IBM System Storage SAN Volume Controller

Software Installation and Configuration Guide

Version 5.1.0

SC23-6628-05

IBM System Storage SAN Volume Controller

Software Installation and Configuration Guide

Version 5.1.0

SC23-6628-05

Note:
Before using this information and the product it supports, read the information in **Notices**.

This edition applies to the IBM System Storage SAN Volume Controller, release 5.1.0, and to all subsequent releases and modifications until otherwise indicated in new editions. This edition replaces SC23-6628-04.

Contents

Chapter 6. Using the CLI 221

Figures

Tables

About this guide

This publication provides information that helps you configure and use the IBM® System Storage™ SAN Volume Controller.

This publication also describes the configuration tools, both command-line and Web-based, that you can use to define, expand, and maintain the storage of the SAN Volume Controller.

Who should use this guide

This guide is intended for system administrators or others who install, configure, and use the IBM System Storage SAN Volume Controller.

Before using the 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.

Summary of changes

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

Summary of changes for SC23-6628-04 and SC23-6628-05, SAN Volume Controller Software Installation and Configuration Guide

This summary of changes provides a list of new, modified, and changed information since the last version of the guide. This topic describes the changes to this guide since the previous edition, SC23-6628-03.

New information

The following sections summarize the changes that have since been implemented from the previous version.

This version includes the following new information:

- Support statements for the SAN Volume Controller 2145-CF8 node and the optional solid-state drive (SSD)
- Support statements for syslog servers to receive error, warning, and informational notifications from the cluster
- Support statements for reverse FlashCopy® mapping with multiple targets
- Support statements for iSCSI 1Gb host attachment using on-board Ethernet ports
- Support statements for remote authentication for users of SAN Volume Controller clusters
- Configuring information for Texas Memory Systems RamSan Solid State Storage systems
- Configuring information for Xiotech Emprise systems

Changed information

The following updates were made in this document:

- Support statements for new managed disk (MDisk) states
- Support statements for viewing quorum disks and setting active quorum disks
- Support statements for SNMP servers to receive error, warning, and informational notifications from the cluster
- Enhancements to e-mail notification and inventory reporting
- Enhancements to statistics collection
- Support statements for second cluster IP and service IP addresses
- Support statements for a second Ethernet port on SAN Volume Controller nodes
- Support statements for extending the capability of creating partnerships among SAN Volume Controller clusters from one to three partners
- Support statements for embedded CIMOM on a cluster

Removed information

The following information was removed from this book:

- The SAN Volume Controller 2145-4F2 node is not supported in this version of SAN Volume Controller

Summary of changes for SC23-6628-03 SAN Volume Controller Software Installation and Configuration Guide

This summary of changes provides a list of new, modified, and changed information since the last version of the guide. This topic describes the changes to this guide since the previous edition, SC23-6628-02.

New information

The following sections summarize the changes that have since been implemented from the previous version:

This version includes the following new information:

- Information about the new SAN Volume Controller hardware model 2145-8A4
- New Capacity Licensing and Physical Disk Licensing options
- Information about the newly supported IBM XIV® systems
- Support details for IBM System Storage DS5000
- Back-end storage requirements for FlashCopy, Metro Mirror and Global Mirror, and space-efficient virtual disks (VDisks)
- Support for N-port virtualization in the host bus adapter (HBA) or SAN switch
- Usage information for the new recover VDisk commands
- New licensing and hardware error codes
- New event and reason codes

Changed information

The following updates were made in this document:

- Change in terminology from storage subsystem to storage system
- Updates to the space-efficient VDisks overview
- Updates to Brocade core-edge fabrics support

Removed information

The following information was removed from this book:

- Information about the master console was moved to the new *IBM System Storage SAN Volume Controller Master Console Guide*, GC27-2223, which is available at the following Web site:

  www.ibm.com/storage/support/2145

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.

SAN Volume Controller library and related publications

Product manuals, other publications, and Web sites contain information that relates to SAN Volume Controller.

SAN Volume Controller Information Center

The IBM System Storage SAN Volume Controller Information Center contains all of the information that is required to install, configure, and manage the SAN Volume Controller. The information center is updated between SAN Volume Controller product releases to provide the most current documentation. The information center is available at the following Web site:

http://publib.boulder.ibm.com/infocenter/svcic/v3r1m0/index.jsp

SAN Volume Controller library

Table 1 on page xviii 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) from the following Web site:

www.ibm.com/storage/support/2145

Table 1. SAN Volume Controller library

Title	Description	Order number
IBM System Storage SAN Volume Controller Planning Guide	This guide introduces the SAN Volume Controller and lists the features that you can order. It also provides guidelines for planning the installation and configuration of the SAN Volume Controller.	GA32-0551
IBM System Storage SAN Volume Controller Model 2145-CF8 Hardware Installation Guide	This guide provides the instructions that the IBM service representative uses to install the hardware for SAN Volume Controller model 2145-CF8.	GC52-1356
IBM System Storage SAN Volume Controller Model 2145-8A4 Hardware Installation Guide	This guide provides the instructions that the IBM service representative uses to install the hardware for SAN Volume Controller model 2145-8A4.	GC27-2219
IBM System Storage SAN Volume Controller Model 2145-8G4 Hardware Installation Guide	This guide provides the instructions that the IBM service representative uses to install the hardware for SAN Volume Controller model 2145-8G4.	GC27-2220
IBM System Storage SAN Volume Controller Models 2145-8F2 and 2145-8F4 Hardware Installation Guide	This guide provides the instructions that the IBM service representative uses to install the hardware for SAN Volume Controller models 2145-8F2 and 2145-8F4.	GC27-2221
IBM System Storage SAN Volume Controller Software Installation and Configuration Guide	This guide provides guidelines for configuring your SAN Volume Controller. Instructions for backing up and restoring the cluster configuration, using and upgrading the SAN Volume Controller Console, using the CLI, upgrading the SAN Volume Controller software, and replacing or adding nodes to a cluster are included.	SC23-6628

Table 1. SAN Volume Controller library (continued)

Title	Description	Order number
IBM System Storage SAN Volume Controller CIM Agent Developer's Guide	This guide describes the concepts of the Common Information Model (CIM) environment. Steps about using the CIM agent object class instances to complete basic storage configuration tasks, establishing new Copy Services relationships, and performing CIM agent maintenance and diagnostic tasks are included.	SC23-6665
IBM System Storage 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).	SC26-7903
IBM System Storage SAN Volume Controller Host Attachment Guide	This guide provides guidelines for attaching the SAN Volume Controller to your host system.	SC26-7905
IBM System Storage SAN Volume Controller Troubleshooting Guide	This guide describes the features of each SAN Volume Controller model, explains how to use the front panel, and provides maintenance analysis procedures to help you diagnose and solve problems with the SAN Volume Controller.	GC27-2227
IBM System Storage SAN Volume Controller Hardware Maintenance Guide	This guide provides the instructions that the IBM service representative uses to service the SAN Volume Controller hardware, including the removal and replacement of parts.	GC27-2226
IBM System Storage SAN Volume Controller Master Console Guide	This guide describes how to install, maintain, and service the master console.	GC27-2223
IBM Systems Safety Notices	This guide contains translated caution and danger statements. Each caution and danger statement in the SAN Volume Controller documentation has a number that you can use to locate the corresponding statement in your language in the *IBM Systems Safety Notices* document.	G229-9054

Other IBM publications

Table 2 lists IBM publications that contain information related to the SAN Volume Controller.

