6** German
h. **7** Italian
i. **8** Spanish
j. **9** Portuguese (Brazilian)

Note: This command does not change the front panel display panel settings.

Related reference

"Advanced function cluster overview using the CLI" on page 203
The command-line interface (CLI) offers advanced functions for your cluster.

Viewing the feature log using the CLI

You can view the feature log using the command-line interface (CLI).

Perform the following steps to view the feature log:

1. Open a command prompt window.
2. Issue the **svcinfo lsfeaturedumps** command to return a list of dumps in the /dumps/feature destination directory. The feature log is maintained by the cluster. The feature log records events that are generated when license parameters are entered or when the current license settings have been breached.
3. Issue the **svcservicemodeinfo lsfeaturedumps** command to return a list of the files that exist of the type specified on the given node.

Related reference

"Advanced function cluster overview using the CLI" on page 203
The command-line interface (CLI) offers advanced functions for your cluster.

Analyzing the error log using the CLI

You can analyze the error log using the command-line interface (CLI).

Perform the following steps to analyze the error log:

1. Open a command prompt window.
2. You can issue any of the following commands to list error log files:
   a. **svcinfo lserrlogbymdisk**
   b. **svcinfo lserrlogbymdiskgroup**
   c. **svcinfo lserrlogbyvdisk**
   d. **svcinfo lserrlogbyhost**
   e. **svcinfo lserrlogbynode**
   f. **svcinfo lserrlogbyiogrp**
   g. **svcinfo lserrlogbyfcconsistgrp**
   h. **svcinfo lserrlogbyfcmap**
   i. **svcinfo lserrlogbyrcconsistgrp**
   j. **svcinfo lserrlogbyrcrelationship**

   These commands will list the error log by type. These commands will return a list of dumps in the appropriate directory. For example, issue the **svcinfo lserrlogbymdisk** command, displays the error log by MDisks.

   You can display the whole log or filter the log so that only errors, events, or unfixed errors are displayed. In addition, you can request the output to be sorted either by error priority or by time. For error priority, the most serious errors are the lowest-numbered errors. They are, therefore, displayed first in the table. For time, either the older or the latest entry can be displayed first in the output.

Related reference

"Advanced function cluster overview using the CLI" on page 203
The command-line interface (CLI) offers advanced functions for your cluster.

Shutting down a cluster or single node using the CLI

You can shut down a cluster using the command-line interface (CLI).

If all input power to a SAN Volume Controller cluster is to be removed for more than a few minutes, (for example, if the machine room power is to be shutdown for maintenance), it is important that the cluster is shutdown before the power is removed. The reason for this is that if the input power is removed from the uninterruptible power supply units without first shutting down the cluster and the uninterruptible power supplies, the uninterruptible power supply units will remain operational and eventually become drained of power.

When input power is restored to the uninterruptible power supplies they will start to recharge but the SAN Volume Controllers will not permit any I/O activity to be performed to the virtual disks until the uninterruptible power supply is charged enough to enable all the data on the SAN Volume Controller nodes to be saved in the event of an unexpected power loss. This might take as long as three hours. Shutting down the cluster prior to removing input power to the uninterruptible power supply units will prevent the battery power being drained and will make it possible for I/O activity to be resumed as soon as input power is restored.

Shutting down the cluster

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

Attention: If you are shutting down the entire cluster, you will lose access to all VDisks being provided by this cluster. Shutting down the cluster will also shut down all nodes. This shutdown will cause the hardened data to be dumped to the internal hard drive.

Attention: Ensure that you have stopped all FlashCopy, Metro Mirror, and data migration operations before you attempt a node or cluster shutdown. Also ensure that all asynchronous deletion operations have completed prior to a shutdown operation.

Begin the process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDisks provided by the cluster.

1. If you are unsure which hosts are using the VDisks provided by the cluster, follow the procedure called, Determining the hosts that a VDisk is mapped to.
2. Repeat the previous step for all VDisks.

Perform the following steps to shut down a cluster:

1. Begin the process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDisks provided by the cluster.
   a. If you are unsure which hosts are using the VDisks provided by the cluster, follow the procedure called, Determining the hosts that a VDisk is mapped to.

b. Repeat the previous step for all VDisks.

2. Open a command prompt window.
3. When all I/O has been stopped, issue the **svctask stopcluster** to shut down a single node or the entire cluster in a controlled manner. If you specify the node ID or node name, you can shut down a single node.

   When you enter this command, you will need to either specify a node ID or node name argument, the node in question is to be shut down. After the command completes, the other node in the I/O group de-stages the contents of its cache and goes into write-through mode until the power to the node is returned and the node rejoins the cluster.

   Attention: If this is the last node in an I/O group, you will lose all access to the virtual disks in the I/O group. Before you enter this command, ensure that this is what you want to do. You must specify the force flag.

   If a shutdown command has been sent to the cluster and both cluster and uninterruptible power supply units have powered off, when input power is restored it will be necessary to restart the uninterruptible power supply units by pressing the power button on the uninterruptible power supply front panel.
4. Close the SSH session if you are using SSH in interactive mode.

Shutting down a single node

Attention: If you are shutting down 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 while this node is shut down. Proceed to 2

Attention: If you are shutting down a single node, and this is the last node in the I/O group, you will lose access to all VDisks being served by this I/O group.

Perform the following steps to shut down a single node:

1. Begin the process of quiescing all I/O to the VDisks being served by this nodes 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 these are mapped to by following the procedure called, Determining the hosts that a VDisk is mapped to.
2. When all I/O has been stopped issue the following command to shut down the node:`svctask stopcluster <nodename/ID>`where <nodename/ID> is the name or ID of the node that you want to shut down.

   Note: If this is the last node in the I/O group you also need to specify the -force parameter. For example to force the shutdown of node1: `svctask stopcluster -force node1`

Related tasks

"Determining the host that a VDisk is mapped to" on page 183
You can determine the host that a VDisk is mapped to using the command-line interface (CLI).

Related reference

"Advanced function cluster overview using the CLI" on page 203
The command-line interface (CLI) offers advanced functions for your cluster.

Chapter 5. Backing up and restoring the cluster configuration

You can back up and restore the cluster configuration.

Maintaining your cluster configuration involves several tasks. This topic lists the tasks that are involved.

- Backing up the cluster configuration
- Restoring the cluster configuration
- Deleting unwanted backup configuration files

Related tasks

"Backing up the cluster configuration"
To back up a cluster configuration, complete these tasks.

"Restoring the cluster configuration" on page 216
You can restore a cluster configuration by using the **svcconfig** command in the command-line interface (CLI).

"Deleting a backup configuration file" on page 220
You can delete a backup cluster configuration by using the Deleting a Cluster Configuration panel or the command-line interface (CLI) **SVCCONFIG** command.

Backing up the cluster configuration

To back up a cluster configuration, complete these tasks.

Plan to regularly back up the business data that is stored on all VDisks using your preferred backup method. It is important that the data on all VDisks is backed up, because it will be lost when the configuration is restored.

- All nodes must be online.
- No object name may begin with an underscore.
- All objects should have non-default names, that is, names that are not assigned by the SAN Volume Controller.

  Note: While it is recommended that objects have non-default names at the time the backup is taken, this is not mandatory. Objects with default names will be renamed when they are restored. The names will appear in the format *name*_r.

The following scenario illustrates the value of configuration backup:

1. Use the Backing up a Cluster Configuration panel in the master console or the **svcconfig** command to create a backup file on the cluster that contains details of the current cluster configuration.
2. Store the backup configuration on some form of tertiary storage.

   Note: You must copy the backup file off the cluster or it will be lost if the cluster crashes.
3. If a severe failure occurs, it will cause the cluster to be lost. Both configuration data (for example, the cluster definitions of hosts, I/O groups, managed disk groups, MDisks) and the application data on the virtualized disks is lost. In this scenario, it is assumed that the application data can be restored from normal customer backup procedures. However, before this can be carried out, it is necessary to reinstate the cluster, as configured at the time of the failure. This

means restoring the same managed disk groups, I/O groups, host definitions, and finally the VDisks that existed prior to the failure. The application data can then be copied back onto these VDisks and operations resumed.

4. Recover the hardware: hosts, SAN Volume Controllers, disk controller systems, disks, and SAN fabric. The hardware and SAN fabric must physically be the same as that used before the failure, although added hardware is permitted.
5. Re-initialize the cluster.
6. Restore your cluster configuration using the backup configuration file generated in step 1 on page 213
7. Restore the data on your virtual disks (VDisks) using your preferred restore solution, or with help from IBM Service.
8. Resume normal operations.

Restrictions: The following restrictions must be observed for the scenario described above, or for a similar scenario, to work:

- The installed hardware must be identical when the backup is taken and when the restore is done, except that new hardware may be present. Otherwise the restore will fail.
- No changes should be made to the fabric of the cluster between backup and restore. If changes are made, you should back up your cluster configuration again.
- There are two phases in the restore process, prepare and execute. Do not make any changes to the fabric or the cluster between the two phases.
- No independent operations that could change the cluster configuration should be running while the backup command is running.

Perform the following steps to backup your cluster configuration.

1. Back up the data that your enterprise uses to run its business using your preferred backup method. It is important that the data on all VDisks is backed up, because it will be lost when the configuration is restored.
2. Back up the cluster configuration using the Backing Up a Cluster Configuration panel or the **svcconfig backup** command.

   Note: Back up the cluster configuration immediately after completing step 1.
3. Ensure that all your SSH keys are available. You need these keys when you restore the cluster configuration.

When the **svcconfig backup** command runs, it produces a file called svc.config.backup.xml, which describes the cluster configuration. This file is stored in /tmp on the configuration node within the cluster. It is important that this file is copied from the cluster to some external storage since, should the configuration node move to another node within the cluster, then the /tmp directory on this node will be inaccessible. (The configuration node might move in response to an error recovery action, or due to some user maintenance activity.)

To copy the svc.config.backup.xml from the node to external storage, use the secure copy command (pscp) on the master console or the secure copy command as described in the example below if you accessing the cluster through your own secure shell installation.

The backup feature of the **svcconfig** command is designed to back up the cluster information, such as VDisks, local Metro Mirror information, MDisk groups, and nodes and not the information that you have written to the VDisks within the cluster.

It is important that any application using the VDisks on the cluster as storage should back up its data as usual using their appropriate backup routines.

All nodes must be online and no object in the cluster can begin with an underscore "_".

To create a backup of your cluster configuration, (your_cluster_name), perform the following steps:

1. Log onto the cluster by issuing the following:

   ```
   ssh -l admin your_cluster_name -p 22
   ```

   This will bring up a session on the cluster where you can issue the **svcconfig** command.
2. Issue the following:

   ```
   svcconfig clear -all
   ```

   This will remove any existing backup files that are on your cluster and ensure a clean directory into which the backup files can be placed.
3. Issue the following:

   ```
   svcconfig backup
   ```

   The cluster will return output similar to the following as the backup runs:

   ```
   CMMVC6112W mdisk mdisk14 ...
   CMMVC6112W node node1 ...
   CMMVC6112W node node2 ...
   ```

   Once the backup has completed and you have returned to the prompt, you will need to exit the cluster and copy the backup files to somewhere off the cluster.
4. Issue the following:

   ```
   exit
   ```
5. To copy the backup files off the cluster, issue the following:

   ```
   scp -P 22 admin@your_cluster:/tmp/svc.config.backup.*
   /offclusterstorage/
   ```

   The following three files will be retrieved from the cluster:
   a. *svc.config.backup.xml* This contains the information about the objects on your cluster.
   b. *svc.config.backup.sh* This contains the **svcinfo** commands that were used to create the backup of the cluster. This can be discarded.
   c. *svc.config.backup.log* This contains the feedback from the backup routine and will contain any error information that may have been reported. This can be discarded after examination.

Because the **svcconfig** command requires that the same cluster configuration be in place before you can restore, it is wise to rename these files with the configuration node name either at the start or end of the file names to make identification easier at restore time. To rename the files with the configuration node name, issue the following:

```
mv /offclusterstorage/svc.config.backup.xml
/offclusterstorage/svc.config.backup.xml_myconfignode
```

Because these files contain details about your cluster, it is advisable to copy them to a location that is under password control to avoid unauthorized access to this configuration information.

Restoring the cluster configuration

You can restore a cluster configuration by using the **svcconfig** command in the command-line interface (CLI).

Note: Data is *not* restored when restoring a configuration from backup. Ensure that you rely on restoring data from traditional data backups once you have restored the cluster configuration.

Ensure that the cluster to which you are restoring the backup configuration files meet the following conditions:

- You have a copy of your configuration files on a server accessible to the cluster (for example, the master console).

  Note: Ensure that you back up the cluster configuration to generate the configuration files.
- You can download the configuration files to your master console by performing the following steps:
  1. Log on to the master console.
  2. Click **Service and Maintenance -> List Dumps**.
  3. Click **Software Dumps**. The Software Dump files window is displayed.
  4. Find the configuration files on the Software Dump files window and download them to your master console.

     Note: When saving the configuration files, ensure that you select the `All Files` option from the drop-down list.
- The installed hardware must be identical when the backup is taken and when the restore is done, except that new hardware can be present. Otherwise the restore will fail. All nodes, controllers and MDisks must be the same. It is permissible to have replaced one or more nodes, provided they present the correct WWNN identifiers. The fabric of the SAN should also be identical.
- The restore must be performed to a single node cluster. You can restore the configuration using any node as the configuration node. However, it is recommended that you use the node that was the configuration node when the cluster was first created, otherwise the SCSI inquiry identifiers of the I/O groups may change. This is known to affect IBM TotalStorage Productivity Center for Fabric and Veritas Volume Manager and may affect other programs that record this information. You can use the same naming convention as described above. All other nodes (or the other single node for a two-node cluster) that were part of the cluster at the time that the backup was taken, must be powered on but not configured to the cluster. These nodes should appear as output from issuing the **svcinfo lsnodecandidate** command or by using the SAN Volume Controller Console.

In addition, the SAN Volume Controller analyzes the backup configuration file and the cluster to verify that the required disk controller systems are available.

Before you begin, hardware recovery must be complete. In particular, the following hardware must be operational: hosts, SAN Volume Controller, disk controller systems, disks, and the SAN fabric.

Before you begin, obtain the featurization settings from the master console before proceeding with the configuration restoration. Perform the following steps to obtain the featurization settings:

- Log on to the SAN Volume Controller
- Click **Service and Maintenance -> Set Features**.
- Write down the virtualization limit (gigabytes) because you will need it in later steps. Also note whether FlashCopy and Metro Mirror are enabled or disabled.

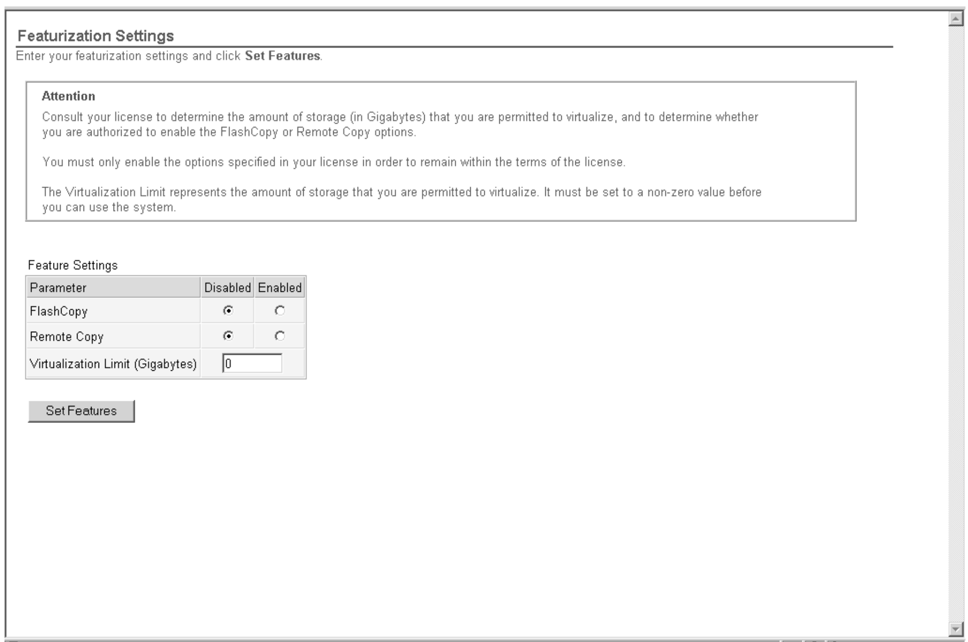

Figure 45. Featurization settings

To restore a cluster configuration, complete these tasks.

1. Perform the following steps to prepare the cluster for configuration restore by removing the cluster using the front panel and using the SAN Volume Controller Console:

   Note: Before removing the cluster, ensure that you have a copy of the cluster and all of its data stored somewhere safe. Also, the following procedure requires that you have superuser administrator authority.

   a. Select delete cluster from the front panel on each node in the cluster.

      Note: Do this for all nodes that are not showing *Cluster :* on the front panel. This means they are already destroyed and are candidates.

   b. Click **Clusters** in the portfolio. The View Clusters panel is displayed.
   c. Select the cluster you want to remove and select **Remove a cluster** from the list. Click **Go**.
2. Create a new cluster from the front panel.
3. Add the cluster using the Adding clusters panel.

4. Connect to the cluster by either using the CLI or the SAN Volume Controller Console. Type the IP address of the cluster and select the Create (Initialize) Cluster check box. When you click OK, the Create a Cluster wizard is displayed.
5. Configure the cluster. During the configuration wizard, the Featurization settings window asks for the feautrization settings. Ensure that you enter the same values for these features that you noted previously.

   Note: You should now be able to connect to the cluster using the master console or the CLI. The PuTTY or SSH keys have been uploaded. You may need to reset the SSH keys after uploading the public key.
6. Copy the file named `svc.config.backup.xml` to the /tmp directory of the cluster using the PuTTY ″pscp″ program from the master console. Perform the following steps to use the PuTTY ″pscp″ program and copy the file:
   a. Open a command prompt from the master console.
   b. Set the path in the command line to use pscp with the following format: `set PATH=C:\path\to\putty\directory;%PATH%`
   c. Issue the following command to specify the location of your private SSH key for authentication. `pscp [options] source [source...] [user@]host:target`
7. Prepare the cluster using the **svcconfig restore -prepare** command. After issuing this command, do not rename the objects in the cluster. The default names are required to successfully complete the restoration. The unique names in the backup configuration file will be restored at the completion of the restoration.

   Note: Copy the output log file to another server accessible to the cluster (the master console) and check for errors. If there are errors, correct the condition which caused the errors and reissue the command. Ensure that you correct all errors before proceeding. If you need assistance, contact IBM Customer Support. Perform the following steps to copy the log file:
   a. Issue the **pscp -i <private key location> [user@]host:source target**
   b. The file name is `/tmp/svc.config.restore.prepare.log`
8. Restore the cluster configuration using the **svcconfig restore -execute** command. Issuing the **svcconfig restore -execute** command on a single node cluster will add the other nodes and hosts to the cluster. This will create a log file named `svc.config.restore.execute.log` in the /tmp directory. Check this log to ensure that no errors or unexpected warnings have been located.

   Note: In some cases, you may receive an error or warning stating that a featurization has not been enabled. This means that the current featurization values did not match the previous featurization values after the recovery process. Nevertheless, the recovery procedure will continue normally and you will be able to input the correct feature values in the SAN Volume Controller Admin Console after the cluster is recovered.

Select the set of backup files you wish the cluster to be recovered to. Use the master console to create a new cluster. Once the new cluster has been created, you can proceed with the restore. Perform the following steps:
1. Log on to the cluster and clear the backup directory of any old backup restore files. Issue the following:

```
ssh -l admin your_cluster_name -p 22
```

```
svcconfig clear -all

exit
```

2. Copy the backup file from your cluster storage by issuing the following:

   ```
   pscp -P 22 /offclusterstorage/svc.config.backup.xml_myconfignode
    admin@your_cluster_name:/tmp/svc.config.backup.xml
   ```

   Log back on to the cluster.

   ```
   ssh -l admin your_cluster_name -p 22
   ```

3. Issue the following:

   ```
   svcinfo lsnode
   ```

   Ensure that only one node is online and identify which one it is. If this node was not the configuration node in the configuration you are trying to restore, either make a cluster where it is or select the appropriate backup file.

id	....	status	IO_group_id	IO_group_name	config_node
1	....	online	0	io_grp0	yes

4. Issue the following:

   ```
   svcinfo lscluster
   ```

   Output similar to the following is displayed:

Id	name	location
0000020066206BE2	your_cluster_name	local

5. Issue the following:

   ```
   svcconfig restore -prepare
   ```

   This will do a comparison of the current cluster configuration and available resource and the backup file you have put onto the cluster. If there are any errors, the command will fail with a CMMVC*nnnn*E error. You will need to fix the error and issue the command again.

   Note: If there has been any change to the fabric since the backup was taken, it will not be possible to restore the chosen configuration.

   When this command has completed, you may have received a number of warning messages. You will need to ensure that the action about to be taken is acceptable. You may need to exit the cluster and copy this log file off the cluster for reading, because it may be quite large.

6. In order to read the log file produced by the -prepare flag to ensure you are aware of all the warnings that have been issued, exit from the cluster. Issue the following:

   ```
   exit

   pscp -P 22 admin@your_cluster_name:/tmp/svc.config.restore.prepare.log
    /offclusterstorage

   cat /offclusterstorage/svc.config.restore.prepare.log|more
   ```

to exit the list.

7. When you are satisfied that the restore will happen as expected, log back on to the cluster and execute the restore command. Issue the following:

```
ssh -l admin your_cluster_name -p 22
```

```
svcconfig restore -execute
```

This will use the svc.config.restore.sh file to attempt to recover your cluster structure onto the available cluster hardware. Once this has completed, you can check the log file to ensure that no errors or unexpected warnings have been issued about the restore. The output is stored in svc.config.restore.execute.log. The following output displays that a successful restore has taken place and that no errors are reported.

....
IBM_2145:admin>

When you have verified the cluster as correct, you may restore your company data from the storage back onto the presented VDisks.

Remove any unwanted configuration backup files from the cluster using the **svcconfig clear** command.

Related tasks

"Backing up the cluster configuration" on page 213
To back up a cluster configuration, complete these tasks.

"Create cluster from the front panel" on page 79
After you have created a pair of nodes, you can now create a cluster from the front panel.

"Configuring a cluster using the SAN Volume Controller Console" on page 88
Once you've created a pair of nodes, you will then need to create and configure a cluster.

"Configuring the cluster using the CLI" on page 158
You can configure the cluster using the command-line interface (CLI).

Related information

Chapter 2, "Preparing to configure the SAN Volume Controller," on page 79
Once you've completed the two phase creation of the cluster, you can begin to configure the SAN Volume Controller.

Deleting a backup configuration file

You can delete a backup cluster configuration by using the Deleting a Cluster Configuration panel or the command-line interface (CLI) **SVCCONFIG** command.

Unneeded backup configuration files and SSH keys can be deleted from the master console or a SAN Volume Controller.

Perform the following steps to delete backup configuration files:

1. Click **Service and Maintenance** from the portfolio.
2. Click **Delete Backup** from the portfolio.
3. Click **OK** to delete the backup configuration file.

Chapter 6. Software upgrade strategy

You can upgrade your software while your day-to-day operations are running.

You must, however, expect performance to be degraded while the software is being installed.

Note: Applying a software update takes approximately one hour. This is in part due to the 30 minute delay which is inserted to allow the multi-pathing software to recover.

Software and microcode for the SAN Volume Controller and its attached adapters is tested and released as a single package. The package number is increased each time a new release is made, although only some of the components might have changed. Included in the package are Linux, Apache, and the SAN Volume Controller software.

If you are upgrading through more than one level; for example, from level 1 to level 3, under some circumstances, you might need to install an intermediate level. For example, if you are upgrading from level 1 to level 3, you might need to install level 2 before you install level 3. Details of any prerequisite levels are provided with the source files.

Attention: Applying a software upgrade while the node is in service mode results in deleting the node from the cluster. Status information stored within the node will be deleted, and this causes data loss if the cluster is dependent solely on this node.

When upgrading cluster software where the cluster participates in one or more intercluster relationships only one cluster should be upgraded at a time. That is, both clusters should not be upgraded concurrently. The software upgrade should be allowed to complete one cluster before it is started on the other cluster. If both clusters are upgraded concurrently, it may lead to a loss of synchronization. It may further lead to a loss of availability.

Attention: Ensure that you have no unfixed errors in the log, and that the Cluster Time/Date is correctly set. Start the Directed Maintenance Procedures and ensure that you fix any outstanding errors before attempting to concurrently upgrade the software.

Related tasks

"Disruptive software upgrade" on page 222
You can perform a disruptive software upgrade using the command-line interface (CLI).

"Upgrading the SAN Volume Controller firmware using the SAN Volume Controller Console" on page 223
You can upgrade the cluster software using the SAN Volume Controller Console.

"Performing the node rescue" on page 226
You can follow the step-by-step instructions to perform the node rescue.

"Manual recovery from software upgrade problems" on page 231
When a revised version of software is committed, you might not be able to return to a previous software version because some data structures might have been changed such that they cannot be used with the previous software version. Therefore, if you have any problems, you must go forward to a later version of the code.

Related reference

"Automatic upgrade" on page 227
When new nodes are added to the cluster, the upgrade packages are usually automatically downloaded to them from the SAN Volume Controller cluster. No manual intervention is needed.

"Automatic recovery from upgrade problems" on page 227
The cluster will automatically terminate the upgrade process if the any of the nodes fail to upgrade to the new software level.

"PuTTY scp" on page 228
PuTTY scp provides a file transfer mechanism for secure shell (SSH) to copy files either between two directories on the SAN Volume Controller configuration node or between the configuration node and another host.

"Installing the upgrade using the CLI" on page 229
You can install upgraded software using the command-line interface (CLI).

"Installing the software" on page 230
The software is delivered to you as a single package.

Disruptive software upgrade

You can perform a disruptive software upgrade using the command-line interface (CLI).

The IBM Total SAN Volume Controller only supports concurrent code upgrades. To ensure that a code upgrade is coordinated across all nodes in the cluster, it is necessary for the nodes to be able to communicate with each other across the fibre-channel SAN. However, some users may prefer to perform a disruptive code upgrade. The following procedure documents how to quiesce I/O to the SAN before performing a concurrent code upgrade to ensure that there is no I/O in progress during the upgrade.

Perform the following steps to complete the disruptive software upgrade process:

1. Stop any host applications and unmount the file-systems that are using storage that is being managed by the SAN Volume Controller. If your hosts are being shutdown then this will occur as the host is shutdown, otherwise it will be necessary to do this manually on each host. This step will ensure that hosts will stop issuing I/O operations and that any data in the filesystem caches is flushed.
2. Shutdown the cluster by issuing the **svctask stopcluster** command. This command will stop the SAN Volume Controllers from issuing I/O to backend controllers and will flush data from the SAN Volume Controller cache.
3. Rezone the switch so that the SAN Volume Controller nodes are in one zone. Ensure that this zone does not include a host HBA or a backend controller (keep the old switch configuration so it can be restored at step 6). This step isolates the SAN Volume Controller from the rest of the SAN.
4. Power on all the SAN Volume Controller nodes and wait for them to reform a cluster.

   Note: Because the IBM Total Storage SAN Volume Controller has been isolated from the backend storage you will get some error logs indicating that this has occurred.
5. Perform the software upgrade in the same manner as for a concurrent code upgrade.
6. Restore the original switch configuration.

7. Clear any error logs produced at step 4 on page 222 indicating that backend storage is unavailable. Check that all backend storage is now online and accessible to the SAN Volume Controllers.
8. Remount file-systems and start host applications.

Related tasks

"Shutting down a cluster or single node using the CLI" on page 210
You can shut down a cluster using the command-line interface (CLI).

Related reference

Chapter 6, "Software upgrade strategy," on page 221
You can upgrade your software while your day-to-day operations are running.

Upgrading the SAN Volume Controller firmware using the SAN Volume Controller Console

You can upgrade the cluster software using the SAN Volume Controller Console.

If you are using Internet Explorer, perform the following:
1. Click on **Tools** in the menu.
2. Select **Internet Options –>Connections** tab.
3. Click on **LAN Settings...** and ensure that the box marked **Use a proxy server** is unchecked. Click **OK** twice to accept the settings.

If you are using Netscape, perform the following:
1. Click on **Edit** in the menu.
2. Click on **Preferences...**. Expand the Advanced section and select **Proxies**.
3. Select the radio button marked **Direct connection to the Internet**. Click **OK** to accept the settings.

Note: You cannot download the upgrade. You need to download the file to your local directory so the package can be uploaded in the process.

Software upgrade files can be quite large, if you experience problems when uploading upgrade files to the cluster reliably, you should disable proxies on the Web browser from where you will upload the file. This should also shorten the file upload time.

Note: If you disable proxies, you may not be able to connect to external Web sites. It is therefore advised that prior to disabling proxies, you make a record of your existing settings in case you need to restore access to other Web sites.

Perform the following steps to upgrade the software:
1. Click **Service and Maintenance** from the portfolio.
2. Click **Upgrade Software** to check the installed software level or to install a new level of software on the cluster. The Software Upgrade panel is displayed.

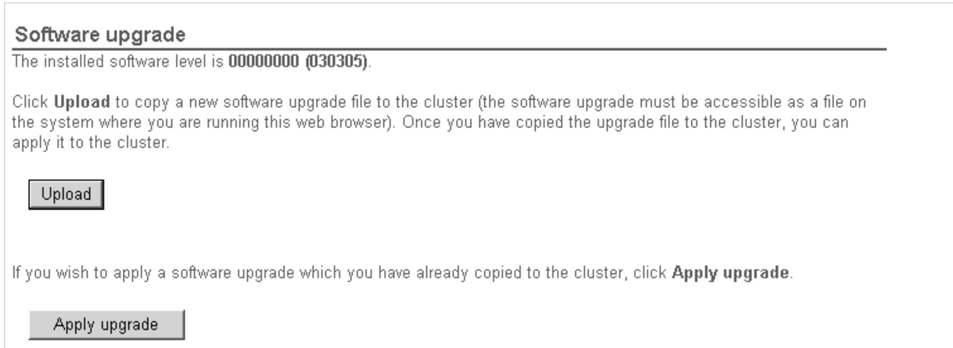

Figure 46. Software upgrade panel

3. Click **Upload** to copy a new software level from your host to the cluster. (This action uses the upload feature of the Web browser.) The Software upgrade - file upload panel is displayed.

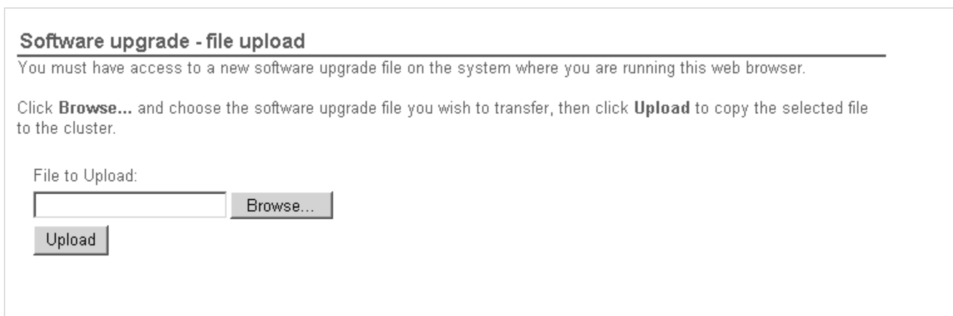

Figure 47. Software upgrade - file upload panel

You can get new software levels from the IBM Product Support Web site, or from an installation CD.

After a successful copy of the file, the install process fails if all the nodes configured into the cluster are not present. This behavior cannot be overridden using the force flag. If any node configured to be a member of the cluster is not present then in order to upgrade the software the node must either be deleted from the cluster or must be brought online. Furthermore, if a node has been deleted from the cluster such that any IO group has only one member then the software upgrade will also fail. This is because the upgrade process will result in loss of access to data. The force flag can be used to override this restriction if you are prepared to loose access to data during the upgrade.

Before you begin the software upgrade, make sure that you are aware of the following:

- The code is distributed to all the nodes in the cluster using fibre channel connections between the nodes.
- Nodes are updated one at a time.
- Nodes will begin executing the new software, concurrently with normal cluster activity.
- The procedure to update a single node takes approximately 5 minutes.
- During the update of a node it does not participate in I/O activity in the I/O group. Thus all I/O activity for the virtual disks in the I/O group is directed to the other node in the I/O group by the host multi-pathing software. During the update of a node the other node in the I/O group will notice that it's partner is

not participating in the cluster and will as a result attempt to flush the writeback cache and set it into write-through mode. This flush is not guaranteed to be successful or to complete and as a result concurrent software update does create a single point of data loss. Should the remaining node in an I/O group experience a failure during a software update of its partner then the only valid copy of dirty data in the writeback cache could be lost.

- All of the nodes connected to one uninterruptible power supply are updated first before any of the nodes connected to the other uninterruptible power supply.
- A 30 minute delay is inserted into the procedure between updating the nodes connected to one uninterruptible power supply and starting to update the nodes on the other uninterruptible power supply. This allows time for the host mult-ipathing software to rediscover paths to the nodes on the first uninterruptible power supply so that when nodes on the second uninterruptible power supply are updated loss of access does not result.
- The update is not committed until all nodes in the cluster have been successfully updated to the new code level. If all nodes successfully restart with the new code the new version is committed. When this happens, the cluster VPD is updated to reflect the new level of code. After this point downgrade to a package with a lower major number is no longer possible.
- New behaviors or functions in the installed software will only be available to be invoked when all member nodes are upgraded and the update is committed.
- Since the software upgrade process takes some time the install command completes as soon as the software package is verified by the cluster. To determine when the upgrade has completed you must either display the software version in the cluster VPD or look for the Software upgrade complete event in the error/event log. If any node fails to restart with the new code level or fails at any other time during the process the code is backed-off.
- During a software upgrade the version number of each node is updated when the software has been installed and that node has been restarted. The cluster software version number is updated when the new version of software is committed.
- When code upgrade starts an entry is made in the error or event log and another entry is made when the upgrade completes or fails.

4. Click **Apply upgrade** to display the Applying Software Upgrade panel. This page enables you to select the upgrade and to apply it to the cluster. This page displays a list of the software levels that you can apply to the cluster.

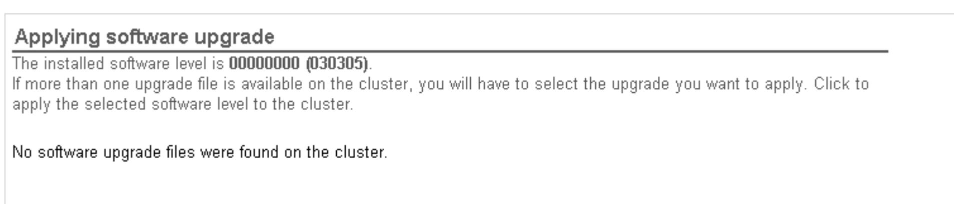

Figure 48. Applying software upgrade panel

When a new code level is applied, it is automatically installed on all the nodes that are in the cluster. Note this process can take up to 30 minutes per node.

Related reference

Chapter 6, “Software upgrade strategy,” on page 221
You can upgrade your software while your day-to-day operations are running.

Performing the node rescue

You can follow the step-by-step instructions to perform the node rescue.

If it is necessary to replace the hard disk drive or if the software on the hard disk drive has become corrupted, you can reinstall the software on the SAN Volume Controller by using the node rescue procedure.

To provide an alternate boot device, a minimal operating system is also available in nonvolatile memory on the service controller. If it is necessary to replace the hard disk drive or the software on the hard disk drive has become corrupted, the SAN Volume Controller cannot boot and the Hardware Boot indicator remains on the front panel display or the boot operation does not progress.

If this occurs, you can reinstall the software on the SAN Volume Controller by using the node rescue procedure. Node rescue works by booting the operating system from the service controller and running a program that will copy all the node software from any other SAN Volume Controller that can be found on the fibre-channel fabric. The following procedure tells you how to run the node rescue procedure.

Perform the following steps to complete the node rescue:

1. Ensure that the fibre-channel cables are connected.
2. Ensure that at least one other SAN Volume Controller node is connected to the fibre-channel fabric.
3. Turn off the SAN Volume Controller.
4. Press and hold the left and right buttons on the front panel.
5. Press the power button.
6. Continue to hold the left and right buttons until the node-rescue-request symbol is displayed on the front panel. Figure 49 shows the node-rescue-request symbol.

Figure 49. Node-rescue-request display

The node rescue request symbol displays on the front panel display until the SAN Volume Controller starts to boot from the service controller. If the node rescue request symbol displays for more than two minutes, check the connection between the service controller and the system board. The service display shows the progress or failure of the node rescue operation.

Note: If the recovered node was part of a cluster, the node will now be offline. Delete the offline node from the cluster and then add the node back into the cluster. If node recovery was used to recover a node that failed during a software upgrade process, the automatic software downgrade process will start but it might not continue until the failed node has been deleted from the cluster. After the failed node is deleted, it is not possible to add the node back into the cluster until the downgrade process has completed. This may take up to two hours.

If the cables are correctly located and the node rescue request symbol still displays, replace the FRUs in the following sequence:

1. System board assembly
2. Service controller

Related reference

Chapter 6, “Software upgrade strategy,” on page 221
You can upgrade your software while your day-to-day operations are running.

Automatic upgrade

When new nodes are added to the cluster, the upgrade packages are usually automatically downloaded to them from the SAN Volume Controller cluster. No manual intervention is needed.

New nodes introduced to the cluster normally have software packages downloaded to them from the cluster without any manual intervention. A new node requiring a code version higher than that currently available on the cluster or a node that already contains a code version higher than that on the cluster will not be configured into the cluster. If a node is added to the network that has no code installed, for example because the disk drive has been replaced, or it has such an old code version installed that it cannot advertise itself to the clusters, a reinstall of the software is forced by using the Node Rescue procedure.

If you add a new SAN Volume Controller node that has a code version that is higher than the one that is available on the cluster, that node is *not* configured into the cluster. It will join the cluster, however the node will be downgraded to the cluster level.

Error counts: During the SAN Volume Controller software upgrade, you can expect to see either I/O error counts displayed by the data-path query adapter, or an increase in the number of **datapath query device** commands if active I/O operations exist between hosts and the SANs. See the *IBM TotalStorage Multipath Subsystem Device Driver: User’s Guide* for more information about **datapath query** commands.

During the software upgrade, each SAN Volume Controller node of a working pair is upgraded sequentially. The SAN Volume Controller node that is being upgraded is temporarily unavailable, and all I/O operations to that SAN Volume Controller fail. As a result, I/O error counts increase. However, failed I/O operations are directed to the other SAN Volume Controller node of the working pair, and applications should not see any I/O failures.

Related reference

Chapter 6, “Software upgrade strategy,” on page 221
You can upgrade your software while your day-to-day operations are running.

Automatic recovery from upgrade problems

The cluster will automatically terminate the upgrade process if the any of the nodes fail to upgrade to the new software level.

In this case, any nodes that have already upgraded to the new software level will downgrade back to the original code level. You should check the error log to determine the reason for the failure before attempting to upgrade the cluster again.

Related reference

Chapter 6, "Software upgrade strategy," on page 221
You can upgrade your software while your day-to-day operations are running.

PuTTY scp

PuTTY scp provides a file transfer mechanism for secure shell (SSH) to copy files either between two directories on the SAN Volume Controller configuration node or between the configuration node and another host.

Overview

You must have appropriate permissions on the source and destination directories on your respective hosts to be able to use pscp. PuTTY scp is available to you when you install an SSH client on your host system. The pscp application can be accessed through a command line.

To start pscp, ensure that it is on your PATH or in your current directory. To add the directory containing pscp to your PATH environment variable, issue the following:

```
set PATH=C:\path\to\putty\directory;%PATH%
```

Open your console window and issue the following: `pscp`. This will bring up a usage message. The message provides the version of pscp you're using and gives you a brief summary of how to use pscp.

If you want to use pscp, perform the following:

1. Start a PuTTY session in order to access your SAN Volume Controller Console cluster.
2. Save your configuration. For example, your saved session is called SVCPUTTY.
3. From a Command Prompt window, set up your path by issuing the following commands:

   ```
   set PATH=C:\path\to\putty\directory;%PATH%
   set PATH=C:\Program Files\Putty;%PATH%
   ```
4. Copy the package onto the node where the CLI is running from the Master console. Ensure you are in the software directory called C:\SVC_Software_Directory.

   ```
   pscp -load SVCPUTTY svc_code_name admin@
    <cluster_ip_address>:/home/admin/upgrade
   ```

   where *<cluster_ip_address>* is your cluster IP address. You are notified of copy failures by error messages from the CLI and the SAN Volume Controller Console. If there is insufficient space on the cluster to store the software upgrade package then the copy operation will fail. If this occurs, issue the **svctask cleardumps** command to make space for the upgrade package, then repeat the copy operation.

   Alternatively, you can transfer error logs to the master console using pscp. To scp the error logs from the cluster to the master console, issue the following:

```
pscp -unsafe -load  SVCPUTTY  admin@<cluster_ip_address>:
 /dumps/elogs/* c:/svc_directory
```

where *<cluster_ip_address>* is your cluster IP address.

Related reference

Chapter 6, “Software upgrade strategy,” on page 221
You can upgrade your software while your day-to-day operations are running.

Installing the upgrade using the CLI

You can install upgraded software using the command-line interface (CLI).

You can use either PuTTY scp (pscp) or the SAN Volume Controller Console to copy the upgrade package to each SAN Volume Controller cluster or issue CLI commands.

If you want to use pscp, perform the following:

1. Start a PuTTY session in order to access your SAN Volume Controller Console cluster.
2. Save your configuration. For example, your saved session is called SVCPUTTY.
3. From a Command Prompt window, set up your path by issuing the following commands:

   ```
   set PATH=C:\path\to\putty\directory;%PATH%
   set PATH=C:\Program Files\Putty;%PATH%
   ```
4. Copy the package onto the node where the CLI is running from the Master console. Ensure you are in the software directory called C:\SVC_Software_Directory.

   ```
   pscp -load SVCPUTTY svc_code_name admin@
    <cluster_ip_address>:/home/admin/upgrade
   ```

   where *<cluster_ip_address>* is your cluster IP address. You are notified of copy failures by error messages from the CLI and the SAN Volume Controller Console. If there is insufficient space on the cluster to store the software upgrade package then the copy operation will fail. If this occurs, issue the **svctask cleardumps** command to make space for the upgrade package, then repeat the copy operation.