Table 2. Other IBM publications

Title	Description	Order number
IBM System Storage Productivity Center Introduction and Planning Guide	This guide introduces the IBM System Storage Productivity Center hardware and software.	SC23-8824
Read This First: Installing the IBM System Storage Productivity Center	This guide describes how to install the IBM System Storage Productivity Center hardware.	GI11-8938
IBM System Storage Productivity Center User's Guide	This guide describes how to configure the IBM System Storage Productivity Center software.	SC27-2336
IBM System Storage Multipath Subsystem Device Driver User's Guide	This guide describes the IBM System Storage Multipath Subsystem Device Driver for IBM System Storage products and how to use it with the SAN Volume Controller.	GC52-1309
Implementing the IBM System Storage SAN Volume Controller V4.3	This IBM Redbooks® publication is a detailed technical guide to the IBM System Storage SAN Volume Controller. It provides a high-level overview of storage virtualization and the SAN Volume Controller architecture, discusses implementing and configuring the SAN Volume Controller, tells you how to migrate existing storage to the SAN Volume Controller, and discusses different supported migration activities.	SG24-6423

IBM documentation and related Web sites

Table 3 lists Web sites that provide publications and other information about the SAN Volume Controller or related products or technologies.

Table 3. IBM documentation and related Web sites

Web site	Address
Support for SAN Volume Controller (2145)	www.ibm.com/storage/support/2145
Support for IBM System Storage and IBM TotalStorage® products	www.ibm.com/storage/support/
IBM Publications Center	www.ibm.com/shop/publications/order/
IBM Redbooks publications	www.redbooks.ibm.com/

Related accessibility information

To view a PDF file, you need Adobe Acrobat Reader, which can be downloaded from the Adobe Web site:

www.adobe.com/support/downloads/main.html

How to order IBM publications

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

The IBM Publications Center offers customized search functions to help you find the publications that you need. Some publications are available for you to view or download at no 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/

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.

SAN Volume Controller installation and configuration overview

The installation and configuration of a SAN Volume Controller cluster requires the completion of various tasks, some of which are normally completed by an IBM service representative.

Additional publications are included with some of the hardware components; however, use the installation and configuration procedures in the documents that are listed here.

When you plan or perform the installation and configuration tasks, have the following SAN Volume Controller publications available:

- *IBM System Storage SAN Volume Controller Planning Guide*
- *IBM System Storage SAN Volume Controller Model 2145-XXX Hardware Installation Guide*, where 2145-*XXX* is a specific node model
- *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*

See the Support for SAN Volume Controller (2145) Web site for access to SAN Volume Controller publications:

www.ibm.com/storage/support/2145

The IBM System Storage Productivity Center (SSPC) is the management environment for SAN Volume Controller clusters. For SSPC planning, installation, and configuration information, see the following publications:

- *IBM System Storage Productivity Center Introduction and Planning Guide*, SC23-8824
- *Read This First: Installing the IBM System Storage Productivity Center*, GI11-8938
- *IBM System Storage Productivity Center User's Guide*, SC27-2336

To access the SSPC publications, go to the **Printable PDFs** section and click the **IBM System Storage Productivity Center** link from the following Web site:

publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp

Note: A master console can be upgraded to support clusters that are running the latest SAN Volume Controller software. For details, see the *IBM System Storage SAN Volume Controller Master Console Guide* at the Support for SAN Volume Controller (2145) Web site:

www.ibm.com/storage/support/2145

Planning tasks to complete before installing the SAN Volume Controller

Before you install the SAN Volume Controller, you must complete the following planning tasks or have them completed by an IBM service representative or IBM Business Partner:

1. **Verify that all the SAN Volume Controller installation requirements have been met.**

Review Chapter 2 of the *IBM System Storage SAN Volume Controller Planning Guide* to make sure that space and power requirements have been met before you begin the installation. SAN Volume Controller nodes and uninterruptible power-supply units are installed in pairs.

2. **Review SAN fabric and zoning guidelines and develop your SAN Volume Controller cluster, host systems, and storage controllers plan.**

   This task helps to assure a seamless configuration. For more information, see Chapters 3 and 4 of the *IBM System Storage SAN Volume Controller Planning Guide*.

3. **Complete all physical planning charts.**

   Chapter 2 of the *IBM System Storage SAN Volume Controller Planning Guide* provides instructions for accessing and completing the following charts and tables:

   - Hardware location chart
   - Cable connection table
   - Configuration data table
   - Redundant ac-power connection chart

   The SAN Volume Controller charts and tables are available at the Support for SAN Volume Controller (2145) Web site:

   www.ibm.com/storage/support/2145

   You can save, edit, and share the charts and tables between members of the installation team.

   For the SSPC, complete the planning worksheet in the Appendix of the *IBM System Storage Productivity Center Introduction and Planning Guide*.

   You can also obtain the planning work sheet from the IBM System Storage Productivity Center Information Center at http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/index.jsp. In the left navigation pane, click **System Storage Productivity Center → Getting started → Planning work sheet**.

Hardware installation tasks that an IBM service representative performs

To install the SAN Volume Controller hardware, an IBM service representative must complete the following tasks:

1. **Verify that you have all of the required parts for the installation.**

   Chapter 2 of the *IBM System Storage SAN Volume Controller Model 2145-XXX Hardware Installation Guide* provides a list of all the parts that are required for an installation. The list includes the SAN Volume Controller nodes, uninterruptible power-supply units, optional redundant ac-power switches, and associated parts.

2. **Install the SAN Volume Controller hardware.**

   Chapter 2 of the *IBM System Storage SAN Volume Controller Model 2145-XXX Hardware Installation Guide* describes the procedures for installing the uninterruptible power-supply units, SAN Volume Controller nodes, and the optional redundant ac-power switches.

3. **Install the SSPC server.**

   Read This First: Installing the IBM System Storage Productivity Center describes how to install the SSPC server.

Configuration tasks

To configure a SAN Volume Controller cluster, you must complete the following tasks or have them completed by an IBM service representative or IBM Business Partner:

1. **Register your product.**

   To receive product support notifications from IBM, you must register your product. To register your product, click **Register** at the Support for IBM System Storage and TotalStorage products Web site:

   www.ibm.com/systems/support/supportsite.wss/brandmain?brandind=5345868

2. **Optionally, check for an updated version of the SAN Volume Controller Console (GUI) software.**

   For the latest information, click **Install/use**, and then click the link for the appropriate recommended software level from the Support for SAN Volume Controller (2145) Web site:

   www.ibm.com/storage/support/2145

   Additionally, preinstalled software on the SSPC console might need to be updated to fully support the latest level of SAN Volume Controller. For the latest information, go to the Support for System Storage Productivity Center (SSPC) Web site:

   www.ibm.com/systems/support/storage/software/sspc

3. **Configure the IBM System Storage Productivity Center.**

   The *IBM System Storage Productivity Center User's Guide* describes how to configure the SSPC for the SAN Volume Controller.

4. **Create a SAN Volume Controller cluster.**

   The *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide* describes this procedure, which is completed in two phases:

   a. Use the Create Cluster option on the front panel of one of the SAN Volume Controller nodes that you have installed to create the cluster.