   Alternatively, you can transfer error logs to the master console using pscp. To scp the error logs from the cluster to the master console, issue the following:

   ```
   pscp -unsafe -load  SVCPUTTY  admin@<cluster_ip_address>:
    /dumps/elogs/* c:/svc_directory
   ```

   where *<cluster_ip_address>* is your cluster IP address.
5. After a successful copy of the file, issue the **svcservicetask applysoftware -file filename** command, where *filename* is the name of the file that you copied the software upgrade package too. This command starts the installation of the code. The installation process will fail if a node is not present and if the node is not paired with another node in an I/O group. You can, however, use the -force option to override this restriction if you are prepared to lose access to data during the upgrade.

   Note: The installation process will *only* fail when some paths between the host systems and the cluster are not available. Data access can be lost temporarily during the upgrading process. You can prevent this if, before you start the installation, you issue a data-path query device on each host system to ensure that all paths are available. See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* for more information about data-path query commands.

Attention: The order in which the nodes are upgraded depends on the following:

- The position of the nodes. The code will be transferred to all the nodes in an I/O group.
- The I/O group ID. The code will be transferred from the lowest I/O group ID that includes nodes on it.

6. To verify that the upgrade was successful, you can perform any one of the following steps:
   - The code level is distributed to all the nodes that are in the cluster. The nodes, in turn, are then restarted. If all the nodes successfully restart with the new code level, the new version is committed and the cluster vital product data (VPD) is updated to new level of code.
   - The software upgrade is complete when the cluster verifies the upgrade package. To determine whether the upgrade has completed, you must either display the software version in the cluster VPD, or look for the `Software upgrade complete` event in the SAN Volume Controller error or event log. If the node does not restart automatically during the upgrade, you should repair or manually delete that node from the cluster to complete the backout process.
   - Alternatively, you can also either perform the following steps:
     a. Issue the **svctask dumperrlog** command to dump the contents of the error log to a text file. You can also use this command to delete unwanted error log dumps from the cluster.
     b. Once you have the contents of the error log dumped into a text file, verify that there were no errors in the text file. If there are no errors, you have successfully upgraded the software and output similar to the following is displayed in the log file:

        ```
        Upgrade completed successfully
        ```

     c. Issue the **svcinfo lsnodevpd** command for each node. You should see that the software version field has been updated.

Related reference

Chapter 6, "Software upgrade strategy," on page 221
You can upgrade your software while your day-to-day operations are running.

"SAN Volume Controller library and related publications" on page xviii
A list of other publications that are related to this product are provided to you for your reference.

Installing the software

The software is delivered to you as a single package.

Software package

Cluster software versions comprise a number of software components that are delivered as a single package. Cluster software versions comprise a number of software components that are delivered as a single package. The size of the software update package depends on the number of components that are being replaced by that upgrade package. The software installation procedure involves copying the new software version to the cluster and then starting an automatic installation process. This installation process might take up to an hour to complete and during the process each of the nodes is restarted in turn. Once all the nodes in the cluster have been successfully restarted with the new software the new software version is automatically committed. While each node is being restarted, there might be some degradation in the maximum input/output rate that can be sustained by the cluster.

Installation operation

The installation operation can normally be performed concurrently with normal user I/O operations. If any restrictions apply to the operations that can be performed during the upgrade, then these restrictions will be documented on the SAN Volume Controller web site from where the upgrade package was obtained. During the upgrade operation, only the following SAN Volume Controller commands will be operational from the time the install process starts to the time that the new software is committed or until the process has been backed-out. All other commands will fail with a message indicating that a software upgrade is in progress. In the following commands, xxxx is the object type.

- **svcinfo lsxxxx**
- **svcinfo lsxxxxcandidate**
- **svcinfo lsxxxxprogress**
- **svcinfo lsxxxxmember**
- **svcinfo lsxxxxextent**
- **svcinfo lsxxxxdumps**
- **svcinfo caterrlog**
- **svcinfo lserrlogbyxxxx**
- **svcinfo caterrlogbyseqnum**
- **svctask rmnode**
- **svcservicetask rmnode**

Because of the operational limitations that occur during the upgrade process, the software installation is a customer task.

Related reference

Chapter 6, “Software upgrade strategy,” on page 221
You can upgrade your software while your day-to-day operations are running.

Manual recovery from software upgrade problems

When a revised version of software is committed, you might not be able to return to a previous software version because some data structures might have been changed such that they cannot be used with the previous software version. Therefore, if you have any problems, you must go forward to a later version of the code.

Attention: This procedure causes a loss of *all* data currently configured in the cluster. This is a last resort and should only be done if you have recently backed-up your data.

In extreme conditions where you cannot wait for a software update and you need to return to the previous software version, you can use the following procedure.

Attention: This procedure, however, causes the total loss of the SAN Volume Controller cluster. This should only be done as a last resort.

Perform the following steps to reset from software upgrade problems:
1. Power-off all but one of the nodes that are in the cluster.
2. Set the powered-on node to the service access mode.
3. Use the service access functions to force the download of the older software package.
4. Repeat the action for each of the failed nodes.
5. From a node that has the new code, create a new cluster.

Related reference

Chapter 6, “Software upgrade strategy,” on page 221
You can upgrade your software while your day-to-day operations are running.

Related information

“Resetting a refused SSH key” on page 146
You can reset a refused SSH key relationship between the SAN Volume Controller Console and the SAN Volume Controller cluster.

Chapter 7. Configuring other SAN devices and SAN switches for use with the SAN Volume Controller

Ensure that you are familiar with configuring disk controllers and switches for use with the SAN Volume Controller.

Related reference

"Configuring and servicing storage subsystems"
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

"Configuring the EMC CLARiiON subsystem" on page 253
This part provides information about configuring the EMC CLARiiON storage system for its attachment to a SAN Volume Controller.

"Configuring the EMC Symmetrix subsystem" on page 265
This part provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

"Configuring the HP StorageWorks subsystem" on page 309
This part provides information about configuring the HP StorageWorks subsystem, which uses the HSG80 controller, so it can attach to a SAN Volume Controller.

Related information

"Configuring the Enterprise Storage Server" on page 273
This part provides information about configuring the Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

"Configuring the FAStT subsystem" on page 278
This part provides information about configuring the FAStT subsystem so it can attach to a SAN Volume Controller.

"Configuring the HDS Lightning subsystem" on page 290
This part provides information about configuring the HDS Lightning subsystem so it can attach to a SAN Volume Controller.

"Configuring the HDS Thunder subsystem" on page 297
This part provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

"Configuring the HPQ Enterprise Virtual Array (EVA) subsystem" on page 326
This part provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Configuring and servicing storage subsystems

Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

Overview

Although virtualization provides many benefits over direct attached or direct SAN attached storage, it is also more liable to and is not restricted by the creation of performance hot-spots. These hot-spots could lead to I/O errors being returned to your hosts, and potentially a loss of access to data may occur.

Related reference

“Configuration guidelines”
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

Related information

“Identifying your storage subsystem”
The serial number presented by the command-line and Web application on the SAN Volume Controller is the serial number of the device.

Identifying your storage subsystem

The serial number presented by the command-line and Web application on the SAN Volume Controller is the serial number of the device.

The serial numbers can be viewed on your storage subsystem. If the serial numbers are not displayed, the WWNN or WWPN will be displayed. The WWNN or WWPN can be used to identify the different subsystems.

Configuration guidelines

Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

Guidelines

- Avoid splitting arrays into multiple logical disks at the storage subsystem layer. Where possible, create a single logical disk from the entire capacity of the array.
- Depending upon the redundancy required, RAID-5 arrays should be created using between 5 and 8 plus parity components. That is 5 + P, 6 + P, 7 + P or 8 + P.
- Ensure that managed disk groups contain managed disks with similar characteristics and approximately the same capacity. Consider the following factors:
  - The underlying RAID type that the storage subsystem is using to implement the managed disk.
  - The number of physical disks in the RAID array and the physical disk type (for example, 10K/15K rpm, FC/SATA).
- If disks are not the same capacity, then they can be specified multiple times when the MDisk group is created. For example, if you have two 400MB disks and one 800MB disk that are identified as MDisk 0, 1, and 2, respectively, then you can create the MDisk group with the candidate IDs of 0:1:2:2. This has the effect of doubling the number of extents on the 800MB drive.
- Do not mix managed disks of greatly differing performance in the same managed disk group. The overall group performance will be limited by the slowest managed disk in the group. Some disk controllers may be able to sustain much higher I/O bandwidths than others, ensure you do not mix managed disks provided by low-end subsystems with those provided by high-end subsystems.
- Avoid leaving virtual disks in image mode. Only use image mode to import existing data into the cluster. This data should be migrated across the other managed disks in the group as soon as possible to optimize the benefits of virtualization.
- Follow the FlashCopy requirements before setting up the storage. Balance the spread of the FlashCopy virtual disks across the managed disk groups and then

the storage subsystems. The I/O characteristics of the application writing to the source virtual disk also effects the impact that FlashCopy operations have on your overall I/O throughput.

- Perform the appropriate calculations to ensure your storage subsystems are configured correctly.

Related reference

"Storage subsystem logical disks"
Most storage subsystems provide some mechanism to create multiple logical disks from a single array. This is useful when the storage subsystem is directly presenting storage to the hosts.

"RAID array configuration"
When using virtualization, ensure that the storage devices are configured to provide some type of redundancy against hard disk failures.

"Optimal managed disk group configurations" on page 236
A managed disk group provides the pool of storage from which virtual disks will be created. It is therefore necessary to ensure that the entire pool of storage provides the same performance and reliability characteristics.

"Considerations for FlashCopy mappings" on page 237
Ensure that you've considered the type of I/O and frequency of update before creating the virtual disks that you wish to use in FlashCopy mappings.

"Image mode and migrating existing data" on page 237
Image mode virtual disks enable you to import and then migrate existing data under the SAN Volume Controller.

Storage subsystem logical disks

Most storage subsystems provide some mechanism to create multiple logical disks from a single array. This is useful when the storage subsystem is directly presenting storage to the hosts.

In a virtualized SAN, however, where possible, there should be a one-to-one mapping between arrays and logical disks. Ensuring that the arrays are configured in this way will make the subsequent load calculations and the managed disk and managed disk group configuration tasks a lot easier.

For example, you have two RAID-5 arrays and both contain 5 + P components. Array A has a single logical disk that is being presented to the SAN Volume Controller cluster. This is seen by the cluster as mdisk0. Array B has three logical disks that are being presented to the SAN Volume Controller cluster. These are seen by the cluster as managed disks 1 through 3. All four managed disks are assigned to the same managed disk group, mdisk_grp0. When a virtual disk is created by striping across this group, what actually takes place is that array A presents the first extent and array B presents the next 3 extents. Therefore, when reading and writing to the virtual disk, the loading is split 25% on the disks in array A and 75% on the disks in array B. The performance of the virtual disk will in general be one third of what array B can sustain.

This example describes the performance degradation and complexity that is introduced by having uneven logical disks in a simple configuration. As stated in the guideline summary, you should aim to create a single logical disk from each array.

RAID array configuration

When using virtualization, ensure that the storage devices are configured to provide some type of redundancy against hard disk failures.

Overview

A failure of a storage device can affect a larger amount of storage being presented to the hosts. To provide redundancy, storage devices should be configured as RAID arrays which use either mirroring or parity to protect against single failures.

When creating RAID arrays with parity protection (for example, RAID-5 arrays) consider how many component disks you want to use in each array. The larger the number of disks, the fewer disks are required to provide availability for the same total capacity (1 per array). However, more disks means a longer time is taken to rebuild a replacement disk after a disk failure, and during this period a second disk failure will cause a loss of all array data. 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 one). In this case, write performance is improved. The number of disk drives required to provide availability may be unacceptable if arrays are too small.

Note:

1. If in doubt, arrays with between 6 and 8 member disks is recommended.
2. When creating RAID arrays with mirroring, the number of component disks in each array does not affect redundancy or performance.

Optimal managed disk group configurations

A managed disk group provides the pool of storage from which virtual disks will be created. It is therefore necessary to ensure that the entire pool of storage provides the same performance and reliability characteristics.

- The performance of a managed disk group will generally be governed by the slowest managed disk in the group.
- The reliability of a managed disk group will generally be governed by the weakest managed disk in the group.
- If a single managed disk in a group fails, access to the entire group will be lost.

These guidelines show that grouping similar disks together is important. The following guidelines should be followed when grouping similar disks:

- Group equally performing managed disks, arrays, in a single group.
- Group similar arrays, for example, all 6 + P RAID-5 arrays, in one group.
- Group managed disks from the same type of storage subsystem in a single group.
- Group managed disks that use the same type of underlying physical disk (for example, fibre-channel, SATA).
- Do not use single disks. Single disks provide no redundancy. Failure of a single disk will result in total data loss of the managed disk group to which it is assigned.

For example, you have two storage subsystems attached behind your SAN Volume Controller. One device is an IBM ESS, which contains ten 6 + P RAID-5 arrays, mdisks 0 through 9. The other device is an IBM FAStT200, which contains a single RAID-1 array, mdisk10, one single JBOD, mdisk11, and a large 15 + P RAID-5 array, mdisk12. If you assigned mdisks 0 through 9 and mdisk11 into a single managed disk group, and the JBOD, mdisk11, fails, you would lose access to all of the ESS arrays, even though they are online. The performance would likely be limited to that available to the JBOD in the FAStT storage subsystem, thus slowing down the ESS arrays.

The ideal configuration with the above components would be to create three groups. One which contained the ESS arrays, MDdisks 0 through 9, another which contained the RAID-1 array, and the third group containing the large RAID-5 array.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Considerations for FlashCopy mappings

Ensure that you've considered the type of I/O and frequency of update before creating the virtual disks that you wish to use in FlashCopy mappings.

FlashCopy will perform in direct proportion to the performance of the source and target disks. That is, if you have a fast source disk and slow target disk, the performance of the source disk will be reduced as it has to wait for the write to happen at the target before it can write to the source.

The FlashCopy implementation provided by the SAN Volume Controller copies at least 256K every time a write is made to the source. This means that *any* write will at a minimum involve a read of 256K from the source, write of the same 256K at the target, and then a write of the original change at the target. Therefore, when an application is performing small 4K writes, this is translated into 256K.

Due to this overhead, consider the type of I/O your application will be performing during a FlashCopy. Ensure that you will not be overloading the storage. The calculations contain a heavy weighting when FlashCopy is active. The weighting itself depends on the type of I/O being performed. Random writes have a much higher overhead than sequential writes, for example, as the sequential write would have copied the entire 256K anyway.

You may spread the FlashCopy source virtual disks and the FlashCopy destination virtual disks between as many managed disk groups as possible. This will limit the potential bottle-necking of a single storage subsystem, (assuming that the managed disk groups contain managed disks from different storage subsystems). However, this may still result in potential bottle-necking if you wish to maintain all your target virtual disks on a single storage subsystem. Ensure that you add the appropriate weighting to your calculations.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Image mode and migrating existing data

Image mode virtual disks enable you to import and then migrate existing data under the SAN Volume Controller.

Ensure that you follow the guidelines when using image mode virtual disks. This may be difficult because a configuration of logical disks and arrays that performed well in a direct SAN-attached environment may contain hot-spots or hot-component disks when connected via the SAN Volume Controller cluster.

If the existing storage subsystems are configured incorrectly with respect to the guidelines you may wish to consider stopping I/O operations on the host systems while migrating the data into the cluster. If I/O operations are continued and the storage subsystem does not follow the guidelines, I/O operations may fail at the hosts and ultimately loss of access to the data will occur.

The procedure for importing many managed disks that contain existing data depends on how much free capacity you have in the SAN Volume Controller cluster.

- You should have the same amount of free space in the cluster as the data you wish to migrate into the cluster.
- If you do not have this amount of capacity available, you can still migrate the data into the cluster, however this is not recommended. The resulting managed disk group will have an uneven distribution of data, with some managed disks being much more heavily loaded then others. Further migration operations will be required to ensure an even distribution of data and, therefore, subsequent I/O loading.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Migrating data with an equivalent amount of free capacity:

If the free capacity in the SAN Volume Controller cluster is equivalent to the capacity of the data being imported, complete these tasks.

1. Stop all I/O operations from the hosts. Un-map the logical disks that contain the data from the hosts.
2. Create one or more managed disk (MDisk) groups with free capacity. Ensure that the MDisk groups have enough free capacity to contain all of the migrating data and that they are configured in a well-balanced manner.
3. Create an empty MDisk group. This will temporarily contain the data being imported.
4. Create an image mode virtual disk from the first unmanaged-mode MDisk that contains the data to be imported. To do this, perform the following steps:
   a. Map one logical disk from the storage subsystem to the SAN Volume Controller ports.
   b. Issue the **svctask detectmdisk** command on the cluster. The new unmanaged-mode MDisk that is found will correspond with the logical disk mapped in the previous step.
   c. Create an image mode virtual disk from this unmanaged-mode MDisk. Assign it to the empty MDisk group just created.
   d. Repeat for all logical disks as required.
5. If you have decided to continue the I/O operations while you migrate the data into SAN Volume Controller, map all the image mode virtual disks to the hosts using the SAN Volume Controller and continue accessing the data through the SAN Volume Controller.
6. Migrate the data to the MDisk groups that you created in step 2. To do this, perform the following:

a. Select the first image mode virtual disk to be migrated.
b. Migrate this virtual disk from its current MDisk group into one of the MDisk groups created in step 2 on page 238 above. This will migrate all the data from the logical disk into the new free space.
c. When this completes, select the next image mode virtual disk and repeat the previous step.
7. When all the virtual disks have been migrated, the MDisk groups created in step 2 on page 238 will contain the data that was on the image mode virtual disks. This data will be striped across the new groups and will be virtualized.
8. You should now go back and destroy the temporary MDisk group that contained the original image mode virtual disks.
9. Go back to the storage subsystem and re-configure the old arrays and logical disks according to the guidelines.
10. Add this storage back under the SAN Volume Controller and use the old storage to create new virtual disks.

Related reference

"Configuration guidelines" on page 234
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

Migrating data with a smaller amount of free capacity:

If the free capacity in the SAN Volume Controller cluster is smaller than the capacity of the data being imported, complete these tasks.

In this scenario, the free capacity in the SAN Volume Controller cluster is smaller than the capacity of the data being imported. The RAID arrays that contain the existing data will still contain the data after importing the arrays to the cluster.

Attention: This procedure will result in an uneven distribution of data across the managed disks in the managed disk group. The impact of this depends on the number of managed disks (MDisks) that are initially in the managed disk group and how many of these have free capacity.

For example, suppose that you have one MDisk in the destination MDisk group. You bring in image mode logical units from an array on the storage subsystem. You migrate these logical units to the destination MDisk group. These logical units are now striped across the one managed-mode disk. Next, you add another logical unit to the destination MDisk group. Now the MDisk contains two managed-mode disks, but all of the data is on the first managed-mode disk. There will always be more data on the first few managed-mode disks than in the last managed-mode disks. Some of the data must be migrated from the overloaded managed-mode disks to the under-utilized ones.

This procedure may require subsequent migration of data within the MDisk group in order to balance the distribution of data across the MDisks in the group.

1. Select a managed disk (MDisk) group that contains enough free capacity to migrate *all* of the logical disks on the first array that is to be migrated to the cluster.
2. Create an empty MDisk group that will temporarily contain the data being imported.

3. Stop all I/O operations to the logical disks that are to be migrated first, and un-map these disks from their hosts.
4. Create an image mode virtual disk from the first unmanaged-mode MDisk that contains the data to be imported. To do this, perform the following steps:
   a. Map one logical disk from the storage subsystem to the SAN Volume Controller ports.
   b. Issue the **svctask detectmdisk** command on the cluster. The new unmanaged-mode MDisk that is found will correspond with the logical disk mapped in the previous step.
   c. Create an image mode virtual disk from this unmanaged-mode MDisk. Assign it to use the empty MDisk group just created.
   d. Repeat for all logical disks as required.
5. If you have decided to continue the I/O operations while you migrate the data to the SAN Volume Controller cluster, map all the image mode virtual disks to the hosts using the SAN Volume Controller and continue accessing the data through the SAN Volume Controller.
6. Migrate the data into the MDisk groups that you created in step 1 on page 239. To do this, perform the following:
   a. Select the first image mode virtual disk to be migrated.
   b. Migrate this virtual disk from its current MDisk group into one of the MDisk groups created in step 1 on page 239 above. This will migrate all the data from the logical disk into the new free space.
   c. When this completes, select the next image mode virtual disk and repeat the previous step.
7. The RAID array that contains the logical disks can now be re-configured and added to the MDisk group selected in step 1 on page 239. To do this, perform the following steps:
   a. Remove the managed disks from the temporary MDisk group.
   b. At the storage subsystem, the logical disks that have been migrated should be unmapped from the SAN Volume Controller cluster and deleted from the array (if more than one existed).
   c. Assuming the array meets the guidelines, a single logical disk should be created using the entire array capacity.
   d. This new logical disk can be mapped to the SAN Volume Controller ports.
   e. Issue the **svctask detectmdisk** command on the cluster. The new managed-mode MDisk that is found will correspond with the new logical disk created.
   f. Add this managed-mode MDisk to the MDisk group selected in step 1 on page 239.
8. Repeat steps 3 through 7 for the next array.

Related reference

"Configuration guidelines" on page 234
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

Configuring a balanced storage subsystem

The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

There are 2 major steps in this process:

1. Setting the characteristics of the SAN Volume Controller to storage connection(s)
2. Mapping logical unit(s) to these connections such that the SAN Volume Controller can access them.

The virtualization features of the IBM Total Storage SAN Volume Controller enable you to choose how your storage is divided up and presented to hosts. While virtualization provides you with a great deal of flexibility, it also offers the potential for setting up a storage subsystem which can be overloaded. A storage subsystem is overloaded if the quantity of I/O transactions that are being issued by the host systems exceeds the capability of the storage to process those transactions. If a storage subsystem is overloaded, then, at best, it causes delays in the host systems and, at worst, it causes I/O transactions to be timed out in the host which leads to errors being logged by the hosts and I/Os being failed back to applications.

As an extreme example of an overloaded storage subsystem it would be possible to use an IBM Total Storage SAN Volume Controller to virtualize a single RAID array and to divide this storage among 64 host systems. Clearly, if all host systems attempt to access this storage at the same time the single RAID array will be overloaded. The following guidelines are provided to help you configure balanced storage subsystems.

1. Calculate the I/O rate for an array. For each RAID array in the storage subsystem use the following table to calculate the approximate number of I/O operations per second that can be processed by the RAID array. Note that the actual number of I/O operations per second that can be processed will vary depending on the location and length of each I/O, whether the I/O is a read or a write operation and on the specifications of the component disks of the RAID array. For example, a RAID-5 array with eight component disks has an approximate I/O rate of 150 * 7 = 1050.

Table 16. Calculate the I/O rate

Type of RAID Array	Number of component disks in the RAID Array	Approximate I/O rate
RAID-1 (mirrored) arrays	2	300
RAID-3, RAID-4, RAID-5 (striped + parity) arrays	N + 1 parity	150 * N
RAID-10, RAID 0+1, RAID 1+0 (striped + mirrored) arrays	N	150 * N

2. Calculate the I/O rate for a managed disk. If there is a one-to-one relationship between backend arrays and managed disks (this is the recommended configuration) then the I/O rate for a managed disk is the same as the I/O rate of the corresponding array. If an array is divided up into multiple managed disks then the I/O rate per managed disk is the I/O rate of the array divided by the number of managed disks that are using that array.
3. Calculate the I/O rate for a managed disk group. The I/O rate for a managed disk group is simply the sum of the I/O rates of the managed disk within that group.

   For example, a managed disk group contains eight managed disks each of which corresponds to a RAID-1 array. From the table above the I/O rate for each managed disks can be calculated as 300. The I/O rate for the managed disk group is therefore 300 * 8 = 2400.

4. Calculate the impact of FlashCopy relationships. If you are using the FlashCopy feature provided by the IBM Total Storage SAN Volume Controller then you need to consider how much additional I/O will be generated by using this feature as this will reduce the rate at which I/O from host systems that can be processed. When a FlashCopy relationship is copying data any write I/Os from host systems to areas of the source or target virtual disk that have not yet been copied will cause extra I/Os to be generated by the IBM Total Storage SAN Volume Controller to copy the data before the write I/O is actually performed. The effect of using FlashCopy depends on the type of I/O workload being generated by an application:

Table 17. Calculate the impact of FlashCopy relationships

Type of application	Impact to I/O rate	Additional Weighting for FlashCopy
Application is doing no I/O	Insignificant impact	0
Application is only reading data	Insignificant impact	0
Application is only issuing random writes	Up to 50 times as much I/O	49
Application is issuing random reads and writes	Up to 15 times as much I/O	14
Application is issuing sequential reads or writes	Up to 2 times as much I/O	1

   For each virtual disk that is either the source or target of an active FlashCopy relationship consider the type of application that will be using that virtual disk and record the additional weighting for that virtual disk.

   For example, a FlashCopy relationship is being used to provide point in time backups. During the FlashCopy process, a host application generates an I/O workload of random reads and writes to the source virtual disk. A second host application reads the target virtual disk and writes the data to tape to create a backup. The additional weighting for the source virtual disk is 14. The additional weighting for the destination virtual disk is 0.

5. Calculate the I/O rate for virtual disks in a managed disk group. Calculate the number of virtual disks in the managed disk group. Add the additional weighting for each virtual disk that is either the source or a target of an active FlashCopy relationship. Divide the I/O rate of the managed disk group by this number to give an I/O rate per VDisk.

   Example 1: A managed disk group has an I/O rate of 2400 and contains 20 virtual disks. There are no FlashCopy relationships. The I/O rate per virtual disk is 2400 / 20 = 120.

   Example 2: A managed disk group has an I/O rate of 5000 and contains 20 virtual disks. There are two active FlashCopy relationships which have source virtual disks in this managed disk group. Both source virtual disks are being accessed by applications issuing random reads and write and hence the additional weighting for each of these virtual disks is 14. The I/O rate per virtual disk is 5000 / (20 + 14 + 14) = 104.

6. Determine whether the storage subsystem is overloaded. The figure determined at step 4 provides some indication of how many I/O operations per second can be processed by each virtual disk in the managed disk group. If you know how many I/O operations per second your host applications generate you can compare these figures to determine if the system is overloaded. If you do not know how many I/O operations per second your host applications generate then

you can either measure this (for example by using the I/O statistics facilities provided by the IBM Total Storage SAN Volume Controller to measure the I/O rate of your virtual disks) or use the following table as a guideline.

Table 18. Determine if the storage subsystem is overloaded

Type of Application	I/O rate per virtual disk
Applications that generate a high I/O workload	200
Applications that generate a medium I/O workload	80
Applications that generate a low I/O workload	10

7. Interpret the result. If the I/O rate generated by the application exceeds the I/O rate you calculated per virtual disk this indicates that you could be overloading your storage subsystem and you should monitor the system carefully to see if the backend storage is actually limiting the overall performance of your system. It is also possible that the calculation above is too simplistic to model your use of storage, for example the calculation assumes that your applications generate the same I/O workload to all virtual disks.

   One method you can use to monitor the performance of your storage subsystem is to use the I/O statistics facilities provided by the IBM Total Storage SAN Volume Controller to measure the I/O rate of your managed disks. Alternatively you could use the performance and I/O statistics facilities provided by your storage subsystems.

   If you find your storage subsystem is over loaded there are several actions that can be take to resolve the problem:

   a. Adding more backend storage to the system will allow you to increase the quantity of I/O that can be processed by your storage subsystem. The virtualization and data migration facilities provided by the IBM Total Storage SAN Volume Controller can be used to redistribute the I/O work load of virtual disks across a greater number of managed disks without having to take the storage offline.
   b. Stop any unessential FlashCopy relationships. This will reduce the amount of I/O operations submitted to the backend storage. If you are making many FlashCopy's in parallel then consider starting less FlashCopy relationships in parallel.
   c. The I/O workload generated by a host can often be limited by adjusting the queue depth (for example, the maximum number of I/O operations that are submitted in parallel). Depending on the type of host and type of host bus adapters it may be possible to limit the queue depth per virtual disk and/or limit the queue depth per host bus adapter. An alternative method of limiting the I/O workload generated by a host would be to use the I/O governing features provided by the IBM Total Storage SAN Volume Controller. These techniques may be particularly applicable if using a mixture of different host systems to prevent one host system from saturating an I/O subsystem to the detriment of the other host systems. Note that although these techniques may be used to avoid I/O time-outs it still means the performance of your system is being limited by the amount of storage.

Expanding a logical unit

A logical unit can be expanded using vendor-specific disk-configuration software.

Some storage subsystems enable you to expand the size of a logical unit (LU) using configuration software that is provided. However, the SAN Volume Controller cannot use extra capacity that is provided in this way. Perform the following task to ensure that this additional capacity is available to the SAN Volume Controller.

The logical unit has increased in size and this additional space must be made available for use.

Perform the following steps to ensure that this additional capacity is available to the SAN Volume Controller:

1. Issue the **svctask migrateexts** command to migrate all the data from the MDisk.

   Note:

   a. For managed mode MDisks, issue the **svctask rmmdisk** command to remove the MDisk from the MDisk group.
   b. For image mode MDisks, issue the **svctask chmdisk** to change the mode of the image mode disk to ″unmanaged″.
2. Issue the **svctask includemdisk <MDisk number>** command. Where *<MDisk number>* is the number of the MDisk that has been expanded.
3. Issue the **svctask detectmdisk** command to re-scan the fibre-channel network for the new managed disk that you have included. This may take a few minutes.
4. Issue the **svcinfo lsmdisk** command to display the additional capacity that has been expanded.

The extra capacity is available for use by the SAN Volume Controller.

Modifying a logical unit mapping

You can modify a logical unit mapping.

The logical unit mapping must be modified, therefore, the logical unit number (LUN) is being changed.

1. Issue the **svctask migrateexts** command to migrate all the data from the MDisk.

   Note:

   a. For managed mode MDisks, issue the **svctask rmmdisk** command to remove the MDisk from the MDisk group.
   b. For image mode MDisks, issue the **svctask chmdisk** to change the mode of the image mode disk to ″unmanaged″.
2. De-configure the mapping on the storage subsystem so that the logical unit is not visible to the SAN Volume Controller.
3. Issue the **svctask includemdisk <MDisk number>** command. Where *<MDisk number>* is the number of the MDisk that you want to modify.
4. Issue the **svctask detectmdisk** command to re-scan the fibre-channel network for the managed disk that you want to rediscover. This may take a few minutes.
5. Issue the **svcinfo lsmdisk** command to verify that the MDisk has been removed. If the MDisk is still displayed, repeat steps 3 and 4. The MDisk should now have been removed from the list of valid candidates.
6. Configure the mapping of the logical unit to the new logical unit number.
7. Issue the **svctask detectmdisk** command.

8. Issue the **svcinfo lsmdisk** command to check that the MDisk candidate(s) now have the correct LUN.

The MDisk candidate(s) now have the correct LUN.

Storage subsystem tasks using the SAN Volume Controller Console

You can determine the storage subsystem name from the SAN Volume Controller name, rename a storage subsystem, and add and remove a storage subsystem using the SAN Volume Controller Console.

Related tasks

"Determining a storage subsystem name from its SAN Volume Controller name using the SAN Volume Controller Console"
You can determine a storage subsystem name from its SAN Volume Controller name.

"Renaming a storage subsystem"
You can rename a storage subsystem from the Renaming a Disk Controller System panel.

"Changing a configuration for an existing storage subsystem" on page 246
You must change the configuration for a storage subsystem in order to delete and replace logical units.

"Adding a new storage controller to a running configuration using the SAN Volume Controller Console" on page 247
You can add a new storage controller to your SAN at any time.

"Removing a storage subsystem using the SAN Volume Controller" on page 248
You can replace or decommission an old storage subsystem.

"Removing managed disks that represent de-configured LUs" on page 249
When de-configuring or removing LUs from your storage subsystem, the managed disks (MDisks) that represent those LUs may still exist in the cluster. Use the following procedure to remove the MDisks.

Determining a storage subsystem name from its SAN Volume Controller name using the SAN Volume Controller Console

You can determine a storage subsystem name from its SAN Volume Controller name.

1. Click **Work with Disk Controllers**.
2. Select the name link for the storage subsystem in question. Write down the WWNN. This can be used to determine the actual storage subsystem by launching the native user interface or using the command line tools it provides to verify the actual storage subsystem that has this WWNN.

Related information

"Storage subsystem tasks using the SAN Volume Controller Console"
You can determine the storage subsystem name from the SAN Volume Controller name, rename a storage subsystem, and add and remove a storage subsystem using the SAN Volume Controller Console.

Renaming a storage subsystem

You can rename a storage subsystem from the Renaming a Disk Controller System panel.

Perform the following steps to rename a storage subsystem:

1. Click **Work with Managed Disks** in the portfolio.

2. Click **Disk Controller Systems** in the portfolio. The Disk Controller Systems panel is displayed.
3. Select the storage subsystem that you want to rename and select **Rename a disk controller system** from the list. Click **Go**. The Renaming a Disk Controller System panel is displayed.

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related information

"Storage subsystem tasks using the SAN Volume Controller Console" on page 245
You can determine the storage subsystem name from the SAN Volume Controller name, rename a storage subsystem, and add and remove a storage subsystem using the SAN Volume Controller Console.

Changing a configuration for an existing storage subsystem

You must change the configuration for a storage subsystem in order to delete and replace logical units.

Perform the following steps to delete existing logical units (LUs) and replace them with new LUs:

1. Delete the managed disks (MDisks), that are associated with the LUs, from their MDisk groups.
2. Delete the existing LUs using the configuration software of the storage subsystem.
3. Delete the associated MDisks from the cluster by running the **svctask detectmdisk** command.
4. Configure the new LUs using the configuration software of the storage subsystem.
5. Add the new LUs to the cluster by running the **svctask detectmdisk** command.

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

"Managed disk groups" on page 26
An *MDisk group* is a collection of MDisks that jointly contain all the data for a specified set of virtual disks (VDisks).

"Managed disks" on page 23
A managed disk (MDisk) is a logical disk (typically a RAID array or partition thereof) that a storage subsystem has exported to the SAN fabric to which the nodes in the cluster are attached.

Related information

"Storage subsystem tasks using the SAN Volume Controller Console" on page 245
You can determine the storage subsystem name from the SAN Volume Controller name, rename a storage subsystem, and add and remove a storage subsystem using the SAN Volume Controller Console.

Adding a new storage controller to a running configuration using the SAN Volume Controller Console

You can add a new storage controller to your SAN at any time.

Follow the switch zoning guidelines and also ensure the controller is setup correctly for use with the SAN Volume Controller.

You should create one or more arrays on the new controller. It is recommend that you use RAID-5, RAID-1 or RAID-0+1 (sometimes called RAID-10) for maximum redundancy and reliability. Generally 5+P arrays are recommended. If your controller provides array partitioning we recommend that you create a single partition from the entire capacity available in the array, remember the LUN number that you assign to each partition as you will need this later. You should also follow the mapping guidelines (if your storage controller requires LUN mapping) to map the partitions or arrays to the SAN Volume Controller ports. You can determine the SAN Volume Controller ports by following the procedure for determining WWPNs..

1. To ensure that the cluster has detected the new storage (MDisks) click **Work with MDisks** and select the **Detect MDisks** task.
2. The controller itself will have automatically been assigned a default name. If you are unsure which controller is presenting the MDisks click **Work with Disk Controllers**. You should see a new controller listed (the one with the highest numbered default name). You must determine the storage controller name to validate that this is the correct controller.
3. Close and reopen the **Work with MDisks** panel using the filter panel to select a mode of **unmanaged** and a controller name that corresponds with the new controller's name. The MDisks shown should correspond with the RAID arrays or partitions you have created. Remember the field controller LUN number, this corresponds with the LUN number you assigned to each of the arrays or partitions.
4. It is recommended that you create a new managed disk group and add only the RAID arrays that belong to the new controller to this MDisk group. You should also avoid mixing RAID types, so for each set of RAID array types (for example, RAID-5, RAID-1) you should create a new MDisk group. Give this MDisk group an appropriate name, so if your controller is called FAST650-fred, and the MDisk group contains RAID-5 arrays, call it something like F600-fred-R5.
5. Click **Work with MDisk group** from the portfolio. Select the **Create MDisk group** task. On the new panel, enter the name you wish to give this group, select the MDisks you wish to add from the list and click **Add**. Select the extent size you wish this group to have and click **OK**.

Related tasks

"Exposing logical units on your existing storage to the cluster via switch zoning" on page 110
You can expose logical units on your existing storage to the cluster via switch zoning.

"Determining the WWPNs for a node using the SAN Volume Controller Console" on page 124
To determine the WWPNs of a node using the SAN Volume Controller Console, complete these steps.

Related reference

"Configuring and servicing storage subsystems" on page 233
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

“Switch zoning for the SAN Volume Controller” on page 65
Consider these constraints when you zone a switch.

Related information

“Storage subsystem tasks using the SAN Volume Controller Console” on page 245
You can determine the storage subsystem name from the SAN Volume Controller name, rename a storage subsystem, and add and remove a storage subsystem using the SAN Volume Controller Console.

Removing a storage subsystem using the SAN Volume Controller

You can replace or decommission an old storage subsystem.

All the MDisks that belong to the storage subsystem that is being decommissioned belong to a single MDisk group. You need to repeat this procedure for each MDisk group in turn before removing the old device.

This procedure takes you through adding the new device, migrating the data off of the old device and removing the old MDisks.

This function can also be performed by migrating all the VDisks that are using storage in this MDisk group to another MDisk group. This procedure has an advantage if you wish to consolidate the VDisks in a single or new group. However, you can only migrate a single VDisk at a time. The procedure outlined below will migrate all the data through a single command. If you wish to migrate the VDisks however, follow the procedure for all VDisks that are using this group. You can determine the relationship between VDisks and MDisks by following the procedure.

This procedure can also be used to remove or replace a single MDisk in a group. If an MDisk has suffered a partial failure, such as a degraded array, and you can still read the data from the disk but cannot write to it, you can follow this procedure to replace just one MDisk. In steps 1 and 3, you can add or remove a single MDisk rather than a list of MDisks.

1. Add new storage.
2. Select the MDisk group that contains the old MDisks you are decommissioning. Select the **Add MDisk** task. On the task dialog, select the new MDisks from the list and click **Add**. Click **OK** to complete the task.
3. You should now have an MDisk group that contains the old MDisks (those to be decommissioned) and the new MDisks (those that are replacing them). Ensure that the capacity of the new MDisks is the same or exceeds that of the old MDisks before proceeding.
4. Force delete the old MDisks from the group. This will migrate all the data from the old MDisks to the new MDisks. Select the **Remove MDisks** task. Selectthe MDisks you wish to remove and click **Add**. Click **OK** to complete the task. When prompted click **Forced Delete**. Depending upon the number and size of the MDisks, and the number and size of the VDisks that are using these MDisks, this operation will take some time to complete although the task will complete immediately.
5. The only way to check progress is by using the command line interface. Issue the following command:

   ```
   svcinfo lsmigrate
   ```
6. When all the migration tasks have completed, for example, the command in step 5 returns no output, you can safely remove the old storage subsystem from the SAN.

7. Once you have removed the old storage subsystem from the SAN, rerun the detect MDisks task to remove the entries for the old MDisks.

Related tasks

"Adding a new storage controller to a running configuration using the CLI" on page 250
You can add a new disk controller system to your SAN at any time using the using the command-line interface (CLI).

Related information

"Storage subsystem tasks using the SAN Volume Controller Console" on page 245
You can determine the storage subsystem name from the SAN Volume Controller name, rename a storage subsystem, and add and remove a storage subsystem using the SAN Volume Controller Console.

Removing managed disks that represent de-configured LUs

When de-configuring or removing LUs from your storage subsystem, the managed disks (MDisks) that represent those LUs may still exist in the cluster. Use the following procedure to remove the MDisks.

MDisks exist in the cluster that are no longer accessible to it. That is because the LUs that these MDisks represent have been deconfigured or removed from the storage subsystem. You should remove these MDisks.

1. Run the **svctask includemdisk** command on all the affected MDisks.
2. Run the **svctask rmmdisk** command on all affected MDisks. This puts the MDisks into the unmanaged mode.
3. Run the **svctask detectmdisk** command. The cluster detects that the MDisks no longer exist in the storage subsystem.

All of the MDisks that represent de-configured LUs are removed from the cluster.

Related tasks

"Discovering MDisks using the CLI" on page 168
You can discover MDisks using the command-line interface (CLI).

Related information

"Storage subsystem tasks using the SAN Volume Controller Console" on page 245
You can determine the storage subsystem name from the SAN Volume Controller name, rename a storage subsystem, and add and remove a storage subsystem using the SAN Volume Controller Console.

Controller tasks using the CLI

You can determine the storage controller name from the SAN Volume Controller name, add and remove a controller using the command-line interface (CLI).

Related tasks

"Determining a storage subsystem name from its SAN Volume Controller name using the CLI" on page 250
You can determine a storage subsystem name from its SAN Volume Controller name using the command-line interface (CLI).

"Adding a new storage controller to a running configuration using the CLI" on page 250
You can add a new disk controller system to your SAN at any time using the using the command-line interface (CLI).

"Removing a storage subsystem using the CLI" on page 251
You can replace or decommission an old storage subsystem using the command-line interface (CLI).

Determining a storage subsystem name from its SAN Volume Controller name using the CLI

You can determine a storage subsystem name from its SAN Volume Controller name using the command-line interface (CLI).

1. List the storage subsystem by issuing the following command:

   ```
   svcinfo lscontroller
   ```

   Remember the name or ID for the storage subsystem you want to determine.
2. For the device in question, issue the following command:

   ```
   svcinfo lscontroller <controllername/id>
   ```

   where *<controllername/id>* is the name or ID. Remember the WWNN for the device. Make a written record of it. The WWNN can be used to determine the actual storage subsystem by launching the native user interface or using the command line tools it provides to verify the actual storage subsystem that has this WWNN.

Adding a new storage controller to a running configuration using the CLI

You can add a new disk controller system to your SAN at any time using the using the command-line interface (CLI).