      This procedure is usually performed by an IBM representative or IBM Business Partner using information that the customer provides.

   b. Use the Add a Cluster function from the SAN Volume Controller Console.

5. **Complete the initial SAN Volume Controller configuration.**

   After you create the SAN Volume Controller cluster, you must perform the configuration procedures that are needed to meet your requirements. You can perform these procedures in stages; for example, add nodes to a cluster, set cluster date and time, and set license features immediately. Later, after your applications are tested and migrated to SAN Volume Controller, create host definitions, assign managed disks (MDisks) to MDisk groups, and set up virtual disks (VDisks) and assign them to hosts.

   You can also set up event notifications, including Call Home e-mails, to immediately notify you and the IBM Support Center if critical problems occur.

   The *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide* describes how to perform these steps by using either the SAN Volume Controller Console or the comand-line interface (CLI).

Chapter 1. SAN Volume Controller overview

The SAN Volume Controller combines software and hardware into a comprehensive, modular appliance that uses symmetric virtualization.

Symmetric virtualization is achieved by creating a pool of managed disks (MDisks) from the attached storage systems. Those storage systems are then mapped to a set of virtual disks (VDisks) for use by attached host systems. System administrators can view and access a common pool of storage on the storage area network (SAN). This functionality helps administrators to use storage resources more efficiently and provides a common base for advanced functions.

A *SAN* is a high-speed fibre-channel network that connects host systems and storage devices. In a SAN, a host system can be connected to a storage device across the network. The connections are made through units such as routers and switches. The area of the network that contains these units is known as the *fabric* of the network.

SAN Volume Controller software

The SAN Volume Controller software performs the following functions for the host systems that attach to SAN Volume Controller:

- Creates a single pool of storage
- Provides logical unit virtualization
- Manages logical volumes
- Mirrors logical volumes

The SAN Volume Controller also provides the following functions:

- Large scalable cache
- Copy Services
  - IBM FlashCopy (point-in-time copy)
  - Metro Mirror (synchronous copy)
  - Global Mirror (asynchronous copy)
  - Data migration
- Space management
  - Mapping that is based on desired performance characteristics
  - Metering of service quality
  - Space-efficient logical volumes (thin provisioning)

SAN Volume Controller hardware

Each SAN Volume Controller node is an individual server in a SAN Volume Controller cluster on which the SAN Volume Controller software runs.

The nodes are always installed in pairs, with a minimum of one and a maximum of four pairs of nodes constituting a *cluster*. 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.

Note: I/O groups take the storage that is presented to the SAN by the storage systems as MDisks and translates the storage into logical disks, known as VDisks, that are used by applications on the hosts. A node resides in only one I/O group and provides access to the VDisks in that I/O group.

The following nodes are supported in SAN Volume Controller 5.1:

- The new SAN Volume Controller 2145-CF8 node is available for purchase, with up to four of the optional solid-state drives (SSDs).
- The SAN Volume Controller 2145-8A4 node remains available for purchase.
- The SAN Volume Controller 2145-8G4 node is no longer available for purchase, but remains supported in SAN Volume Controller 5.1.
- The SAN Volume Controller 2145-8F4 node is no longer available for purchase, but remains supported in SAN Volume Controller 5.1.
- The SAN Volume Controller 2145-8F2 node is no longer available for purchase, but remains supported in SAN Volume Controller 5.1.

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 systems. These systems can be supplied from various vendors. The pool can be split into virtual disks (VDisks) that are visible to the host systems that use them. Therefore, VDisks 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 the following four main levels:

At the server level
Manages 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.

At the storage device level
Uses striping, mirroring and RAIDs to create disk systems. 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.

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.

At the file system level
Provides the highest benefit because data is shared, allocated, and protected at the data level rather than the volume level.

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 system level. Multiple RAID controllers of different types require knowledge of, and software that is specific to, the given hardware. Virtualization provides a central point of control for disk creation and maintenance.

One problem area that virtualization addresses is unused capacity. Before virtualization, individual host systems each had their own storage, which wasted unused storage capacity. Using virtualization, storage is pooled so that jobs from any attached system that need large amounts of storage capacity can use it as needed. Virtualization makes it easier to regulate the amount of available storage without having to use host system resources or to turn storage devices off and on to add or remove capacity. Virtualization also provides the capability to move storage between storage systems transparently to host systems.

Types of virtualization

Virtualization can be performed either asymmetrically or symmetrically. Figure 1 on page 4 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 and presents disks to the hosts, but hides 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, the benefits of those levels can also be combined. For example, you can combine benefits by attaching a RAID controller to a virtualization engine that provides virtual volumes for 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.

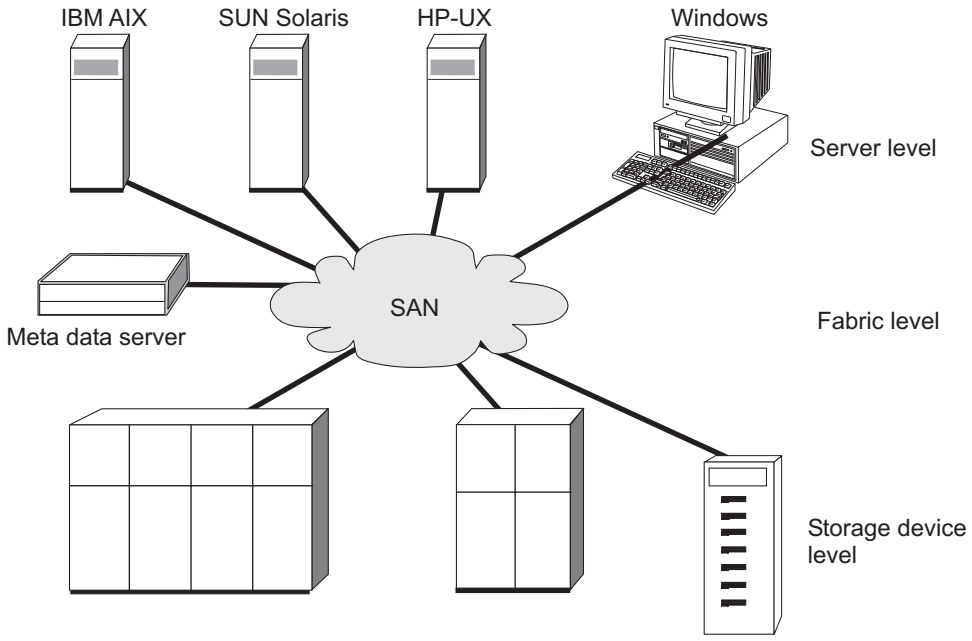

Figure 1. Levels of virtualization

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 2 on page 5, 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. However, there are disadvantages to asymmetric virtualization.

Asymmetric virtualization can have the following disadvantages:

- 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.