You can add a new disk controller system to your SAN at any time. Follow the switch zoning guidelines in the section about switch zoning. Also, ensure the controller is setup correctly for use with the SAN Volume Controller.

You should create one or more arrays on the new controller. It is recommend that you use, RAID-5, RAID-1 or RAID-0+1 (sometimes called RAID-10) for maximum redundancy and reliability. Generally 5+P arrays are recommend. If your controller provides array partitioning we recommend that you create a single partition from the entire capacity available in the array, remember the LUN number that you assign to each partition as you will need this later. You should also follow the mapping guidelines (if your disk controller system requires LUN mapping) to map the partitions or arrays to the SAN Volume Controller ports.

1. To ensure that the cluster has detected the new storage (MDisks) issue the following command:

   ```
   svctask detectmdisk
   ```
2. The controller itself will have automatically been assigned a default name. If you are unsure which controller is presenting the MDisks, list the controllers by issuing the following command:

   ```
   svcinfo lscontroller
   ```

   You should see a new controller listed (the one with the highest numbered default name). Remember the controller name and follow the instructions in the section about determining a disk controller system name.
3. You should give this controller a name that you can easily use to identify it. Issue the following command:

```
svctask chcontroller -name <newname> <oldname>
```

4. List the unmanaged MDisks by issuing the following command:

   ```
   svcinfo lsmdisk -filtervalue mode=unmanaged:controller_name=<new_name>
   ```

   These MDisks should correspond with the RAID arrays or partitions you have created. Remember the field controller LUN number. This corresponds with the LUN number you assigned to each of the arrays or partitions.
5. It is recommended that you create a new managed disk group and add only the RAID arrays that belong to the new controller to this MDisk group. You should also avoid mixing RAID types, so for each set of RAID array types (for example, RAID-5, RAID-1) you should create a new MDisk group. Give this MDisk group an appropriate name, so if your controller is called FAST650-fred, and the MDisk group contains RAID-5 arrays, call it something like F600-fred-R5). Issue the following command:

   ```
   svctask mkmdiskgrp -ext 16 -name <mdisk_grp_name>
   -mdisk <colon separated list of RAID-x mdisks returned
   in step 4.
   ```

   This will create a new MDisk group with an extent size of 16MB.

Related tasks

"Determining the WWPNs of a node using the CLI" on page 181
You can determine the WWPNs of a node using the command-line interface (CLI).

Related reference

"Configuring and servicing storage subsystems" on page 233
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

"Switch zoning for the SAN Volume Controller" on page 65
Consider these constraints when you zone a switch.

Removing a storage subsystem using the CLI

You can replace or decommission an old storage subsystem using the command-line interface (CLI).

All the MDisks that belong to the storage subsystem that is being decommissioned belong to a single MDisk group. You need to repeat this procedure for each MDisk group in turn before removing the old device.

This procedure takes you through adding the new device, migrating the data off of the old device and removing the old MDisks.

This function can also be performed by migrating all the VDisks that are using storage in this MDisk group to another MDisk group. This procedure has an advantage if you wish to consolidate the VDisks in a single or new group. However, you can only migrate a single VDisk at a time. The procedure outlined below will migrate all the data through a single command.

This procedure can also be used to remove or replace a single MDisk in a group. If an MDisk has suffered a partial failure, such as a degraded array, and you can still read the data from the disk but cannot write to it, you can follow this procedure to replace just one MDisk. In the following procedure, add or remove a single MDisk rather than a list of MDisks.

1. Add the new storage subsystem to your cluster configuration.
2. Issue the following command:

   ```
   svctask addmdisk -mdisk <mdiskx:mdisky:mdiskz...> <mdisk_grp_name>
   ```

   Where *<mdiskx:mdisky:mdiskz...>* are the names of new MDisks that have a total capacity that is larger than the MDisks that are being decommissioned. *<mdisk_grp_name>* is the name of the MDisk group that contains the MDisks that are being decommissioned.
3. You should now have an MDisk group that contains the old MDisks (those to be decommissioned) and the new MDisks (those that are replacing them). Ensure that the capacity of the new MDisks is the same or exceeds that of the old MDisks before proceeding.
4. Force delete the old MDisks from the group. This will migrate all the data from the old MDisks to the new MDisks. Issue the following command:

   ```
   svctask rmmdisk -force -mdisk <colon separated
    mdisk list of all the old mdisks> <mdisk_grp_name>
   ```

   Depending upon the number and size of the MDisks, and the number and size of the VDisks that are using these MDisks, this operation will take some time to complete although the command will return immediately.
5. Check progress by issuing the following command:

   ```
   svcinfo lsmigrate
   ```
6. When all the migration tasks have completed, for example, the command in step 5 returns no output, you can safely remove the old device from the SAN.
7. Once you have removed the old device from the SAN, rerun the **svctask detectmdisk** command to remove the entries for the old MDisks.

Related tasks

"Removing a storage subsystem using the CLI" on page 251
You can replace or decommission an old storage subsystem using the command-line interface (CLI).

"Determining the relationship between VDisks and MDisks using the CLI" on page 183
You can determine the relationship between VDisks and MDisks using the command-line interface (CLI).

"Migrating VDisks between MDisk groups using the CLI" on page 200
You can migrate VDisks between MDisk groups using the command-line interface (CLI).

Creating a quorum disk

A quorum disk is used to resolve tie-break situations when the ″voting set″ of nodes disagree on the current cluster state.

Quorum disk creation and extent allocation

During quorum disk discovery, the system assesses each logical unit (LU) to determine its potential use as a quorum disk. From the set of eligible LUs, the system nominates three candidates and selects one.

To be a candidate for a quorum disk, an LU must meet the following criteria:

- It must be in managed space mode.
- It must be visible to all nodes in the cluster.
- It must be presented by a storage subsystem that is an approved host for quorum disks.
- It must have sufficient free extents to hold the cluster state and the configuration metadata.

If possible, the quorum disk candidates will be presented by different devices. Once the system has selected a quorum disk, it does not attempt to ensure that the other candidates are presented by different devices. The set of quorum disk candidates can be updated by configuration activity if other eligible LUs are available.

Quorum disk candidates may not exist. If so, the following will occur:

- If LUs do not exist in managed space mode, then there are no quorum disk candidates. An error will be logged.
- If LUs exist in managed space mode and quorum disk candidates do not exist because they do not meet the eligibility criteria, an error will be logged.

Manual discovery

When creating or removing LUNs on a storage subsystem, the MDisk view will not be automatically updated. It must be updated manually.

To ensure that the MDisk view is updated, type **svctask detectmdisk** to start a manual discovery.

Servicing storage subsystems

When servicing your storage subsystems, it is imperative that you follow the service instructions contained in the vendor documentation.

If the instructions state that all I/O operations be stopped for a particular service action, ensure that the SAN Volume Controller has terminated all FlashCopy activity and that all data migration requests are complete.

Configuring the EMC CLARiiON subsystem

This part provides information about configuring the EMC CLARiiON storage system for its attachment to a SAN Volume Controller.

Related concepts

"Managed disk groups" on page 26
An *MDisk group* is a collection of MDisks that jointly contain all the data for a specified set of virtual disks (VDisks).

Related tasks

"Configuring the EMC CLARiiON controller with Access Logix installed" on page 255
To configure the EMC CLARiiON controller that has Access Logix installed on it, complete these tasks.

"Configuring the EMC CLARiiON controller (Access Logix not installed)" on page 257
To configure an EMC CLARiiON controller that does not have Access Logix installed, complete these steps.

“Creating managed disk groups” on page 107
You can create a new managed disk (MDisk) group using the Create a Managed Disk Group wizard.

Related reference

“Access Logix”
Access Logix is an optional feature of the firmware code. This provides the functionality that is known as LUN Mapping or LUN Virtualization.

“Supported models of the EMC CLARiiON” on page 258
The following table lists the supported models of the EMC CLARiiON. See the following Web site for the latest supported models. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

“Supported firmware levels for the EMC CLARiiON” on page 258
See the following Web site for specific firmware levels and the latest supported hardware. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

“Concurrent maintenance on the EMC CLARiiON” on page 258
Concurrent maintenance of the EMC CLARiiON is supported with the SAN Volume Controller.

“User interface on CLARiiON” on page 259
Ensure that you are familiar with the user interface that supports the CLARiiON subsystem.

“Sharing the EMC CLARiiON between a host and the SAN Volume Controller” on page 259
The EMC CLARiiON can be shared between a host and a SAN Volume Controller.

“Switch zoning limitations for the EMC CLARiiON” on page 260
There are limitations in switch zoning for the SAN Volume Controller and EMC CLARiiON.

“Quorum disks on the EMC CLARiiON” on page 260
This topic provides information about quorum disks on the EMC CLARiiON.

“Advanced functions for the EMC CLARiiON” on page 261
Some EMC CLARiiON advanced functions are not compatible with the SAN Volume Controller.

“Logical unit creation and deletion on the EMC CLARiiON” on page 261
Binding an LU to a RAID group can take a significant amount of time on the EMC CLARiiON.

“Configuring settings for the EMC CLARiiON” on page 261
A number of settings and options are available through the EMC CLARiiON configuration interface.

Access Logix

Access Logix is an optional feature of the firmware code. This provides the functionality that is known as LUN Mapping or LUN Virtualization.

It is possible to determine whether Access Logix is installed using the SAN Volume Controller Console via the software tab in the storage subsystem properties page. Once Access Logix is installed it can be disabled but not removed. This gives two modes of operation, as follows:

- **Access Logix not installed:** In this mode, all LUNs are accessible from all target ports by any host. Therefore, the SAN fabric must be zoned to ensure that only the SAN Volume Controller can access the target ports.

- **Access Logix enabled:** In this mode of operation, a storage group can be formed from a set of LUNs. Only the hosts assigned to the storage group are allowed access to these LUNs.

Configuring the EMC CLARiiON controller with Access Logix installed

To configure the EMC CLARiiON controller that has Access Logix installed on it, complete these tasks.

It is assumed at this point that the EMC CLARiiON controller is not connected to the SAN Volume Controller. It is also assumed that you already have a RAID controller with logical units and you've identified which LUs will be presented to the SAN Volume Controller.

With Access Logix installed, the SAN Volume Controller will not have access to the storage controllers LUs. In order to give SAN Volume Controller access to a LU, the SAN Volume Controller and LU must be associated using the EMC CLARiiON configuration tools. The association is formed by creating a Storage Group containing the LU and the SAN Volume Controller.

Related tasks

"Configuring your storage groups" on page 257
To allow the SAN Volume Controller to access LUs configured on the EMC CLARiiON controller, complete these steps.

Registering the SAN Volume Controller ports with the EMC CLARiiON

To register SAN Volume Controller ports with an EMC CLARiiON controller that has Access Logix installed, complete these steps.

At this point, the EMC CLARiiON controller is not connected to the SAN Volume Controller. It is assume you already have a RAID controller with logical units (LUs) and you've identified which LUs will be presented to the SAN Volume Controller.

Each initiator port (WWPN) must be registered against a host name and against a target port to which access is granted. These associations are listed in the following table.

If a host has multiple initiator ports, there will be multiple table entries with the same host name. If a host is allowed access via multiple target ports, there will be multiple table entries. For SAN Volume Controller hosts, it is recommended that all WWPN entries carry the same host name.

Option	EMC CLARiiON Default Setting	SAN Volume Controller Required Setting
WWPN	n/a	Any
WWN	n/a	Any
Host name	n/a	Any
SP port	n/a	Any
Initiator type	3	3
ArrayCommPath	enable	disable
Failover mode	0	2
Unit Serial Number	array	Any

1. Connect the fibre channel and zone the fabric as required.

2. Issue an **svctask detectmdisk** command .
3. From the Enterprise Storage window, right-click on the storage subsystem.
4. Select **Connectivity Status**. The Connectivity Status window is displayed.
5. Click **New**. The Create Initiator Record window is displayed.
6. Wait for the list of SAN Volume Controller ports to appear in the dialog box. (Identify them from the WWPN.) This can take several minutes.
7. Click **Group Edit**.
8. Select all instances of all the SAN Volume Controller ports in the Available dialog box. Then click the right arrow to move them to the selected box.
9. Fill in the HBA WWN field. You must know the following information:
   a. WWNN of each SAN Volume Controller in the cluster
   b. WWPN of each port ID for each node on the cluster

   The HBA WWN field is made up of the WWNN and the WWPN for the SAN Volume Controller port. The following output shows an example:

   ```
   50:05:07:68:01:00:8B:D8:50:05:07:68:01:20:8B:D8
   ```
10. Select A in the field marked SP and 0 in the SP Port field.
11. For the Initiator Type field, select **CLARiiON Open** in the drop down list.
12. Deselect the ArrayCommPath checkbox if it has been selected.
13. Select **2** from the Failover Mode field drop down list.

    Note: Failure to select failover mode 2 will prevent SVC from being able to failover I/O. Your data may become unavailable in the event of a single failure.

    Note:
    a. If this is the first time that a port has been registered, ensure that you select the New Host option. Otherwise, select Existing Host.
    b. Ensure that the same host name is entered for each port that is registered.
14. Assign a host name in the Host Name field.
15. Click **OK**.
16. Specify the IP address of your switch. The EMC CLARiiON will not use this IP address. However it must be unique (within the EMC CLARiiON) to prevent errant behavior by Navisphere.
17. Perform step 10 for all possible combinations. The following example shows the different combinations of a subsystem with four ports:
    a. SP: A SP Port: 0
    b. SP: A SP Port: 1
    c. SP: B SP Port: 0
    d. SP: B SP Port: 1
18. Refer to steps 1 on page 255 through 17 to register the rest of your SAN Volume Controller WWPNs.

All your WWPNs are registered against the host name that you specified.

Related tasks

“Adding nodes to a cluster using the CLI” on page 163
You can add nodes to a cluster using the CLI.

Configuring your storage groups

To allow the SAN Volume Controller to access LUs configured on the EMC CLARiiON controller, complete these steps.

LUN mapping is provided by Access Logix in the following manner:

Note: Storage groups can only be configured if Access Logix is installed and enabled.

- A subset of LUs can form a storage group.
- An LU can be in multiple storage groups.
- A host can be added to a storage group. This host has access to all LUs in the storage group.
- A host *cannot* be added to a second storage group.

1. From the Enterprise Storage window, right-click on the storage subsystem.
2. Select **Create Storage Group**. The Create Storage Group window is displayed.
3. Choose a name for your storage group. Enter this name in the Storage Group Name field.
4. Select **Dedicated** in the Sharing State field.
5. Click **OK**. The storage group has been created.
6. From the Enterprise Storage window, right-click the storage group that was just created.
7. Select **Properties**. The Storage Group Properties window is displayed.
8. From the Storage Group Properties window, perform the following steps:
   a. Select the **LUNs** tab.
   b. Select the LUNs you want SAN Volume Controller to manage in the Available LUNs table.

      Attention: Ensure that the logical units that you have selected are not used by another storage group.
   c. Click the forward arrow button.
   d. Click **Apply**. A Confirmation window is displayed.
   e. Click **Yes** to continue. A Success window is displayed.
   f. Click **OK**.
   g. Select the **Hosts** tab.
   h. Select the host you created from the steps in "Registering the SAN Volume Controller Ports with your EMC CLARiiON."

      Attention: Ensure only SAN Volume Controller hosts (initiator ports) are in the storage group.
   i. Click the forward arrow button.
   j. Click **OK**. The Confirmation window is displayed.
   k. Click **Yes** to continue. A Success window is displayed.
   l. Click **OK**.

Configuring the EMC CLARiiON controller (Access Logix not installed)

To configure an EMC CLARiiON controller that does not have Access Logix installed, complete these steps.

If Access Logix is not installed on an EMC CLARiiON controller, all LUs that were created on the controller may be used by the SAN Volume Controller. No further configuration of the EMC CLARiiON controller is necessary.

Configure the switch zoning such that no hosts can access these LUs.

Related reference

"Switch zoning for the SAN Volume Controller" on page 65
Consider these constraints when you zone a switch.

Supported models of the EMC CLARiiON

The following table lists the supported models of the EMC CLARiiON. See the following Web site for the latest supported models. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Table 19. Supported models of the EMC CLARiiON

Model
FC4700-1
FC4700-2
CX200
CX300
CX400
CX500
CX600
CX700

Supported firmware levels for the EMC CLARiiON

See the following Web site for specific firmware levels and the latest supported hardware. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Concurrent maintenance on the EMC CLARiiON

Concurrent maintenance of the EMC CLARiiON is supported with the SAN Volume Controller.

Concurrent maintenance is the ability to perform I/O operations to a controller while simultaneously performing maintenance on it. This topic lists the components of the EMC CLARiiON that can be maintained in this way.

The EMC CLARiiON model FC series allows for concurrent replacement of the following components:

- Disk drives
- Controller fans (Fans must be replaced within 2 minutes or controllers will shut down.)
- Disk enclosure fans (Fans must be replaced within 2 minutes or controllers will shut down.)
- Controller (Service Processor: cache must be disabled first)
- Fibre Channel Bypass cards (LCC)
- Power supplies (Fans must be removed first.)
- UPS battery (SPS)

The EMC CLARiiON model CX series allows for concurrent replacement of the following components:

- Disk drives

- Controller (Service processor or drawer controller)
- Power/cooling modules (Modules must be replaced within 2 minutes or controllers will shut down.)
- UPS battery (SPS)

Note:

1. EMC CLARiiON procedures for concurrent upgrade must be followed in all cases.
2. The CX Series also has a feature called Data In Place Upgrade which allows you to upgrade from one model to another (for example, from the CX200 to the CX600) with no data loss or migration required. This is **not** supported by the SAN Volume Controller.

User interface on CLARiiON

Ensure that you are familiar with the user interface that supports the CLARiiON subsystem.

NaviSphere or Navicli

The following user interface applications are available with the EMC CLARiiON:

- The Web-based application NaviSphere can be accessed from any Web browser.
- The command-line application Navicli is installed as part of the NaviSphere Agent software (the host software).

Note: Some options and features are only accessible through the CLI.

Communication with the controller in both cases in out-of-band. Therefore, the host does not need to be connected to the storage over fibre-channel and moreover must not be connected without Access Logix.

Web Server

A Web server is running on each of the controllers on the subsystem. During normal operation, the user interface only allows basic monitoring of the subsystem and displays an error log. If a controller is put into diagnostic mode by pressing the reset button on the controller, the user interface allows for firmware upgrades and subsystem configuration resets.

Sharing the EMC CLARiiON between a host and the SAN Volume Controller

The EMC CLARiiON can be shared between a host and a SAN Volume Controller.

- Split controller access is only supported when Access Logix is installed and enabled.
- A host cannot be connected to both the SAN Volume Controller and EMC CLARiiON at the same time.
- LUs must not be shared between a host and a SAN Volume Controller.
- Partitions in a RAID group must not be shared between a host and a SAN Volume Controller.

Related concepts

"Storage subsystems" on page 54
Follow these rules when you are planning the configuration of storage subsystems in the SAN fabric.

Switch zoning limitations for the EMC CLARiiON

There are limitations in switch zoning for the SAN Volume Controller and EMC CLARiiON.

The number of connections (process logins) consumed by the SAN Volume Controller cluster and the EMC CLARiiON must be carefully considered. For a single fabric, the number of connections is:

- Number of SAN Volume Controller nodes * the number of initiator ports * the number of target ports

If this exceeds the subsystem capabilities, the solution is to reduce the number of initiator or target ports in the configuration, without introducing a single point of failure.

- When reducing the number of *initiator* ports, use only two of the four ports on each SAN Volume Controller node (one per HBA) and configure two fabrics, or fabric zones, such that these are the only initiator ports visible to each target port.
- When reducing the number of *target* ports, ensure that ports from more than one controller are used.

The EMC CLARiiON CX200 provides two ports and supports 30 connections. Using a single SAN fabric, a four-node cluster requires 32 connections (4 * 4 * 2). This exceeds the CX200 capability and exposes the SAN Volume Controller cluster integrity. Since there are only two target ports available, the solution is to reduce the number of initiator ports. This consumes only 16 of the available 30 connections.

Note: The EMC CLARiiON CX200 cannot be used in an eight-node cluster configuration because the number of initiator ports cannot be fewer than 16 (two per node) and the number of target ports cannot be fewer than two. This consumes 32 connections and still exceeds the subsystem limit.

EMC CLARiiON FC4700 and CX400 systems provide four target ports and support 64 connections. Using a single SAN fabric, a four-node cluster requires 64 connections (4 * 4 * 4). This equals the EMC CLARiiON capabilities and is therefore only a problem if split support with other hosts is required. Reducing either the number of initiator ports or target ports consumes 32 of the available 64 connections.

EMC CLARiiON CX600 provides eight target ports and supports 128 connections. A four-node cluster consumes all 128 connections (4 * 4 * 8). An eight-node cluster exceeds the connection limit an done of the reduction schemes must be used.

Related reference

"Switch zoning for the SAN Volume Controller" on page 65
Consider these constraints when you zone a switch.

Quorum disks on the EMC CLARiiON

This topic provides information about quorum disks on the EMC CLARiiON.

The EMC CLARiiON supports quorum disks.

A SAN Volume Controller configuration that only includes the EMC CLARiiON is permitted

Related reference

"Configuring and servicing storage subsystems" on page 233
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

Related information

"Creating a quorum disk" on page 252
A quorum disk is used to resolve tie-break situations when the ″voting set″ of nodes disagree on the current cluster state.

Advanced functions for the EMC CLARiiON

Some EMC CLARiiON advanced functions are not compatible with the SAN Volume Controller.

Flash Copy and SnapView

The EMC CLARiiON's form of FlashCopy called SnapView is not supported by the SAN Volume Controller. In a split controller configuration, SnapView is not supported even for the LUs that are controlled by the host.

Metro Mirror and MirrorView

The EMC CLARiiON's form of Metro Mirror called MirrorView is not supported by the SAN Volume Controller. In a split controller configuration, MirrorView is not supported even for the LUs that are controlled by the host.

SAN Copy

The EMC CLARiiON provides a form of FlashCopy called SAN Copy that is not supported by the SAN Volume Controller. In a split controller configuration, SAN Copy is not supported even for the LUs that are controlled by the host.

MetaLUN

MetaLUN allows an LU to be expanded using LUs in other RAID groups. The SAN Volume Controller supports MetaLUN for migration purposes in Image Mode only.

Logical unit creation and deletion on the EMC CLARiiON

Binding an LU to a RAID group can take a significant amount of time on the EMC CLARiiON.

The LU must not be added to a storage group until binding is complete. As a safeguard, the SAN Volume Controller will not discover the LU if binding is in progress. A subsequent manual discovery is required.

Related tasks

"Discovering MDisks using the CLI" on page 168
You can discover MDisks using the command-line interface (CLI).

Configuring settings for the EMC CLARiiON

A number of settings and options are available through the EMC CLARiiON configuration interface.

Some of the options and settings are supported by the SAN Volume Controller. These options and settings cover the following:

- Subsystem
- Port
- Logical unit

Related reference

“Global settings for the EMC CLARiiON”
Global settings apply across an EMC CLARiiON subsystem.

“Controller settings for the EMC CLARiiON” on page 263
Controller settings are settings that apply across one EMC CLARiiON subsystem.

“Port settings for the EMC CLARiiON” on page 263
Port settings are configurable at the port level.

“LU settings for the EMC CLARiiON” on page 264
LU settings are configurable at the LU level.

Global settings for the EMC CLARiiON

Global settings apply across an EMC CLARiiON subsystem.

Table 20. EMC CLARiiON global settings supported by the SAN Volume Controller

Option	EMC CLARiiON Default Setting	SAN Volume Controller Required Setting
Access Controls (Access Logix installed)	Not installed	Either Installed or Not Installed
Subsystem Package Type	3	3
Queue Full Status	Disable	Disable
Recovered Errors	Disable	Disable
Target Negotiate	Displays the state of the target negotiate bit.	Displays the state of the target negotiate bit.
Mode Page 8 Info	Disable	Disable
Base UUID	0	0
Write Cache Enabled	Enabled	Enabled
Mirrored Write Cache	Enabled	Enabled
Write Cache Size	600 MB	Default recommended
Enable Watermarks	Enabled	Enabled
Cache High Watermark	96%	Default
Cache Low Watermark	80%	Default
Cache Page Size	4 Kb	4 Kb
RAID3 Write Buffer Enable	Enable	Default recommended
RAID3 Write Buffer	0 MB	Default recommended

Related concepts

“Storage subsystems” on page 22
A storage subsystem 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.

Related tasks

“Configuring a balanced storage subsystem” on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Controller settings for the EMC CLARiiON

Controller settings are settings that apply across one EMC CLARiiON subsystem.

Table 21 describes the options that can be set by the EMC CLARiiON.

Table 21. EMC CLARiiON controller settings supported by the SAN Volume Controller

Option	EMC CLARiiON Default Setting	SAN Volume Controller Required Setting
Read Cache Enabled	Enable	Enable
Read Cache Size	200 MB	Enable
Statistics Logging	Disable	Either Enable or Disable

Note: The SAN Volume Controller cannot obtain or change the configuration options listed above. It is therefore your responsibility to configure the options as recommended.

Related concepts

“Storage subsystems” on page 22
A storage subsystem 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.

Related tasks

“Configuring a balanced storage subsystem” on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Port settings for the EMC CLARiiON

Port settings are configurable at the port level.

This topic lists port settings, the EMC CLARiiON defaults, and the required settings for the SAN Volume Controller.

Table 22. EMC CLARiiON port settings supported by the SAN Volume Controller

Option	EMC CLARiiON Default Setting	SAN Volume Controller Required Setting
Port speed	2 GB	Either 1 or 2 GB

Note: The SAN Volume Controller cannot obtain or change the configuration options listed above. It is therefore your responsibility to configure the options as recommended.

Related concepts

“Storage subsystems” on page 22
A storage subsystem 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.

“Virtual disk-to-host mapping” on page 32
Virtual disk-to-host mapping is the process of controlling which hosts have access to specific virtual disks (VDisks) within the SAN Volume Controller.

Related tasks

“Configuring a balanced storage subsystem” on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

LU settings for the EMC CLARiiON

LU settings are configurable at the LU level.

This topic lists EMC CLARiiON defaults and required settings are for the SAN Volume Controller.

Table 23 describes the options that must be set for each logical unit that is accessed by the SAN Volume Controller. LUs that are accessed by hosts can be configured differently.

Table 23. EMC CLARiiON LU settings supported by the SAN Volume Controller

Option	EMC CLARiiON Default Setting	SAN Volume Controller Required Setting
LU ID	Auto	N/A
RAID Type	5	Any RAID Group
RAID Group	Any available RAID Group	Any available RAID Group
Offset	0	Any setting
LU Size	ALL LBAs in RAID Group	Any setting
Placement	Best Fit	Either Best Fit or First Fit
UID	N/A	N/A
Default Owner	Auto	N/A
Auto Assignment	Disabled	Disabled
Verify Priority	ASAP	N/A
Rebuild Priority	ASAP	N/A
Strip Element Size	128	N/A
Read Cache Enabled	Enabled	Enabled
Write Cache Enabled	Enabled	Enabled
Idle Threshold	0–254	0–254
Max Prefetch Blocks	0–2048	0–2048
Maximum Prefetch IO	0–100	0–100
Minimum Prefetch Size	0–65534	0–65534
Prefetch Type	0, 1, or 2	0, 1, or 2
Prefetch Multiplier	0 to 2048 or 0 to 324	0 to 2048 or 0 to 324
Retain prefetch	Enabled or Disabled	Enabled or Disabled
Prefetch Segment Size	0 to 2048 or 0 to 32	0 to 2048 or 0 to 32
Idle Delay Time	0 to 254	0 to 254
Verify Priority	ASAP, High, Medium, or Low	Low

Table 23. EMC CLARiiON LU settings supported by the SAN Volume Controller (continued)

Option	EMC CLARiiON Default Setting	SAN Volume Controller Required Setting
Write Aside	16 to 65534	16 to 65534

Note: The SAN Volume Controller cannot obtain or change the configuration options listed above. It is therefore your responsibility to configure the options as recommended.

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Configuring the EMC Symmetrix subsystem

This part provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

Related reference

"Supported models of the EMC Symmetrix controller" on page 266
The following table lists the supported models of the EMC Symmetrix controller. See the following Web site for the latest supported models. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

"Supported firmware levels for the EMC Symmetrix controller" on page 266
See the following Web site for specific firmware levels and the latest supported hardware. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

"Concurrent maintenance on the EMC Symmetrix" on page 266
Concurrent maintenance is the capability to perform I/O operations to an EMC Symmetrix while simultaneously performing maintenance operations on it.

"User interface on Symmetrix" on page 267
Ensure that you are familiar the user interface application that supports the Symmetrix subsystem.

"Sharing the EMC Symmetrix controller between a host and the SAN Volume Controller" on page 268
The EMC Symmetrix can be shared between a host and a SAN Volume Controller.

"Switch zoning limitations for the EMC Symmetrix" on page 268
Consider the following limitations when planning switch zoning and connecting to the SAN.

"Quorum disks on EMC Symmetrix" on page 269
Managed disks presented by the EMC Symmetrix will be chosen by the SAN Volume Controller as quorum disks.

"Advanced functions for EMC Symmetrix" on page 269
Advanced copy functions for Symmetrix (for example, SRDF and TimeFinder)

are *not* supported for disks that are managed by the SAN Volume Controller because the copy does not extend to the SAN Volume Controller cache.

“Logical unit creation and deletion on EMC Symmetrix” on page 269
An LU exported by Symmetrix, meaning it is visible to a host, is either a *Symmetrix device* or a *Meta device*.

“Configuring settings for the EMC Symmetrix” on page 270
A number of settings and options are available through the EMC Symmetrix configuration interface.

Supported models of the EMC Symmetrix controller

The following table lists the supported models of the EMC Symmetrix controller. See the following Web site for the latest supported models. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Table 24. Supported models of the EMC Symmetrix

Model
DMX 800
DMX 1000
DMX 2000
DMX 3000
Symmetrix 8000 (Symm 5)-8130
Symmetrix 8000 (Symm 5)-8230
Symmetrix 8000 (Symm 5)-8430
Symmetrix 8000 (Symm 5)-8530
Symmetrix 8000 (Symm 5)-8730
Symmetrix 8000 (Symm 5)-8830

Related reference

“Configuring the EMC Symmetrix subsystem” on page 265
This part provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

Supported firmware levels for the EMC Symmetrix controller

See the following Web site for specific firmware levels and the latest supported hardware. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Related reference

“Configuring the EMC Symmetrix subsystem” on page 265
This part provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

Concurrent maintenance on the EMC Symmetrix

Concurrent maintenance is the capability to perform I/O operations to an EMC Symmetrix while simultaneously performing maintenance operations on it.

This topic lists the components of the EMC Symmetrix that can be maintained in this way. Symmetrix supports non-distruptive microcode upgrade procedures.

The EMC Symmetrix is an Enterprise class device that supports non-disruptive replacement of the following components:

- Channel Director
- Disk Director
- Cache card
- Disk drive
- Cooling fan
- Comms card
- EPO card
- Operator panel
- PSU
- Service Processor
- Batteries
- Ethernet hub

Service actions and upgrade procedures may only be performed by an EMC Customer Engineer. Consequently, concurrent maintenance of Symmetrix is not supported by the SAN Volume Controller. Concurrent code upgrade under the SAN Volume Controller is also not supported.

Related reference

"Configuring the EMC Symmetrix subsystem" on page 265
This part provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

User interface on Symmetrix

Ensure that you are familiar the user interface application that supports the Symmetrix subsystem.

EMC Control Center

A basic Symmetrix configuration is performed by an EMC Customer Engineer (CE) via the Symmetrix service processor. After the initial configuration, you can configure and control the exported storage. The CE defines the storage device types and sets the configurable options. The user is then able to configure and control the exported storage as described below.

You use the EMC Control Center to manage and monitor the Symmetrix storage system.

You use Volume Logix for volume configuration management. You can use it to control access rights to the storage when multiple hosts share target ports.

Command-line interface

The Symmetrix Command Line Interface (SYMCLI) allows the server to monitor and control Symmetrix.

Related reference

"Configuring the EMC Symmetrix subsystem" on page 265
This part provides information about configuring the EMC Symmetrix for

attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

Sharing the EMC Symmetrix controller between a host and the SAN Volume Controller

The EMC Symmetrix can be shared between a host and a SAN Volume Controller.

This topic briefly discusses the restrictions.

- Target ports must not be shared between the SAN Volume Controller and other hosts.
- A single host must not be connected to a SAN Volume Controller and a Symmetrix because the multi-pathing drivers (for example, Subsystem Device Driver (SDD) and PowerPath), do not coexist.
- Other hosts may be connected directly to Symmetrix at the same time as SAN Volume Controller, under the following conditions:
  – The fabric must be zoned such that other hosts cannot access the target ports used by the SAN Volume Controller.
  – Symmetrix must be configured such that other hosts cannot access the LUs that are managed by the SAN Volume Controller.

Related concepts

"Storage subsystems" on page 54
Follow these rules when you are planning the configuration of storage subsystems in the SAN fabric.

Related reference

"Configuring the EMC Symmetrix subsystem" on page 265
This part provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

Switch zoning limitations for the EMC Symmetrix

Consider the following limitations when planning switch zoning and connecting to the SAN.

Switch zoning

The SAN Volume Controller switch zone must include at least one target port on two or more fibre-channel adapters (FAs) in order to have no single point of failure.

Connecting to the SAN

The Symmetrix connects to the SAN via a fibre-channel director. Directors are installed in pairs and each consists of two boards, one of which is a fibre-channel adapter (FA). The FA provides 2 - 12 target ports. Symmetrix assigns a WWNN per target port and SAN Volume Controller can resolve up to four WWNN's per subsystem. In order to connect more than four target ports to a SAN Volume Controller, the following procedure must be performed:

1. Divide the set of target ports into groups of 2 - 4.
2. Define a discrete set of logical units for each group.
3. Map the logical units to each target port in their group.

The SAN Volume Controller views each group of target ports as a separate subsystem. Ensure that no LUs are a member of more than one group.

Related reference

"Configuring the EMC Symmetrix subsystem" on page 265
This part provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

"Switch zoning for the SAN Volume Controller" on page 65
Consider these constraints when you zone a switch.

Quorum disks on EMC Symmetrix

Managed disks presented by the EMC Symmetrix will be chosen by the SAN Volume Controller as quorum disks.

This topic discusses the implications.

The SAN Volume Controller uses a logical unit (LU) presented by an EMC Symmetrix as a quorum disk. In addition, it will provide a quorum disk, even if the connection is by a single port.

Related reference

"Configuring the EMC Symmetrix subsystem" on page 265
This part provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

"Configuring and servicing storage subsystems" on page 233
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

Advanced functions for EMC Symmetrix

Advanced copy functions for Symmetrix (for example, SRDF and TimeFinder) are *not* supported for disks that are managed by the SAN Volume Controller because the copy does not extend to the SAN Volume Controller cache.

Related reference

"Configuring the EMC Symmetrix subsystem" on page 265
This part provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

Logical unit creation and deletion on EMC Symmetrix

An LU exported by Symmetrix, meaning it is visible to a host, is either a *Symmetrix device* or a *Meta device*.

Symmetrix device is an EMC term for an LU that is hosted by a Symmetrix. These are all emulated devices and have exactly the same characteristics:

- N cylinders
- 15 tracks per cylinder
- 64 logical blocks per track
- 512 bytes per logical block

Symmetrix devices can be created using the **create dev** command from the Symmetrix Command Line Interface (SYMCLI). The configuration of an LU can be changed using the **convert dev** command from the Symmetrix Command Line Interface (SYMCLI). Each physical storage device in a Symmetrix is partitioned into

1 - 128 hyper-volumes or hypers. Each hyper may be up to 16GB. A Symmetrix device maps to one or more hypers, depending on how it is configured. For example:

- Hypers may be mirrored (2-way, 3-way, 4-way)
- Hypers may be formed into RAID-S groups

Meta device is an EMC term for a concatenated chain of Symmetrix devices. This enables Symmetrix to provide logical units that are larger than a hyper. Up to 255 hypers may be concatenated to form a single meta device. Meta devices can be created using the **form meta** and **add dev** commands from the Symmetrix Command Line Interface (SYMCLI).

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Related reference

"Configuring the EMC Symmetrix subsystem" on page 265
This part provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

Configuring settings for the EMC Symmetrix

A number of settings and options are available through the EMC Symmetrix configuration interface.

These options and settings can have a scope of a:

- Subsystem
- Port
- Logical unit

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Related reference

"Configuring the EMC Symmetrix subsystem" on page 265
This part provides information about configuring the EMC Symmetrix for attachment to a SAN Volume Controller. Except where explicitly noted, all information in this section also applies to EMC DMX subsystems.

"Global settings for the EMC Symmetrix"
Global settings apply across an EMC Symmetrix subsystem.

"Port settings for the EMC Symmetrix"
Target port characteristics can be set using the **set port** command.

"LU settings for the EMC Symmetrix" on page 272
LU settings are configurable at the LU level.

"Mapping and virtualization settings for the EMC Symmetrix" on page 273
LUs can be mapped to a particular director or target port using the **map dev** command from the Symmetrix Command Line Interface (SYMCLI).

Related information

"Servicing storage subsystems" on page 253
When servicing your storage subsystems, it is imperative that you follow the service instructions contained in the vendor documentation.

Global settings for the EMC Symmetrix

Global settings apply across an EMC Symmetrix subsystem.

Subsystem characteristics can be set using the **set Symmetrix** command. The characteristics can be viewed using the **symconfigure** command from the Symmetrix Command Line Interface (SYMCLI).

Table 25. EMC Symmetrix global settings supported by the SAN Volume Controller

Option	EMC Symmetrix Default Setting	SAN Volume Controller Required Setting
max_hypers_per_disk		any
dynamic_rdf	disable	disable
fba_multi_access_cache	disable	n/a
Raid_s_support	disable	enable, disable

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Port settings for the EMC Symmetrix

Target port characteristics can be set using the **set port** command.

The characteristics can be viewed using the **symcfg** command from the Symmetrix Command Line Interface (SYMCLI).

Table 26. EMC Symmetrix port settings supported by the SAN Volume Controller

Option	EMC Symmetrix Default Setting	SAN Volume Controller Required Setting
Disk_Array	enabled	disabled
Volume_Set_Addressing	enabled	disabled
Hard_Addressing	enabled	enabled

Table 26. EMC Symmetrix port settings supported by the SAN Volume Controller (continued)

Option	EMC Symmetrix Default Setting	SAN Volume Controller Required Setting
Non_Participating	disabled	disabled
Global_3rdParty_Logout	enabled	enabled
Tagged_Commands	enabled	enabled
Common_Serial_Number		enabled
Disable_Q_Reset_on_UA	disabled	disabled
Return_busy_for_abort	disabled	disabled
SCSI-3	disabled	disabled
Environ_Set	disabled	disabled
Unique_WWN	enabled	enabled
Point_to_Point	disabled	enabled
VCM_State	disabled	either
OpenVMS	disabled	disabled

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

"Virtual disk-to-host mapping" on page 32
Virtual disk-to-host mapping is the process of controlling which hosts have access to specific virtual disks (VDisks) within the SAN Volume Controller.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

LU settings for the EMC Symmetrix

LU settings are configurable at the LU level.

This topic lists the LU settings, the EMC Symmetrix defaults settings and the settings that are required for the SAN Volume Controller. LU characteristics can be set using the **set device** command from the Symmetrix Command Line Interface (SYMCLI).

Table 27. EMC Symmetrix LU settings supported by the SAN Volume Controller

Option	EMC Symmetrix Default Setting	SAN Volume Controller Required Setting
emulation		FBA
attribute		RAD

Related concepts

“Storage subsystems” on page 22
A storage subsystem 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.

Related tasks

“Configuring a balanced storage subsystem” on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Mapping and virtualization settings for the EMC Symmetrix

LUs can be mapped to a particular director or target port using the **map dev** command from the Symmetrix Command Line Interface (SYMCLI).

They can be unmapped using the **unmap dev** command. Mapping a logical unit to a host is a function of the EMC Control Center.

Related reference

“Switch zoning for the SAN Volume Controller” on page 65
Consider these constraints when you zone a switch.

Configuring the Enterprise Storage Server

This part provides information about configuring the Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Related tasks

“Configuring the Enterprise Storage Server (ESS)” on page 274
To configure the Enterprise Storage Server (ESS), perform these steps.

Related reference

“Supported models of the ESS” on page 275
The following table lists the supported models of the ESS. See the following Web site for the latest supported models. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

“Supported firmware levels for the ESS” on page 275
See the following Web site for specific firmware levels and the latest supported hardware. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

“Concurrent maintenance on the ESS” on page 275
Concurrent maintenance is the capability to perform I/O operations to an ESS while simultaneously performing maintenance operations on it.

“User interface on ESS” on page 276
Ensure that you are familiar with the user interface application the supports the ESS subsystem.

“Sharing the ESS between a host and the SAN Volume Controller” on page 276
The ESS can be shared between a host and a SAN Volume Controller.

“Switch zoning limitations for the ESS” on page 276
Consider the following limitations when planning switch zoning and connection to the SAN.

“Quorum disks on the ESS” on page 277
Managed disks presented by the ESS will be chosen by the SAN Volume Controller as quorum disks.

“Advanced functions for the ESS” on page 277
ESS provides advanced functions that are not compatible with the SAN Volume Controller.

"Logical unit creation and deletion on the ESS" on page 278
Before creating or deleting a logical unit on the ESS, consider these constraints.

Configuring the Enterprise Storage Server (ESS)

To configure the Enterprise Storage Server (ESS), perform these steps.

1. Use a Web browser to access the ESS Specialist by entering the IP address of the ESS.
2. Login using the user name and password.
3. Click **ESS Specialist**.
4. Click **Storage Allocation**.
5. Click **Open System Storage**.
6. Click **Modify Host Systems**.
7. Create a host entry for every initiator port on every SAN Volume Controller node in your cluster. Complete the following fields:

   Nickname
   Type a unique name for each port (for example, `knode` or `lnode`).

   Host Type
   Select **IBM SAN Volume Controller** or **RS/6000** if that is not available.

   Host Attachment
   Select **Fibre Channel attached**.

   Hostname/IP address
   Leave this field blank.