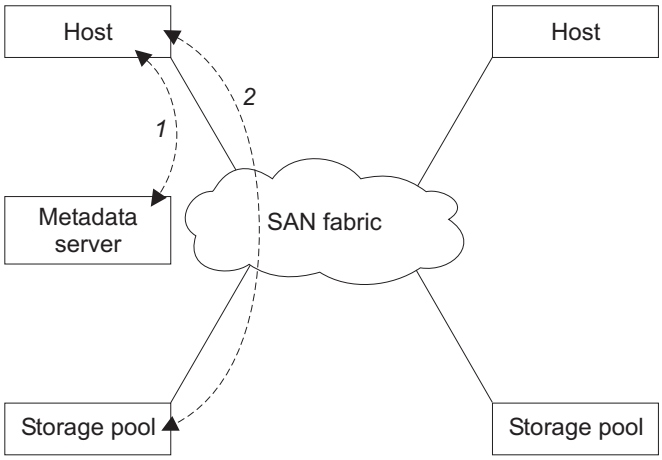

Figure 2. Asymmetrical virtualization

Symmetric virtualization

The SAN Volume Controller provides symmetric virtualization.

Virtualization splits the storage that is presented by the storage systems into smaller chunks that are known as extents. These extents are then concatenated, using various policies, to make virtual disks (VDisks). With symmetric virtualization, host systems can be isolated from the physical storage. Advanced functions, such as data migration, can run without the need to reconfigure the host. With symmetric virtualization, the virtualization engine is the central configuration point for the SAN.

Figure 3 shows that the storage is pooled under the control of the virtualization engine, because the separation of the control from the data occurs in the data path. The virtualization engine performs the logical-to-physical mapping.

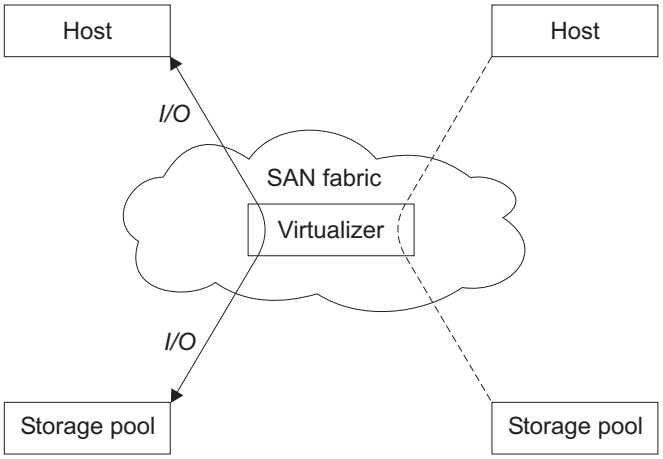

Figure 3. Symmetrical virtualization

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. Therefore, the virtualization engine is a central point of control for device and advanced function management. Symmetric virtualization allows you to build a firewall in the storage network. Only the virtualization engine can grant access through the firewall.

Symmetric virtualization can cause some problems. The main problem that is associated with symmetric virtualization is scalability. Scalability can cause poor performance because all input/output (I/O) must flow through the virtualization engine. To solve this problem, you can use an *n-way* cluster of virtualization engines that has failover capacity. You can scale the additional processor power, cache memory, and adapter bandwidth to achieve the desired level of performance. Additional memory and processing power are needed to run advanced services such as Copy Services and caching.

The 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.

SAN Volume Controller operating environment

To use the SAN Volume Controller, you must meet the minimum hardware and software requirements and ensure that other operating environment criteria are met.

Minimum requirements

You must set up your SAN Volume Controller operating environment according to the following requirements:

- Minimum of one pair of SAN Volume Controller nodes
- Minimum of two uninterruptible power supply units
- One IBM System Storage Productivity Center or one master console per SAN installation for configuration

SAN Volume Controller 2145-CF8 node features

The SAN Volume Controller 2145-CF8 node has the following features:

- A 19-inch rack-mounted enclosure
- One 4-port 8 Gbps fibre-channel adapter
- 24 GB memory
- One quad-core processor
- Dual, redundant power supplies
- Supports up to four optional solid-state drives (SSDs)

SAN Volume Controller 2145-8A4 node features

The SAN Volume Controller 2145-8A4 node has the following features:

- A 19-inch rack-mounted enclosure
- One 4-port 4 Gbps fibre-channel adapter
- 8 GB cache memory
- One dual-core processor

Solid-state drive (SSD) features

Support for solid-state drives (SSDs) is an optional feature of the SAN Volume Controller 2145-CF8. SSDs include the following features:

- Up to four SSDs can be installed on each SAN Volume Controller 2145-CF8 node. An IBM PCIe SAS host bus adapter (HBA) is required on each node that contains an SSD.
- Each SSD is a 2.5-inch Serial Attached SCSI (SAS) drive.
- Each SSD provides up to 146 GB of real capacity.
- SSDs are hot-pluggable and hot-swappable.

Supported hosts

In a SAN environment, host systems are application servers that access data from the storage controllers that are connected to the SAN. Hosts that are running in a number of operating environments can connect to the storage using the SAN Volume Controller. Host connections to the SAN Volume Controller are either SCSI over the fibre-channel SAN or iSCSI over an Ethernet network.

For a list of the supported host operating systems, go to the IBM System Storage SAN Volume Controller Web site:

www.ibm.com/servers/storage/software/virtualization/svc

From the Web site, take the following steps:

1. In the **Learn more** column, click **Interoperability**.
2. Click **Recommended software levels** for your SAN Volume Controller code version.
3. Click **Multipathing / Host Drivers, Clustering and SAN Boot Support - By Host Operating System** to view a list of supported operating systems and to access host attachment scripts.

Multipathing software

For the most current information, go to the following Web site:

www.ibm.com/servers/storage/software/virtualization/svc

From the Web site, take the following steps:

1. In the **Learn more** column, click **Interoperability**.
2. Click **Recommended software levels** for your SAN Volume Controller code version.
3. Click **Multipathing / Host Drivers, Clustering and SAN Boot Support - By Host Operating System** to view a list of supported operating systems and to access multipath drivers. You can also view **Multipath Driver Co-existence with SDD** information.

User interfaces

The SAN Volume Controller software provides the following user interfaces:

- The 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) that uses Secure Shell (SSH)

Application programming interfaces

The SAN Volume Controller software provides an application programming interface called the Common Information Model (CIM) agent, which supports the Storage Management Initiative Specification (SMI-S) of the Storage Network Industry Association.

SAN Volume Controller objects

The SAN Volume Controller solution is based on a group of virtualization concepts. Before setting up your SAN Volume Controller environment, you should understand the concepts and the objects in the environment.

Each SAN Volume Controller is a single processing unit called a *node*. Nodes are deployed in pairs to make up a cluster. A cluster can consist of one to four pairs of nodes. Each pair of nodes is known as an *I/O group* and each node can be in only one I/O group.

Virtual disks (VDisks) are logical disks that are presented by the clusters. Each VDisk is associated with a particular 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 VDisk, it can access the VDisk with either of the nodes in the I/O group. Because each I/O group has only two nodes, the distributed cache is only two-way.

Each node does not contain any internal battery backup units and therefore must be connected to an *uninterruptible power supply*, which provides data integrity in the event of a cluster wide power failure. In such situations, the uninterruptible power supply maintains 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 that is presented by back-end *disk controllers* as a number of disks, known as *managed disks (MDisks)*.

Each MDisk is divided into a number of *extents* which are numbered, from 0, sequentially from the start to the end of the MDisk. MDisks are collected into groups, known as MDisk groups.

Each VDisk is made up of one or two VDisk copies. Each VDisk copy is an independent physical copy of the data that is stored on the VDisk. A VDisk with two copies is known as a *mirrored VDisk*. VDisk copies are made out of MDisk extents. All the MDisks that contribute to a particular VDisk copy must belong to the same MDisk group.

A VDisk can be space-efficient. This means that the capacity of the VDisk as seen by host systems, called the virtual capacity, can be different from the amount of storage that is allocated to the VDisk from MDisks, called the real capacity. Space-efficient VDisks can be configured to automatically expand their real capacity by allocating new extents.

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

For a SCSI over fibre-channel connection, the nodes detect the fibre-channel ports that are connected to the SAN. These correspond to the worldwide port names (WWPNs) of the fibre-channel host bus adapters (HBAs) that are present in the application servers. You can create logical host objects that group WWPNs that belong to a single application server or to a set of them.

For an iSCSI over Ethernet connection, the iSCSI qualified name (IQN) identifies the iSCSI target (destination) adapter. Host objects can have both IQNs and WWPNs.

SAN Volume Controller hosts are virtual representations of the physical host systems and application servers that are authorized to access the cluster VDisks. Each SAN Volume Controller host definition specifies the connection method (SCSI over fibre-channel or iSCSI over Ethernet), the fibre-channel port or Ethernet IP address, and the VDisks that the host applications can access.

The cluster provides block-level aggregation and volume management for disk storage within the SAN. The cluster manages a number of back-end storage controllers and maps the physical storage within those controllers into logical disk images that can be seen by application servers and workstations in the SAN. The SAN is configured in such a way that the application servers cannot see the back-end physical storage. This prevents any possible conflict between the cluster and the application servers both trying to manage the back-end storage.

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.

Clusters

All your configuration, monitoring, and service tasks 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.

A cluster can consist of between two and eight SAN Volume Controller nodes.

All configuration settings are replicated across all nodes in the cluster. Because configuration is performed at the cluster level, management IP addresses are assigned to the cluster instead of to each node. The cluster is configured using the SAN Volume Controller Console, the command-line interface (CLI) or an application developed to access the SAN Volume Controller CIMOM. Each interface accesses the cluster remotely through the Ethernet cluster-management address.

Each node has two Ethernet ports that can be used for management. Ethernet port 1 must be configured and connected on the configuration node. Ethernet port 1 must be connected on all cluster nodes. The use of Ethernet port 2 is optional. At any point in time, only one node in the cluster can operate as the focal point for configuration and monitoring requests. This node is called the *configuration node*. It is the only node that activates the cluster IP addresses. You can use one or more of these addresses to access the cluster through the SAN Volume Controller graphical user interface or the command-line interface (CLI).

Each SAN Volume Controller cluster can have one to four management IP addresses. You can assign up to two IPv4 addresses and up to two IPv6 addresses. When a node has been assigned to a cluster, you can display the cluster IP addresses on the front panel by selecting **Cluster** from the menu.

Each SAN Volume Controller cluster can have optional Small Computer System Interface over Internet Protocol (iSCSI IP) addresses.

Note: Management IP addresses that are assigned to a cluster must be different from the iSCSI IP addresses and are used for different purposes. If iSCSI is used, iSCSI addresses are assigned to individual node ports. On the configuration node, a port will have multiple IP addresses active at the same time.

Cluster management:

A cluster is managed using a command-line session or the SAN Volume Controller console over an Ethernet connection.

Each SAN Volume Controller node has two Ethernet ports that can be used for management. Ethernet port 1 must be configured and connected on the configuration node. Ethernet port 1 must be connected on all cluster nodes. The use of Ethernet port 2 is optional. At any point in time, only one node in the cluster can operate as the focal point for configuration and monitoring requests. This node is called the *configuration node*. It is the only node that activates the cluster IP addresses. You can use one or more of these addresses to access the cluster through the SAN Volume Controller graphical user interface or the command-line interface (CLI).

At any point in time, only one node in the cluster is assigned the role of configuration node. The configuration node is the only node that has active cluster IP addresses, and is the only node that receives cluster management requests.

You can assign IPv4 addresses, IPv6 addresses, or both to a cluster. When a node has been added to a cluster, you can display the cluster IP addresses on the front panel by selecting **Cluster** from the menu.

Each SAN Volume Controller cluster has one or two cluster IP addresses, as well as Small Computer System Interface over Internet Protocol (iSCSI IP) addresses.

Note: Cluster IP addresses that are assigned to a cluster are different from iSCSI IP addresses and are used for different purposes. If iSCSI is used, iSCSI addresses are assigned to node ports. On the configuration node, a port has multiple IP addresses active at the same time.

Cluster IP failover:

If the configuration node fails, the cluster IP addresses are transferred to a new node. The cluster services are used to manage the transfer of the cluster IP addresses from the failed configuration node to the new configuration node.

The following changes are performed by the cluster service:

- If software on the failed configuration node is still operational, the software shuts down the cluster IP interfaces. If the software cannot shut down the cluster IP interfaces, the hardware service forces the node to shut down.

- When the cluster IP interfaces shut down, all remaining nodes choose a new node to host the configuration interfaces.
- The new configuration node initializes the configuration daemons, including sshd and httpd, and then binds the cluster IP interfaces to its Ethernet ports.
- The router is configured as the default gateway for the new configuration node.
- The routing tables are established on the new configuration node for the cluster IP addresses. The new configuration node sends five unsolicited address resolution protocol (ARP) packets for each IP address to the local subnet broadcast address. The ARP packets contain the cluster IP and the media access control (MAC) address for the new configuration node. All systems that receive ARP packets are forced to update their ARP tables. Once the ARP tables are updated, these systems can connect to the new configuration node.

  Note: Some Ethernet devices might not forward ARP packets. If the ARP packets are not forwarded, connectivity to the new configuration node cannot be established automatically. To avoid this problem, configure all Ethernet devices to pass unsolicited ARP packets. You can restore lost connectivity by logging into the SAN Volume Controller and starting a secure copy to the affected system. Starting a secure copy forces an update to the ARP cache for all systems connected to the same switch as the affected system.

Ethernet link failures

If the Ethernet link to the SAN Volume Controller cluster fails because of an event unrelated to the SAN Volume Controller, such as a cable being disconnected or an Ethernet router failure, the SAN Volume Controller does not attempt to failover the configuration node to restore cluster IP access. SAN Volume Controller provides the option for two Ethernet ports, each with its own management IP address, to protect against this type of failure. If you cannot connect through one IP address, attempt to access the cluster through the alternate IP address.

Note: IP addresses that are used by hosts to access the cluster over an Ethernet connection are different from cluster IP addresses.

Routing considerations for event notification

SAN Volume Controller supports the following protocols that make outbound connections from the cluster:

- E-mail
- Simple Network Mail Protocol (SNMP)
- Syslog
- Network Time Protocol (NTP)

One or more of these protocols can be configured on the cluster to receive event notifications. When making outbound connections, the SAN Volume Controller uses the following routing decisions:

- If the destination IP address is in the same subnet as one of the cluster IP addresses, then SAN Volume Controller sends the packet immediately.
- If the destination IP address is not in the same subnet as either of the cluster IP addresses, then SAN Volume Controller sends the packet to the default gateway for Ethernet port 1.

- If the destination IP address is not in the same subnet as either of the cluster IP addresses and Ethernet port 1 is not connected to the Ethernet network, then SAN Volume Controller sends the packet to the default gateway for Ethernet port 2.