   WWPN
   Either select the WWPN from the list, or type it manually. A configuration command will fail if you use WWPN 0 in the command string.
8. After you are finished adding all of the ports, click **Perform Configuration Update**.
9. Click **Add Volumes** to add the volumes on which you want the SAN Volume Controller to run.
10. From the Add Volumes window, perform the following actions:
    a. Select any of the SAN Volume Controller host ports that you created earlier.
    b. Select the necessary ESS adapter to create the volumes.
    c. Click **Next**.
    d. Create volumes using your desired size, placement, and RAID level.
    e. After you are done creating all the volumes, click **Perform Configuration Update**.
11. Map the volumes to all of your SAN Volume Controller ports by performing the following steps:
    a. Click **Modify Volume Assignments**.
    b. Select all of the volumes that you created earlier.
    c. Click **Assigning selected volumes to target hosts**.
    d. Select all of the remaining SAN Volume Controller host ports that you created earlier.
    e. Select the **Use same ID/LUN in source and target** check box.
    f. Click **Perform Configuration Update**.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

"Expanding a logical unit" on page 243
A logical unit can be expanded using vendor-specific disk-configuration software.

Related reference

"Configuring and servicing storage subsystems" on page 233
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

Related information

"Configuring the Enterprise Storage Server" on page 273
This part provides information about configuring the Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Supported models of the ESS

The following table lists the supported models of the ESS. See the following Web site for the latest supported models. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Table 28. Supported models of the Enterprise Storage Server

Model
2105-F20
2105-750 (See Note.)
2105-800
2107-900 (See Note.)
1750-500 (See Note.)
Note: Support for these models is dependent on the product availability date.

Related information

"Configuring the Enterprise Storage Server" on page 273
This part provides information about configuring the Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Supported firmware levels for the ESS

See the following Web site for specific firmware levels and the latest supported hardware. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Related information

"Configuring the Enterprise Storage Server" on page 273
This part provides information about configuring the Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Concurrent maintenance on the ESS

Concurrent maintenance is the capability to perform I/O operations to an ESS while simultaneously performing maintenance operations on it.

IBM supports all ESS concurrent maintenance procedures.

Related information

“Configuring the Enterprise Storage Server” on page 273
This part provides information about configuring the Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

User interface on ESS

Ensure that you are familiar with the user interface application the supports the ESS subsystem.

Web Server

A Web server is running on each of the controllers on the subsystem. During normal operation, the user interface application only allows basic monitoring of the subsystem and displays an error log. If a controller is put into diagnostic mode by pressing the reset button on the controller, the user interface application allows for firmware upgrades and subsystem configuration resets.

Related information

“Configuring the Enterprise Storage Server” on page 273
This part provides information about configuring the Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Target port groups for the ESS

The ESS model, 1750-500 uses the SCSI Target Port Groups feature to indicate a preferred path for each LU.

Sharing the ESS between a host and the SAN Volume Controller

The ESS can be shared between a host and a SAN Volume Controller.

IBM supports sharing an ESS between a SAN Volume Controller and other hosts. However, if an ESS port is in the same zone as a SAN Volume Controller port, that same ESS port should not be in the same zone as another host.

A single host can have both ESS direct-attached and SAN Volume Controller virtualized disks configured to it. If a LUN is managed by the SAN Volume Controller, it should never be mapped to another host.

See the following Web site for the latest supported configurations:
http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Related concepts

“Storage subsystems” on page 54
Follow these rules when you are planning the configuration of storage subsystems in the SAN fabric.

Related information

“Configuring the Enterprise Storage Server” on page 273
This part provides information about configuring the Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Switch zoning limitations for the ESS

Consider the following limitations when planning switch zoning and connection to the SAN.

The minimum number of cables recommended for redundancy is 2 cables from 2 separate adapter bays. Up to 16 cables can be used to connect to an ESS. Only 1 or 2 Gb fibre-channel attached is supported.

Note: ESCON, FICON, and Ultra SCSI attachment is not supported with the SAN Volume Controller.

Related reference

"Switch zoning for the SAN Volume Controller" on page 65
Consider these constraints when you zone a switch.

Related information

"Configuring the Enterprise Storage Server" on page 273
This part provides information about configuring the Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Quorum disks on the ESS

Managed disks presented by the ESS will be chosen by the SAN Volume Controller as quorum disks.

Related reference

"Configuring and servicing storage subsystems" on page 233
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

Related information

"Configuring the Enterprise Storage Server" on page 273
This part provides information about configuring the Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

"Creating a quorum disk" on page 252
A quorum disk is used to resolve tie-break situations when the ″voting set″ of nodes disagree on the current cluster state.

Advanced functions for the ESS

ESS provides advanced functions that are not compatible with the SAN Volume Controller.

Only open system storage is supported by the SAN Volume Controller.

FlashCopy and Concurrent copy

FlashCopy and Concurrent copy are not supported on any LUN that is managed by the SAN Volume Controller.

Metro Mirror and Global Copy

Metro Mirror and Global Copy (formerly known as extended distance Remote Copy) are not supported on any LUN that is managed by the SAN Volume Controller.

Related information

"Configuring the Enterprise Storage Server" on page 273
This part provides information about configuring the Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Logical unit creation and deletion on the ESS

Before creating or deleting a logical unit on the ESS, consider these constraints.

Certain ESS types are supported for use with the SAN Volume Controller.

Before you delete or un-map a logical unit from the SAN Volume Controller, remove the logical unit from the MDisk group. The following is supported:

- The supported logical unit size is 1GB to 2TB.
- RAID 5 and RAID 10 logical units are supported.
- Logical units can be added dynamically. When adding a new logical unit, you *must* select the **Use same ID/LUN in source and target** checkbox.

  Note: Failure to select the **Use same ID/LUN in source and target** checkbox can cause loss in redundancy or a loss of data integrity.

  The detect MDisks action in the SAN Volume Controller Console or the **lsinfo detectmdisks** command must be run in order for the SAN Volume Controller to detect the new disks.

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Related information

"Configuring the Enterprise Storage Server" on page 273
This part provides information about configuring the Enterprise Storage Server (ESS) so it can attach to a SAN Volume Controller.

Configuring the FAStT subsystem

This part provides information about configuring the FAStT subsystem so it can attach to a SAN Volume Controller.

Related tasks

"Configuring FAStT disk controllers for the storage server" on page 279
FAStT disk controllers provide functionality that is compatible with the SAN Volume Controller.

"Support actions for the FAStT controller" on page 281
The FAStT disk controllers provide functionality that can be used with the SAN Volume Controller.

Related reference

"Supported models of the IBM FAStT controller" on page 282
The following table lists the supported models of the IBM FAStT controller that are supported by the SAN Volume Controller. See the following Web site for the latest supported models. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

"Supported firmware levels for the FAStT" on page 282
See the following Web site for specific firmware levels and the latest supported hardware. It includes the maximum number of LUNs per partition, depending on

the firmware level. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

"Concurrent maintenance on the IBM FAStT" on page 282
Concurrent maintenance is the capability to perform I/O operations to an IBM FAStT controller while simultaneously performing maintenance operations on it. Refer to your FAStT documentation for information about concurrent maintenance.

"User interface on FAStT" on page 283
Ensure that you are familiar with the user interface that supports the FAStT subsystem.

"Sharing the IBM FAStT controller between a host and the SAN Volume Controller" on page 283
The IBM FAStT controller can be shared between a host and a SAN Volume Controller.

"Quorum disks on the IBM FAStT" on page 283
Managed disks presented by the IBM FAStT controller will be chosen by the SAN Volume Controller as quorum disks.

"Advanced functions for the IBM FAStT" on page 284
FlashCopy and Metro Mirror facilities are provided as advanced functions on the IBM FAStT controller. These controller-delivered Copy Services are not supported by the SAN Volume Controller.

"Logical unit creation and deletion on the IBM FAStT" on page 285
To create or delete logical units on the IBM FAStT, complete the following tasks.

"Configuration interface for the IBM FAStT" on page 285
The IBM FAStT provides a configuration application.

"Controller settings for the IBM FAStT" on page 286
Controller settings are settings that apply across one FAStT controller.

Configuring FAStT disk controllers for the storage server

FAStT disk controllers provide functionality that is compatible with the SAN Volume Controller.

Attention: The SAN Volume Controller does not concurrently support I/O operations with the download of ESM (Environmental Services Monitor) firmware. You must quiesce all I/O operations from the hosts that are using storage provided by the FAStT controllers you wish to update before installing new ESM firmware.

The FAStT storage server has many options and actions. The following list provides the supported actions and its impact on the SAN Volume Controller and its configuration.

1. host type:
   a. You must set either the default host type of your FAStT or the host type of the chosen partition to:

      `IBM TS SAN VCE`

      1) Click **Storage Subsystem –> Change –> Default Host Type**, or
      2) For each host port you can specify the host type of that port or modify existing ports.
2. WWNN:
   a. Set the subsystem so that both controllers have the same WWNN. Scripts are available from the FAStT support Web site to change the set up of the FAStT if required.

www.storage.ibm.com

3. auto volume transfer (AVT):
   a. Make sure the auto volume transfer is enabled. The host type selection should have enabled this function already.
   b. View the storage subsystem profile data to confirm that you have the AVT function enabled. This storage profile is presented as a text view in a separate window.
   c. Scripts are available from the FAStT Web site to enable AVT if required.

      www.storage.ibm.com
4. limitations:
   a. Only one FAStT storage partition can be created that contains any of the ports of any of the nodes in a single SAN Volume Controller cluster.
   b. You must not map more than one partition to any of the ports on any of the nodes in the same SAN Volume Controller cluster. Otherwise, unexpected behavior might result. For example, there will not be any warning messages, however, there will be errors logged in the SAN Volume Controller error log and access to storage may be lost.
5. access LUN:
   a. The access LUN, also known as the Universal Transport Mechanism (UTM) LUN, might not be in a partition that contains the SAN Volume Controller ports. It is not required by the SAN Volume Controller. The UTM LUN is a special LUN that allows the SAN Volume Controller to be configured through suitable software over the Fibre channel connection. However, the SAN Volume Controller does not require the UTM LUN, therefore does not generate errors either way.
   b. The FAStT *must not* have the Access (UTM) LUN presented as Logical Unit Number 0 (zero).
6. logical unit:
   a. The SAN Volume Controller attempts to follow the FAStT specified preferred ownership. You can specify which controller (A or B) is used to do I/O operations to a given Logical Unit. If the SAN Volume Controller can see the ports of the preferred controller and no error conditions exist, then it will access that Logical Unit through one of the ports on that controller.
   b. Under error conditions, the ownership is ignored. Meaning, the SAN Volume Controller has found a given path through the fabric to be errant, or there is no connection to a given port.
7. Copy services (FlashCopy and Metro Mirror):
   a. FAStT copy services must *not* be used when the SAN Volume Controller is attached to the FAStT. Partitioning might allow copy services to be used on other host platforms.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

"Configuring the Enterprise Storage Server (ESS)" on page 274
To configure the Enterprise Storage Server (ESS), perform these steps.

"Expanding a logical unit" on page 243
A logical unit can be expanded using vendor-specific disk-configuration software.

Related reference

"Configuring and servicing storage subsystems" on page 233
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

Related information

"Configuring the FAStT subsystem" on page 278
This part provides information about configuring the FAStT subsystem so it can attach to a SAN Volume Controller.

Support actions for the FAStT controller

The FAStT disk controllers provide functionality that can be used with the SAN Volume Controller.

The FAStT storage manager has many options and actions. The following shows the supported actions and their impact on the SAN Volume Controller and its configuration.

1. controller run diagnostics:
   a. The diagnostics should be automatically recovered by the SAN Volume Controller software.
   b. Check your MDisks to make sure that they have not been set to degraded mode after this action.
2. controller disable data transfer:
   a. This option is not supported when a SAN Volume Controller is attached to the FAStT. Loss of availability and redundancy may occur if data transfer is disabled.
3. setting an array Offline:
   a. Do not set an array Offline. If you use this setting, you might lose access to the MDisk group.
4. array increase capacity:
   a. Increasing capacity is supported but the new capacity is not usable until the MDisk is removed from an MDisk group and then added again. You might have to migrate data to increase the capacity.
5. redistribute logical drives or change ownership of the preferred path:
   a. These actions are supported but might not take effect until a cluster rediscovery is initiated on the SAN Volume Controller cluster. This can be achieved using the **svctask detectmdisk** command.
6. controller reset
   a. Controller reset should only be performed if directed to do so by service personnel, the alternate controller is functional and available to the SAN. The SAN Volume Controller reset should be automatically recovered by the SAN Volume Controller software.
   b. Check your MDisks to make sure that they have not been set to degraded state during this operation. You can issue the **svctask includemdisk** to repair degraded MDisks.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

"Configuring the Enterprise Storage Server (ESS)" on page 274
To configure the Enterprise Storage Server (ESS), perform these steps.

Related reference

"Configuring and servicing storage subsystems" on page 233
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

Related information

"Configuring the FAStT subsystem" on page 278
This part provides information about configuring the FAStT subsystem so it can attach to a SAN Volume Controller.

Supported models of the IBM FAStT controller

The following table lists the supported models of the IBM FAStT controller that are supported by the SAN Volume Controller. See the following Web site for the latest supported models. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Table 29. Supported models of the IBM FAStT controller

Model
1724 FAStT model 100
3542 FAStT model 200
3552 FAStT model 500
1722 FAStT model 600
1742/1RU FAStT model 700
1742/90U FAStT model 900

Note: Some levels of FAStT microcode support a maximum of 32 LUNs per host partition, newer versions allow up to 256 LUNs per host partition.

Related information

"Configuring the FAStT subsystem" on page 278
This part provides information about configuring the FAStT subsystem so it can attach to a SAN Volume Controller.

Supported firmware levels for the FAStT

See the following Web site for specific firmware levels and the latest supported hardware. It includes the maximum number of LUNs per partition, depending on the firmware level. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Related information

"Configuring the FAStT subsystem" on page 278
This part provides information about configuring the FAStT subsystem so it can attach to a SAN Volume Controller.

Concurrent maintenance on the IBM FAStT

Concurrent maintenance is the capability to perform I/O operations to an IBM FAStT controller while simultaneously performing maintenance operations on it. Refer to your FAStT documentation for information about concurrent maintenance.

Related information

"Configuring the FAStT subsystem" on page 278
This part provides information about configuring the FAStT subsystem so it can attach to a SAN Volume Controller.

User interface on FAStT

Ensure that you are familiar with the user interface that supports the FAStT subsystem.

Web Server

A Web server is running on each of the controllers on the subsystem. During normal operation, the user interface only allows basic monitoring of the subsystem and displays an error log. If a controller is put into diagnostic mode by pressing the reset button on the controller, the user interface allows for firmware upgrades and subsystem configuration resets.

Related information

"Configuring the FAStT subsystem" on page 278
This part provides information about configuring the FAStT subsystem so it can attach to a SAN Volume Controller.

Sharing the IBM FAStT controller between a host and the SAN Volume Controller

The IBM FAStT controller can be shared between a host and a SAN Volume Controller.

Attention: The use of the FAStT term ″partitioning″ does not hold the same meaning as used by IBM.

The FAStT function known as partitioning must be used to separate groups of logical units that are directly attached to hosts or groups of hosts from the SAN Volume Controller accessed logical units.

Note: The SAN Volume Controller partition must either contain all the ports of the SAN Volume Controller cluster that are connected to the SAN, or are zoned to have access to the FAStT ports. At least one port from each FAStT controller must be visible by the SAN Volume Controller cluster.

Related concepts

"Storage subsystems" on page 54
Follow these rules when you are planning the configuration of storage subsystems in the SAN fabric.

Related information

"Configuring the FAStT subsystem" on page 278
This part provides information about configuring the FAStT subsystem so it can attach to a SAN Volume Controller.

Quorum disks on the IBM FAStT

Managed disks presented by the IBM FAStT controller will be chosen by the SAN Volume Controller as quorum disks.

Note: The FAStT 200 does not support quorum disks.

Related reference

"Configuring and servicing storage subsystems" on page 233
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

Related information

“Configuring the FAStT subsystem” on page 278
This part provides information about configuring the FAStT subsystem so it can attach to a SAN Volume Controller.

“Creating a quorum disk” on page 252
A quorum disk is used to resolve tie-break situations when the ″voting set″ of nodes disagree on the current cluster state.

Advanced functions for the IBM FAStT

FlashCopy and Metro Mirror facilities are provided as advanced functions on the IBM FAStT controller. These controller-delivered Copy Services are not supported by the SAN Volume Controller.

Related concepts

“Metro Mirror” on page 46
Metro Mirror (formerly known as 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 on the other virtual disk.

Related reference

“Data migration on an existing FAStT installation that contains partitions”
You can migrate data on an existing FAStT installation that contains partitions.

Related information

“Configuring the FAStT subsystem” on page 278
This part provides information about configuring the FAStT subsystem so it can attach to a SAN Volume Controller.

Data migration on an existing FAStT installation that contains partitions

You can migrate data on an existing FAStT installation that contains partitions.

You can enable the SAN Volume Controller to be introduced to an existing SAN environment, so that you have the option of utilizing image mode LUNs to import the existing data into the virtualization environment without requiring a backup and restore cycle. For example, each FAStT partition may contain up to 32 LUNs. Each partition can only access a unique set of HBA ports (as defined by WWPNs). That is, for a single host to access multiple partitions, unique host fibre ports (WWPNs) need to be assigned to each partition. All LUNs within a partition are surfaced to assigned host fibre ports (no sub-partition LUN mapping).

Host A is mapped to LUN 0, 1, 2 in Partition 0
Host B is mapped to LUN 0, 1, 2, 3, 4, 5 in Partition 1
Host C is mapped to LUN 0, 1, 2 in Partition 2

To allow Host A to access the LUNs in partition B, it is necessary to remove one of the HBAs (for example, A1) from the access list for partition 0 and add it to partition 1 (A1 cannot be on the access list for more than one partition).

To add a SAN Volume Controller into this configuration without save and restore cycles would require that a set of unique SAN Volume Controller HBA port WWPNs for each partition. This would allow the FAStT to surface the LUNs (with your data) to the SAN Volume Controller, which would then configure these LUNs as image-mode LUNs and surface them to the required hosts. Unfortunately, this violates a requirement that all SAN Volume Controller nodes be able to see all backend storage. To work around this problem, change the FAStT to allow more than 32 LUNs in 1 storage partition, so that you can move all the LUNs from all the other partitions into 1 partition and map to the SAN Volume Controller cluster.

For example, lets say the FAStT has 8 partitions with 30 LUNs in each. Perform the following:

1. Change the mappings for the first 4 partitions on the FAStT such that each partition is mapped to 1 port on each node, this maintains redundancy across the cluster.
2. Create a new partition on the FAStT that is mapped to all 4 ports on all the SAN Volume Controllers (actually not a partition at all)
3. Gradually migrate the data into the MDisks in the target partition, as storage is freed from the source partitions this can be reused as new storage in the target partition. As partitions are deleted new partitions that need to be migrated can be mapped and migrated in the same way. The host side data access and integrity would be maintained throughout this process.

Logical unit creation and deletion on the IBM FAStT

To create or delete logical units on the IBM FAStT, complete the following tasks.

Certain IBM FAStT controller types are supported for use with the SAN Volume Controller.

To create a logical disk, you must set either the default host type of your FAStT or the host type of the chosen partition to:

```
IBM TS SAN VCE
```

You can set the host type in 2 ways:

1. Click **Storage Subsystem –> Change –> Default Host Type**, or
2. For each host port you can specify the host type of that port or modify existing ports.

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Related information

"Configuring the FAStT subsystem" on page 278
This part provides information about configuring the FAStT subsystem so it can attach to a SAN Volume Controller.

Configuration interface for the IBM FAStT

The IBM FAStT provides a configuration application.

The access LUN, also known as the Universal Transport Mechanism (UTM) LUN, is the configuration interface for the IBM FAStT controller.

The access LUN might not be in a partition that contains the SAN Volume Controller ports. It is not required by the SAN Volume Controller. The UTM LUN is a special LUN that allows the SAN Volume Controller to be configured through suitable software over the Fibre channel connection. However, the SAN Volume Controller does not require the UTM LUN, therefore does not generate errors either way. The FAStT *must not* have the Access (UTM) LUN presented as Logical Unit Number 0 (zero).

It is possible to use in-band (over fibre-channel) and out-of-band (Ethernet) to allow the FAStT configuration software to communicate with more than one FAStT. If using in-band configuration, the ″Access″ logical unit will need to be configured in a partition that does not include any logical units being accessed by the SAN Volume Controller cluster.

Note: In-band is not supported via access to the LUN while in the SAN Volume Controller partition.

Related information

"Configuring the FAStT subsystem" on page 278
This part provides information about configuring the FAStT subsystem so it can attach to a SAN Volume Controller.

Controller settings for the IBM FAStT

Controller settings are settings that apply across one FAStT controller.

For restrictions on controller settings, see the following:

- You must set either the default host type of your FAStT or the host type of the chosen partition to:

  ```
  IBM TS SAN VCE
  ```

  You can set the host type in 2 ways:

  1. Click **Storage Subsystem –> Change –> Default Host Type**, or
  2. For each host port you can specify the host type of that port or modify existing ports.
- Set the subsystem so that both controllers have the same WWNN. Scripts are available from the FAStT support Web site to change the set up of the FAStT if required.

  www.storage.ibm.com
- Make sure the auto volume transfer is enabled. The host type selection should have enabled this function already. View the storage subsystem profile data to confirm that you have the AVT function enabled. This storage profile is presented as a text view in a separate window. Scripts are available from the FAStT Web site to enable AVT if required.

  www.storage.ibm.com
- Ensure that you have the following enabled on any logical units mapped to the :
  - read caching
  - write caching
  - write cache mirroring

  Caching without batteries must not be enabled.

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Related reference

"Configuring settings for the IBM FAStT"
This topic provides information about configuring settings for the IBM FAStT.

"Global settings for the IBM FAStT" on page 288
Global settings apply across an IBM FAStT controller.

"LU settings for the IBM FAStT" on page 288
LU settings are configurable at the LU level.

"Miscellaneous settings for the IBM FAStT" on page 289
There are options of a miscellaneous nature that must be set properly in order for the IBM FAStT controller to work with the SAN Volume Controller. Refer to your FAStT documentation for information about other settings.

"Mapping and virtualization settings for IBM FAStT" on page 289
Review the following information to understand LUN mapping (or masking) and virtualization from the point of view of the IBM FAStT controller and their use in a SAN Volume Controller environment.

Related information

"Configuring the FAStT subsystem" on page 278
This part provides information about configuring the FAStT subsystem so it can attach to a SAN Volume Controller.

Configuring settings for the IBM FAStT

This topic provides information about configuring settings for the IBM FAStT.

The IBM FAStT controller configuration interface provides configuration settings and options. This topic and its subtopics discusses those options and settings that are supported with the SAN Volume Controller.

These options and settings can have a scope of a:

- Subsystem
- Logical unit, see the following:
  - The SAN Volume Controller attempts to follow the FAStT specified preferred ownership. You can specify which controller (A or B) is used to do I/O operations to a given Logical Unit. If the SAN Volume Controller can see the ports of the preferred controller and no error conditions exist, then it will access that Logical Unit through one of the ports on that controller. Under error conditions, the ownership is ignored. Meaning, the SAN Volume Controller has found a given path through the fabric to be errant, or there is no connection to a given port.
  - Ensure that you have the following enabled on any logical units mapped to the SAN Volume Controller:
    - read caching
    - write caching
    - write cache mirroring

Caching without batteries must not be enabled.

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

"Configuring FAStT disk controllers for the storage server" on page 279
FAStT disk controllers provide functionality that is compatible with the SAN Volume Controller.

Related information

"Servicing storage subsystems" on page 253
When servicing your storage subsystems, it is imperative that you follow the service instructions contained in the vendor documentation.

Global settings for the IBM FAStT

Global settings apply across an IBM FAStT controller.

Table 30. IBM FAStT controller global settings supported by the SAN Volume Controller

Option	FAStT Default Setting
Start flushing	80%
Stop flushing	80%
Cache block size	4 Kb

These settings can be adjusted depending on the performance requirements. Modification of these settings is not recommended unless directed to do so by support personnel.

If partitioning is not being used, meaning all the FAStT logical units are visible to the SAN Volume Controller, then the default host type for the FAStT subsystem can be set. See step 1 on page 279. If partitioning is being used to group SAN Volume Controller ports and host ports apart then the host type for each partition, or group of SAN Volume Controller ports, must be defined. When defining host ports, the host type must be set to:`IBM TS SAN VCE`

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

LU settings for the IBM FAStT

LU settings are configurable at the LU level.

LUs that are accessed by hosts can be configured differently. This topic lists those settings, what the IBM FAStT controller defaults are, and what the required settings are for the SAN Volume Controller.

Read ahead cache multiplier is typically set to 0 or 1. Modification of these settings is not recommended unless directed to do so by support personnel.

Ensure that you have the following enabled on any logical units mapped to the SAN Volume Controller:

- read caching
- write caching
- write cache mirroring

Caching without batteries must not be enabled.

When creating a new logical unit set the host type for that logical unit to the host type `IBM TS SAN VCE`.

Note: `IBM TS SAN VCE` is set as the default if the default type was already displayed.

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Miscellaneous settings for the IBM FAStT

There are options of a miscellaneous nature that must be set properly in order for the IBM FAStT controller to work with the SAN Volume Controller. Refer to your FAStT documentation for information about other settings.

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Mapping and virtualization settings for IBM FAStT

Review the following information to understand LUN mapping (or masking) and virtualization from the point of view of the IBM FAStT controller and their use in a SAN Volume Controller environment.

Related reference

"Switch zoning for the SAN Volume Controller" on page 65
Consider these constraints when you zone a switch.

Configuring the HDS Lightning subsystem

This part provides information about configuring the HDS Lightning subsystem so it can attach to a SAN Volume Controller.

Related reference

"Supported models of the HDS Lightning"
The following table lists the models of the HDS Lightning that are supported by the SAN Volume Controller. See the following Web site for the latest supported models. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

"Supported firmware levels for HDS Lightning" on page 291
See the following Web site for specific HDS Lightning firmware levels and the latest supported hardware. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

"Concurrent maintenance on the HDS Lightning 99xxV" on page 291
Concurrent upgrade of the HDS Lightning 99xxV firmware is supported with the SAN Volume Controller.

"User interface on Lightning" on page 291
Ensure that you are familiar with the user interface application that supports supports the Lightning subsystem.

"Sharing the HDS Lightning 99xxV between host and the SAN Volume Controller" on page 291
The HDS Lightning 99xxV can be shared between a host and a SAN Volume Controller according to certain restrictions.

"Quorum disks on HDS Lightning 99xxV" on page 292
HDS Lightning 99xxV is not an approved host for quorum disks. Therefore, configurations with only Lightning are not possible.

"Advanced functions for HDS Lightning" on page 293
The advanced functions of the HDS Lightning 99xxV may or may not be compatible with the SAN Volume Controller. Advanced copy functions are not supported for LUs managed by SAN Volume Controller since they do not extend to the SAN Volume Controller cache.

"LU configuration for Lightning" on page 294
Logical unit configuration for Lightning supports both RAID 1 and RAID 5 arrays.

"Configuring settings for Lightning" on page 295
The Lightning configuration interface provides functionality for configuration.

Supported models of the HDS Lightning

The following table lists the models of the HDS Lightning that are supported by the SAN Volume Controller. See the following Web site for the latest supported models. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Table 31. Supported HDS Lightning models

Model
Lightning 9910
Lightning 9960
Lightning 9970
Lightning 9980

Related information

“Configuring the HDS Lightning subsystem” on page 290
This part provides information about configuring the HDS Lightning subsystem so it can attach to a SAN Volume Controller.

Supported firmware levels for HDS Lightning

See the following Web site for specific HDS Lightning firmware levels and the latest supported hardware. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Note: Concurrent upgrade of the controller firmware is **not** supported with the SAN Volume Controller.

Related information

“Configuring the HDS Lightning subsystem” on page 290
This part provides information about configuring the HDS Lightning subsystem so it can attach to a SAN Volume Controller.

Concurrent maintenance on the HDS Lightning 99xxV

Concurrent upgrade of the HDS Lightning 99xxV firmware is supported with the SAN Volume Controller.

Related information

“Configuring the HDS Lightning subsystem” on page 290
This part provides information about configuring the HDS Lightning subsystem so it can attach to a SAN Volume Controller.

User interface on Lightning

Ensure that you are familiar with the user interface application that supports supports the Lightning subsystem.

SVP

Lightning has a laptop in the controller frame. The laptop runs the primary configuration user interface, which is called SVP. Using SVP, it is possible to do almost every configuration task as well as monitor the controller.

HiCommand

HiCommand is a graphical user interface that allows basic creation of storage and system monitoring. It communicates with Lightning through Ethernet.

Related information

“Configuring the HDS Lightning subsystem” on page 290
This part provides information about configuring the HDS Lightning subsystem so it can attach to a SAN Volume Controller.

Sharing the HDS Lightning 99xxV between host and the SAN Volume Controller

The HDS Lightning 99xxV can be shared between a host and a SAN Volume Controller according to certain restrictions.

Sharing ports

The HDS Lightning 99xxV can be shared between a host and a SAN Volume Controller according to certain restrictions. The following restriction applies:

- The same host cannot be connected to both a SAN Volume Controller and a Lightning at the same time because HDLM and the Subsystem Device Driver do not coexist.
- A controller port cannot be shared between a host and a SAN Volume Controller. In other words, if a controller port is used by a SAN Volume Controller it must not be present in a switch zone which allows a host to access the port.
- Logical units (LUs) cannot be shared between a host and a SAN Volume Controller.

Supported Topologies

SAN Volume Controller supports connection to the Lightning according to the following restrictions:

- The SAN Volume controller resolves up to four WWNNs per subsystem and allows up to 512 LUs per WWNN. Lightning assigns a WWNN per port, therefore the SAN Volume Controller could be a limitation to both capacity (2048 LUs) and bandwidth (4 ports). You can use the following procedure for Lightning subsystems with 8 ports if more capacity or bandwidth is required:
  1. Divide the set of ports into groups of between 2 and 4.
  2. Assign a discreet set of logical units to each group.

  SAN Volume Controller then interprets each group as a separate subsystem.
- If a logical unit is mapped to the SAN Volume Controller port as LUN*x*, it must appear as LUN*x* to all the SAN Volume Controller ports in the cluster and must also appear as LUN*x* through all of the controller ports that it is mapped to.
- Command LUNs must not be mapped to the SAN Volume Controller.
- LUN Expansion (LUSE) and Virtual LVI/LUN operations cannot be run on a disk managed by the SAN Volume Controller. LUNs created using LUSE and Virtual LVI/LUN can be mapped to the SAN Volume Controller after they have been created.
- Only disks with open emulation can be mapped to the SAN Volume Controller. S/390 disks cannot be used with the SAN Volume Controller. Only fibre channel connections can be used to connect the SAN Volume Controller to the Lightning.

Related concepts

"Storage subsystems" on page 54
Follow these rules when you are planning the configuration of storage subsystems in the SAN fabric.

Related information

"Configuring the HDS Lightning subsystem" on page 290
This part provides information about configuring the HDS Lightning subsystem so it can attach to a SAN Volume Controller.

Quorum disks on HDS Lightning 99xxV

HDS Lightning 99xxV is not an approved host for quorum disks. Therefore, configurations with only Lightning are not possible.

Related information

"Configuring the HDS Lightning subsystem" on page 290
This part provides information about configuring the HDS Lightning subsystem so it can attach to a SAN Volume Controller.

"Creating a quorum disk" on page 252
A quorum disk is used to resolve tie-break situations when the ″voting set″ of nodes disagree on the current cluster state.

Advanced functions for HDS Lightning

The advanced functions of the HDS Lightning 99xxV may or may not be compatible with the SAN Volume Controller. Advanced copy functions are not supported for LUs managed by SAN Volume Controller since they do not extend to the SAN Volume Controller cache.

ShadowImage

ShadowImage is functionally similar to FlashCopy. ShadowImage is not supported when the disk controller system is being used with the SAN Volume Controller. Even when a Lightning 99xxV is shared between a host and a SAN Volume Controller, ShadowImage is not supported on the ports that are zoned directly with the host.

LU Expansion

The Lightning 99xxV supports Logical Unit expansion (LUSE). LUSE is a non-concurrent operation. LUSE is accomplished by concatenating between 2 and 26 existing logical units together. Before LUSE can be performed on a logical unit, it must be removed from an mdisk group and unmapped from the SAN Volume Controller.

Attention: This procedures will destroy all data that exists on the logical unit, except on a Windows system.

TrueCopy

TrueCopy is functionally similar to Metro Mirror. TrueCopy is not supported when the disk controller system is being used with the SAN Volume Controller. Even when a Lightning 99xxV is shared between a host and a SAN Volume Controller, TrueCopy is not supported on the ports that are zoned directly with the host.

Virtual LVI

The Lightning 99xxV supports Virtual LVI/LUNs. You can use this method to modify a LUN size that the Lightning uses by dividing it into several smaller virtual LUNs. This is a non-concurrent procedure that requires you to first create existing LUNs into free space, then you must define their own LUNs using that free space. Virtual LVI/LUNs must not be managed or mapped to a SAN Volume Controller.

LUNs that are set up using either LUSE or Virtual LVI/LUNs appear as normal LUNs after they have been created. Therefore, LUNs set up using LUSE or Virtual LVI/LUNs can be used by the SAN Volume Controller after they have been created.

Write protect

Logical units (LUs) cannot be explicitly set to be write-protected. However, some of the advanced features, such as Metro Mirror, can be used to write-protect a LU as part of the function. Metro Mirror must not be used for LUs in use by a SAN Volume Controller.

Related concepts

"FlashCopy" on page 34
FlashCopy is a copy service that is available with the SAN Volume Controller.

"Metro Mirror" on page 46
Metro Mirror (formerly known as 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 on the other virtual disk.

Related information

"Configuring the HDS Lightning subsystem" on page 290
This part provides information about configuring the HDS Lightning subsystem so it can attach to a SAN Volume Controller.

LU configuration for Lightning

Logical unit configuration for Lightning supports both RAID 1 and RAID 5 arrays.

The Lightning subsystem can have up to 8192 logical units defined, however, only 256 logical units can be mapped to a single port. Report LUNs is supported by LUN 0, so the SAN Volume Controller can detect all LUNs.

In the event that no LUN 0 is configured, the Lightning subsystem will present a pseudo-LUN at LUN 0. The inquiry data for this psuedo-LUN differs slightly from the inquiry data of normal LUNs, allowing the SAN Volume Controller to recognize it and exclude it from I/O. The puesdo-LUN will accept the report LUNs command.

The Lightning subsystem supports both open-mode attachment and S/390 attachment. The emulation mode is set when the LU is defined. All LUNs presented to an SAN Volume Controller must use open emulation. All LUNs with open emulation use a standard 512 byte block size.

The Lightning subsystem can only have certain sized logical units defined. These logical units can be expanded by merging 2 - 36 of these LUs together using the LU Size Expansion (LUSE) feature. They can also be made into several, smaller ″virtual″ LUNs by using the Virtual LVI/LUN feature.

Special LUs

When a LU is mapped to a host, there is the option of making it a *command LUN*. These LUNs support in-band configuration commands, but not I/O. Therefore, no LUNs mapped to the SAN Volume Controller can be command LUNs.

Related reference

"Logical unit creation and deletion on HDS Lightning"
Before you create or delete a logical unit on HDS Lightning, consider the following prerequisites.

Related information

"Configuring the HDS Lightning subsystem" on page 290
This part provides information about configuring the HDS Lightning subsystem so it can attach to a SAN Volume Controller.

Logical unit creation and deletion on HDS Lightning

Before you create or delete a logical unit on HDS Lightning, consider the following prerequisites.

The SAN Volume Controller supports Logical Unit expansion (called LUSE). LUSE is accomplished by integrating between 2 and 26 existing logical units together.

Before LUSE can be performed on a logical unit, it must be unmounted from a host, and have no available paths. The LUSE procedure will destroy all data that exists on the logical unit, except for logical units on a Windows system. LUSE must not be performed on any disk that is managed by the SAN Volume Controller. If data exists on a disk and you wish to bring in the data via image mode, LUSE should not be used on the disk prior to you importing the data.

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

"Configuring FAStT disk controllers for the storage server" on page 279
FAStT disk controllers provide functionality that is compatible with the SAN Volume Controller.

Configuring settings for Lightning

The Lightning configuration interface provides functionality for configuration.

These options and settings can have a scope of a:

- Subsystem
- Port
- Logical unit

Related reference

"Global settings for Lightning"
Global settings apply across a Lightning disk controller system.

"Controller settings for Lightning" on page 296
Controller settings are settings that apply across the entire Lightning controller.

"Port settings for Lightning" on page 296
Port settings are configurable at the port level.

"LU settings for Lightning" on page 296
Logical unit (LU) settings apply to individual LUs configured in the Lightning controller.

Related information

"Configuring the HDS Lightning subsystem" on page 290
This part provides information about configuring the HDS Lightning subsystem so it can attach to a SAN Volume Controller.

Global settings for Lightning

Global settings apply across a Lightning disk controller system.

The following table lists the global settings for Lightning.

Table 32. Lightning global settings supported by the SAN Volume Controller

Option	Lightning Default Setting	SAN Volume Controller Required Setting
Spare disk recover	Interleave	Interleave
Disk copy place	Medium	Medium
Copy operation	Correction copy and dynamic sparing	Correction copy and dynamic sparing

Table 32. Lightning global settings supported by the SAN Volume Controller (continued)

Option	Lightning Default Setting	SAN Volume Controller Required Setting
Read configuration data mode	Selected	Selected
PS off timer	Not selected	Not selected

Related reference

"Configuring settings for Lightning" on page 295
The Lightning configuration interface provides functionality for configuration.

Controller settings for Lightning

Controller settings are settings that apply across the entire Lightning controller.

Table 33. Lightning controller settings supported by the SAN Volume Controller

Option	Lightning Default Setting	SAN Volume Controller Required Setting
PCB mode	Standard	Standard

Related reference

"Configuring settings for Lightning" on page 295
The Lightning configuration interface provides functionality for configuration.

Port settings for Lightning

Port settings are configurable at the port level.

This topic lists the port settings, the Lightning defaults, and the settings that are required for the SAN Volume Controller. Port settings for clearing LUN reservations are also listed.

There are no options available with the scope of a single controller.

- The ports are included in switch zones
- The switch zones only present the ports directly to the hosts and not to a SAN Volume Controller.

Table 34. Lightning port settings supported by the SAN Volume Controller

Option	Lightning Default Setting	SAN Volume Controller Required Setting
Address	AL/PA	AL/PA
Fabric	On	On
Connection	Point-to-Point	Point-to-Point
Security switch	On	On or off
Host type	Default	Windows

Related reference

"Configuring settings for Lightning" on page 295
The Lightning configuration interface provides functionality for configuration.

LU settings for Lightning

Logical unit (LU) settings apply to individual LUs configured in the Lightning controller.

Lightning LUs must be configured as described in Table 35 if the Logical Unit Number (LUN) is associated with ports in a switch zone that is accessible to the SAN Volume Controller.

Table 35. Lightning LU settings for the SAN Volume Controller

Option	Required Values	SAN Volume ControllerDefault Setting
Command device	Off	Off
Command security	Off	Off

Note: These settings only apply to LUs accessible by SAN Volume Controller.

Related reference

"Configuring settings for Lightning" on page 295
The Lightning configuration interface provides functionality for configuration.

Configuring the HDS Thunder subsystem

This part provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

Related tasks

"Setting up a Thunder with more than four ports" on page 300
Perform the following steps to set up a Thunder with more than four ports.

Related reference

"Supported models of the HDS Thunder" on page 298
The following table lists the models of the Thunder that are supported by the SAN Volume Controller. See the following Web site for the latest supported models. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

"Supported firmware levels for HDS Thunder" on page 298
See the following Web site for specific HDS Thunder firmware levels and the latest supported hardware. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

"Concurrent maintenance on the HDS Thunder" on page 298
Concurrent maintenance of the HDS Thunder 9200 and 9500V is not supported with the SAN Volume Controller.

"User interface on the HDS Thunder" on page 299
Ensure that you are familiar with user interface application that supports the HDS Thunder subsystem.

"Sharing the HDS Thunder between host and the SAN Volume Controller" on page 299
The HDS Thunder 9200 and 95xxV can be shared between a host and a SAN Volume Controller according to certain restrictions.

"Quorum disks on HDS Thunder" on page 301
Managed disks presented by the Thunder 9200 and 95xxV may be chosen by the SAN Volume Controller as quorum disks.

"Advanced functions for HDS Thunder" on page 301
The HDS Thunder provides advanced functions that may or may not be compatible with the SAN Volume Controller.

"Logical unit creation and deletion on HDS Thunder" on page 302
Before you create or delete a logical unit on HDS Thunder, consider the following constraints.

"Configuring settings for HDS Thunder" on page 303
The Thunder configuration interface provides functionality for configuration.

Supported models of the HDS Thunder

The following table lists the models of the Thunder that are supported by the SAN Volume Controller. See the following Web site for the latest supported models. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Table 36. Supported Thunder 9200 models

Model	Description
Thunder 9200 Rackmount	Up to 100 disks
Thunder 9200 Deskside 20	Maximum of 20 disks
Thunder 9200 Deskside 10	Maximum of 10 disks

Table 37. Supported Thunder 95xxV models

Model	Description
Thunder 9530V Deskside/Rackmount	Supports 2 - 14 disks
Thunder 9531V Deskside	Pre-configured with 5 disks
Thunder 9532V Deskside	Pre-configured with 9 disks
Thunder 9533V Deskside	Pre-configured with 13 disks
Thunder 9570V Deskside/Rackmount	Supports 2 - 224 disks
Thunder 9580V Rackmount	Supports 5 - 449 disks
Thunder 9585V Rackmount	Supports 5 - 449 disks

Related information

"Configuring the HDS Thunder subsystem" on page 297
This part provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

Supported firmware levels for HDS Thunder

See the following Web site for specific HDS Thunder firmware levels and the latest supported hardware. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Note: Concurrent upgrade of the controller firmware is supported with the SAN Volume Controller.

Related information

"Configuring the HDS Thunder subsystem" on page 297
This part provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

Concurrent maintenance on the HDS Thunder

Concurrent maintenance of the HDS Thunder 9200 and 9500V is not supported with the SAN Volume Controller.

Concurrent upgrade of the controller firmware is supported with the SAN Volume Controller.