When configuring any of these protocols for event notifications, use these routing decisions to ensure error notification works correctly in the event of a network failure.

Cluster configuration backup:

Cluster 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 your cluster configuration in the event that it is lost. Only the data that describes the cluster configuration is backed up. You must back up your application data using the appropriate backup methods.

Configuration restore:

Configuration restore is the process of using a backup cluster configuration data file or files to restore a specific cluster configuration.

Restoring the cluster configuration is an important part of a complete backup and disaster recovery solution. You must also regularly back up your application data using appropriate backup methods because you might need to restore your application data after you have restored your cluster configuration.

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. 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 4 describes the operational states of a node.

Table 4. Node state

State	Description
Adding	The node was added to the cluster but is not yet synchronized with the cluster state (see Note). The node state changes to Online after synchronization is complete.
Deleting	The node is in the process of being deleted from the cluster.

Table 4. Node state (continued)

State	Description
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: A node can stay in the Adding state for a long time. You should wait for at least 30 minutes before taking further action, but if after 30 minutes, the node state is still Adding, you can delete the node and add it again. If the node that has been added is at a lower code level than the rest of the cluster, the node is upgraded to the cluster code level, which can take up to 20 minutes. During this time, the node is shown as adding.	

Configuration node:

A *configuration node* is a single node that manages configuration activity of the cluster.

If the configuration node fails, the cluster chooses a new configuration node. This action is called configuration node failover. The new configuration node takes over the cluster IP addresses. Thus you can access the cluster through the same IP addresses 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.

Figure 4 shows an example cluster containing four nodes. Node 1 has been designated the configuration node. User requests (*1*) are handled by Node 1.

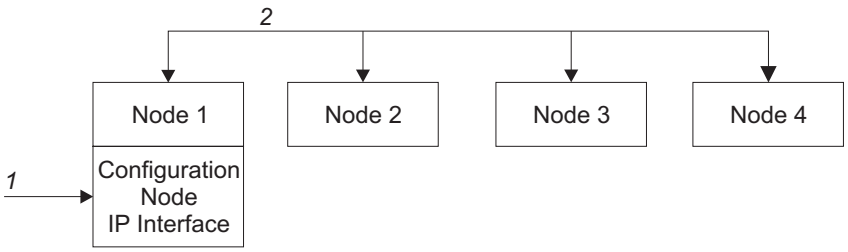

Figure 4. Configuration node

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 can only be in one I/O group.

Virtual disks (VDisks) are logical disks that are presented to the SAN by SAN Volume Controller nodes. VDisks 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.

I/O groups

Each pair of nodes is known as an *input/output (I/O) group*. An I/O group is defined during the cluster configuration process.

Each node can only be in one I/O group. The I/O groups are connected to the SAN so that all backend storage and all application servers are visible to all of the I/O groups. Each pair of nodes has the responsibility to serve I/O operations on a particular set of virtual disks (VDisks).

VDisks are logical disks that are presented to the SAN by SAN Volume Controller nodes. VDisks are also associated with an I/O group. Nodes do 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. The uninterruptible power supply only provides power long enough to enable the SAN Volume Controller cluster to shutdown and save cache data. The uninterruptible power supply is not intended to maintain power and keep the nodes running during an outage.

When an application server performs I/O to a VDisk, it can access the VDisk with either of the nodes in the I/O group. When you create a VDisk, you can specify a preferred node. Many of the multipathing driver implementations that SAN Volume Controller supports use this information to direct I/O to the preferred node. The other node in the I/O group is used only if the preferred node is not accessible.

If you do not specify a preferred node for a VDisk, the node in the I/O group that has the fewest VDisks is selected by the SAN Volume Controller to be the preferred node.

After the preferred node is chosen, it can be changed only when the VDisk is moved to a different I/O group.

Attention: Moving a VDisk to a different I/O group can be disruptive to host I/O.

To view the current preferred node assignment, run the svcinfo lsvdisk command.

An I/O group consists of two nodes. When a write operation is performed to a VDisk, the node that processes the I/O duplicates the data onto the partner node that is in the I/O group. After the data is protected on the partner node, the write operation to the host application is completed. The data is physically written to disk later.

Read I/O is processed by referencing the cache in the node that receives the I/O. If the data is not found, it is read from the disk into the cache. The read cache can provide better performance if the same node is chosen to service I/O for a particular VDisk.

I/O traffic for a particular VDisk is, at any one time, managed exclusively by the nodes in a single I/O group. Thus, although a cluster can have eight nodes within it, the nodes manage I/O in independent pairs. This means that the I/O capability of the SAN Volume Controller scales well, because additional throughput can be obtained by adding additional I/O groups.

Figure 5 on page 15 shows a write operation from a host (*1*), that is targeted for VDisk A. This write is targeted at the preferred node, Node 1 (*2*). The write is

cached and a copy of the data is made in the partner node, Node 2's cache (*3*). The host views the write as complete. At some later time, the data is written, or de-staged, to storage (*4*). Figure 5 also shows two uninterruptible power supply units correctly configured so that each node is in a different power domain.

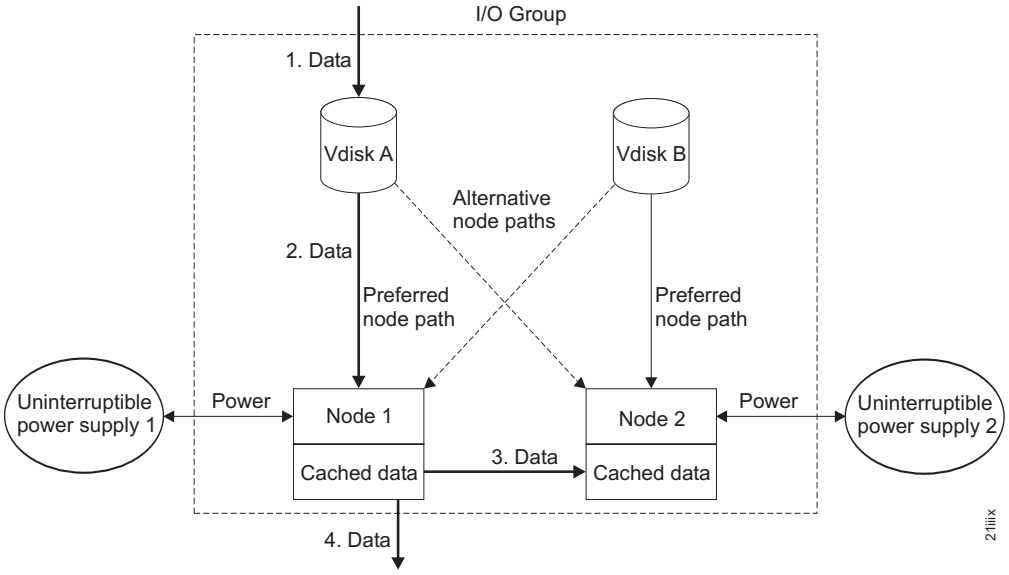

Figure 5. I/O group and uninterruptible power supply

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

If only one node is assigned to an I/O group or if 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 VDisks that are assigned to this I/O group are not cached; they are sent directly to the storage device. If both nodes in an I/O group go offline, the VDisks that are assigned to the I/O group cannot be accessed.

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

The cluster also provides a virtual recovery I/O group that can be used for certain service actions. You can move the VDisks to the recovery I/O group and then into a working I/O group. I/O access is not possible when VDisks are assigned to a recovery I/O group.

I/O 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 and Global Mirror primary and secondary VDisks as follows:

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

2145 UPS-1U

A 2145 UPS-1U is used exclusively to maintain data that is held in the SAN Volume Controller dynamic random access memory (DRAM) in the event of an unexpected loss of external power. This use differs from the traditional uninterruptible power supply that enables continued operation of the device that it supplies when power is lost.

With a 2145 UPS-1U, data is saved to the internal disk of the SAN Volume Controller node. The uninterruptible power supply units are required to power the SAN Volume Controller nodes even when the input power source is considered uninterruptible.

Note: The uninterruptible power supply maintains continuous SAN Volume Controller-specific communications with its attached SAN Volume Controller nodes. A SAN Volume Controller node cannot operate without the uninterruptible power supply. The uninterruptible power supply must be used in accordance with documented guidelines and procedures and must not power any equipment other than a SAN Volume Controller node.

2145 UPS-1U operation:

Each SAN Volume Controller node monitors the operational state of the uninterruptible power supply to which it is attached.

If the 2145 UPS-1U reports a loss of input power, the SAN Volume Controller node stops all I/O operations and dumps the contents of its dynamic random access memory (DRAM) to the internal disk drive. When input power to the 2145 UPS-1U is restored, the SAN Volume Controller node restarts and restores the original contents of the DRAM from the data saved on the disk drive.

A SAN Volume Controller node is not fully operational until the 2145 UPS-1U battery state indicates that it has sufficient charge to power the SAN Volume Controller node long enough to save all of its memory to the disk drive. In the event of a power loss, the 2145 UPS-1U has sufficient capacity for the SAN Volume Controller to save all its memory to disk at least twice. For a fully charged 2145 UPS-1U, even after battery charge has been used to power the SAN Volume Controller node while it saves dynamic random access memory (DRAM) data, sufficient battery charge remains so that the SAN Volume Controller node can become fully operational as soon as input power is restored.

Important: Do not shut down a 2145 UPS-1U without first shutting down the SAN Volume Controller node that it supports. Data integrity can be compromised by pushing the 2145 UPS-1U on/off button when the node is still operating. However, in the case of an emergency, you can manually shut down the 2145 UPS-1U by pushing the 2145 UPS-1U on/off button when the node is still operating. Service actions must then be performed before the node can resume normal operations. If

multiple uninterruptible power supply units are shut down before the nodes they support, data can be corrupted.

Storage systems and MDisks

The nodes in a cluster detect the storage exported by SAN-attached storage systems as a number of disks, known as managed disks (MDisks). The SAN Volume Controller does not attempt to provide recovery from physical disk failures within a storage system. An MDisk is usually, but not necessarily, a RAID array.

An MDisk is either a disk from a storage system or an internal solid-state drive (SSD).

Storage systems

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

Storage systems provide the storage that a SAN Volume Controller cluster detects as one or more managed disks (MDisks).

SAN Volume Controller supports both RAID and non-RAID storage systems. RAID storage systems provide redundancy at the disk level, which prevents a single physical disk failure from causing an MDisk, MDisk group or associated VDisk failure. To minimize data loss, only virtualize the following RAID storage systems: RAID 1, RAID 10, RAID 0+1 and RAID 5.

Many storage systems can be used for storage that is provided by a RAID to be divided up into many Small Computer System Interface (SCSI) logical units (LUs) that are presented on the SAN. With the SAN Volume Controller, ensure that the storage systems are configured to present each RAID as a single SCSI LU that are recognized by the SAN Volume Controller as a single MDisk. The virtualization features of the SAN Volume Controller can then be used to divide up the storage into VDisks.

The exported storage devices are detected by the cluster and reported by the user interfaces. The cluster can also determine which MDisks each storage system is presenting and can provide a view of MDisks that is filtered by the storage system. Therefore, you can associate the MDisks with the RAID that the system exports.

The storage system can have a local name for the RAID 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 system. The storage system makes the storage devices visible with a unique ID, called the logical unit number (LUN). This ID, along with the storage system serial number or numbers (there can be more than one controller in a storage system), can be used to associate the MDisks in the cluster with the RAID exported by the system.

The size of an MDisk cannot be changed once it becomes a managed MDisk by adding it to an MDisk group. If the size of the LUN that is presented by the storage system is reduced to below the size of the managed MDisk, the MDisk is taken offline by the SAN Volume Controller. If the size of the LUN that is presented by the storage system is increased, the SAN Volume Controller does not use the additional space. To increase the storage capacity that is managed on a storage system, create a new LU on the storage system and add the MDisk that represents this LU to the MDisk group.

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.

MDisks

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

An MDisk might, therefore, consist of multiple physical disks that are presented as a single logical disk to the SAN. An MDisk 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 MDisk is divided into a number of extents, which are numbered, from 0, sequentially from the start to the end of the MDisk. The extent size is a property of MDisk groups. When an MDisk is added to an MDisk group, the size of the extents that the MDisk is 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. The following list provides the three types of possible access 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 an MDisk that contains existing data to an MDisk group while the MDisk is in unmanaged or managed mode, you lose the data that it contains. The *image mode* is the only mode that preserves this data.

Figure 6 on page 19 shows physical disks and MDisks.

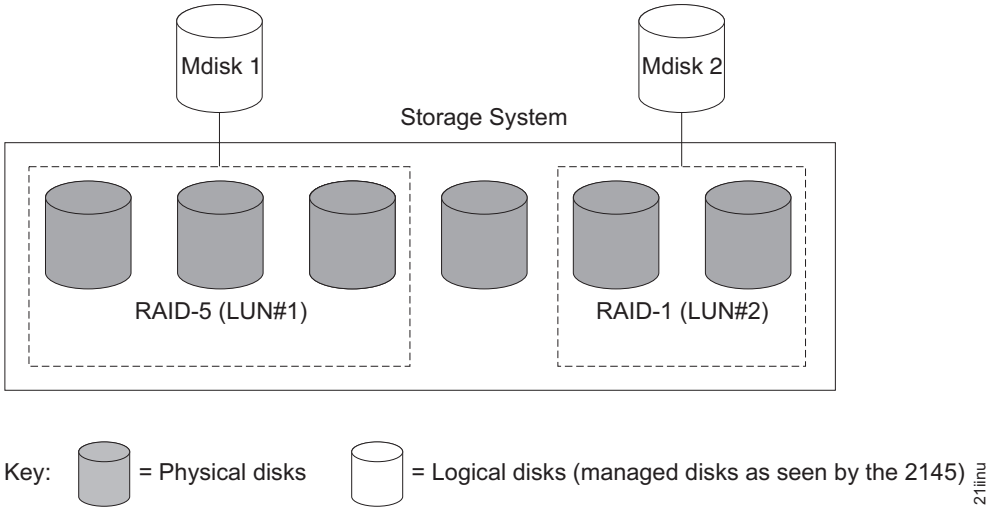

Figure 6. Controllers and MDisks

Table 5 describes the operational states of an MDisk.