Related information

"Configuring the HDS Thunder subsystem" on page 297
This part provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

User interface on the HDS Thunder

Ensure that you are familiar with user interface application that supports the HDS Thunder subsystem.

In-band configuration

Ensure that you have the Thunder command LUN disabled when using the Thunder user interface application.

DAMP

Disk Array Management Program (DAMP) is the primary user interface application for configuration. Use DAMP for firmware upgrades, changing settings, and creating and monitoring storage. Use DAMP for the initial configuration of the subsystem.

DAMP supports an Ethernet connection to Thunder. An out-of-band command-line interface is available with DAMP. It provides the majority of the functions that are provided in DAMP.

HiCommand

Ensure that you have access to DAMP in order to configure the settings using HiCommand.

HiCommand is another configuration user interface available for the Thunder, which also uses Ethernet as a transport. HiCommand is more restricted than DAMP. It allows basic creation of storage and some monitoring features. HiCommand works for both Thunder and Lightning subsystems.

Web Server

A Web server is running on each of the controllers on the subsystem. During normal operation, the user interface only allows basic monitoring of the subsystem and displays an error log. If a controller is put into diagnostic mode by pressing the reset button on the controller, the user interface allows for firmware upgrades and subsystem configuration resets.

Related information

"Configuring the HDS Thunder subsystem" on page 297
This part provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

Sharing the HDS Thunder between host and the SAN Volume Controller

The HDS Thunder 9200 and 95xxV can be shared between a host and a SAN Volume Controller according to certain restrictions.

- The same host cannot be connected to both a SAN Volume Controller and a Thunder at the same time because Hitachi Dynamic Link Manager (HDLM) and the Subsystem Device Driver do not coexist.

- For Thunder 9200 only, a target port cannot be shared between a host and a SAN Volume Controller. In other words, if a target port is used by a SAN Volume Controller it must not be present in a switch zone which allows a host to access the port.
- Logical units (LUs) cannot be shared between a host and a SAN Volume Controller. Thus, Thunder 9200 must be set into M-TID M-LUN mode and Mapping Mode enabled on Thunder 95xx. No LU can have a LUN Number associated with a port which is zoned for host use while also having a LUN Number associated with a port which is zoned for a SAN Volume Controller.

Related concepts

"Storage subsystems" on page 54
Follow these rules when you are planning the configuration of storage subsystems in the SAN fabric.

Related information

"Configuring the HDS Thunder subsystem" on page 297
This part provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

Setting up a Thunder with more than four ports

Perform the following steps to set up a Thunder with more than four ports.

1. Set the Mapping Mode to **Enabled**.
2. Divide the ports into groups of four (or two). For redundancy, at least one port from each controller should be in each group.
3. Make a note of all of the LUNs currently on the array. Each LUN that you want managed by the San Volume Controller should be in one group.
4. Divide the LUNs into groups: One group of LUNs for each group of ports.
5. From the **Host Groups** view:
   a. Select the first port in the first port group.
   b. Select **Option**

      Set the port options.

      Select **Logical Unit**.

      From the menu, select **Modify Mapping**.

      From the Modify Mapping panel:

      1) Select a LUN from the first LUN group from the ″LUN″ column
      2) Select ″Host LUN″ 0, and click **Add**.

         This repositions the mapping to the ″reserved configuration″ column.
      3) Select the next LUN from the first group
      4) Select ″Host LUN″ 1, and click **Add**.

         Repeat the previous step for all ports in the first port group. Ensure that the LUN and Host LUN ids are identical for all ports. Failure to make identical will result in I/O failures.
      5) Repeat the previous two steps for all port groups.

Related information

"Configuring the HDS Thunder subsystem" on page 297
This part provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

Quorum disks on HDS Thunder

Managed disks presented by the Thunder 9200 and 95xxV may be chosen by the SAN Volume Controller as quorum disks.

Managed disks presented by the Thunder 9200 and 95xxV may be chosen by the SAN Volume Controller as quorum disks during initialization of the cluster. The selection made can be changed by the following methods:

- **Set quorum disk** command
- Setting a Quorum Disk panel

Related information

"Configuring the HDS Thunder subsystem" on page 297
This part provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

"Creating a quorum disk" on page 252
A quorum disk is used to resolve tie-break situations when the ″voting set″ of nodes disagree on the current cluster state.

Advanced functions for HDS Thunder

The HDS Thunder provides advanced functions that may or may not be compatible with the SAN Volume Controller.

ShadowImage

ShadowImage is functionally similar to FlashCopy. ShadowImage is not supported when the disk controller system is being used with the SAN Volume Controller. Even when a HDS Thunder is shared between a host and a SAN Volume Controller, ShadowImage is not supported on the ports that are zoned directly with the host.

TrueCopy

TrueCopy is functionally similar to Metro Mirror. TrueCopy is not supported when the disk controller system is being used with the SAN Volume Controller. Even when a HDS Thunder is shared between a host and a SAN Volume Controller, TrueCopy is not supported on the ports that are zoned directly with the host.

LUN Security

LUN Security enables LUN masking by the WWN of the initiator port. This function is not supported for logical units (LUs) used by the SAN Volume Controller.

Partitioning

Thunder supports Partitioning. Partitioning is splitting a RAID array into up to 128 smaller LUs, each of which behaves as an independent disk like entity. The SAN Volume Controller fully supports this function.

Dynamic array expansion

The Thunder allows the last LU defined in a RAID group to be expanded. This function is not supported with the SAN Volume Controller attachment. It must **not** be used for LUs in use by a SAN Volume Controller.

Note: Use in this context means that the LU has an LUN number that is associated with a fibre-channel port, and this fibre-channel port is contained in a switch zone that also contains SAN Volume Controller fibre-channel ports.

Host storage domains (HSD) and virtual fibre-channel ports for Thunder 95xxV

The Thunder 95xxV supports host storage domains (HSD) and virtual fibre-channel ports. Each fibre-channel port may support multiple HSDs. Each host in a given HSD is essentially presented with a virtual target port and a unique set of LUNs.

The Thunder 9200 does not support HSD and virtual fibre-channel ports.

Related concepts

"FlashCopy" on page 34
FlashCopy is a copy service that is available with the SAN Volume Controller.

"Metro Mirror" on page 46
Metro Mirror (formerly known as 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 on the other virtual disk.

Related information

"Configuring the HDS Thunder subsystem" on page 297
This part provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

Logical unit creation and deletion on HDS Thunder

Before you create or delete a logical unit on HDS Thunder, consider the following constraints.

The Thunder configuration interface enables you to create and delete logical unit number (LUNs). You must avoid certain creation and deletion scenarios to prevent data corruption. This topic discusses those scenarios.

Creation and deletion scenarios

The Thunder configuration interface enables you to create and delete LUNs. Certain creation and deletion scenarios must be avoided to prevent data corruption. For example, the configuration interface enables you to create LUN A, delete LUN A, and then create LUN B with the same unique ID as LUN A. Doing this with a SAN Volume Controller attached could cause data corruption because the SAN Volume Controller might not realize that LUN B is different than LUN A.

Attention: Before you delete a LUN using the Thunder configuration interface, the LUN must first be removed from the managed disk group that contains it.

Dynamic addition of LUNs

Perform the following procedure to add LUNs dynamically. Using this procedure prevents the existing LUNs from rejecting I/O and returning a status of unavailable during dynamic addition of LUNs.

1. Create the new LUNs using the Disk Array Management Program (DAMP), which is the Thunder configuration tool.
2. Quiesce all I/O.

3. Perform either an offline format or an online format of all new LUNs on the controller using DAMP. Wait for the format to complete.
4. Go into the LUN mapping function of DAMP. Add mapping for the new LUN to all of the controller ports that are available to the SAN Volume Controller on the fabric.
5. Restart the controller. (Model 9200 only)
6. After the controller has restarted, restart I/O.

LUN mapping considerations

If LUN mapping is used as described in the LUN mapping topic, the controller must be restarted to pick up the new LUN mapping configuration. For each managed disk group (MDisk) group that contains an MDisk that is supported by an LU on the Thunder disk controller, all virtual disks in those MDisk groups will go offline.

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

"Managed disks" on page 23
A managed disk (MDisk) is a logical disk (typically a RAID array or partition thereof) that a storage subsystem has exported to the SAN fabric to which the nodes in the cluster are attached.

Related reference

"Mapping and virtualization settings for HDS Thunder" on page 308
Thunder supports different modes of operation. These modes affect LUN mapping or masking and virtualization.

Related information

"Configuring the HDS Thunder subsystem" on page 297
This part provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

Configuring settings for HDS Thunder

The Thunder configuration interface provides functionality for configuration.

These options and settings can have a scope of a:

- Subsystem
- Port
- Logical unit

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Related reference

“Global settings for the HDS Thunder”
Global settings apply across a Thunder disk controller system.

“Controller settings for HDS Thunder” on page 305
Controller settings are settings that apply across the entire HDS Thunder controller. There are no options available with the scope of a single controller.

“Port settings for the HDS Tunder” on page 305
Port settings are configurable at the port level.

“LU settings for the HDS Thunder” on page 307
Logical unit (LU) settings apply to individual LUs configured in the Thunder controller.

“Mapping and virtualization settings for HDS Thunder” on page 308
Thunder supports different modes of operation. These modes affect LUN mapping or masking and virtualization.

Related information

“Configuring the HDS Thunder subsystem” on page 297
This part provides information about configuring the HDS Thunder subsystem so it can attach to a SAN Volume Controller.

“Servicing storage subsystems” on page 253
When servicing your storage subsystems, it is imperative that you follow the service instructions contained in the vendor documentation.

Global settings for the HDS Thunder

Global settings apply across a Thunder disk controller system.

The following table lists the global settings for HDS Thunder.

Table 38. Thunder global settings supported by the SAN Volume Controller

Option	Thunder Default Setting	SAN Volume Controller Required Setting
Start attribute	Dual active mode	Dual active mode
SCSI ID/Port takeover mode	N/A	N/A
Default controller	N/A	N/A
Data-share mode	Used	Used
Serial number		Same as the Thunder default setting
Delay planned shutdown	0	0
Drive detach mode	False	False
Multipath controller (Thunder 9200 only)	False	False
PROCOM mode	False	False
Report status	False	False
Multipath (Array unit)	False	False
Turbu LU warning	False	False
NX mode	False	False
Auto reconstruction mode	False	False
Forced write-through mode	False	False
Changing logical unit mode 1	False	False
Multiple stream mode (Thunder 9200 only)	False	False

Table 38. Thunder global settings supported by the SAN Volume Controller (continued)

Option	Thunder Default Setting	SAN Volume Controller Required Setting
Multiple stream mode (write) (Thunder 95xxV only)	False	False
Multiple stream mode (read) (Thunder 95xxV only)	False	False
RAID 3 mode (Thunder 9200 only)	False	False
Target ID (9200 only) Mapping mode on 95xx	S-TID, M-LUN	M-TID, M-LUN (if sharing controller, otherwise S-TID, M-LUN)
Data striping size	16K; 32K; 64K	Any (Thunder 9200) 64K (Thunder 95xxV)
Operation if processor failure occurs	Reset the fault	Reset the fault
Command queuing	True	True
ANSI Version	N/A	N/A
Vendor ID	HITACHI	HITACHI
Product ID (Thunder 9200)	DF500F	DF500F
Product ID (Thunder 95xxV)	DF500F	DF600F
ROM microprogram version	<Empty>	<Empty>
RAM microprogram version	<Empty>	<Empty>
Web title	<Empty>	Any setting supported
Cache mode (Thunder 9200 only)	All off	All off
Link separation (Thunder 9200 only)	False	False
ROM Pseudo-response command processing (Thunder 9200 only)	N/A	N/A
Save data pointer response (Thunder 9200 only)	N/A	N/A
Controller identifier	False	False
RS232C error information outflow mode	Off	Any
Execute write and verify mode	True	True

Controller settings for HDS Thunder

Controller settings are settings that apply across the entire HDS Thunder controller. There are no options available with the scope of a single controller.

Port settings for the HDS Tunder

Port settings are configurable at the port level.

This topic lists those per port settings, what the Thunder defaults are, and what the required settings are for the SAN Volume Controller. Port settings for clearing LUN reservations are also listed.

The settings listed in the table Table 39 apply to those HDS Thunder 9200 disk controllers that are in a switch zone that contains SAN Volume Controllers. If the Thunder disk controller is shared between a SAN Volume Controller and another host, you can configure with different settings than shown if both of the following conditions are true:

There are no options available with the scope of a single controller.

- The ports are included in switch zones
- The switch zones only present the ports directly to the hosts and not to a SAN Volume Controller

Table 39. Thunder port settings supported by the SAN Volume Controller

Option	**Thunder Default Setting**	**SAN Volume Controller Required Setting**
Host connection mode 1	Standard	Standard
VxVM DMP mode (Thunder 9200 only)	False	False
HP connection mode	False	False
Report inquiry page 83H (Thunder 9200 only)	False	True
UA (06/2A00) suppress mode	False	False
HISUP mode	False	False
CCHS mode	False	False
Standard inquiry data expand (Thunder 9200 only)	False	False
Host connection mode 2	False	False
Product ID DF400 mode	False	False
HBA WWN report mode (Thunder 9200 only)	False	False
NACA mode	False	False
SUN cluster connection mode	False	False
Persistent RSV cluster mode	False	False
ftServer connection mode 1 (Thunder 9200 only)	False	False
ftServer connection mode 2	False	False
SRC Read Command reject	False	False
Reset/LIP mode (signal)	False	False
Reset/LIP mode (progress)	False	False
Reset ALL LIP port mode	False	False
Reset target (reset bus device mode)	False	True
Reserve mode	False	True
Reset logical unit mode	False	True

Table 39. Thunder port settings supported by the SAN Volume Controller (continued)

Option	Thunder Default Setting	SAN Volume Controller Required Setting
Reset logout of third party process mode	False	False
Read Frame minimum 128 byte mode (Thunder 950xxV only)	False	False
Topology	Point-to-point	Fabric

LU settings for the HDS Thunder

Logical unit (LU) settings apply to individual LUs configured in the Thunder controller.

Logical unit (LU) settings apply to individual LUs configured in the Thunder controller. Thunder LUs must be configured as described in Table 40 if the Logical Unit Number (LUN) is associated with ports in a switch zone that is accessible to the SAN Volume Controller.

Table 40. Thunder LU settings for the SAN Volume Controller

Option	Required Values	Default Setting
LUN default controller	Controller 0 or Controller 1	Any

Note: These settings only apply to LUs accessible by SAN Volume Controller.

Data corruption scenarios to avoid

Scenario 1: The configuration application enables you to change the serial number for an LU. Changing the serial number also changes the unique user identifier (UID) for the LU. Since the serial number is also used to determine the WWPN of the controller ports, two LUNs cannot have the same unique ID on the same SAN because two controllers cannot have the same WWPN on the same SAN.

Scenario 2: The serial number is also used to determine the WWPN of the controller ports. Therefore, two LUNs must not have the same ID on the same SAN, because this results in two controllers having the same WWPN on the same SAN. This is not a valid configuration.

Attention: Do not change the serial number for an LU that is managed by a SAN Volume Controller because this could result in data loss or undetected data corruption.

Scenario 3: The configuration application enables you to create LUN A, delete LUN A, and create LUN B with the same unique ID as LUN A. Doing this with a LUN managed by a SAN Volume Controller could result in data corruption because the SAN Volume Controller might not recognize that LUN B is different than LUN A.

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Mapping and virtualization settings for HDS Thunder

Thunder supports different modes of operation. These modes affect LUN mapping or masking and virtualization.

SAN Volume Controller supports the S-TID M-LUN and M-TID M-LUN modes on Thunder 9200, and Mapping Mode enabled or disabled on Thunder 95xx.

Attention: Thunder does not provide an interface that enables a SAN Volume Controller to detect and ensure that the mapping or masking and virtualization options are set properly. Therefore, you or someone else in your organization must ensure that these options are set as outlined in this topic.

S-TID M-LUN modes

In S-TID M-LUN mode all LUs are accessible via all ports on Thunder with the same LUN Number on each port. This is the simplest mode and it should be used for all situations, except, where a Thunder subsystem is being shared between a host and a SAN Volume Controller.

M-TID M-LUN modes

If a Thunder is being shared between a host and a SAN Volume Controller, then you must use M-TID M-LUN mode. Configure your Thunder so that all SAN Volume Controllers accessible LUs have the same LUN Number on all ports through which they can be accessed.

A SAN Volume Controller can access controller ports x and y. The SAN Volume Controller also sees an LU on port x that has LUN number p. In this situation the following conditions must be met:

- The SAN Volume Controller must see either the same LU on port y with LUN number p or it must not see the LU at all on port y.
- The LU should not appear as any other LUN number on port y.
- The LU must not be mapped to any Thunder port which is zoned for use directly by a host in a configuration where the Thunder is shared between a host and a SAN Volume Controller.

M-TID M-LUN mode enables LU virtualization by target port. In this mode, a single LU can be seen as different LUN numbers across all of the controller ports. For example, LU A may be LUN 0 on controller port 1, LUN 3 on controller port 2, and not visible at all on controller ports 3 and 4.

Note: The SAN Volume Controller does not support this.

In addition, M-TID M-LUN mode enables a single LU to be seen as multiple LUN numbers on the same controller port. For example, LU B may be LUN 1 and LUN 2 on controller port 1.

Note: The SAN Volume Controller does not support this.

Note: Thunder 9200 controllers require a reboot in order for changes to LUN mapping to take effect.

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Configuring the HP StorageWorks subsystem

This part provides information about configuring the HP StorageWorks subsystem, which uses the HSG80 controller, so it can attach to a SAN Volume Controller.

Support for HSG80 based products is limited in SAN Volume Controller version 1.1.1 to a single port connection. Version 1.2.0 enables multi-port connection but this introduces some restrictions on LUN partitioning.

Managed Disk Groups and MDisks

Attention: This topic only applies to SAN Volume Controller version 1.1.1 when only a single port connection was supported.

An MDisk group should contain either no HSG80 LUNs or LUNs that are only from a single HSG80 subsystem. **No other configuration is supported.**An MDisk group that consists of LUNs from HSG80 storage and non-HSG80 storage would potentially contain a single point of failure, if the HSG80 subsystem was connected to the cluster by a single port. Consequently, any virtual disks created from such a MDisk group would potentially contain a single point of failure.

Related concepts

"Managed disk groups" on page 26
An *MDisk group* is a collection of MDisks that jointly contain all the data for a specified set of virtual disks (VDisks).

Related tasks

"Configuring the HP StorageWorks controller" on page 313
To configure the HP StorageWorks HSG80 controller, complete these steps.

"Creating managed disk groups" on page 107
You can create a new managed disk (MDisk) group using the Create a Managed Disk Group wizard.

Related reference

"Supported models of the HP StorageWorks controller" on page 316
The following table lists the models of the HP StorageWorks controller that are supported by the SAN Volume Controller. See the following Web site for the latest supported models. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

"Supported firmware levels for the HP StorageWorks controller" on page 317
See the following Web site for specific HP StorageWorks firmware levels and the latest supported hardware. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

“Concurrent maintenance on the HP StorageWorks” on page 317
Concurrent maintenance is the capability to perform I/O operations to a HP StorageWorks controller while simultaneously performing maintenance operations on it.

“Configuration interface for the HP StorageWorks” on page 317
The Command Console configuration and service utility is the configuration interface for the HSG80 subsystem.

“Sharing the HP StorageWorks controller between a host and the SAN Volume Controller” on page 318
The HP StorageWorks controller can be shared between a host and a SAN Volume Controller.

“Switch zoning limitations for the HP StorageWorks subsystem” on page 318
Consider the following limitations when planning switch zoning and connection to the SAN.

“Quorum disks on HP StorageWorks” on page 319
Managed disks presented by the HP StorageWorks controllers will be chosen by the SAN Volume Controller as quorum disks.

“Support for HP StorageWorks advanced functions” on page 320
HP StorageWorks provides advanced functions that may or may not be compatible with the SAN Volume Controller.

“HP StorageWorks advanced functions” on page 320
VDisks that are created from MDisks presented by an HSG80 controller may be used in SAN Volume Controller FlashCopy mappings or SAN Volume Controller Metro Mirror relationships.

“Logical unit creation and deletion on the HP StorageWorks” on page 321
Ensure you are familiar with the HSG80 container types for logical unit configuration.

“Configuring settings for the HP StorageWorks” on page 321
The HP StorageWorks HSG80 configuration interface provides configuration settings and options that are supported with the SAN Volume Controller.

Related information

“HP StorageWorks definitions”
The following terms are used in the IBM and HP documentation and have different meanings.

HP StorageWorks definitions

The following terms are used in the IBM and HP documentation and have different meanings.

IBM term	IBM definition	HP term	HP definition
container	A visual user-interface component that holds objects.	**container**	(1) Any entity that is capable of storing data, whether it is a physical device or a group of physical devices. (2) A virtual, internal controller structure representing either a single disk or a group of disk drives linked as a storageset. Stripesets and mirrorsets are examples of storageset containers that the controller uses to create units.
device	A piece of equipment that is used with the computer. A device does not generally interact directly with the system, but is controlled by a controller.	**device**	In its physical form, a magnetic disk that can be attached to a SCSI bus. The term is also used to indicate a physical device that has been made part of a controller configuration; that is, a physical device that is known to the controller. Units (virtual disks) can be created from devices, once the devices have been made known to the controller.
just a bunch of disks (JBOD)	See *non-RAID*.	**just a bunch of disks (JBOD)**	A group of single-device logical units not configured into any other container type.
mirrorset	See *RAID 1*.	**mirrorset**	A RAID storageset of two or more physical disks that maintains a complete and independent copy of the entire virtual disk's data. This type of storageset has the advantage of being highly reliable and extremely tolerant of device failure. Raid level 1 storagesets are referred to as mirrorsets.
non-RAID	Disks that are not in a redundant array of independent disks (RAID).	**non-RAID**	See *just a bunch of disks*.

IBM term	IBM definition	HP term	HP definition
RAID 0	RAID 0 allows a number of disk drives to be combined and presented as one large disk. RAID 0 does not provide any data redundancy. If one drive fails, all data is lost.	**RAID 0**	A RAID storageset that stripes data across an array of disk drives. A single logical disk spans multiple physical disks, allowing parallel data processing for increased I/O performance. While the performance characteristics of RAID level 0 is excellent, this RAID level is the only one that does not provide redundancy. Raid level 0 storagesets are referred to as stripesets.
RAID 1	A form of storage array in which two or more identical copies of data are maintained on separate media. Also known as mirrorset.	**RAID 1**	See *mirrorset*.
RAID 5	A form of parity RAID in which the disks operate independently, the data strip size is no smaller than the exported block size, and parity check data is distributed across the disks in the array.	**RAID 5**	See *RAIDset*.
RAIDset	See *RAID 5*.	**RAIDset**	A specially developed RAID storageset that stripes data and parity across three or more members in a disk array. A RAIDset combines the best characteristics of RAID level 3 and RAID level 5. A RAIDset is the best choice for most applications with small to medium I/O requests, unless the application is write intensive. A RAIDset is sometimes called parity RAID. RAID level 3/5 storagesets are referred to as RAIDsets.
partition	A logical division of storage on a fixed disk.	**partition**	A logical division of a container represented to the host as a logical unit.
stripeset	See *RAID 0*.	**stripeset**	See *RAID 0*.

Configuring the HP StorageWorks controller

To configure the HP StorageWorks HSG80 controller, complete these steps.

Assume that the HP StorageWorks subsystem is not in use.

Note: An eight-node SAN Volume Controller cluster cannot be connected to an HSG80-based subsystem via all four target ports. This would exceed the limit of 96 process logins. In this type of configuration, only two of the target ports can be in the same switch zone as the SAN Volume Controller initiator ports.

1. Verify that the SAN Volume Controller front panel is clear of errors.
2. Ensure that the HP StorageWorks Operator Control Panel (OCP) on each HSG80 controller is clear of errors. The Operator Control Panel consists of seven green LED's at the rear of each HSG80 controller.
3. Ensure that you can use an HP StorageWorks command-line interface (CLI) to configure the HSG80 controllers.
4. Issue the **SHOW THIS** command and **SHOW OTHER** command to verify the following:
   a. That the controller software is at the supported level. See http://www-1.ibm.com/servers/storage/support/virtual/2145.html.
   b. Ensure that the controllers are configured for MULTIBUS FAILOVER with each other.
   c. Ensure that the controllers are running in SCSI-3 mode.
   d. Ensure that MIRRORED_CACHE is enabled.
   e. Ensure that the Host Connection Table is *not* locked.
5. Issue the **SHOW DEVICES FULL** command to verify the following:
   a. Ensure that none of the LUNs are TRANSPORTABLE.
   b. Ensure that all LUNs are configured. For example, the LUNs report their serial numbers and TRANSFER_RATE_REQUESTED correctly.
6. Issue the **SHOW FAILEDSET** command to verify that there are no failing disks.

   Note: To verify, there should be no orange lights on any disks in the subsystem.
7. Issue the **SHOW UNITS FULL** command to verify the following:
   a. Ensure that all LUNs are set to RUN and NOWRITEPROTECT
   b. Ensure that all LUNs are ONLINE to either THIS or OTHER controller.
   c. Ensure that all LUNs that are to be made available to the SAN Volume Controller have ALL access.
   d. Ensure that all LUNs have Host Based Logging NOT specified.
8. If you have partitioned LUNs, refer to the HP StorageWorks controllers topic.
9. Issue the **SHOW CONNECTIONS FULL** command to verify that you have enough spare entries for all combinations of SAN Volume Controller ports and HP StorageWorks ports.
10. Connect up to four known good fibre-channel cables between your fibre-chanel switches and your HP StorageWorks subsystem.
11. Ensure that your fibre-channel switches are zoned such that the SAN Volume Controller and the HP StorageWorks subsystem are in a zone. Refer to zoning a switch.

12. Issue the **SHOW THIS** command and **SHOW OTHER** command to verify that each connected port is running. Something similar to the following output should be displayed:`PORT_1_TOPOLOGY=FABRIC` (fabric up).
13. Issue the **SHOW CONNECTIONS FULL** command to verify that the new connections have appeared for each SAN Volume Controller port and HP StorageWorks port combination.
14. Verify that `No rejected hosts` is displayed at the end of the SHOW CONNECTIONS output.
15. On the SAN Volume Controller, issue the **svctask detectmdisk** command to discover the controller.
16. Issue the **svcinfo lscontroller** command to verify that the two HSG80 serial numbers appear under the `ctrl s/n`.
17. Issue the **svcinfo lsmdisk** command to verify that the additional MDisks that correspond to the UNITS shown in the HP StorageWorks subsystem.

You can now use the SAN Volume Controller commands to create an MDisk group. You can also create and map VDisks from these MDisk groups. Check your SAN Volume Controller front panel has no errors. Once you ensure that your host has reloaded its fibre-channel driver then you should be able to perform I/O to the VDisks. See the *IBM TotalStorage SAN Volume Controller: Host Attachment Guide* for detailed information.

Related tasks

"Creating managed disk (MDisk) groups using the CLI" on page 169
You can create an MDisk group using the command-line interface (CLI).

"Create virtual disks (VDisks)" on page 172
You can create a VDisk using the command-line interface (CLI).

Related reference

"HP StorageWorks controllers"
The SAN Volume Controller supports HP StorageWorks Controllers.

"Switch zoning for the SAN Volume Controller" on page 65
Consider these constraints when you zone a switch.

HP StorageWorks controllers

The SAN Volume Controller supports HP StorageWorks Controllers.

The SAN Volume Controller version 1.2.0 and later supports up to four fibre-channel connections per HP StorageWorks subsystem. The support for partitioned LUNs is restricted to a single fibre-channel connection.

HP StorageWorks does not recognize more than eight logical units per port. In order to accommodate this, the SAN Volume Controller associates a ″type″ with a host. This can be set using the **svctask mkhost** or **svctask chhost** command. The type can be set to generic, which is the default, or HPUX.

Attention: Before making any changes to your system, consider backing up important data.

If you issue the HP StorageWorks command SHOW UNITS, any units that are partitioned are shown. The following table shows an example.

Table 41. Determining partition usage

HSG80 ″SHOW UNITS″ LUN	Uses	Used by
D1	R50	

Table 41. Determining partition usage (continued)

HSG80 ″SHOW UNITS″ LUN	Uses	Used by
D2	R52	
D3	R53	(partition)
D4	R54	
D5	DISK50000	(partition)
D6	D51	
D7	DISK30300	(partition)
D8	DISK10000	(partition)
D9	R55	

Here *D3, D5, D7* and *D8* are partitioned units.

Scenario 1

This scenario assumes that you have no partitioned units on any HP StorageWorks controllers.

If there are no partitions on any of the HP StorageWorks controllers that are or will be connected to the SAN Volume Controller, then simply ensuring that the SAN Volume Controller code, version 1.2.0, is installed on each of the SAN Volume Controller clusters is all that is necessary. Once this code level is present and running correctly, then additional fibre channel connections can be zoned (and physically connected).

Scenario 2

This scenario assumes that you are using HP StorageWorks controllers on the SAN Volume Controller code, version 1.1.1, with a single fibre channel attached or zoned in. If partitions are present on the HP StorageWorks controllers, then there are two options available:

Option 1: Migrate data from the partitioned units

Migrate data residing on partitioned units and then delete the partitioned units. Perform the following steps to migrate your data:

1. Perform a concurrent code load to get the code to version 1.2.0.
2. Migrate data residing on partitioned units. This can be done in two ways:
   a. Migrate all virtual disks, using the **svctask migratevdisk**, that are in groups that include at least one partitioned unit to groups that contain no partitioned units. You can use the **svcinfo lsmdisk** command and the ″SHOW UNITS FULL″ command to correlate which HP StorageWorks units correspond with which MDisks on the SAN Volume Controller by comparing the unit identifiers (UIDs).
   b. Or, ensure that the managed disks groups have enough unused space on the MDisks that correspond to unpartitioned units for a copy of all the data on MDisks that correspond to the partitioned units. Then, delete the MDisks using the **svctask rmmdisk** command, which may require the use of the force flag.
3. Rezone to take advantage of the extra ports on the HP StorageWorks controller.

Option 2: Retain partitioned units

Perform a concurrent code load to get the SAN Volume Controller code to version 1.2.0. Retain the partitioned units and continue using a single fibre channel attachment.

Note: You must not zone in any additional fibre channel ports on the HP StorageWorks controller because MDisks based on partitioned units will be taken offline. If you have partitioned LUNs that are not allocated to unit numbers and you subsequently add these to your configuration, then these units must be online to the controller that has the zoned in fibre channel port. You can accomplish this by pressing the reset button on the other controller. This is only necessary for unmanaged MDisks.

Scenario 3

This scenario assumes that you have partitions present on a HP StorageWorks controller that is to be connected to the SAN Volume Controller already running version 1.2.0.

In this case, you must initially zone in a single fibre channel connection to one of the HP StorageWorks controllers, and ensure that all the units are online to this controller. You can accomplish this by pressing the reset button on the other controller. You can then choose from the two options shown in Scenario 2. You won't need to do the concurrent code load as the code is already at version 1.2.0.

Supported models of the HP StorageWorks controller

The following table lists the models of the HP StorageWorks controller that are supported by the SAN Volume Controller. See the following Web site for the latest supported models. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Attention: The SAN Volume Controller only supports configurations in which the HSG80 cache is enabled in writeback mode. Running with only a single controller in the HSG80 subsystem results in a single point of data loss.

Note: Transportable disks are not supported for any models.

Table 42. Supported models of the HP StorageWorks HSG80

Model	Description
MA8000	1 controller enclosure (one or two HSG80 controllers), 3 dual bus 14 bay drive enclosures, 22U modular storage cabinet
EMA12000 D14	3 controller enclosures (each with one or two HSG80 controllers), 9 dual bus 14 bay drive enclosures, 42U modular storage cabinet
EMA12000 S14	1 controller enclosure (with one or two HSG80 controllers), 6 single bus 14 bay drive enclosures, 42U modular storage cabinet
EMA12000 Blue	1 controller enclosure (with one or two HSG80 controllers), 3 dual bus 14 bay drive enclosures, 41U modular storage cabinet

Table 42. Supported models of the HP StorageWorks HSG80 (continued)

Model	Description
EMA16000 S14	2 controller enclosures with dual HSG80 controllers, 12 single bus 14 bay drive enclosures, wide 41U storage cabinet
EMA16000 D14	4 controller enclosures with dual HSG80 controllers, 12 dual bus 14 bay drive enclosures, wide 41U storage cabinet

Supported firmware levels for the HP StorageWorks controller

See the following Web site for specific HP StorageWorks firmware levels and the latest supported hardware. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Note: Concurrent upgrade of the controller firmware is also not supported with the SAN Volume Controller.

Concurrent maintenance on the HP StorageWorks

Concurrent maintenance is the capability to perform I/O operations to a HP StorageWorks controller while simultaneously performing maintenance operations on it.

Note: HP StorageWorks maintenance documentation uses the phrase 'rolling upgrade' in place of 'concurrent maintenance'. Please refer to this documentation as, in some instances, the level of I/O must be reduced before performing the maintenance procedure.

The HP StorageWorks Controller allows for concurrent replacement of the following components:

- Drive
- EMU
- Blower
- Dual power supply (one unit can be removed and replaced. The fans' speed increases when only one power supply unit is present.)

The following components are hot-pluggable, but maintenance concurrent with the SAN Volume Controller I/O is not supported.

- Controller

The HP StorageWorks Controller does **not** allow for concurrent replacement of the following components:

- Single power supply (in a single power supply configuration, the enclosure is disabled when the power supply fails.)
- SCSI bus cables
- I/O module
- Cache

Configuration interface for the HP StorageWorks

The Command Console configuration and service utility is the configuration interface for the HSG80 subsystem.

The configuration and service utility can connect to the subsystem in the following ways:

- RS232
- In band over fibre channel
- Over TCP/IP to a proxy agent which then communicates with the HSG80 in band over fibre channel.

In band

Attention: POTENTIAL DATA CORRUPTION

In order for the Command Console to communicate with the HSG80 controllers, the host that runs the service utility must be able to access the HSG80 ports over the SAN. This host could therefore also access LUs that are visible to SVC and cause data corruption. To avoid this, set the UNIT_OFFSET option to 199 for all connections to this host. This ensures that the host is only able to see the CCL.

Related reference

"Connection settings for the HP StorageWorks" on page 325
The HP StorageWorks HSG80 includes options that are configurable at the connection level.

Sharing the HP StorageWorks controller between a host and the SAN Volume Controller

The HP StorageWorks controller can be shared between a host and a SAN Volume Controller.

- A host must not be connected to both a SAN Volume Controller and a HP StorageWorks HSG80 subsystem at the same time.
- Target ports must not be shared between a host and a SAN Volume Controller. Specifically, if a controller port is being used by a SAN Volume Controller it must not be present in a switch zone which enables a host to access the port.
- LU's and RAID arrays must not be shared between a host and a SAN Volume Controller.
- Partitions in the same container must all be either on the SAN Volume Controller or on hosts.

Related concepts

"Storage subsystems" on page 54
Follow these rules when you are planning the configuration of storage subsystems in the SAN fabric.

Switch zoning limitations for the HP StorageWorks subsystem

Consider the following limitations when planning switch zoning and connection to the SAN.

Attention: HSG80 based subsystems are supported with a single controller or dual controllers in the subsystem. Since the SAN Volume Controller only supports configurations in which the HSG80 cache is enabled in write-back mode, running with only a single controller in the HSG80 subsystem results in a single point of data loss.

Switch zoning

For code version 1.1.1, regardless of whether the subsystem uses one or two HSG80 controllers, only a single fibre-channel (FC) port attached to the subsystem can be present in a switch zone that contains SAN Volume Controller FC ports. This guarantees that the nodes in the cluster can access at most one port on an HSG80 subsystem.

For code version 1.2.0, switches should be zoned so that ports on the HSG80 are in the switch zone which contains all ports of the SAN Volume Controller nodes.

Connecting to the SAN

Multiple ports from an HSG80 should be physically connected to the fibre-channel SAN in order to enable servicing of the subsystem. However, switch zoning must be used in the manner outlined in this topic.

Note: If the HPQ Command Console is not able to access an FC port on each of the controllers in a two-controller subsystem, there is a risk of an undetected single point of failure.

If you're using partitioned containers then there are restrictions on the way the HSG80 should be connected to the SAN Volume Controller nodes.

Related reference

"Switch zoning for the SAN Volume Controller" on page 65
Consider these constraints when you zone a switch.

Quorum disks on HP StorageWorks

Managed disks presented by the HP StorageWorks controllers will be chosen by the SAN Volume Controller as quorum disks.

The SAN Volume Controller uses a logical unit (LU) presented by an HSG80 controller as a quorum disk. It will provide a quorum disk even if the connection is by a single port, although this is not recommended. If you are connecting the HSG80 subsystem with a single fibre, then you should ensure that you have another subsystem on which to put your quorum disk. Use the command line **svctask setquorum** to move quorum disks to another subsystem.

Managed disks provided by HSG80 may be chosen by the SAN Volume Controller software as quorum disks and can be set as quorum disks using the command-line interface. This means that clusters attached only to the HSG80 controllers are supported.

Related reference

"Configuring and servicing storage subsystems" on page 233
Follow the guidelines and procedures outlined in this section to make the most of the performance available from your storage subsystems and to avoid potential I/O problems.

Related information

"Creating a quorum disk" on page 252
A quorum disk is used to resolve tie-break situations when the ″voting set″ of nodes disagree on the current cluster state.

Support for HP StorageWorks advanced functions

HP StorageWorks provides advanced functions that may or may not be compatible with the SAN Volume Controller.

FlashCopy and Metro Mirror facilities are provided as advanced functions on the HSG80 controller, however, these controller-delivered copy services are not supported by SAN Volume Controller.

Partitioning

HSG80 supports Partitioning. A partition is a logical division of a container, represented to the host as a logical unit (LU). A container can be a RAID array or a JBOD. Any container type is a candidate for partitions. Any non-transportable disk or storageset can be divided into a maximum of 8 partitions.

Use of this feature is subject to the following restrictions:

- Partitioned containers are fully supported if the HSG80 subsystem is connected to the SAN by a single port.
- Partitioned containers will not be configured by the SAN Volume Controller if the HSG80 subsystem is connected to the SAN by multiple ports.
- Partitioned containers will be removed from the configuration if a single port connection becomes a multi-port connection.
- Partitioned containers will be configured if a multi-port connection becomes a single port connection.

It is recommended that containers are partitioned such that no spare capacity exists because there is no way to detect this 'unused' partition. With a multi-port connection, subsequent attempts to use this capacity will remove all partitions on the container from the configuration.

Dynamic array expansion (LU expansion)

HSG80 does not provide this feature.

Write protection of LUNs

This feature is not supported for use with the SAN Volume Controller.

HP StorageWorks advanced functions

VDisks that are created from MDisks presented by an HSG80 controller may be used in SAN Volume Controller FlashCopy mappings or SAN Volume Controller Metro Mirror relationships.

That is, SAN Volume Controller copy services fully supports the use of MDisks presented by an HSG80 controller.

Related concepts

"Metro Mirror" on page 46
Metro Mirror (formerly known as 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 on the other virtual disk.

Logical unit creation and deletion on the HP StorageWorks

Ensure you are familiar with the HSG80 container types for logical unit configuration.

Table 43 lists the valid container types.

Table 43. HSG80 container types for logical unit configuration

Container	Number of Members	Maximum Size
JBOD - non-transportable **Attention:** Provides no redundancy at the physical disk drive level, that is, a single disk failure may result in the loss of an entire managed disk group and its associated virtual disks.	1	disk size minus metadata
Mirrorset	2 to 6	smallest member
RAIDset	3 to 14	1.024 terabytes
Stripeset	2 to 24	1.024 terabytes
Striped Mirrorset	2 to 48	1.024 terabytes

Note: Logical units can be created and deleted on an HSG80 subsystem while I/O operations are performed to other LUs. You do not need to reboot the HSG80 subsystem.

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Configuring settings for the HP StorageWorks

The HP StorageWorks HSG80 configuration interface provides configuration settings and options that are supported with the SAN Volume Controller.

These options and settings can have a scope of a:

- Subsystem (global)
- Controller
- Port
- Logical unit
- Connection

Related concepts

“Storage subsystems” on page 22
A storage subsystem 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.

Related tasks

“Configuring a balanced storage subsystem” on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Related reference

“Global settings for the HP StorageWorks”
Global settings apply across a HP StorageWorks HSG80 subsystem.

“Controller settings for the HP StorageWorks”
Controller settings are settings that apply across one HSG80 controller.

“Port settings for the HP StorageWorks” on page 323
Port settings are configurable at the port level.

“LU settings for the HP StorageWorks” on page 324
LU settings are configurable at the LU level.

“Connection settings for the HP StorageWorks” on page 325
The HP StorageWorks HSG80 includes options that are configurable at the connection level.

“Mapping and virtualization settings for the HP StorageWorks” on page 326
Consider these constraints for LUN mapping or masking and virtualization from the point of view of the HSG80 controller and their use in a SAN Volume Controller environment.

Related information

“Servicing storage subsystems” on page 253
When servicing your storage subsystems, it is imperative that you follow the service instructions contained in the vendor documentation.

Global settings for the HP StorageWorks

Global settings apply across a HP StorageWorks HSG80 subsystem.

Table 44. HSG80 global settings supported by the SAN Volume Controller

Option	HSG80 Default Setting	SAN Volume Controller Required Setting
DRIVE_ERROR_THRESHOLD	800	Default
FAILEDSET	Not defined	n/a

Related concepts

“Storage subsystems” on page 22
A storage subsystem 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.

Related tasks

“Configuring a balanced storage subsystem” on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Controller settings for the HP StorageWorks

Controller settings are settings that apply across one HSG80 controller.

Table 45 describes the options that can be set by HSG80 command-line interface (CLI) commands for each controller.

Table 45. HSG80 controller settings supported by the SAN Volume Controller

Option	HSG80 Default Setting	SAN Volume Controller Required Setting
ALLOCATION_CLASS	0	Any value
CACHE_FLUSH_TIME	10	Any value
COMMMAND_CONSOLE_LUN	Not defined	Any value
CONNECTIONS_UNLOCKED	CONNECTIONS_ UNLOCKED	CONNECTIONS_ UNLOCKED
NOIDENTIFIER	Not defined	No identifier
MIRRORED_CACHE	Not defined	Mirrored
MULTIBUS_FAILOVER	Not defined	MULTIBUS_FAILOVER
NODE_ID	Worldwide name as on the label	Default
PROMPT	None	Any value
REMOTE_COPY	Not defined	Any value
SCSI_VERSION	SCSI-2	SCSI-3
SMART_ERROR_EJECT	Disabled	Any value
TERMINAL_PARITY	None	Any value
TERMINAL_SPEED	9600	Any value
TIME	Not defined	Any value
UPS	Not defined	Any value

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Port settings for the HP StorageWorks

Port settings are configurable at the port level.

Only one port per HSG80 pair may be used with the SAN Volume Controller.

Note: The port settings are set using the following commands:

- SET THIS PORT_1_TOPOLOGY=FABRIC
- SET THIS PORT_2_TOPOLOGY=FABRIC
- SET OTHER PORT_1_TOPOLOGY=FABRIC
- SET OTHER PORT_2_TOPOLOGY=FABRIC

These values can be checked using the following commands:

- SHOW THIS

- SHOW OTHER

Table 46. HSG80 port settings supported by the

Option	HSG80 Default Setting	HSG80 SAN Volume Controller Required Setting
PORT_1/2-AL-PA	71 or 72	n/a
PORT_1/2_TOPOLOGY	Not defined	FABRIC

Note: The HSG80 supports LUN masking using the ″SET <unit number> ENABLE_ACCESS_PATH command. When used with a SAN Volume Controller, ENABLE_ACCESS_PATH must be set to all (″SET <unit number> ENABLE_ACCESS_PATH=ALL″) and all LUN masking handled exclusively by the SAN Volume Controller. The access rights can be checked using the ″SHOW CONNECTIONS FULL″ and any UNIT_OFFSETs using the ″SHOW CONNECTIONS FULL″ command.

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

"Virtual disk-to-host mapping" on page 32
Virtual disk-to-host mapping is the process of controlling which hosts have access to specific virtual disks (VDisks) within the SAN Volume Controller.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

LU settings for the HP StorageWorks

LU settings are configurable at the LU level.

Table 47 describes the options that must be set for each logical unit that is accessed by the SAN Volume Controller. LUs that are accessed by hosts can be configured differently.

Table 47. HSG80 LU settings supported by the SAN Volume Controller

Option	HSG80 Default Setting	SAN Volume Controller Required Setting
TRANSFER_RATE_REQUESTED	20MHZ	n/a
TRANSPORTABLE/ NOTRANSPORTABLE	NOTRANSPORTABLE	NOTRANSPORTABLE
ENABLE_ACCESS_PATH	ENABLE_ACCESS_ PATH=ALL	ENABLE_ACCESS_ PATH=ALL
DISABLE_ACCESS_PATH (See Note.)	NO DEFAULT	NO DEFAULT
IDENTIFIER/ NOIDENTIFIER	NOIDENTIFIER	n/a
MAX_READ_CACHE_SIZE	32	n/a
MAX_WRITE_CACHE_SIZE	32	64 or higher
MAX_CACHED_TRANSFER_SIZE	32	n/a

Table 47. HSG80 LU settings supported by the SAN Volume Controller (continued)

Option	HSG80 Default Setting	SAN Volume Controller Required Setting
PREFERRED_PATH/ NOPREFERRED_PATH	NOPREFERRED_PATH is set	n/a
READ_CACHE/ NOREAD_CACHE	READ_CACHE	n/a
READAHEAD_CACHE/ NOREADAHEAD_CACHE	READAHEAD_CACHE	n/a
RUN/ NORUN	RUN	RUN
WRITE_LOG/NOWRITE_LOG	NOWRITE_LOG	NOWRITE_LOG
WRITE_PROTECT/ NOWRITE_PROTECT	NOWRITE_PROTECT	NOWRITE_PROTECT
WRITEBACK_CACHE/ NOWRITEBACK_CACHE	WRITEBACK_CACHE	WRITEBACK_CACHE
Note: DISABLE_ACCESS_PATH can be used to disable access from specific hosts. It should always be over ridden by using ENABLE_ACCESS_PATH=ALL on all connections to the SAN Volume Controller nodes.		

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Connection settings for the HP StorageWorks

The HP StorageWorks HSG80 includes options that are configurable at the connection level.

Table 48. HSG80 connection options and their required settings for the

Option	HSG80 Default Setting	HSG80 Required Setting
OPERATING_SYSTEM	Not defined	WINNT
RESERVATION_STYLE	CONNECTION_BASED	n/a
UNIT_OFFSET	0	0 or 199

Related concepts

"Storage subsystems" on page 22
A storage subsystem 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.

Related tasks

"Configuring a balanced storage subsystem" on page 240
The attachment of a storage subsystem to a SAN Volume Controller requires that some specific settings be applied to the device, some limitations are also listed for each storage type.

Related reference

“Configuration interface for the HP StorageWorks” on page 317
The Command Console configuration and service utility is the configuration interface for the HSG80 subsystem.

Mapping and virtualization settings for the HP StorageWorks

Consider these constraints for LUN mapping or masking and virtualization from the point of view of the HSG80 controller and their use in a SAN Volume Controller environment.

The HSG80 configuration interface requires you assign a unit number to each logical unit when it is defined. By default the LUN is the unit number. It is therefore possible for gaps to exist in the LUN range if the unit numbers used in the configuration commands are not contiguous. By default, each LUN is visible on all controller ports on both controllers.

LUN masking

The HSG80 supports the concept of connection names. It supports a maximum of 96 connection names which contain the following parameters:

- HOST_ID
- ADAPTER_ID
- CONTROLLER
- PORT
- REJECTED_HOST

Note: The SAN Volume Controller ports should not be in the REJECTED_HOSTS list. This list can be seen with SHOW CONNECTIONS FULL.

LUN masking must not be used on LUs which are in use by the SAN Volume Controller to restrict either the initiator ports or the Target ports which the SAN Volume Controller uses to access the LUs. Configurations that use LUN masking in this way are not supported. LUN masking can be used to prevent other initiators on the SAN from accessing LUs which are in use by the SAN Volume Controller but the preferred method for this is to use SAN zoning.

LU Virtualization

The HSG80 also provides LU virtualization by the port and by the initiator. This is achieved by specifying a UNIT_OFFSET for the connection. The use of LU Virtualization for connections between the HSG80 target ports and initiator ports on the SAN Volume Controller is not supported.

Related reference

“Switch zoning for the SAN Volume Controller” on page 65
Consider these constraints when you zone a switch.

Configuring the HPQ Enterprise Virtual Array (EVA) subsystem

This part provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Related reference

“Supported models of the HPQ EVA” on page 327
The following table lists the models of the HPQ EVA that are supported by the SAN Volume Controller. See the following Web site for the latest supported models. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

“Supported firmware levels for HPQ EVA” on page 328
See the following Web site for specific HPQ EVA firmware levels and the latest supported hardware. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

“User interface on HPQ EVA” on page 328
Ensure that you are familiar with the user interface which supports the HPQ EVA subsystem.

“Sharing the HPQ EVA controller between a host and the SAN Volume Controller” on page 328
The HPQ EVA controller can be shared between a host and a SAN Volume Controller.

“Switch zoning limitations for the HPQ EVA subsystem” on page 328
Consider the following limitations when planning switch zoning and connection to the SAN.

“Quorum disks on HPQ EVA” on page 329
Managed disks presented by the HPQ EVA controllers will be chosen by the SAN Volume Controller as quorum disks.

“Support for HPQ EVA advanced functions” on page 329
HPQ EVA provides advanced functions that may or may not be compatible with the SAN Volume Controller.

“Logical unit configuration on the HPQ EVA” on page 329
An EVA logical unit is referred to as a virtual disk (VDisk). An EVA subsystem can support up to 512 VDisks. VDisks are created within a set of physical disk drives, referred to as a disk group. A VDisk is striped across all the drives in the group.

“Logical unit presentation” on page 330
A VDisk must be explicitly presented to a host before it can be used for I/O operations.

“Configuration interface for the HPQ EVA” on page 330
The HPQ EVA systems are configured, managed and monitored via a Storage Management Appliance. The Storage Management Appliance is a server that runs a software agent called Command View EVA. The agent is accessed via a graphical user interface that is provided by a standard Web browser.

“Configuring settings for the HPQ EVA” on page 331
The HPQ EVA configuration interface provides configuration settings and options that are supported with the SAN Volume Controller.

Supported models of the HPQ EVA

The following table lists the models of the HPQ EVA that are supported by the SAN Volume Controller. See the following Web site for the latest supported models. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Table 49. Supported HPQ EVAmodels

Model
EVA 3000
EVA 5000

Related information

“Configuring the HPQ Enterprise Virtual Array (EVA) subsystem” on page 326
This part provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Supported firmware levels for HPQ EVA

See the following Web site for specific HPQ EVA firmware levels and the latest supported hardware. http://www-1.ibm.com/servers/storage/support/virtual/2145.html

Note: Concurrent upgrade of the controller firmware is **not** supported with the SAN Volume Controller.

Related information

"Configuring the HPQ Enterprise Virtual Array (EVA) subsystem" on page 326
This part provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

User interface on HPQ EVA

Ensure that you are familiar with the user interface which supports the HPQ EVA subsystem.

Storage Management Appliance

EVA systems are configured, managed and monitored via a Storage Management Appliance. This is a PC server that runs a software agent called Command View EVA. The agent is accessed via a user interface that is provided by a standard web browser.

Command View EVA communicates in-band with the HSV controllers.

Related information

"Configuring the HPQ Enterprise Virtual Array (EVA) subsystem" on page 326
This part provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Sharing the HPQ EVA controller between a host and the SAN Volume Controller

The HPQ EVA controller can be shared between a host and a SAN Volume Controller.

- A host must not be connected to both a SAN Volume Controller and an HPQ EVA subsystem at the same time.
- LU's and RAID arrays must not be shared between a host and a SAN Volume Controller.

Related information

"Configuring the HPQ Enterprise Virtual Array (EVA) subsystem" on page 326
This part provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Switch zoning limitations for the HPQ EVA subsystem

Consider the following limitations when planning switch zoning and connection to the SAN.

Fabric zoning

The SAN Volume Controller switch zone must include at least one target port from each HSV controller in order to have no single point of failure.

Related information

"Configuring the HPQ Enterprise Virtual Array (EVA) subsystem" on page 326
This part provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Quorum disks on HPQ EVA

Managed disks presented by the HPQ EVA controllers will be chosen by the SAN Volume Controller as quorum disks.

Related information

"Configuring the HPQ Enterprise Virtual Array (EVA) subsystem" on page 326
This part provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Support for HPQ EVA advanced functions

HPQ EVA provides advanced functions that may or may not be compatible with the SAN Volume Controller.

FlashCopy and Metro Mirror facilities are provided as advanced functions on the HPQ EVA controller, however, these controller-delivered copy services are not supported by SAN Volume Controller.

SAN Volume Controller FlashCopy and Metro Mirror functions are fully supported for LUs presented by EVA.

EVA advanced functions VSnap and SnapClone are not supported for LUs managed by the SAN Volume Controller.

Related information

"Configuring the HPQ Enterprise Virtual Array (EVA) subsystem" on page 326
This part provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Logical unit configuration on the HPQ EVA

An EVA logical unit is referred to as a virtual disk (VDisk). An EVA subsystem can support up to 512 VDisks. VDisks are created within a set of physical disk drives, referred to as a disk group. A VDisk is striped across all the drives in the group.

The minimum size of a disk group is eight physical drives. The maximum size of a disk group is all available disk drives.

EVA VDisks are created and deleted using the Command View EVA utility.

Note: A VDisk is formatted during the creation process, therefore, the capacity of the VDisk will determine the length of time it takes to be created and formatted. Ensure that you wait until the VDisk is created before you present it to the SAN Volume Controller.

A single VDisk can consume the entire disk group capacity or the disk group may be used for multiple VDisks. The amount of disk group capacity consumed by a VDisk depends on the VDisk capacity and the selected redundancy level. There are three redundancy levels:

- Vraid 1 - High redundancy (mirroring)
- Vraid 5 - Moderate redundancy (parity striping)
- Vraid 0 - No redundancy (striping)

Related reference

"Logical unit creation and deletion on the HPQ EVA"
EVA VDisks are created and deleted using the Command View EVA utility.

Related information

"Configuring the HPQ Enterprise Virtual Array (EVA) subsystem" on page 326
This part provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Logical unit creation and deletion on the HPQ EVA

EVA VDisks are created and deleted using the Command View EVA utility.

VDisks are formatted during creation. The time it takes to format the VDisks depends on the capacity.

Note: Selecting a host for presentation at creation time is not recommended. Ensure that you wait until the VDisk has been created before presenting it to the SAN Volume Controller.

Logical unit presentation

A VDisk must be explicitly presented to a host before it can be used for I/O operations.

The SAN Volume Controller supports LUN masking on an EVA controller. When presenting a VDisk, the LUN can be specified or allowed to default to the next available value.

The SAN Volume Controller supports LUN virtualization on an EVA controller. The LUN-host relationship is set on a per-host basis.

Note: All nodes and ports in the SAN Volume Controller cluster must be represented as one host to the EVA.

HPQ EVA will not recognize more than eight LUNs per port using the generic SCSI behavior. HP StorageWorks does not recognize more than eight logical units per port. In order to accommodate this, the SAN Volume Controller associates a ″type″ with a host. This can be set using the **svctask mkhost** or **svctask chhost** command. The type can be set to generic, which is the default, or HPUX.

Special LUs

The Console LU is a special VDisk that represents the SCSI target device. It is presented to all hosts as LUN 0.

Related information

"Configuring the HPQ Enterprise Virtual Array (EVA) subsystem" on page 326
This part provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Configuration interface for the HPQ EVA

The HPQ EVA systems are configured, managed and monitored via a Storage Management Appliance. The Storage Management Appliance is a server that runs a software agent called Command View EVA. The agent is accessed via a graphical user interface that is provided by a standard Web browser.

In band

The Command View EVA subsystem communicates in-band with the HSV controllers.

Related information

"Configuring the HPQ Enterprise Virtual Array (EVA) subsystem" on page 326
This part provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Configuring settings for the HPQ EVA

The HPQ EVA configuration interface provides configuration settings and options that are supported with the SAN Volume Controller.

These options and settings can have a scope of a:

- Subsystem (global)
- Logical unit
- Host

Related reference

"Global settings for the HPQ EVA"
Global settings apply across a HPQ EVA subsystem.

"LU settings for the HPQ EVA"
LU settings are configurable at the LU level.

"Host settings for the HPQ EVA" on page 332
Host settings are configurable.

Related information

"Configuring the HPQ Enterprise Virtual Array (EVA) subsystem" on page 326
This part provides information about configuring the HPQ Enterprise Virtual Array (EVA) disk controller system so it can attach to a SAN Volume Controller.

Global settings for the HPQ EVA

Global settings apply across a HPQ EVA subsystem.

The following subsystem options can be accessed via the Command View EVA.

Table 50. HPQ EVA global settings supported by the SAN Volume Controller

Option	HPQ EVA Default Setting	SAN Volume Controller Required Setting
Console LUN ID	0	Any
Disk replacement delay	1	Any

LU settings for the HPQ EVA

LU settings are configurable at the LU level.

Table 51 on page 332 describes the options that must be set for each logical unit that is accessed by other hosts. LUs that are accessed by hosts can be configured differently.

Table 51. HPQ EVA LU settings supported by the SAN Volume Controller

Option	HPQ EVA Default Setting	SAN Volume Controller Required Setting
Capacity	None	Any
Write cache	Mirrored Write-back	Mirrored
Read cache	On	On
Redundancy	Vraid0	Any
Preferred path	No preference	No preference
Write protect	Off	Off

Host settings for the HPQ EVA

Host settings are configurable.

The following host options can be accessed via the Command View EVA.

Table 52. HPQ EVA host settings supported by the SAN Volume Controller

Option	HPQ EVA Default Setting	SAN Volume Controller Required Setting
OS type		Windows
Direct eventing	Disabled	Disabled

Chapter 8. IBM TotalStorage support for Microsoft Volume Shadow Copy service

The SAN Volume Controller provides support for the Microsoft Volume Shadow Copy service. The Microsoft Volume Shadow Copy service can provide a point-in-time (″shadow″) copy of a Windows host volume while the volume is mounted and files are in use.

The following components make up the support for the service:

- SAN Volume Controller
- SAN Volume Controller master console
- IBM TotalStorage hardware provider, known as ″IBM TotalStorage Support for Microsoft Volume Shadow Copy Service″
- Microsoft Volume Shadow Copy service

The IBM TotalStorage hardware provider is installed on the Windows host.

To provide the point-in-time shadow copy, the components complete the following process:

1. A backup application on the Windows host initiates a snapshot backup.
2. The Volume Shadow Copy service notifies the IBM TotalStorage hardware provider that a copy is needed.
3. The SAN Volume Controller prepares volumes for a snapshot.
4. The Volume Shadow Copy service quiesces applications (the software applications that are writing data on the host) and flushes file system buffers to prepare for the copy.
5. The SAN Volume Controller makes the shadow copy with the FlashCopy Copy Service.
6. The Volume Shadow Copy service notifies the writing applications that I/O operations can resume, and it notifies the backup application that the backup was successful.

The Volume Shadow Copy service maintains a “free pool” of virtual disks (VDisks) for use as a FlashCopy target and a “reserved pool″ of (VDisks). These pools are implemented as virtual host systems in the SAN Volume Controller.

Installation overview

The steps for implementing the IBM TotalStorage Support for Microsoft Volume Shadow Copy Service must be completed in the correct sequence.

Before you begin, you must have experience with or knowledge of administering a Windows server operating system.

You must have experience with or knowledge of administering a SAN Volume Controller .

Complete the following tasks.

1. Verify that the system requirements are met.
2. If not previously installed, install the SAN Volume Controller master console.
3. Install the IBM TotalStorage hardware provider.

4. Verify the installation.
5. Create a free pool of volumes and a reserved pool of volumes on the SAN Volume Controller.

System requirements for the IBM TotalStorage support for Microsoft Volume Shadow Copy service

Ensure that your system satisfies the following requirements before you install the IBM TotalStorage hardware provider on a Windows Server 2003 operating system.

The following software is required:

- IBM TotalStorage Support for Microsoft Volume Shadow Copy Service version 2.3 or later.
- SAN Volume Controller version 1.2.1.0 or later with the FlashCopy feature enabled.
- SAN Volume Controller master console version 1.2.1.2 or later. You must install the master console *before* you install the IBM TotalStorage hardware provider.
- Windows Server 2003 operating system. The following editions of Windows Server 2003 are supported:
  - Standard Server Edition
  - Enterprise Edition, 32-bit version
  - Datacenter Edition, 32-bit version

Installing the IBM TotalStorage hardware provider

This section includes the steps to install the IBM TotalStorage hardware provider on a Windows server.

You must satisfy all of the prerequisites that are listed in the system requirements section before starting the installation.

During the installation, you will be prompted to enter information about the SAN Volume Controller master console, including the location of the truststore file. The truststore file is generated during the installation of the master console. You must copy this file to a location that is accessible to the IBM TotalStorage hardware provider on the Windows server.

When the installation is complete, the installation program might prompt you to restart the system.

Complete the following steps to install the IBM TotalStorage hardware provider on the Windows server.

1. Log on to the Windows server as an administrator.
2. Run the InstallShield Wizard by inserting the *IBM TotalStorage Support for Microsoft Shadow Copy Service* CD into the CD-ROM drive.
3. The Welcome window opens. Click **Next** to continue with the InstallShield Wizard. You can click **Cancel** at any time to exit the installation. To move back to previous screens while using the wizard, click **Back**.

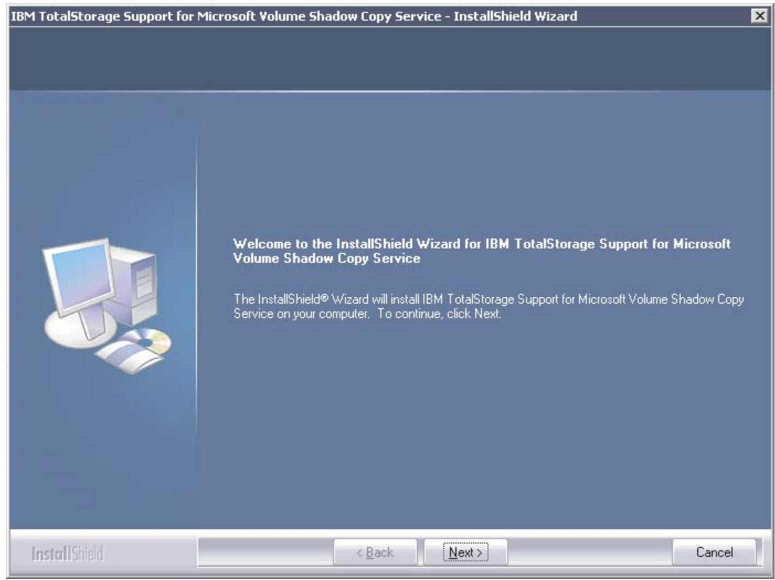

4. The License Agreement window opens. Read the license agreement information. Select whether you accept the terms of the license agreement, and click **Next**. If you do not accept, you cannot continue with the installation.

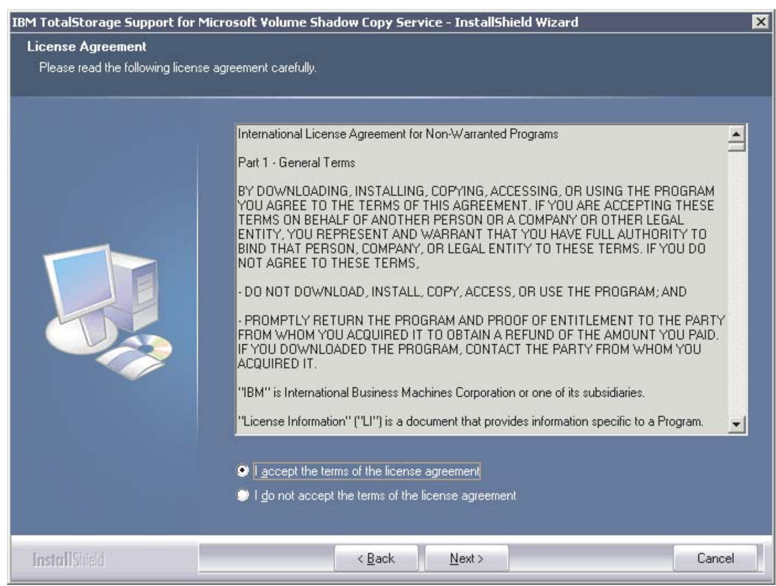

5. The Choose Destination Location Window opens. Click **Next** to accept the default directory where the setup program will install the files, or click **Change** to select a different directory. Click **Next**.

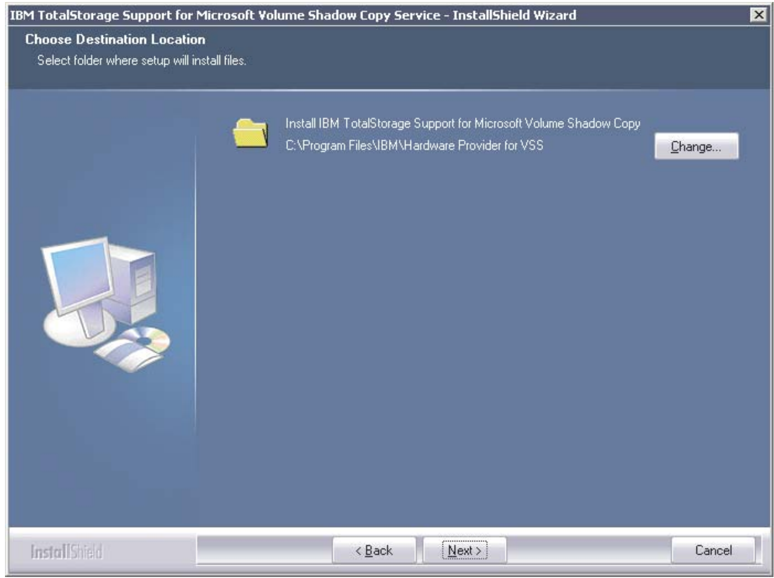

6. The Select Components window opens. Select **Volume Shadow Copy Service for SVC**. Click **Next**.

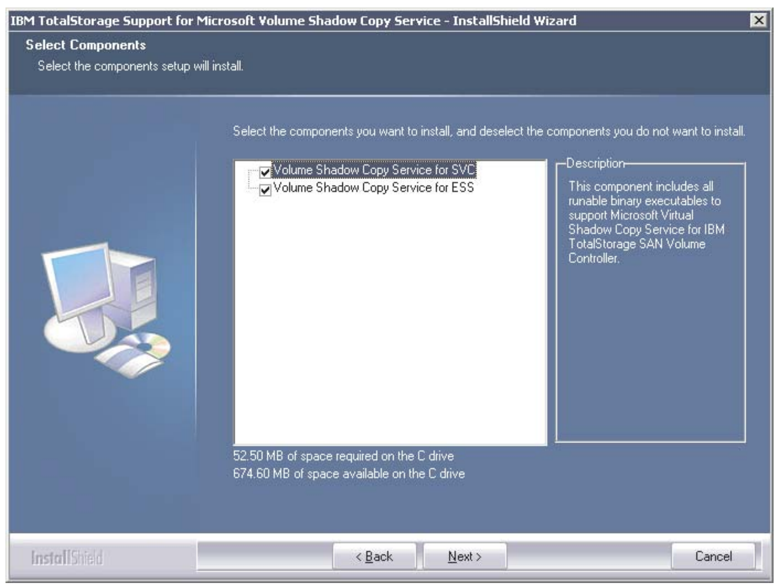

7. The Edit Data window opens.
   - In the **CIM Host** box, type the name of the server where the SAN Volume Controller master console is installed.
   - In the **CIM Port** box, type the port number of the master console. The default value is `5999`.
   - In the **CIM User** box, type the user name that the IBM TotalStorage hardware provider will use to gain access to the master console.

   Click **Next**.

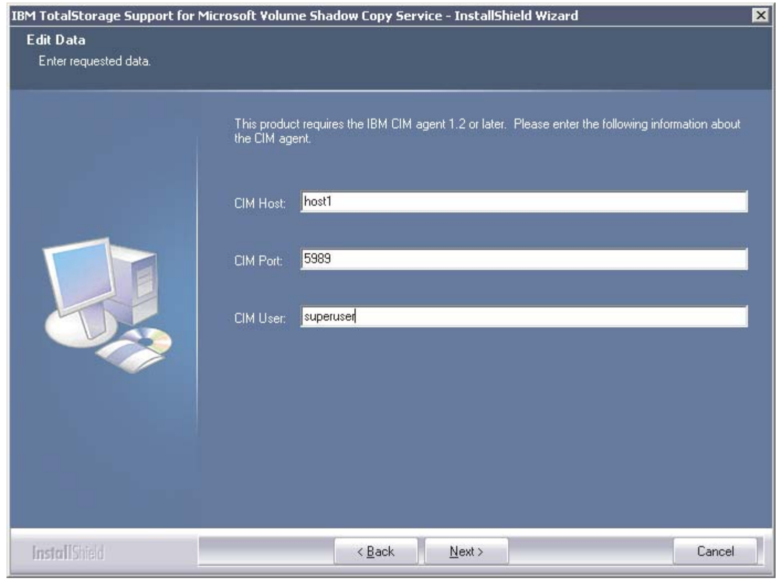

8. The Password window opens. Enter the password for the user name that the IBM TotalStorage hardware provider will use to gain access to the master console, and click **Next**.

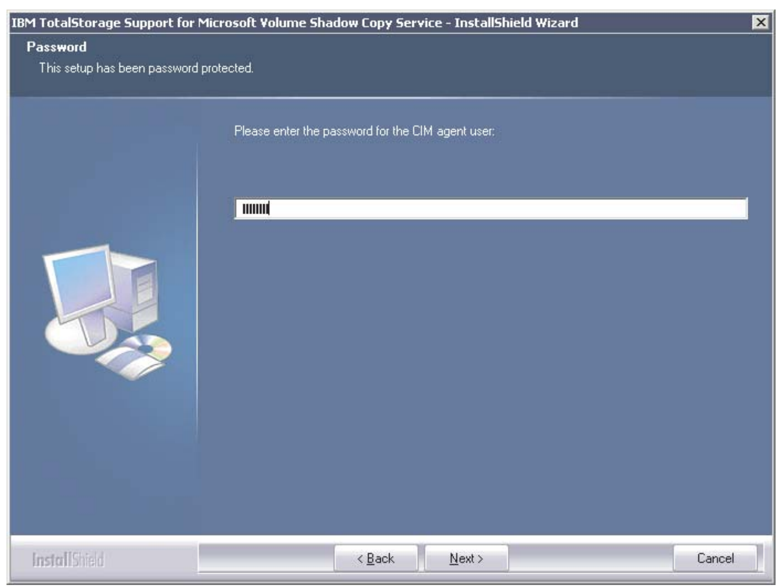

9. The Setup Type window opens. Select whether you want to use the Secure Sockets Layer (SSL) protocol to connect to the master console, and click **Next**. The default communication protocol for the master console is SSL.

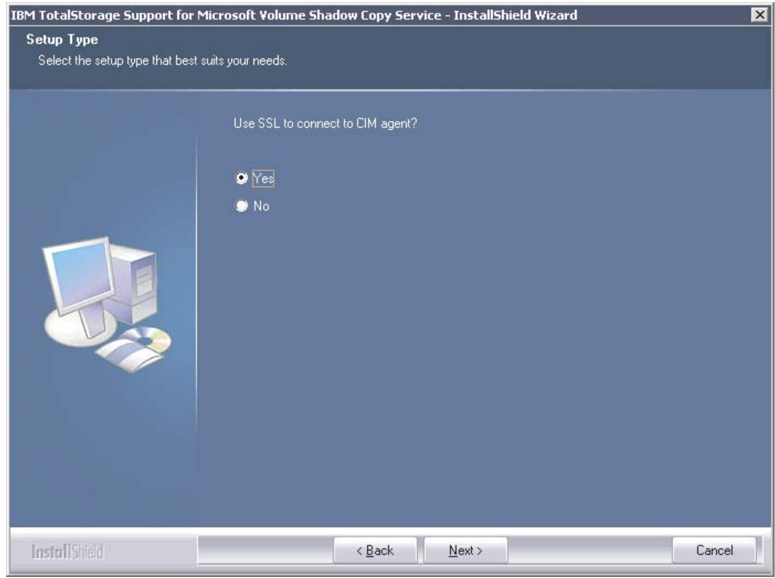

10. The Edit Data window opens.
    - In the **Truststore** box, enter the path to the master console truststore file.
    - In the **Trustpassword** box, enter the truststore password. The default truststore password is “ibmstore”.

    Click **Next**.

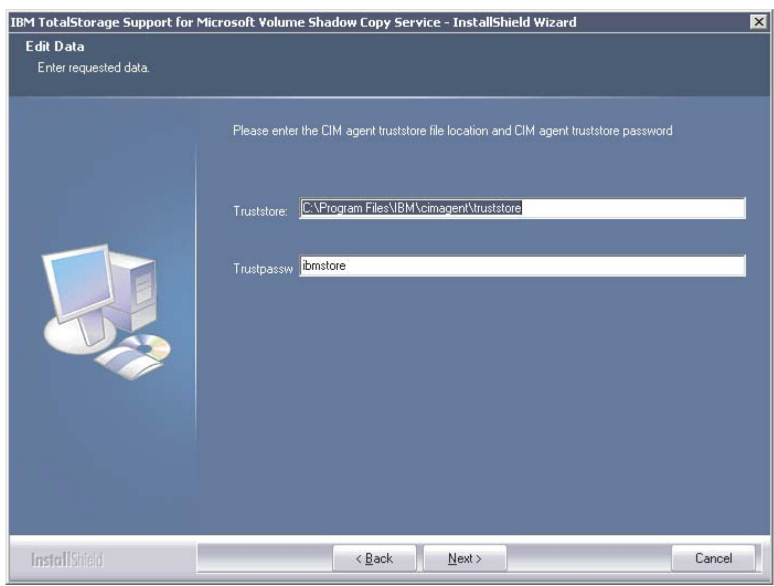

11. The Ready to Install the Program window opens. Click **Install** to begin the installation.

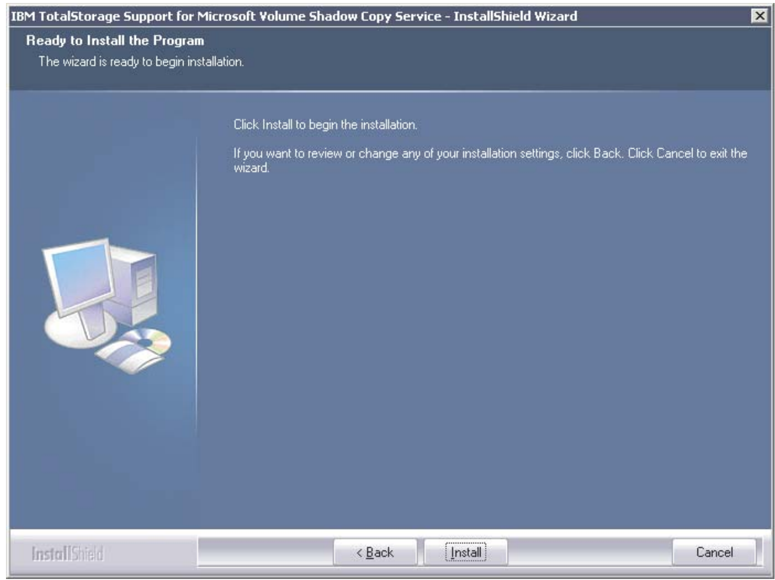

The Setup Status window opens. Wait for the setup to complete.

12. The InstallShield Wizard Complete window opens.
    - If you are prompted to restart the system, you must restart the system to finish the installation.
      – Click **Yes**, and then click **Finish** to restart the system.
      – If you want to restart the system later, click **No**, and then click **Finish** to close the wizard.
    - If you are not prompted to restart the system, click **Finish** to close the wizard.

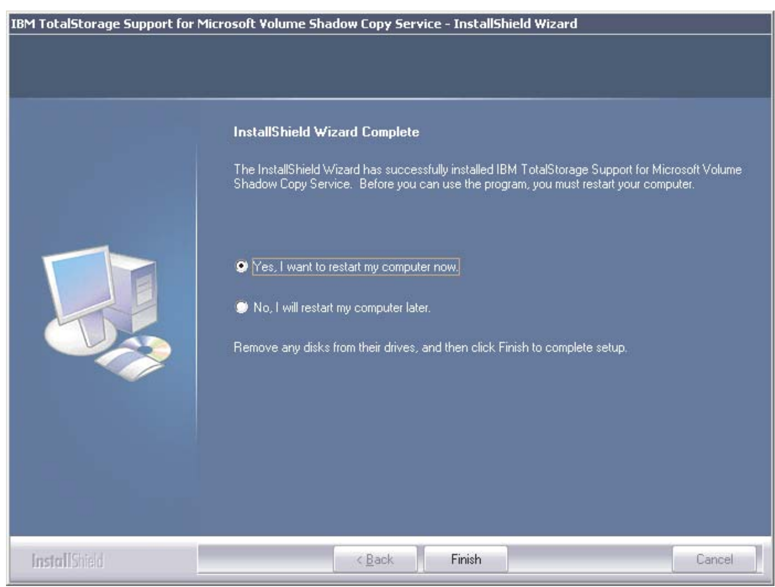

Validating the truststore certificate expiration

In order to successfully log onto the master console, you must ensure that you have a valid truststore certificate.

When signing onto the master console, you might receive a message similar to the following:

```
CMMUI8304E The Administrative server is unable to find a valid
certificate in the truststore file.
```

This message is displayed when a certificate in the truststore file expires. The Administrative server uses the certificates in the truststore file to create a secure connection with the CIM agent. Because the Administrative server cannot find a valid certificate for the CIM agent in the truststore file, no authentication can occur.

To resolve the problem, you must verify that the truststore file was created correctly. If you have any problems, contact your service representative.

Perform the following steps to regenerate a truststore certificate:

1. Go to the C:\Program Files\IBM\svcconsole\cimom directory.
2. Double-click on the **mkcertificate.bat** file. A ″Generating Certificates″ message is displayed. The new certificate is generated and stored in the C:\Program Files\IBM\svcconsole\cimom directory.
3. Copy the truststore file to the following sub directories:

   Note: Each directory begins with C:\Program Files\IBM\svcconsole\console\embeddedWAS...

   C:\...\config\cells\DefaultNode\applications\
   ICAConsole.ear\deployments\ICAConsole\ICAConsole.war\
   WEB-INF

   C:\...\config\cells\DefaultNode\applications\
   SVCConsole.ear\deployments\SVCConsole\SVCConsole.war\
   WEB-INF

   C:\...\config\installedApps\DefaultNode\
   ICAConsole.ear\ICAConsole.war\WEB-INF

   C:\...\config\installedApps\DefaultNode\
   SVCConsole.ear\SVCConsole.war\WEB-INF
4. Stop and then restart the following applications. The following services are located in **Start ▸ Settings ▸ Control Panel ▸ Administrative Tools ▸ Component Services**.
   - IBM CIM Object Manager
   - IBM WebSphere Application Server V5 - SVC

   To stop and then restart the services, right-click on the application and select **Stop**, then **Start**.

   Note: If the stop command times-out in the IBM WebSphere application, you can restart the master console because this restarts the application, as well.
5. Ensure that both applications are running again. Launch the SAN Volume Controller Console and logon.

Verifying the installation

This task verifies that the IBM TotalStorage Support for Microsoft Volume Shadow Copy Service is correctly installed on the Windows server.

Perform the following steps to verify the installation.

1. On the Windows server on the **Start** menu, select **All Programs ▸ Administrative Tools ▸ Services**. The **Services** window opens.
2. Ensure that the service named “IBM TotalStorage Support for Microsoft Volume Shadow Copy Service” appears and that the Status is “Started” and the Startup Type is “Automatic.”

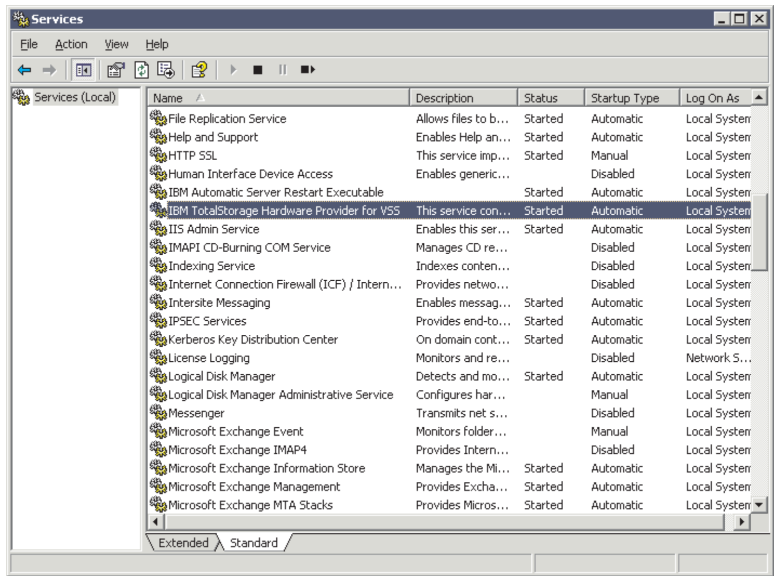

3. Open a command prompt window and type the following command:

   ```
   vssadmin list providers
   ```

   Ensure that the service named “IBM TotalStorage Support for Microsoft Volume Shadow Copy Service” is listed as a provider.

If you are able to perform all of these verification tasks successfully, then the IBM TotalStorage hardware provider was successfully installed on the Windows server.

Creating the free and reserved pools of volumes

The IBM TotalStorage hardware provider maintains a “free pool” of volumes and a “reserved pool” of volumes. Because these objects do not exist in the SAN Volume Controller, the free and reserved pools are implemented as virtual host systems. You must define these two hosts on the SAN Volume Controller.

When a shadow copy is made, the IBM TotalStorage hardware provider selects a volume in the free pool, assigns it to the reserved pool, and then removes it from the free pool. This protects the volume from being overwritten by other Volume Shadow Copy Service users.

To successfully perform a Volume Shadow Copy Service operation, there must be enough VDisks mapped to the free pool. The VDisks must be the same size as the source VDisks.

Use the SAN Volume Controller master console or the SAN Volume Controller command-line interface to perform the following steps.

1. Create a host for the ″free pool″ of virtual disks.
   - You can use the default name VSS_FREE or specify a different name.
   - Associate the host with the worldwide port name 5000000000000000 (15 zeroes).
2. Create a virtual host for the ″reserved pool″ of volumes.
   - You can use the default name VSS_RESERVED or specify a different name.
   - Associate the host with the worldwide port name 5000000000000001 (14 zeroes).
3. Map the logical units (virtual disks, or VDisks) to the free pool of volumes.
   - If you already have VDisks created for the free pool of volumes, you must assign those VDisks to the free pool.
   - Any VDisk can be used.
   - The VDisks must not be mapped to any other hosts.
4. To add the VDisks to the free pool, create VDisk to host mappings between the VDisks selected in step 3 and the VSS_FREE host. Alternatively, you can use the ibmvcfg add command to add vdisks to the free pool.
5. Verify that the VDisks have been mapped. Use one of the following methods:
   - The SAN Volume Controller master console Viewing Virtual Disk-to-Host Mappings panel.
   - The SAN Volume Controller command-line interface `svcinfo lsvdiskhostmap` command.

   Verify that the VDisks are listed with the worldwide port names that they are assigned to.

If you do not use the default WWPNs 5000000000000000 and 5000000000000001, then you must configure the IBM TotalStorage hardware provider with the WWPNs.

Configuration commands

After installation, you can change or correct the parameters that you defined when you installed the IBM TotalStorage hardware provider. To change the parameters, use the utility ibmvcfg.exe.

Table 53 shows the commands.

Table 53. Configuration commands

Command	Description	Example
ibmvcfg showcfg	Lists the current settings.	`ibmvcfg showcfg`
ibmvcfg set username <username>	Sets the user name to access the SAN Volume Controller master console.	`ibmvcfg set username johnny`
ibmvcfg set password <password>	Sets the password of the user name that will access the master console.	`ibmvcfg set password mypassword`

Table 53. Configuration commands (continued)

Command	Description	Example
ibmvcfg set targetSVC <ipaddress>	Specifies the IP address of the SAN Volume Controller on which the VDisks are located when VDisks are moved to and from the free pool with the ibmvcfg add and ibmvcfg rem commands. The IP address is overridden if you use the -s flag with the ibmvcfg add and ibmvcfg rem commands.	`set targetSVC 64.157.185.191`
set backgroundCopy	Sets the background copy rate for FlashCopy.	`set backgroundCopy 80`
ibmvcfg set trustpassword <trustpassword>	Sets the password for the truststore file. The default value is `ibmstore`.	`ibmvcfg set trustpassword ibmstore`
ibmvcfg set truststore <path>	Specifies the truststore file location.	`ibmvcfg set truststore c:\truststore`
ibmvcfg set usingSSL	Specifies whether to use Secure Sockets Layer protocol to connect to the master console.	`ibmvcfg set usingSSL yes`
ibmvcfg set cimomPort <portnum>	Specifies the master console port number. The default value is `5999`.	`ibmvcfg set cimomPort 5999`
ibmvcfg set cimomHost <server name>	Sets the name of the server where the master console is installed.	`ibmvcfg set cimomHost cimomserver`
ibmvcfg set namespace <namespace>	Specifies the namespace value that master console is using. The default value is `\root\ibm`.	`ibmvcfg set namespace \root\ibm`
ibmvcfg set vssFreeInitiator <WWPN>	Specifies the WWPN of the host. The default value is `5000000000000000`. Modify this value only if there is a host already in your environment with a WWPN of 5000000000000000.	`ibmvcfg set vssFreeInitiator 5000000000000000`
ibmvcfg set vssReservedInitiator <WWPN>	Specifies the WWPN of the host. The default value is `5000000000000001`. Modify this value only if there is a host already in your environment with a WWPN of 5000000000000001.	`ibmvcfg set vssFreeInitiator 5000000000000001`

Adding and removing VDisks

To perform the pool management tasks of adding, removing, and listing virtual disks (VDisks), you can use the SAN Volume Controller master console, the SAN Volume Controller command-line interface, or the utility ibmvcfg.exe.

The Microsoft Volume Shadow Copy service maintains a “free pool” of volumes and a “reserved pool” of volumes. These pools are implemented as virtual host systems on the SAN Volume Controller.

Table 54 shows the ibmvcfg.exe commands for adding or removing volumes from the free pool of volumes.

Table 54. Pool management commands

Command	Description	Example
ibmvcfg listvols	Lists all VDisks, including information about size, location, and VDisk to host mappings.	`ibmvcfg listvols`
ibmvcfg listvols all	Lists all VDisks, including information about size, location, and VDisk to host mappings.	`ibmvcfg listvols all`
ibmvcfg listvols free	Lists the volumes that are currently in the free pool.	`ibmvcfg listvols free`
ibmvcfg listvols unassigned	Lists the volumes that are currently not mapped to any hosts.	`ibmvcfg listvols unassigned`

ibmvcfg add -s <ipaddress>	Adds one or more volumes to the free pool of volumes. Use the -s parameter to specify the IP address of the SAN Volume Controller where the VDisks are located. The -s parameter overrides the default IP address that is set with the ibmvcfg set targetSVC command.	`ibmvcfg add vdisk12`
		`ibmvcfg add 600507`
		`68018700035000000`
		`0000000BA`
		`-s 66.150.210.141`

ibmvcfg rem -s <ipaddress>	Removes one or more volumes from the free pool of volumes. Use the -s parameter to specify the IP address of the SAN Volume Controller where the VDisks are located. The -s parameter overrides the default IP address that is set with the ibmvcfg set targetSVC command.	`ibmvcfg rem vdisk12`
		`ibmvcfg rem 600507`
		`68018700035000000`
		`0000000BA`
		`-s 66.150.210.141`

Error codes

The IBM TotalStorage hardware provider for the Microsoft Volume Shadow Copy service logs errors in the Windows Event Monitor and in private log files.

Error messages can be viewed in the following locations on the Windows server where the IBM TotalStorage hardware provider is installed:

- The Windows Event Viewer in Application Events. Check this log first.
- The log file ibmVSS.log, which is located in the directory where the IBM TotalStorage hardware provider is installed.

Table 55 lists the errors that are reported by the IBM TotalStorage hardware provider.

Table 55. Error messages for the IBM TotalStorage hardware provider for the Microsoft Volume Shadow Copy services

Code	Message	Symbolic name
1000	JVM Creation failed.	ERR_JVM
1001	Class not found: %1.	ERR_CLASS_NOT_FOUND
1002	Some required parameters are missing.	ERR_MISSING_PARAMS
1003	Method not found: %1.	ERR_METHOD_NOT_FOUND
1004	A missing parameter is required. Use the configuration utility to set this parameter: %1.	ERR_REQUIRED_PARAM
1600	The recovery file could not be created.	ERR_RECOVERY_FILE_ CREATION_FAILED
1700	ibmGetLunInfo failed in AreLunsSupported.	ERR_ARELUNSSUPPORTED_ IBMGETLUNINFO
1800	ibmGetLunInfo failed in FillLunInfo.	ERR_FILLLUNINFO_IBMGETLUNINFO
1900	Failed to delete the following temp files: %1	ERR_GET_TGT_CLEANUP
2500	Error initializing log.	ERR_LOG_SETUP
2501	Unable to search for incomplete Shadow Copies. Windows Error: %1.	ERR_CLEANUP_LOCATE
2502	Unable to read incomplete Shadow Copy Set information from file: %1.	ERR_CLEANUP_READ
2503	Unable to cleanup snapshot stored in file: %1.	ERR_CLEANUP_SNAPSHOT
2504	Cleanup call failed with error: %1.	ERR_CLEANUP_FAILED
2505	Unable to open file: %1.	ERR_CLEANUP_OPEN
2506	Unable to create file: %1.	ERR_CLEANUP_CREATE
2507	HBA: Error loading hba library: %1.	ERR_HBAAPI_LOAD
3000	An exception occurred. Check the ESSService log.	ERR_ESSSERVICE_EXCEPTION
3001	Unable to initialize logging.	ERR_ESSSERVICE_LOGGING
3002	Unable to connect to the CIM agent. Check your configuration.	ERR_ESSSERVICE_CONNECT

Table 55. Error messages for the IBM TotalStorage hardware provider for the Microsoft Volume Shadow Copy services (continued)

Code	Message	Symbolic name
3003	Unable to get the Storage Configuration Service. Check your configuration.	ERR_ESSSERVICE_SCS
3004	An internal error occurred with the following information: %1.	ERR_ESSSERVICE_INTERNAL
3005	Unable to find the VSS_FREE controller.	ERR_ESSSERVICE_FREE_CONTROLLER
3006	Unable to find the VSS_RESERVED controller. Check your configuration.	ERR_ESSSERVICE_RESERVED_ CONTROLLER
3007	Unable to find suitable targets for all volumes.	ERR_ESSSERVICE_INSUFFICIENT_ TARGETS
3008	The assign operation failed. Check the CIM agent log for details.	ERR_ESSSERVICE_ASSIGN_FAILED
3009	The withdraw FlashCopy operation failed. Check the CIM agent log for details.	ERR_ESSSERVICE_WITHDRAW_ FAILED

Uninstalling the IBM TotalStorage hardware provider

Use Windows to uninstall the IBM TotalStorage hardware provider from the Windows server.

1. Log on to the Windows server as the local administrator.
2. On the **Start** menu, click **Control Panel**. The Control Panel window opens.
3. Double-click **Add or Remove Programs**. The Add or Remove Programs window opens.
4. Select **IBM TotalStorage Support for Microsoft Volume Shadow Copy Service**, and then click **Remove**.
5. When you are prompted to verify that you want to completely remove the program and all of its components, click **Yes**.
6. When the Finish window opens, click **Finish**. The removal is complete.

Reference

This topic provides reference information for the SAN Volume Controller.

Installing or upgrading the IBM TotalStorage SAN Volume Controller Console for Windows

This part provides an overview of the installation process and instructions for installing or upgrading and configuring the IBM TotalStorage SAN Volume Controller Console on a Windows 2000 Server operating system.

Note: Installing the SAN Volume Controller Console on your host system is optional. The SAN Volume Controller Console comes preinstalled on the master console.

Related tasks

"Installing or upgrading the SAN Volume Controller Console in graphical mode" on page 351
If you choose to install or upgrade the IBM TotalStorage SAN Volume Controller Console in unattended mode skip this section. You must satisfy all prerequisites before starting the installation.

"Installing or upgrading the SAN Volume Controller Console in unattended (silent) mode" on page 356
The unattended (silent) mode install or upgrade option enables you to run the installation or upgrade installation unattended.

"Verifying the Windows services associated with the SAN Volume Controller Console" on page 361
To verify that the Windows services associated with your IBM TotalStorage SAN Volume Controller Console are correctly installed and started, complete these steps.

"Post installation tasks" on page 361
To get started using the SAN Volume Controller Console, complete these steps.

"Removing the SAN Volume Controller Console" on page 364
You can remove the IBM TotalStorage SAN Volume Controller Console from your Windows system.

Related information

"Installation overview for the SAN Volume Controller Console" on page 348
Before you install or upgrade and configure the IBM TotalStorage SAN Volume Controller Console on a Windows 2000 Server operating system, consider these prerequisites.

"SAN Volume Controller Console hardware installation requirements" on page 350
Before starting the installation, ensure that your system satisfies the following hardware installation prerequisites for installing the IBM TotalStorage SAN Volume Controller Console on a Windows 2000 Server or Windows Server 2003 operating system.

"SAN Volume Controller Console workstation space requirements" on page 350
Before starting the installation, ensure that your system satisfies the following workstation space prerequisites for installing the IBM TotalStorage SAN Volume Controller Console on a Windows 2000 Server operating system.

"SAN Volume Controller Console software installation requirements" on page 350

Before starting the installation, ensure that your system satisfies the following software installation prerequisites for installing the IBM TotalStorage SAN Volume Controller Console on a Windows 2000 Server operating system.

Installation overview for the SAN Volume Controller Console

Before you install or upgrade and configure the IBM TotalStorage SAN Volume Controller Console on a Windows 2000 Server operating system, consider these prerequisites.

You should have some knowledge of how to administer a Windows 2000 Server operating system. You should also become familiar with the command that you use during the installation or upgrade installation of the SAN Volume Controller Console.

You must be aware of the following list of installation and configuration tasks *before* you install or upgrade the SAN Volume Controller Console:

1. Check the hardware and software requirements.

   Because the software is preinstalled, it is not normally necessary to replace the software on a node. However, if the software is lost for some reason, for example if the hard disk drive in the node fails, it is possible to copy all the software from another node connected to the same fibre-channel fabric. This process is known as node rescue.

   If the SAN Volume Controller detects software errors an error code is generated. The additional data logged with the error will indicate the source of the software error. The additional data might look like this:

   Assert File /build/lodestone/030129_nd/src/user/vg/vgagentevt.c Line 1234

   To view the additional data you will need to access the SAN Volume Controller web pages and select the Analyze error log option for the software error that you are investigating. Report the error code and the additional data to your IBM Product Support Center.

   If this problem is known for your version of software, the customer will be advised to upgrade to the latest software level. If the problem is not known to the Support Center you might be asked to provide additional information for this error. In most cases a dump will automatically be taken when the software error is detected.

   If requested to do so by your Support Center, you can use the SAN Volume Controller Console application on the master console to list and save dump data. If more than one dump file exists, select the dump file with a time stamp closest to the time stamp on the software error report and save this file for use by the Support Center.
2. If the SSH client software called PuTTY is not yet installed on your system, you must install the SSH client software. You can get more information about PuTTY from the PuTTY Web site home page:

   http://www.chiark.greenend.org.uk/~sgtatham/putty/

   and download PuTTY from the following Web site download page:

   http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html

   Note: For your convenience the PuTTY installation program (putty-o.53b-installer.exe) is on the SAN Volume Controller Console installation CD-ROM in the SSHClient/PuTTY directory.

3. Install or upgrade the SAN Volume Controller Console either in graphical mode with the help of an installation wizard or in unattended mode (also known as silent mode), which involves customizing a response file and issuing a command.

   Cluster software versions comprise a number of software components that are delivered as a single package. The size of the software update package depends on the number of components that are being replaced by that upgrade package. The software installation procedure involves copying the new software version to the cluster and then starting an automatic installation process. This installation process might take up to an hour to complete and during the process each of the nodes is restarted in turn. Once all the nodes in the cluster have been successfully restarted with the new software the new software version is automatically committed. While each node is being restarted there might be some degradation in the maximum input/output rate that can be sustained by the cluster.
4. Installation or upgrade installation operation. Because of the operational limitations that occur during the software upgrade process, the software installation is a customer task. The installation or upgrade installation operation can normally be performed concurrently with normal user I/O operations. If any restrictions apply to the operations that can be performed during the upgrade, then these restrictions will be documented on the SAN Volume Controller Web site from where the upgrade package was obtained. During the upgrade operation, only the SAN Volume Controller commands will be operational from the time the installation process starts to the time that the new software is committed or until the process has been backed-out. For a complete list of SAN Volume Controller commands, enter the following command:`svcinfo -?`
5. Verify the following Windows services associated with the SAN Volume Controller Console are installed and started:
   - Service Location Protocol
   - IBM CIM Object Manager - SVC
   - IBM Websphere Application Server V5 - SVC
6. Get started using the SAN Volume Controller Console. Use a Web browser to access the SAN Volume Controller Console. You will identify the clusters to be managed to the SAN Volume Controller Console as well as complete the creation (initialization) of the SAN Volume Controller clusters.

   To allow nodes to operate as a cluster, you must run all nodes at the same version of software. This rule is enforced by the cluster software itself. When you attempt to add a node to a cluster its software version is examined, and if it is not running the same version of the software as the other nodes in the cluster, the software revisions are automatically copied from one of the other nodes in the cluster before the add operation is completed. If for some reason it is not possible to update the software on the node that you are adding, the operation fails and the cluster logs an error to explain the cause of the failure.
7. Remove the SAN Volume Controller Console. You only need to perform this optional task if you get errors during installation verification.

   Related information

   "Installing or upgrading the IBM TotalStorage SAN Volume Controller Console for Windows" on page 347

SAN Volume Controller Console hardware installation requirements

Before starting the installation, ensure that your system satisfies the following hardware installation prerequisites for installing the IBM TotalStorage SAN Volume Controller Console on a Windows 2000 Server or Windows Server 2003 operating system.

Hardware prerequisites

The following hardware is required:

- Any Intel®-based PC running Windows 2000 Server SP 3 or Windows Server 2003
- Intel Pentium® processor at 1 GHz, or faster
- Support for a communications adapter
- CD-ROM drive
- Minimum 2 GB RAM recommended

Related information

"Installing or upgrading the IBM TotalStorage SAN Volume Controller Console for Windows" on page 347

SAN Volume Controller Console workstation space requirements

Before starting the installation, ensure that your system satisfies the following workstation space prerequisites for installing the IBM TotalStorage SAN Volume Controller Console on a Windows 2000 Server operating system.

Workstation space

The following space on your workstation is required:

- 350 MB of disk space

  Note: You might need to increase the total available disk space on your hard drives if the IBM TotalStorage SAN Volume Controller Console and other associated products are split between more than one logical drive. Also, the IBM TotalStorage SAN Volume Controller Console might require additional memory to operate if you configure it to manage many devices or devices with large configurations.
- Up to 65 MB of temporary disk space for installation purposes

Related information

"Installing or upgrading the IBM TotalStorage SAN Volume Controller Console for Windows" on page 347

SAN Volume Controller Console software installation requirements

Before starting the installation, ensure that your system satisfies the following software installation prerequisites for installing the IBM TotalStorage SAN Volume Controller Console on a Windows 2000 Server operating system.

Software

The following software is required:

- Operating systems:
  - Windows 2000 Server SP3

- If the SSH client software called PuTTY is not yet installed on your system, you must install the SSH client software. You can get more information about PuTTY from the PuTTY Web site home page:

  http://www.chiark.greenend.org.uk/~sgtatham/putty/

  and download PuTTY from the following Web site download page:

  http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html

  For your convenience the PuTTY installation program (putty-o.53b-installer.exe) is on the SAN Volume Controller Console installation CD-ROM in the SSHClient/PuTTY directory.
- IBM TotalStorage SAN Volume Controller Console. This is on the IBM TotalStorage SAN Volume Controller Console CD.
- Transmission Control Protocol/Internet Protocol (TCP/IP)
- Adobe Acrobat Reader version 4.0 or later (optional)

  You need the Adobe Acrobat Reader to read License Agreement and product information from the SAN Volume Controller Console LaunchPad. You can download the Adobe Acrobat Reader from the following Web site:
  - http://www.adobe.com/support/downloads/main.html

Related information

"Installing or upgrading the IBM TotalStorage SAN Volume Controller Console for Windows" on page 347

Installing or upgrading the SAN Volume Controller Console in graphical mode

If you choose to install or upgrade the IBM TotalStorage SAN Volume Controller Console in unattended mode skip this section. You must satisfy all prerequisites before starting the installation.

1. Log onto your system as a local system administrator.
2. Insert the IBM TotalStorage SAN Volume Controller Console CD into the CD drive.

   The IBM TotalStorage SAN Volume Controller Console program should start within 15 - 30 seconds if you have **autorun** mode set on your system. If the LaunchPad panel does not open, perform either one of the following steps:
   a. Use a Command Prompt to change to the W2K directory on the CD. Type:
      ```
      LaunchPad
      ```
   b. Using Windows Explorer, (**Start->Programs->Accessories->Windows Explorer**), go to the W2K directory located on the CD drive. Then double-click on the **LaunchPad.bat** file.

      Note: If you are viewing the folder using the Windows Explorer with the option selected to *Hide file extensions for known file types*, find the LaunchPad file with the file type of MS-DOS Batch File.
3. The following options are displayed when the LaunchPad panel opens:

   SVC Console overview
   Offers information about the IBM TotalStorage SAN Volume Controller Console.

Readme file
Offers any last minute product information that did not make it into sections concerning the installation of the IBM TotalStorage SAN Volume Controller Console.

Configuration guide
Contains instructions about how to install the IBM TotalStorage SAN Volume Controller Console (a softcopy of this document).

License agreement
Offers information about the license for the IBM TotalStorage SAN Volume Controller Console.

SAN Volume Controller Web site
Offers information from the product Web site.

Installation wizard
Starts the IBM TotalStorage SAN Volume Controller Console installation program.

Post installation tasks
Details information about validating the installation, accessing the SAN Volume Controller Console URL and adding the SAN Volume Controller Console cluster to the SAN Volume Controller Console management facility.

Exit Exits the IBM TotalStorage SAN Volume Controller Console LaunchPad program.

4. Click **Readme file** from the LaunchPad panel or from the **README.txt** file located in the doc or W2K directory on the IBM TotalStorage SAN Volume Controller Console CD to check for information that might supersede the information in this guide.
5. Click **Installation wizard** from the LaunchPad panel to start the installation.

   Note: The LaunchPad panel remains open behind the installation wizard so that you can access product information during the installation process. Click **Exit** if you want to close the LaunchPad.
6. There might be a slight delay while the software loads on your system. After the software loads a DOS prompt window opens to display the following message:

```
Initializing InstallShield Wizard...
Preparing Java <tm> Virtual Machine ..................
.....................................................
.....................................................
```

7. The Welcome panel opens suggesting what documentation you should review prior to installation. Click **Next** to continue, or click **Cancel** to exit the installation.
8. The License Agreement panel opens. Read the license agreement information. Select **I accept the terms of the license agreement**, then click **Next** to accept the license agreement. Otherwise, keep the selection **I do not accept the terms of the license agreement** (it is the default) and click **Cancel** to exit the installation.
9. The installation wizard verifies that your workstation meets the installation requirements.
   a. If you have a Service Location Protocol (SLP) service that is different from the SLP that the IBM TotalStorage SAN Volume Controller Console

requires, the installation wizard displays an error and asks you to stop the installation and remove this SLP service from the system.

b. The installation wizard checks if the PuTTY SSH client is installed on your workstation.

c. The installation wizard determines whether this is a new installation, reinstallation or upgrade installation of the SAN Volume Controller Console. If the installation wizard determines that the SAN Volume Controller Console was previously installed on the system, it does a comparison of the current version, release, modification, and fix code level with that of the code currently installed on the system. If the level is the same, this is a reinstallation. If the new code has a higher level, it is an upgrade. If the new code level is lower than the level on the system, the installation is invalid. In the case of reinstallation or upgrade installation, the installation wizard performs the following actions:

   1) Checks if the Service Location Protocol (SLP), the IBM CIM Object Manager (CIMOM) service, and WebSphere Application Server V5 - SVC are started. If any of these services are started, the program asks if you want to continue the installation process by clicking **Next**. If you want to exit the installation program click **Cancel**. If you choose to continue, you must stop all the applications that use these services.

   2) Presents a panel with the check-box option to Preserve Configuration. If you chose to preserve the existent Configuration, the installation program skips the next steps and goes directly to the Installation Confirmation panel discussed below.

10. The Destination Directory panel opens. Select one of the following options:

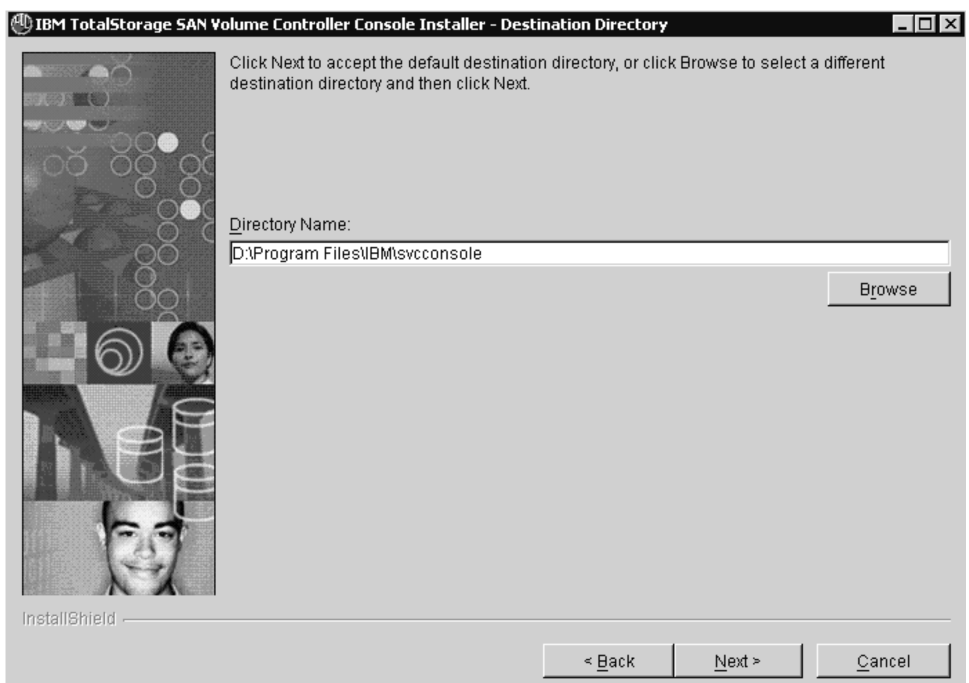

Figure 50. Destination Directory panel

a. Click **Next** to accept the default directory.

b. Click **Browse** to select a different directory for installation and then click **Next** to continue the installation process.

c. Click **Cancel** to exit the installation process.

Note:

a. The directory name, including the drive letter, must be a maximum of 44 characters.

b. If the program detects insufficient space for the IBM TotalStorage SAN Volume Controller Console installation in the chosen destination, an error message is displayed. You can free some space on the destination drive and then click **Next** or you can stop the installation program by clicking **Cancel**. You can also go back by clicking **Back**, and choosing another destination directory for the product.

11. The PuTTY configuration panel opens when the product space check is completed.

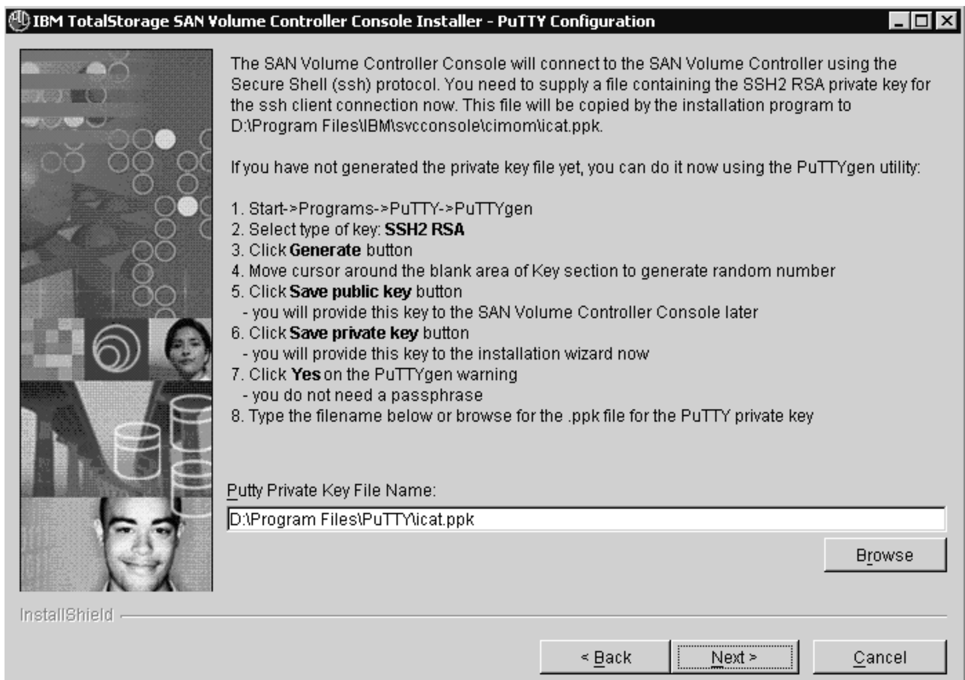

Figure 51. PuTTY Configuration panel

12. Enter the name and location on your system of your PuTTY SSH2 RSA private key file or click **Browse** to select the key file. If you have not prepared a PuTTY private key file yet, the steps on this panel tell you how to generate the PuTTY private and public key. Click **Next** to continue.

13. The Updating Embedded WAS Ports panel is displayed.

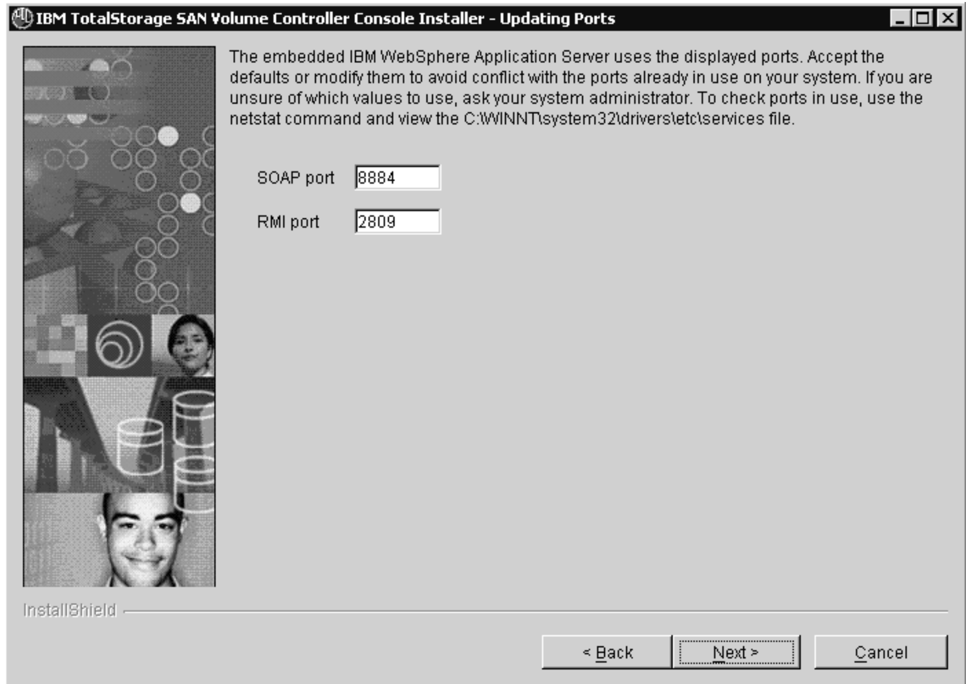

Figure 52. Updating Embedded WAS Ports panel

14. Update the default ports assignments by typing unique port numbers for the products that have been registered on your system. To check ports in use, use the **netstat -a** command and view the C:\WINNT\system32\drivers\etc\services file. Click **Next** to continue.
15. The Updating CIMOM ports panel is displayed. Update the default port assignments and the default communication protocol by typing the unique port numbers and choosing the desired communication protocol for the products that have been registered on your system. To check ports in use, use the **nestat -a** command and view the C:\WINNT\system32\drivers\etc\services file. Click **Next** to continue.
16. The Installation Confirmation panel opens. Click **Install** to confirm the installation location and file size and to start the final installation, reinstallation or upgrade installation. Click **Cancel** to exit the installation wizard or click **Back** to go to the previous panel.
17. The Installation Progress panel opens indicating how much of the installation has been completed. Installation usually takes 3 - 10 minutes depending on the configuration of your workstation.

    Note: If you click **Cancel** a popup panel opens asking you to confirm the cancellation of the installation wizard: `Cancel the current operation?`. You may confirm the cancellation by clicking **Yes** or continue the installation by selecting **No**. If you confirm the cancellation, the information you entered or selected in the previous panel is not saved. You must start the installation again from the beginning.

    After the completion of the successful installation of the IBM TotalStorage SAN Volume Controller Console, the installer attempts to start the following services:

    - Service Location Protocol
    - The IBM CIM Object Manager

- The IBM WebSphere Application Server V5 - SVC

18. When the Installation Progress panel closes, the Finish panel opens. Before proceeding, you might want to review the log file for any possible error messages. The log file is located in xxx\logs\install.log, where *xxx* is the destination directory where the IBM TotalStorage SAN Volume Controller Console for Windows was installed. The install.log contains a trace of the installation actions.

    Note: At the bottom of the Finish panel is a checkbox labeled **View post installation tasks**. If you check this box and then click **Finish**, the wizard will exit and the post installation tasks text file is displayed. The LaunchPad panel Post installation tasks link also displays this same text file. You can avoid the display of the text file by unchecking the **View post installation tasks** box before you click the **Finish** button.

19. Click **Finish** to exit the installation wizard.

    Note: Ordinarily, you do not need to restart your system during or after the installation of the IBM TotalStorage SAN Volume Controller Console. However, the installation wizard might determine that a restart is necessary. Restart your system if required. After you restart the system, the installation wizard continues with the installation.

20. If you have not yet reviewed the post-installation tasks from the installation Finish panel, review the post installation tasks from the LaunchPad program.
    a. Click **Post installation tasks** on the LaunchPad panel which opens the same file available from the Installation Finish panel.
    b. Continue with the post-installation tasks for the SAN Volume Controller by following the instructions in this file.
21. Exit the LaunchPad program by clicking **Exit** on the LaunchPad panel.
22. Verify that the Windows services associated with your SAN Volume Controller Console are correctly installed and started.

Related information

"Installing or upgrading the IBM TotalStorage SAN Volume Controller Console for Windows" on page 347

Installing or upgrading the SAN Volume Controller Console in unattended (silent) mode

The unattended (silent) mode install or upgrade option enables you to run the installation or upgrade installation unattended.

Use this method of installation to customize a response file and issue a command from a command prompt window. The response file is a template on the IBM TotalStorage SAN Volume Controller CD. You can also create a standard response file to ensure that the product is installed consistently on multiple systems. You must satisfy all prerequisites before starting the installation.

The installation wizard determines whether this is a reinstallation or upgrade of the SAN Volume Controller. If the installation wizard determines that the SAN Volume Controller was previously installed on the system, it does a comparison of the current version, release, modification, and fix code level with that of the code currently installed on the system. If the level is the same, this is a reinstallation. If the new code has a higher level, it is an upgrade. If the new code level is lower than the level on the system, the installation is invalid.

1. Log on to the system as a local system administrator.

2. Insert the IBM TotalStorage SAN Volume Controller Console CD into the CD drive.
3. If you have autorun mode set on your system, the IBM TotalStorage SAN Volume Controller Console program will start within 15-30 seconds. Click **Exit** from the LaunchPad.
4. Locate the response file named, responsefile, on your IBM TotalStorage SAN Volume Controller Console CD in the W2K directory.
5. Using Windows Explore or a command prompt, copy the response file to your hard drive.
6. The SAN Volume Controller Console will connect to the SAN Volume Controller using the Secure Shell (SSH) protocol. You need to supply a file containing the SSH2 RSA private key for the SSH client connection. This file will be copied by the installation program to <inst_dir>\cimom\icat.ppk, for example C:\ProgramFiles\IBM\svcconsole\cimom\icat.ppk. If you have not generated the private key file before, you can do it now using the PuTTYgen utility. To generate the private key using the PuTTYgen utility, perform the following steps:
   a. Click **Start -> Programs -> PuTTY -> PuTTYgen**.
   b. Select the type of key: **SSH RSA**.
   c. Click **Generate**.
   d. Move the cursor around the blank area of the Key section to generate a random number.
   e. Click **Save public key**. You will provide this key to the SAN Volume Controller Console later.
   f. Click **Save private key**. You will provide this key to the installation wizard using the option below in the response file.
   g. Click **Yes** on the PuTTYgen warning. You do not need a passphrase.
   h. Ensure that you set the value of the <-W puttyConfiguration.puttyPrivateKeyFile> option in the response file, to the name of the file containing the PuTTY private key.
7. Using a text editor modify the default options in the response file with the values you want to supply to the installation program:
   a. Remove the # character from the beginning of a line if you do not want to use the default value. Change the default value to the value that you want for that option. You *must* enclose all values in double quotation marks (″″).
   b. Depending on whether you are doing a new installation, reinstallation or upgrade, certain response file lines must be active as detailed below. If a response file line is active but inappropriate to the mode (new, reinstall or upgrade) it will be ignored.

   New Installation:
   - The <*-P product.installLocation*> option defines the default directory where the product is to be installed. To specify a destination directory other than the default, remove the # character from the corresponding line and replace the default directory with the desired directory.
   - The <*-G checkPrerequisite*> option checks the prerequisites. If you want to disable this option, remove the # character from the corresponding line and change the value of the option to `no`.
   - Change the default ports values for the embedded WebSphere Application Server - V5 SVC using the update ports variables options. If you want to change a specific port used for a particular WebSphere service, remove the

character from the beginning of the line containing the value of the option and set it to the value you desire. The following are the embedded WebSphere ports options:

– <-W ports.portSOAP=″8884″>
– <-W ports.portRMI=″2809″>
– <-W ports.portHTTP=″9080″>
– <-W ports.portHTTPS=″9443″>

- Change the default ports values and default server communication type for the IBM CIM Object Manager server using variables options below. If you want to change a specific port or the default server communication type, remove the # character from the beginning of the line containing the option's value and set it to the value you desire. The following are the CIM Object Manager server options:
  – <-W cimObjectManagerPorts.port=″5999″>
  – <-W cimObjectManagerPorts.indicationPort=″5990″>
  – <-W cimObjectManagerPorts.serverCommunication=″HTTPS″>
- The <-W puttyConfiguration.puttyPrivateKeyFile> options specifies the name and location of the PuTTY private key file that the SAN Volume Controller Console software should use to connect to the SAN Volume Controller clusters. Remove the # character from the corresponding line and add the fully qualified location of the PuTTY private key file. Save the responsefile *without* a file extension such as .txt.

Reinstallation or Upgrade:

- The *<-G startUpgrade>* option must be enabled to permit the new SAN Volume Controller Console to be reinstalled (having the same version) or upgraded (installed at a higher version). Enable this option be removing the # character from the corresponding line and changing the value of the option to yes.
- The *<-G stopProcessesResponse>* option tells the install program whether or not to automatically stop SLP, IBM CIM Object Manager (CIMOM), and WebShpere Application Server - V5 SAN Volume Controller services when reinstalling or upgrading the product. By default, this option is set to `no`. If you do not change this default value, the reinstallation or upgrade stops when these services are running If you want to automatically stop the SLP and IBM CIM Object Manager (CIMOM), remove the # character from the corresponding line and change its value to yes.
- The *<-G saveConfiguration>* option specifies whether or not to save the configuration files when reinstalling or upgrading the product. If you do not want to save the configuration files when reinstalling or upgrading, remove the # character from the corresponding line and change the value of the option to `no`. If you do not chose to save the configuration, you will have to make the following active or accept the default values.
  – Change the default ports values for the embedded WebSphere Application Server - V5 SAN Volume Controller using the update ports variables options. If you want to change a specific port used for a particular WebSphere service, remove the # character from the beginning of the line containing the value of the option and set it to the value you desire. The following are the embedded WebSphere ports options:
    - <-W ports.portSOAP=″8884″>
    - <-W ports.portRMI=″2809″>
    - <-W ports.portHTTP=″9080″>

   - <-W ports.portHTTPS=″9443″>
 – Change the default ports values and the default server communication type for the CIM Object Manager server using variables options bellow. If you want to change a specific port or the default server communication type, remove the # character from the beginning of the line containing the option's value and set it to the value you desire. The following are the CIM Object Manager server options:
   - <-W cimObjectManagerPorts.port=″5999″>
   - <-W cimObjectManagerPorts.indicationPort=″5990″>
   - <-W cimObjectManagerPorts.serverCommunication=″HTTPS″>
 – The <*-W puttyConfiguration.puttyPrivateKeyFile*> options specifies the name and location of the PuTTY private key file that the SAN Volume Controller Console software should use to connect to the SAN Volume Controller clusters. Remove the # character from the corresponding line and add the fully qualified location of the PuTTY private key file. Save the response file without a file extension such as .txt.

8. From a command prompt window, type the following command:

   a. `<CD drive path>\W2K\install -options <response file path>\responsefile`

   where <*CD drive path*> is the path of your CD drive. <*response file path*> is the path of the responsefile file that you copied in step 5 on page 357 and customized in step 7 on page 357.

   Note: The directory name, including the drive letter, must be a maximum of 44 characters.

9. During the installation, dotted lines are displayed across the screen. When the installation program ends, control returns to the Command Prompt.
10. Check for installation errors in the install.log file. This file is initially created in the system temporary file under the subdirectory named, cimagent. After all the prerequisites checks have been performed, the log file is copied to the <dest-path>\logs directory. The following is an example of an install.log file:

```
(May 15, 2003 9:36:06 AM), This summary log is an overview of the
 sequence of the installation of the IBM TotalStorage SAN Volume
Controller Console 1.0.0.12
(May 15, 2003 9:38:22 AM), IBM TotalStorage SAN Volume Controller
 Console installation process started with the following install
parameters:
Target Directory: C:\Program Files\IBM\svcconsole
SOAP port: 8884
RMI port: 2809
(May 15, 2003 9:38:28 AM), Copying Service Location Protocol Files ...
(May 15, 2003 9:38:29 AM), Service Location Protocol successfully installed
(May 15, 2003 9:38:29 AM), Copying CIM Object Manager Files ...
(May 15, 2003 9:39:26 AM), The PuTTY private key successfully copied
 into file C:\Program Files\IBM\svcconsole\cimom\icat.ppk
(May 15, 2003 9:39:51 AM), The file setupCmdLine.bat successfully updated.
(May 15, 2003 9:39:51 AM), Compile MOF files started ...
(May 15, 2003 9:40:06 AM), MOF files successfully compiled.
(May 15, 2003 9:40:06 AM), Generate a certificate store started ...
(May 15, 2003 9:40:19 AM), Certificate store called truststore
successfully generated.
(May 15, 2003 9:40:20 AM), IBM CIM Object Manager successfully installed
(May 15, 2003 9:40:20 AM), Installing embedded version of IBM WebSphere
 Application Server ...
(May 15, 2003 9:41:42 AM), Websphere Application Server - SVC
successfully installed.
(May 15, 2003 9:43:20 AM), Copying SAN Volume Controller Console Ear Files...
(May 15, 2003 9:46:11 AM), The ICAConsole application successfully installed.
(May 15, 2003 9:47:24 AM), The SVCConsole application successfully installed.
(May 15, 2003 9:48:06 AM), The help application successfully installed.
(May 15, 2003 9:48:27 AM), The ""C:\Program Files\IBM\svcconsole\console\
embeddedWAS\bin\expressPorts\UpdateExpressMultiPorts.bat"  -soap  8884
 -boot  2809  -remove" command updated successfully embedded WAS ports
in configuration files.
(May 15, 2003 9:48:27 AM), Command to be executed : net start cimomsrv
(May 15, 2003 9:48:49 AM), Command to be executed : net start
"IBMWAS5Service - SVC"
(May 15, 2003 9:50:15 AM), The following services started successfully:
Service Location Protocol
IBM CIM Object Manager
IBM WebSphere Application Server V5 - SVC
(May 15, 2003 9:50:15 AM), INSTSUCC: The IBM TotalStorage SAN Volume
 Controller Console has been successfully installed.
```

11. Close the command prompt window by entering a command, for example **exit**.
12. After the completion of the successful installation of the IBM TotalStorage SAN Volume Controller Console, the installer attempts to start the following services:
    a. Service Location Protocol
    b. The IBM CIM Object Manager
    c. IBM WebSphere Application Server V5 - SVC
13. Continue with the post installation tasks for the IBM TotalStorage SAN Volume Controller Console using the instructions in the following section. You can also view the post installation tasks using the following option:
    a. From a Command Prompt, change directory into the W2K directory on the CD drive. Open the LaunchPad by typing:

       ```
       LaunchPad
       ```
    b. Click **Post installation tasks** on the LaunchPad window. Continue with the post installation tasks for the IBM TotalStorage SAN Volume Controller Console by following the instructions in this file.
14. Verify that the Windows services associated with your SAN Volume Controller Console are correctly installed and started.

Related information

"Installing or upgrading the IBM TotalStorage SAN Volume Controller Console for Windows" on page 347

Verifying the Windows services associated with the SAN Volume Controller Console

To verify that the Windows services associated with your IBM TotalStorage SAN Volume Controller Console are correctly installed and started, complete these steps.

Perform the following steps to verify your Service Location Protocol (SLP), IBM CIM Object Manager (CIMOM), and IBM WebSphere Application Server V5 - SVC services were correctly installed

1. Verify the installation of the Service Location Protocol (SLP).
   a. Verify that the Service Location Protocol is started. Select **Start -> Settings -> Control Panel**. Double-click the **Administrative Tools** icon. Double-click the **Services** icon.
   b. Find **Service Location Protocol** in the **Services** list. For this component, the **Status** column should be marked `Started`.
   c. If the Service Location Protocol is not started, right-click on **Service Location Protocol** and select **Start** from the pop-up menu. Wait for the **Status** column to change to `Started`.
   d. Do not close the Services window because you will also use it to verify the CIM Object Manager (CIMOM) service.
2. Verify the installation of the SAN Volume Controller Console.
   a. Find the **IBM CIM Object Manager - SVC** in the **Services** list. For this component, the **Status** column should be marked `Started`.
   b. If the IBM CIM Object Manager is not started, right click on **IBM CIM Object Manager - SVC** and select **Start** from the pop-up menu. Wait for the **Status** column to change to `Started`.
   c. Do not close the Services window because you will also use it to verify the IBM WebSphere Application Server V5 - SVC service.
3. Verify the installation of the IBM WebSphere Application Server V5 - SVC service.
   a. Find the **IBM WebSphere Application Server V5 - SVC** in the **Services** list. For this component, the **Status** column should be marked `Started`.
   b. If the **IBM WebSphere Application Server V5 - SVC** service is not started, right click on **IBM WebSphere Application Server V5 - SVC** and select **Start** from the pop-up menu. Wait for the **Status** column to change to `Started`.
   c. Close the Services window.
   d. Close the Administrative Tools window.

Related information

"Installing or upgrading the IBM TotalStorage SAN Volume Controller Console for Windows" on page 347

Post installation tasks

To get started using the SAN Volume Controller Console, complete these steps.

Once you have installed the IBM TotalStorage SAN Volume Controller and the services (IBM CIM Object Manager, IBM WebSphere Application Server V5 - SVC, Service Location Protocol) have started, you will use a browser to access the Web pages of the Console for purposes of administering the SAN Volume Controller as well as configuring SAN Volume Controller clusters.

Each time you wish to add a SAN Volume Controller cluster to the collection of clusters managed by the IBM TotalStorage SAN Volume Controller, you must store the PuTTY SSH client public key which is located on the SAN Volume Controller system on the SAN Volume Controller cluster.

Attention: If you do not store the SSH public key on the SAN Volume Controller cluster, the SAN Volume Controller Console software cannot connect to the cluster.

When you installed the SAN Volume Controller Console, you provided the name and location of the PuTTY SSH client private key. At the time you used PuTTYGen to generate the PuTTY SSH private key, you also generated an SSH public key. Familiarize yourself with the name and location of the PuTTY SSH public key on the SAN Volume Controller Console system.

Note: This is a long term administrative task and not just a post-installation task.

To enter config mode:

```
switch#config-t
```

To enable ssh:

```
switch (config)#ssh server enable
```

This document has an overview of the steps necessary to get to the web page where you identify the PuTTY public key to the clusters. These steps are documented in more detail in other sections of this manual and references are included to the relevant section titles.

1. Start your Web browser to access the SAN Volume Controller Console. It is recommended that you log onto the SAN Volume Controller Console system from a browser on which the SAN Volume Controller Console is installed to complete uploading the client public SSH key for each cluster that you want to manage. You can access the SAN Volume Controller Console by typing the following:

   ```
   http://localhost:9080/ica
   ```

   Note: `9080` is the default HTTP port. If a different port number for HTTP was assigned during the installation process, then you must substitute that port number in the URL.
2. Log onto the SAN Volume Controller Console using the default super user name and password. The default super user name is `superuser` and the default super user password is `passw0rd`. The first time you log onto the SAN Volume Controller Console using the default super user name and password, you will be prompted to change the default password.
3. Access user assistance. This is an optional step.

   You can access help for the specific task on which you are working by clicking the small information icon just below the banner in the upper right section of the Web page. The help assistant panel opens on the right-hand side of the page.

   You can also launch a separate user assistance panel by clicking the small question mark icon just below the banner in the upper right section of the Web page. A secondary browser window opens which has icons in the frame labeled **Contents** for you to select to make extensive user assistance information available to you.

4. Identify the SAN Volume Controller clusters to the SAN Volume Controller Console. The steps you might need to perform to add SAN Volume Controller clusters to the SAN Volume Controller Console collection of managed clusters, depends on the current status of the cluster in which you are interested.

   Choose one of the following two steps, depending on whether the cluster has completed the cluster creation (initialization)process:

   a. Uninitialized SAN Volume Controller cluster.

      If you have not yet created a SAN Volume Controller cluster using the front panel of the SAN Volume Controller cluster, you will need to perform that phase of the cluster creation first. You will be given a special password by the customer engineer (CE) to be used in later steps of initializing the SAN Volume Controller Console.

      After you create the SAN Volume Controller cluster using the front panel of cluster, you will need to complete the creation of the cluster by using the SAN Volume Controller Console Web pages.

      Enter the IP address of the cluster and check **Create (Initialize) Cluster**. When you click the **OK** button, the Create a Cluster wizard will take over and present you with the panels you need to complete initializing the cluster.

      The browser will then prompt you to enter the network password. Enter the user name `admin` and the password provided to you by the customer engineer (CE) during the cluster front panel creation phasewhich is configured for the cluster.

      During the initializing of the cluster, using the SAN Volume Controller Console, you will be taken to a Web page to provide the PuTTY SSH client public key to upload the key to the cluster. Step 5 below continues with the SSH public key input description. This PuTTY SSH client public key is the other key of the key pair you provided to the SAN Volume Controller Console during the installation program.

   b. Previously initialized SAN Volume Controller cluster.

      If the SAN Volume Controller cluster has completed the initialization (creation) process but is not yet registered with the SAN Volume Controller Console, click the **Add SAN Volume Controller Cluster** button and then add the cluster IP address but *do not* check **Create (Initialize) Cluster**, which is above the **OK** button. When you click the **OK** button, you will be taken to the Web page to provide the PuTTY SSH client public key to upload to the cluster. Step 5 below continues with the SSH key input description.

      The browser will then prompt you to enter the network password. Enter the user name `admin` and the password which is configured for the cluster. Then Click **OK**.

5. Store the SAN Volume Controller Console system SSH public key on the SAN Volume Controller Console. This PuTTY client SSH public key is the other key in the key pair you provided to the SAN Volume Controller Console during the installation program. Each key is associated with an ID string that you define that can consist of up to 30 characters. Up to 100 keys can be stored on a cluster. You can add keys to provide either *administrator* access or *service* access. Perform the following steps to store the SSH public key on the cluster:

   a. Enter the SSH public key name and directory location on your local browser system in the field labeled **Public Key (file upload)** or click **Browse** to identify the key on the local system. Alternatively, you can paste the SSH key into the **Public Key (direct input)** field.

   b. Enter an ID string in the field labeled **ID**. This is a unique ID to distinguish the key and is not related to a user name.

c. Select the *administrator* **Access Level** radio button.
d. Click **Add Key** to store this SSH public key on the cluster.

6. Launch the secondary Web browser window to manage your specific cluster.

   Once you have identified the SAN Volume Controller clusters to the SAN Volume Controller Console you can see a summary of all of the clusters. From this point, you can select the specific cluster in which you are interested and then launch the browser window specifically for the cluster. Perform the following steps to launch the browser window:

   a. Click **Clusters** in the portfolio section of your browser window in the left-hand frame. A new view will be displayed in the work area.
   b. Check the small box in the Select column left of the cluster in which you are interested to select that cluster. Select **Launch the SAN Volume Controller application** in the drop down list of the work area and click **Go**. A secondary browser window opens to the SAN Volume Controller Web application. Now you can work with the specific SAN Volume Controller cluster which you selected.

      Note: The ClusterName parameter in the browser location URL, identifies the cluster with which you are working.
      For example:

      http://9.43.147.38:9080/svc/Console?Console.login
      Token=79334064:f46d035f31:-7ff1&Console.
      ClusterName=9.43.225.208

      Select **Manage Cluster** and click **View Cluster Properties**in the portfolio section.

This completes the verification of the connection to the SAN Volume Controller.

Related information

"Installing or upgrading the IBM TotalStorage SAN Volume Controller Console for Windows" on page 347

Removing the SAN Volume Controller Console

You can remove the IBM TotalStorage SAN Volume Controller Console from your Windows system.

1. Log onto the system as a local system administrator.
2. Stop the IBM CIM Object Manager (CIMOM), IBM WebSphere Application Server V5 - SVC, and the Service Location Protocol (SLP) services if they are started.
   a. Click **Start -> Settings -> Control Panel**. In the Control Panel window, double-click on the **Administrative Tools** icon and then double-click the **Services** icon. The Services window opens.
   b. Stop the IBM CIM Object Manager (CIMOM) service:
      1) In the Services window, scroll to `IBM CIM Object Manager (CIMOM)`. Click on the service to select it.
      2) If the **Status** column shows `Started`, right-click the service, then click **Stop** on the menu.
   c. Stop the IBM WebSphere Application Server V5 - SVC service:
      1) In the Services window, scroll to `IBM WebSphere Application Server V5 - SVC`. Click on the service to select it.

2) If the **Status** column shows `Started`, right-click the service, then click **Stop** on the menu.
3) Wait for the service to stop.

d. Stop the Service Location Protocol (SLP) service:

Note: You must be careful if you have other applications that use the Service Location Protocol (SLP) service. In this case, you must stop these applications before stopping Service Location Protocol (SLP) service, because during the removal process the Service Location Protocol (SLP) service will be deleted. You must also stop the configuration utilities for the IBM TotalStorage SAN Volume Controller Console, if they are running.

1) In the Services window, scroll to `Service Location Protocol`. Click on this service to select it.
2) If it is running (the **Status** column shows `Started`), right-click the service, then click **Stop** on the menu.

   (If you did not stop the IBM CIM Object Manager (CIMOM) service, the system now asks if you want to stop the IBM CIM Object Manager (CIMOM) . Because the IBM CIM Object Manager (CIMOM) service is dependent on the Service Location Protocol service which you just stopped, you must click **Yes** to stop the IBM CIM Object Manager (CIMOM).)
3) Wait for the services to stop.
4) Close the Services window.
5) Close the Administrative Tools window.

3. Use the Windows Add/Remove Programs facility to remove the IBM TotalStorage SAN Volume Controller Console and the Service Location Protocol components.
   a. From the Windows menu bar, click **Start -> Settings -> Control Panel**. Double-click **Add/Remove Programs**.
   b. Click **IBM TotalStorage SAN Volume Controller Console** from the list of currently installed programs and click **Remove** to remove the product.
4. The Welcome panel for the Uninstaller opens. Click **Next** to continue or click **Cancel** to stop the removal of the IBM TotalStorage SAN Volume Controller Console.
5. The program detects whether the Service Location Protocol, IBM CIM Object Manager (CIMOM), and the IBM WebSphere Application Server V5 - SVC services are running.
   a. If any of these services are found to be running, the uninstaller will stop these services before proceeding with the uninstallation. You should consider at this point whether applications other than the IBM TotalStorage SAN Volume Controller Console are dependent on the services. You can either:
      - Click **Next** to have the program stop the services for you.
      - Click **Cancel** to exit the removal process if you wish to manually stop the services and any dependent applications. Instructions for stopping the services are described in step 2 on page 364. You must then restart the removal process from the Windows Add/Remove facility.
6. The Confirmation panel opens. Click **Remove** to continue or click **Cancel** to stop the removal of the IBM TotalStorage SAN Volume Controller Console. Click **Back** to return to the previous panel.

7. The Uninstallation Progress panel opens. Wait for the program to remove the IBM TotalStorage SAN Volume Controller Console product.
8. The Finish panel for the Uninstaller opens. This panel indicates the result of the removal process (successful or failed). Click **Finish** to complete the removal process and exit the wizard.

   Note: If the Uninstaller could not remove some information from the system, you will see a **Next** button instead of a **Finish** button. Click **Next** to open the Reboot panel. If the reboot panel opens, you can chose to either restart your computer now or restart your computer at a later time. Then click **Finish** to complete the removal process and exit the wizard.
9. Close the Add/Remove Programs window.

Perform the following steps to complete the removal process:
1. If the system has not been restarted since IBM TotalStorage SAN Volume Controller Console was removed, do so now.
2. Log onto the system as a local system administrator.
3. The removal process saves files uniquely related to the configuration in a backup directory under the destination path where you installed the IBM TotalStorage SAN Volume Controller Console. You may want those files if you intend to reinstall the product. Otherwise you can remove the backup folder and files. An example of the default destination path is: C:\Program Files\IBM\svcconsole.
4. Perform other cleanup tasks:
   - Empty your Windows Recycle Bin to reclaim the disk space that was made available during the removal process.

Related information

"Installing or upgrading the IBM TotalStorage SAN Volume Controller Console for Windows" on page 347

Valid combinations of FlashCopy and Metro Mirror functions

The following table outlines the combinations of FlashCopy and Metro Mirror functions that are valid for a single virtual disk (VDisk).

Table 56. Valid combinations of FlashCopy and Metro Mirror interactions

FlashCopy	Metro Mirror Primary	Metro Mirror Secondary
FlashCopy source	Supported	Supported
FlashCopy target	Not supported	Not supported

Moving data between MDisk groups with Copy Services

You cannot use the SAN Volume Controller data migration function to move a VDisk between MDisk groups that have different extent sizes. However, you can use Copy Services to move the data by copying a VDisk between MDisk groups that have different extent sizes.

To copy a VDisk between MDisk groups that have different extent sizes, you have the following options:
- Use FlashCopy to copy a VDisk between a source and a destination MDisk group.

- Use intracluster Metro Mirror to copy a VDisk between a source and a destination MDisk group.

Using FlashCopy

This option is available if you have licensed the FlashCopy feature.

Use the following guidelines for using FlashCopy to copy a VDisk between a source and a destination MDisk group that have different extent sizes:
- The VDisk must not be in another FlashCopy or Metro Mirror relationship.
- Stop all I/O operations from the hosts while the VDisk is being copied.
- Once the copy is complete, configure host mappings for the new VDisk and configure the hosts to access the destination VDisk rather than the source VDisk.

Using Metro Mirror

This option is available if you have licensed the Metro Mirror feature.

Use the following guidelines for using Metro Mirror to copy a VDisk between a source and a destination MDisk group that have different extent sizes:
- The VDisk must not be in another FlashCopy or Metro Mirror relationship.
- Create an intracluster Metro Mirror relationship for the VDisks.
- I/O operations from the host can continue while the copy is being performed, however there will be some performance degradation while write I/O operations are mirrored.
- Once the copy is complete, stop all I/O operations from the hosts. Configure host mappings for the new VDisk and configure the hosts to access the destination VDisk rather than the source VDisk.

Setting up SNMP traps

You can set up SNMP traps if the master console has been installed on a separate machine.

Prerequisites

There are two steps required to enable the Call-Home process:
1. Set up the SAN Volume Controller SNMP Trap destination, a specific machine (IP Address)
2. Set up IBM Director to send a correctly formatted e-mail

Overview

To set up the SAN Volume Controller SNMP trap destination, the destination is normally set up as part of the SAN Volume Controller Installation process, but can also be done through the SAN Volume Controller Web pages, by using a browser to log on to the SAN Volume Controller cluster and selecting the option Error notification. See *IBM TotalStorage SAN Volume Controller: Installation Guide* for more information.

Configuring IBM Director overview

You can configure IBM Director for Call-Home and e-mail if it has been installed on a separate machine or is reinstalled on the master console.

1. Set up an Event Action Plan
2. Set up a correctly formatted e-mail

Related tasks

"Setting up an event action plan"
You can set up event action plans if the IBM Director has been installed on a separate machine or is reinstalled on the master console.

"Setting up an e-mail notification to IBM" on page 369
You can set up an e-mail call home to IBM if the IBM Director has been installed on a separate machine or is reinstalled on the master console.

"Setting up an e-mail user notification" on page 370
You can set up the e-mails if the IBM Director has been installed on a separate machine or is reinstalled on the master console.

Setting up an event action plan

You can set up event action plans if the IBM Director has been installed on a separate machine or is reinstalled on the master console.

In order for IBM Director to present the correct SAN Volume Controller information to enable an action plan to be configured, it has to have received a trap from the SAN Volume Controller.

1. Create a SAN Volume Controller trap by removing the ac power from one of the uninterruptible power supply units that are supplying the cluster. Replace the power after 30 seconds.
2. Click **Event Log (ALL)** from the IBM Director Console and check that the trap from the SAN Volume Controller has been received.
3. Click **Tasks –> Event Action Plan Builder** from the IBM Director Console.
4. Right-click **Simple Event Filter**.
5. Click **New**.
6. Click the **Event type** tab from the Simple Event Filter Builder window.
7. Clear the **Any** check box.
8. In the list, select the following items in this sequence:
   a. SNMP
   b. 1 (iso)
   c. 2 (org)
   d. 6 (dod)
   e. 1 (internet)
   f. 4 (private)
   g. 1 (enterprise)
   h. 2 (ibm)
   i. 6 (ibmprod)
   j. 190
   k. 1
9. Click the **Category** tab.
10. Clear the **Any** check box.

11. Click **Alert**.
12. On the menu bar, click **File**and save the file with the name `2145 Error`.
13. From the Event Filter List, select the newly created **2145 Error** filter, and drag and drop it on to the **Log All Events** icon in the Event Action Plan column. This action causes the **2145 Error** filter to be called upon when any event is logged.
14. Perform steps 4 through 11 again (do not do step 8k). On the menu bar, click **File** and save the file with the name `2145 Event`.
15. From the Event Filter List, select the newly created **2145 Event** filter, and drag and drop it on to the **Log All Events** icon in the Event Action Plan column. This action causes the **2145 Event** filter to be called upon when any event is logged.

Related tasks

"Configuring IBM Director overview" on page 368
You can configure IBM Director for Call-Home and e-mail if it has been installed on a separate machine or is reinstalled on the master console.

Setting up an e-mail notification to IBM

You can set up an e-mail call home to IBM if the IBM Director has been installed on a separate machine or is reinstalled on the master console.

1. From the IBM Director Console menu bar, select **Tasks –> Event Action Plan Builder**.
2. In the **Actions** column, right-click on **Send an Internet (SMTP) E-mail** and select **Customize**.
3. In the resulting **Customize Action: Send an Internet (SMTP) E-mail** panel fill-in:

   Internet E-mail Address

   - Enter the IBM Retain E-mail address
     - CALLHOME1@de.ibm.com for customers in North America, Latin America, South America and Caribbean Islands
     - CALLHOME0@de.ibm.com for customer's outside of the USA.

   Reply to

   - Enter the E-mail address that you require any replies to be directed

   SMPT E-mail Server

   - Enter the address of your E-mail server

   SMPT Port

   - Change this, if required to your SMTP Server port number

   Subject of E-mail Message

   - Fill in `2145 Error Notification`.

   Body of the E-mail Message

   - Fill in the following information:
     - Contact name....................not required in the E-mail to Admin

       Note: There is a limitation of 72 characters per field.
     - Contact phone number........not required in the E-mail to Admin
     - Offshift phone number........not required in the E-mail to Admin
     - Machine location
     - Record Type = 1

&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.1
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.2
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.3
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.4
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.5
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.6
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.7
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.8
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.9
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.10
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.11
&iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.12

4. Click **Save** to save the information, using the name **2145CallHome**.
5. From the **Send an Internet (SMTP) E-mail** list select the newly created **2145CallHome** E-mail and Drag and Drop it on to the **2145 Error** action plan icon in the **Event Action Plan** column. This action causes the **2145CallHome** to be call when the **2145 Error** filter is satisfied.

Related tasks

"Configuring IBM Director overview" on page 368
You can configure IBM Director for Call-Home and e-mail if it has been installed on a separate machine or is reinstalled on the master console.

Setting up an e-mail user notification

You can set up the e-mails if the IBM Director has been installed on a separate machine or is reinstalled on the master console.

1. From the IBM Director Console menu bar, select **Tasks –> Event Action Plan Builder**.
2. In the **Actions** column, right-click on **Send an Internet (SMTP) E-mail** and select **Customize**.
3. In the resulting **Customize Action: Send an Internet (SMTP) E-mail** panel fill-in :

   Internet E-mail Address
   - Enter the E-mail address you require for notification

   Reply to
   - Enter the E-mail address that you require any replies to be directed

   SMPT E-mail Server
   - Enter the address of your E-mail server

   SMPT Port
   - Change this, if required to your SMTP Server port number

   Subject of E-mail Message
   - Fill in `2145 Error Notification`.

   Body of the E-mail Message
   - Fill in the following information:
     – # Machine location = xxxx

iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.1
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.2
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.3
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.4
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.5

iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.6
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.7
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.8
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.9
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.10
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.11
iso.org.dod.internet.private.enterprises.ibm.ibmProd.190.4.12

Where *xxxx* is information relevant to your organization.

4. Click **Save** to save the information, using the name **2145ErrorNot**.
5. From the **Send an Internet (SMTP) E-mail** list select the newly created **2145ErrorNot** E-mail and Drag and Drop it on to the **2145 Event** action plan icon in the **Event Action Plan** column. This action causes the **2145ErrorNot** to be call when the **2145 Event** filter is satisfied.

Related tasks

"Configuring IBM Director overview" on page 368
You can configure IBM Director for Call-Home and e-mail if it has been installed on a separate machine or is reinstalled on the master console.

Object types

The following table lists the object codes and corresponding object types.

Table 57. Object types

Object code	Object type
0	IC_TYPE_Unknown
1	IC_TYPE_Vlun
2	IC_TYPE_Vlungrp
3	IC_TYPE_Hlun
4	IC_TYPE_Node
5	IC_TYPE_Host
6	IC_TYPE_Hostgrp
7	IC_TYPE_Hws
8	IC_TYPE_Fcgrp
9	IC_TYPE_Rcgrp
10	IC_TYPE_Fcmap
11	IC_TYPE_Rcmap
12	IC_TYPE_Wwpn
13	IC_TYPE_Cluster
15	IC_TYPE_Hba
16	IC_TYPE_Device
17	IC_TYPE_SCSILun
18	IC_TYPE_Quorum
19	IC_TYPE_TimeSeconds
20	IC_TYPE_ExtSInst
21	IC_TYPE_ExtInst
22	IC_TYPE_Percentage

Table 57. Object types (continued)

Object code	Object type
23	IC_TYPE_VPD_SystemBoard
24	IC_TYPE_VPD_Processor
25	IC_TYPE_VPD_Processor_Cache
26	IC_TYPE_VPD_Memory_Module
27	IC_TYPE_VPD_Fan
28	IC_TYPE_VPD_FC_Card
29	IC_TYPE_VPD_FC_Device
30	IC_TYPE_VPD_Software
31	IC_TYPE_VPD_Front_Panel
32	IC_TYPE_VPD_UPS
33	IC_TYPE_VPD_Port
34	IC_TYPE_FC_Adapter
35	IC_TYPE_Migrate

Event codes

The system generates information and configuration event codes.

There are two different types of event codes:
- Information event codes
- Configuration event codes

Information event codes provide information on the status of an operation. Information event codes are recorded in the error log and an SNMP trap and sometimes an e-mail is generated if the corresponding management flag is set in the Preference cache.

Configuration event codes are generated when configuration parameters are set. Configuration event codes are recorded in a separate log and do not generate SNMP traps or e-mails and their error fixed flags are ignored.

Related reference

"Information event codes"
The information event codes provide information on the status of a particular operation.

"Configuration event codes" on page 374
Configuration event codes are generated when configuration parameters are set.

Information event codes

The information event codes provide information on the status of a particular operation.

Information event codes are recorded in the error log and an SNMP trap and sometimes an e-mail is generated if the corresponding management flag is set in the Preference cache.

Information event codes generate information type (I) descriptions or warning type (W) descriptions.

Table 58. Information event codes

Event code	Type	Description
980221	I	Error log cleared.
980310	I	Degraded or offline Managed Disk group is now online.
980435	W	Failed to obtain directory listing from remote node
980440	W	Failed to transfer file from remote node
980446	I	Secure Delete complete
980500	W	Featurization Violation
981001	W	Cluster Fabric View has been updated by a multiphase discovery
981007	W	Preferred port is not being used for Managed Disk access
981014	W	LUN Discovery failed. Cluster has a connection to a device through this node but this node cannot discovery the Managed Disks associated LUN correctly.
981020	W	Managed Disk error count warning threshold met.
982003	W	Insufficient Virtual Extents.
982004	W	Migration suspended due to insufficient Virtual Extents or too many media errors on the source MDisk.
982007	W	Migration Stopped.
982009	I	Migrate Complete
982010	W	Copied disk I/O medium error.
983001	I	FlashCopy prepared
983002	I	FlashCopy complete
983003	W	FlashCopy stopped
984001	W	First customer data being pinned in a Virtual Disk working set
984002	I	All customer data in a Virtual Disk working set now unpinned
984003	W	Virtual Disk working set cache mode being changed to synchronous destage because too must pinned data has now been unpinned for that Virtual Disk working set.
984004	I	Virtual Disk working set cache mode now allows asynchronous destage because enough customer data has now been unpinned for that Virtual Disk working set.
985001	I	Metro Mirror, background copy complete
985002	I	Metro Mirror ready to restart
985003	W	Unable to find path to disk in remote cluster within timeout
987102	W	Node power-off requested from power switch
987103	W	Coldstart
987301	W	Connection to a configured remote cluster has been lost.
987400	W	The node unexpectedly lost power but has now been restored to the cluster.

Table 58. Information event codes (continued)

Event code	Type	Description
988100	W	An overnight maintenance procedure has failed to complete. Resolve any hardware and configuration problems that you are experiencing on the SAN Volume Controller cluster. If the problem persists, contact your IBM service representative for assistance.

Related reference

"Event codes" on page 372
The system generates information and configuration event codes.

Configuration event codes

Configuration event codes are generated when configuration parameters are set.

Configuration event codes are recorded in a separate log and do not generate SNMP traps or e-mails and their error fixed flags are ignored.

Table 59. Configuration event codes

Event code	Description
990101	Modify cluster (attributes in the **svctask chcluster** command)
990105	Delete node from cluster (attributes in the **svctask rmnode** command)
990106	Create host (attributes in the **svctask mkhost** command)
990112	Cluster config dumped to file (attributes in the **svctask dumpconfig** command)
990117	Create cluster (attributes in the **svctask mkcluster** command)
990118	Modify node (attributes in the **svctask chnode** command)
990119	Configure set controller name
990120	Shut down node (attributes in the **svctask stopcluster** command)
990128	Modify host (attributes in the **svctask chhost** command)
990129	Delete node (attributes in the **svctask rmnode** command)
990138	Virtual Disk Modify (attributes in the **svctask chvdisk** command)
990140	Virtual Disk Delete (attributes in the **svctask rmvdisk** command)
990144	Modify Managed Disk group (attributes in the **svctask chmdiskgrp** command)
990145	Delete Managed Disk group (attributes in the **svctask rmdiskgrp** command)
990148	Create Managed Disk group (attributes in the **svctask mkmdiskgrp** command)
990149	Modify Managed Disk (attributes in the **svctask chmdisk** command)
990158	VLUN included
990159	Quorum created
990160	Quorum Destroy
990168	Modify the HWS a Virtual Disk is assigned to
990169	Create a new Virtual Disk (attributes in the **svctask mkvdisk** command)

Table 59. Configuration event codes (continued)

Event code	Description
990173	Add a Managed Disk to Managed Disk group (attributes in the **svctask addmdisk** command)
990174	Delete a Managed Disk from Managed Disk group (attributes in the **svctask rmmdisk** command)
990178	Add a port to a Host (attributes in the **svctask addhostport** command)
990179	Delete a port from a Host (attributes in the **svctask rmhostport** command)
990182	Create a Virtual Disk to Host SCSI mapping (attributes in the **svctask mkvdiskhostmap** command)
990183	Delete an Virtual Disk to Host SCSI mapping (attributes in the **svctask rmdiskhostmap** command)
990184	Create a FlashCopy mapping (attributes in the **svctask mkfcmap** command)
990185	Modify a FlashCopy mapping (attributes in the **svctask chfcmap** command)
990186	Delete a FlashCopy mapping (attributes in the **svctask rmfcmap** command)
990187	Prepare a FlashCopy mapping (attributes in the **svctask prestartfcmap** command)
990188	Prepare a FlashCopy consistency group (attributes in the **svctask prestartfcconsistgrp** command)
990189	Trigger a FlashCopy mapping (attributes in the **svctask startfcmap** command)
990190	Trigger a FlashCopy consistency group (attributes in the **svctask startfcconsistgrp** command)
990191	Stop a FlashCopy mapping (attributes in the **svctask stopfcmap** command)
990192	Stop a FlashCopy consistency group (attributes in the **svctask stopfcconsistgrp** command)
990193	FlashCopy set name
990194	Delete a list of ports from a Host (attributes in the **svctask rmhostport** command)
990196	Shrink a Virtual Disk.
990197	Expand a Virtual Disk (attributes in the **svctask expandvdisksize** command)
990198	Expand single extent a Virtual Disk
990199	Modify govern a Virtual Disk
990203	Initiate manual Managed Disk discovery (attributes in the **svctask detectmdisk** command)
990204	Create FlashCopy consistency group (attributes in the **svctask mkfcconsistgrp** command)
990205	Modify FlashCopy consistency group (attributes in the **svctask chfcconsistgrp** command)
990206	Delete FlashCopy consistency group (attributes in the **svctask rmfcconsistgrp** command)

Table 59. Configuration event codes (continued)

Event code	Description
990207	Delete a list of Hosts (attributes in the **svctask rmhost** command)
990213	Change the HWS a node belongs to (attributes in the **svctask chiogrp** command)
990216	Apply software upgrade (attributes in the **svcservicetask applysoftware** command)
990219	Analyze error log (attributes in the **svctask finderr** command)
990220	Dump error log (attributes in the **svctask dumperrlog** command)
990222	Fix error log entry (attributes in the **svctask cherrstate** command)
990223	Migrate a single extent (attributes in the **svctask migrateexts** command)
990224	Migrate a number of extents
990225	Create Metro Mirror relationship (attributes in the **svctask mkrcrelationship** command)
990226	Modify Metro Mirror relationship (attributes in the **svctask chrcrelationship** command)
990227	Delete Metro Mirror relationship (attributes in the **svctask rmrcrelationship** command)
990229	Start Metro Mirror relationship (attributes in the **svctask startrcrelationship** command)
990230	Stop Metro Mirror relationship (attributes in the **svctask stoprcrelationship** command)
990231	Switch a Metro Mirror relationship (attributes in the **svctask switchrcrelationship** command)
990232	Start Metro Mirror consistency group (attributes in the **svctask startrcconsistgrp** command)
990233	Stop Metro Mirror consistency group (attributes in the **svctask stoprcconsistgrp** command)
990234	Switch a Metro Mirror consistency group (attributes in the **svctask switchrcconsistgrp** command)
990235	Managed disk migrated to a Managed Disk group
990236	Virtual Disk migrated to a new Managed Disk
990237	Create partnership with remote cluster (attributes in the **svctask mkpartnership** command)
990238	Modify partnership with remote cluster (attributes in the **svctask chpartnership** command)
990239	Delete partnership with remote cluster (attributes in the **svctask rmpartnership** command)
990240	Create Metro Mirror consistency group (attributes in the **svctask mkrcconsistgrp** command)
990241	Modify Metro Mirror consistency group (attributes in svctask chrcconsistgrp)
990242	Delete Metro Mirror consistency group (attributes in the **svctask rmrcconsistgrp** command)
990245	Node pend

Table 59. Configuration event codes (continued)

Event code	Description
990246	Node remove
990247	Node unpend
990380	Time zone changed (attributes in the **svctask settimezone** command)
990383	Change cluster time (attributes in the **svctask setclustertime** command)
990385	System time changed
990386	SSH key added (attributes in the **svctask addsshkey** command)
990387	SSH key removed (attributes in the **svctask rmsshkey** command)
990388	All SSH keys removed (attributes in the **svctask rmallsshkeys** command)
990390	Add node to the cluster
990395	Shutdown or reset node
990410	Software Install started
990415	Software Install completed
990420	Software Install failed
990430	Planar Serial Number changed
990501	The featurization has changed. See feature log for details.
991024	IO tracing has finished, trigger occurred for given Managed Disk.

Related reference

"Event codes" on page 372
The system generates information and configuration event codes.

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 master console:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. The following screen readers have been tested: JAWS v4.5 and IBM Home Page Reader v3.0.
- 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 SAN Volume Controller Console 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 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-1.ibm.com/servers/storage/support/virtual/2145.html

Related reference

"SAN Volume Controller library and related publications" on page xviii
A list of other publications that are related to this product are provided to you for your reference.

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 reference

"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
- Tivoli
- TotalStorage
- xSeries

Intel and Pentium are trademarks of Intel 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.

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

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

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

Glossary

Ensure you are familiar with the list of terms and their definitions used in this guide.

A

application server
A host that is attached to the storage area network (SAN) and that runs applications.

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.

Call Home
A communication service that links a machine to a service provider. The machine can use this link to place a call to IBM or to another service provider when service is required. With access to the machine, service personnel can perform service tasks, such as viewing error and problem logs or initiating trace and dump retrievals.

cluster
In SAN Volume Controller, a pair of nodes that provides a single configuration and service interface.

concurrent maintenance
Service that is performed on a unit while it is operational.

configuration node
A node that acts as the focal point for configuration commands and manages the data that describes the cluster configuration.

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

consistent copy
: In a Global Mirror 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.

container

- IBM definition: A visual user-interface component that holds objects.
- HP definition:
  1. Any entity that is capable of storing data, whether it is a physical device or a group of physical devices.
  2. A virtual, internal controller structure representing either a single disk or a group of disk drives linked as a storageset. Stripesets and mirrorsets are examples of storageset containers that the controller uses to create units.

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.

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.

D

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

degraded
: Pertaining to a valid configuration that has suffered a failure but continues to be supported and legal. Typically, a repair action can be performed on a degraded configuration to restore it to a valid configuration.

dependent write operations
: A set of write operations that must be applied in the correct order to maintain cross-volume consistency.

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

device

- In the CIM Agent, the storage server that processes and hosts client application requests.
- IBM definition: A piece of equipment that is used with the computer and does not generally interact directly with the system, but is controlled by a controller.
- HP definition: In its physical form, a magnetic disk that can be attached to a SCSI bus. The term is also used to indicate a physical device that has been made part of a controller configuration; that is, a physical device that is known to the controller. Units (virtual disks) can be created from devices after the devices have been made known to the controller.

directed maintenance procedures
: The set of maintenance procedures that can be run for a cluster. These procedures are documented in the service guide.

disconnected
: In a Global Mirror relationship, pertains to two clusters when they cannot communicate.

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).

disk zone
: A zone defined in the storage area network (SAN) fabric in which the SAN Volume Controller can detect and address the logical units that the disk controllers present.

E

error code
: A value that identifies an error condition.

ESS See *IBM TotalStorage Enterprise Storage Server®*.

IBM TotalStorage Enterprise Storage Server (ESS)
: An IBM product that provides an intelligent disk-storage subsystem across an enterprise.

exclude
: To remove a managed disk (MDisk) from a cluster because of certain error conditions.

excluded
: In SAN Volume Controller, 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

failover
: In SAN Volume Controller, 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.

fibre-channel extender
: A long-distance communication device that interconnects storage area network (SAN) fabric components.

FC See *fibre channel*.

FlashCopy service
: In SAN Volume Controller, 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.*

FlashCopy mapping
A relationship between two virtual disks.

FlashCopy relationship
See *FlashCopy mapping*.

H

HBA See *host bus adapter*.

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

host An open-systems computer that is connected to the SAN Volume Controller through a fibre-channel interface.

host ID
In SAN Volume Controller, 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).

host zone
A zone defined in the storage area network (SAN) fabric in which the hosts can address the SAN Volume Controllers.

I

IBM TotalStorage Multipath 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.

illegal configuration
A configuration that will not operate and will generate an error code to indicate the cause of the problem.

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*.

image VDisk
A virtual disk (VDisk) in which there is a direct block-for-block translation from the managed disk (MDisk) to the VDisk.

inconsistent
In a Global Mirror 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.

integrity
The ability of a system to either return only correct data or respond that it cannot return correct data.

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.

I/O See *input/output.*

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

I/O throttling rate
The maximum rate at which an I/O transaction is accepted for this virtual disk (VDisk).

IP See *Internet Protocol*.

J

JBOD (just a bunch of disks)
IBM definition: See *non-RAID*. HP definition: A group of single-device logical units not configured into any other container type.

L

LBA See *logical block address*.

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

local/remote fabric interconnect
The storage area network (SAN) components that are used to connect the local and remote fabrics together.

logical block address (LBA)
The block number on a disk.

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).

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*.

mirrorset
IBM definition: See *RAID-1*. HP definition: A RAID storageset of two or more physical disks that maintain a complete and independent copy of the data from the virtual disk. This type of storageset has the advantage of being highly reliable and extremely tolerant of device failure. Raid level 1 storagesets are referred to as mirrorsets.

N

node One SAN Volume Controller. Each node provides virtualization, cache, and Copy Services to the storage area network (SAN).

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

non-RAID
Disks that are not in a redundant array of independent disks (RAID). IBM definition: Disks that are not in a redundant array of independent disks (RAID). HP definition: See *JBOD*.

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

partition
- IBM definition: A logical division of storage on a fixed disk.
- HP definition: A logical division of a container represented to the host as a logical unit.

partnership
In Global Mirror, the relationship between two clusters. In a cluster partnership, one cluster is defined as the local cluster and the other cluster as the remote cluster.

paused
In SAN Volume Controller, the process by which the cache component quiesces all ongoing I/O activity below the cache layer.

pend To cause to wait for an event.

port The physical entity within a host, SAN Volume Controller, or disk controller system that performs the data communication (transmitting and receiving) over the fibre channel.

primary virtual disk
In a Global Mirror 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

RAID See *redundant array of independent disks*.

RAID 0

- IBM definition: RAID 0 allows a number of disk drives to be combined and presented as one large disk. RAID 0 does not provide any data redundancy. If one drive fails, all data is lost.
- HP definition: A RAID storageset that stripes data across an array of disk drives. A single logical disk spans multiple physical disks, allowing parallel data processing for increased I/O performance. While the performance characteristics of RAID level 0 is excellent, this RAID level is the only one that does not provide redundancy. Raid level 0 storagesets are referred to as stripesets.

RAID 1
SNIA dictionary definition: A form of storage array in which two or more identical copies of data are maintained on separate media. IBM definition: A form of storage array in which two or more identical copies of data are maintained on separate media. Also known as mirrorset. HP definition: See *mirrorset*.

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.

RAID 5

- SNIA definition: A form of parity RAID in which the disks operate independently, the data strip size is no smaller than the exported block size, and parity check data is distributed across the array's disks. (S)
- IBM definition: See above.
- HP definition: A specially developed RAID storageset that stripes data and parity across three or more members in a disk array. A RAIDset combines the best characteristics of RAID level 3 and RAID level 5. A RAIDset is the best choice for most applications with small to medium I/O requests, unless the application is write intensive. A RAIDset is sometimes called parity RAID. RAID level 3/5 storagesets are referred to as RAIDsets.

RAID 10
A type of RAID that optimizes high performance while maintaining fault tolerance for up to two failed disk drives by striping volume data across several disk drives and mirroring the first set of disk drives on an identical set.

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.*

rejected
A status condition that describes a node that the cluster software has removed from the working set of nodes in the cluster.

relationship
In Global Mirror, the association between a master 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.*

Global Mirror
In SAN Volume Controller, 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.

S

SAN See *storage area network*.

SAN Volume Controller fibre-channel port fan in
The number of hosts that can see any one SAN Volume Controller port.

SCSI See *Small Computer Systems Interface*.

sequential VDisk
A virtual disk that uses extents from a single managed disk.

Small Computer System Interface (SCSI)
A standard hardware interface that enables a variety of peripheral devices to communicate with one another.

secondary virtual disk
In Global Mirror, 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*.

stand-alone relationship
In FlashCopy and Global Mirror, relationships that do not belong to a consistency group and that have a null consistency group attribute.

stop A configuration command that is used to stop the activity for all copy relationships in a consistency group.

stopped
The status of a pair of virtual disks (VDisks) that have a copy relationship that the user has temporarily broken because of a problem.

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)

stripeset
See *RAID 0*.

superuser authority
The level of access required to add users.

suspended
The status of a pair of virtual disks (VDisks) that have a copy relationship that has been temporarily broken because of a problem.

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 Global Mirror, the status condition that exists when both virtual disks (VDisks) of a pair that has a copy relationship contain the same data.

T

trigger
To initiate or reinitiate copying between a pair of virtual disks (VDisks) that have a copy relationship.

U

unconfigured mode
A mode in which I/O operations cannot be performed. See also *image mode* and *managed space mode*.

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.

unit identifiers (UIDs)
A unit identifier can be one of the following:

1. an integer expression whose value must be zero or positive
2. an * (asterisk) that corresponds to unit 5 for input or unit 6 for output
3. the name of a character array, character array element, or character substring for an internal file

V

valid configuration
A configuration that is supported.

VDisk See *virtual disk*.

virtual disk (VDisk)
In SAN Volume Controller, a device that host systems attached to the storage area network (SAN) recognize as a Small Computer System Interface (SCSI) disk.

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.

vital product data (VPD)
Information that uniquely defines system, hardware, software, and microcode elements of a processing system.

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

D

E

F

G

H

T

U

V

W

Z

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

IBM TotalStorage SAN Volume Controller
Configuration Guide
Version 2.1.0

Publication No. SC26-7543-04

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

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

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

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

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. ______________________

Readers' Comments — We'd Like to Hear from You
SC26-7543-04

Cut or Fold
Along Line

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
Information Development
Department 61C
9032 South Rita Road
Tucson, Arizona
USA 85775-4401

Fold and Tape **Please do not staple** Fold and Tape

SC26-7543-04

Cut or Fold
Along Line

Part Number: 64P8291

Printed in USA

(1P) P/N: 64P8291

SC26-7543-04

Spine information:

IBM TotalStorage SAN Volume Controller

SAN Volume Controller Configuration Guide

Version 2.1.0