Table 5. MDisk 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. • Logical unit number (LUN) inventory of the target ports correctly reported the MDisk. • Discovery of this LUN completed successfully. • All of the MDisk target ports report this LUN as available with no fault conditions.
Degraded paths	The MDisk is not accessible to one or more nodes in the cluster. Degraded path status is most likely the result of incorrect configuration of either the disk controller or the fibre-channel fabric. However, hardware failures in the disk controller, fibre-channel fabric, or node could also be a contributing factor to this state. Complete the following actions to recover from this state: 1. Verify that the fabric configuration rules for storage systems are correct. 2. Ensure that you have configured the storage system properly. 3. Correct any errors in the error log.

Table 5. MDisk status (continued)

Degraded ports	The MDisk has one or more 1220 errors in the error log. The 1220 error indicates that the remote fibre-channel port has been excluded from the MDisk. This error might cause reduced performance on the storage controller and usually indicates a hardware problem with the storage controller. To fix this problem you must resolve any hardware problems on the storage controller and fix the 1220 errors in the error log. To resolve these errors in the log, select **Service and Maintenance** → **Run Maintenance Procedures** in the SAN Volume Controller Console. On the Maintenance Procedures panel, select **Start Analysis**. This action displays a list of unfixed errors that are currently in the error log. For these unfixed errors, select the error name to begin a guided maintenance procedure to resolve them. Errors are listed in descending order with the highest priority error listed first. Resolve highest priority errors first.
Excluded	The MDisk has been excluded from use by the cluster after repeated access errors. Run the Directed Maintenance Procedures to determine the problem.
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, storage system, or one or more physical disks connected to the storage system. The MDisk is 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 that provides the logical connection between MDisks and VDisk copies.

Attention: If you have observed intermittent breaks in links or if you have been replacing cables or connections in the SAN fabric or LAN configuration, you might 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, 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 **Work with Managed Disks** → **Managed Disk** → **Include an MDisk** in the SAN Volume Controller Console, or by issuing the following command in the command-line interface (CLI):

```
svctask includemdisk mdiskname/id
```

Where *mdiskname/id* is the name or ID of your MDisk.

MDisk path

Each MDisk has an online path count, which is the number of nodes that have access to that MDisk; this represents a summary of the I/O path status between the cluster nodes and the 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. If the current path count is not equal to the maximum path count, the

MDisk might be degraded. That is, one or more nodes might not see the MDisk on the fabric.

MDisk groups and VDisks

Managed disks (MDisks) are collected into groups known as *managed disk groups*. Virtual disks (VDisks) are logical disks that are presented to the SAN by SAN Volume Controller nodes. VDisks, like nodes, are associated with an I/O group.

VDisk copies are created from the extents of MDisks.

MDisk groups

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

Figure 7 shows an MDisk group containing four MDisks.

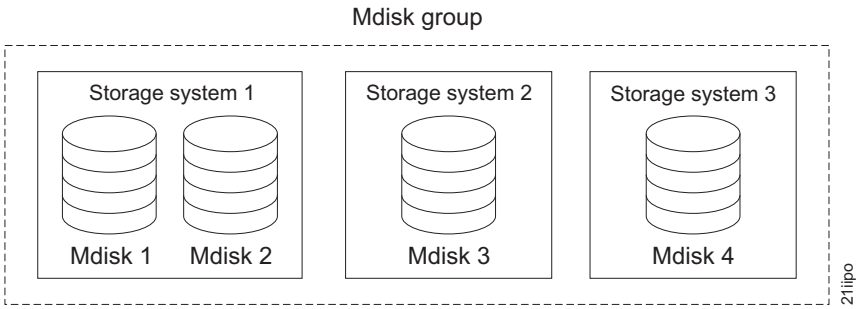

Figure 7. MDisk group

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 either to increase the number of extents that are available for new VDisk copies or to expand existing VDisk copies.

You can specify a warning capacity for an MDisk group. A warning event is generated when the amount of space that is used in the MDisk group exceeds the warning capacity. This is especially useful in conjunction with space-efficient VDisks that have been configured to automatically consume space from the MDisk group.

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.
- If the VDisk is mirrored and the synchronized copies of the VDisk are all in the MDisk group, the mirrored VDisk is destroyed when the MDisk group is deleted.
- If the VDisk is mirrored and there is a synchronized copy in another MDisk group, the VDisk remains after the MDisk group is deleted.

Table 6 describes the operational states of an MDisk group.

Table 6. MDisk group status

Status	Description
Online	The MDisk group is online and available. All the MDisks in the group are available.
Degraded paths	This status indicates that one or more nodes in the cluster cannot access all the MDisks in the group. Degraded path state is most likely the result of incorrect configuration of either the disk controller or the fibre-channel fabric. However, hardware failures in the disk controller, fibre-channel fabric, or node could also be a contributing factor to this state. Complete the following actions to recover from this state: 1. Verify that the fabric configuration rules for storage systems are correct. 2. Ensure that you have configured the storage system properly. 3. Correct any errors in the error log.

Table 6. MDisk group status (continued)

Status	Description
Degraded ports	This status indicates that one or more 1220 errors have been logged against the MDisks in the MDisk group. The 1220 error indicates that the remote fibre-channel port has been excluded from the MDisk. This error might cause reduced performance on the storage controller and usually indicates a hardware problem with the storage controller. To fix this problem you must resolve any hardware problems on the storage controller and fix the 1220 errors in the error log. To resolve these errors in the log, select **Service and Maintenance → Run Maintenance Procedures** in the SAN Volume Controller Console. On the Maintenance Procedures panel, select **Start Analysis**. This action displays a list of unfixed errors that are currently in the error log. For these unfixed errors, select the error name to begin a guided maintenance procedure to resolve them. Errors are listed in descending order with the highest priority error listed first. Resolve highest priority errors first.
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 VDisk copies 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:

- 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 system 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 is reduced to the lowest performer in the group.
- If you intend to keep the VDisk allocation within the storage exported by a storage system, ensure that the MDisk group that corresponds with a single storage system is presented by that storage system. This also enables nondisruptive migration of data from one storage system to another storage system 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.
- In general, MDisk groups that consist of single-port attached systems are not supported by the SAN Volume Controller. However, in some cases, specifically on HP StorageWorks MA and EMA systems that contain RAID partitions, the only way these systems can be attached to the SAN Volume Controller is through single-port attach mode.

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, 512, 1024, or 2048 MB.

Table 7 compares the maximum VDisk capacity for each extent size. The maximum is different for space-efficient VDisks.

Table 7. Maximum VDisk capacity by extent size

Extent size (MB)	Maximum VDisk capacity in GB (not space-efficient VDisks)	Maximum VDisk capacity in GB (space-efficient VDisks)
16	2048 (2 TB)	2000
32	4096 (4 TB)	4000
64	8192 (8 TB)	8000
128	16,384 (16 TB)	16,000
256	32,768 (32 TB)	32,000
512	65,536 (64 TB)	65,000
1024	131,072 (128 TB)	130,000
2048	262,144 (256 TB)	260,000

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.

You cannot use the SAN Volume Controller data migration function to migrate VDisks between MDisk groups that have different extent sizes. However, you can use the following SAN Volume Controller functions to move data to an MDisk that has a different extent size:

- FlashCopy to copy a VDisk between a source and a destination MDisk group that have different extent sizes.
- Intracluster Metro Mirror or Global Mirror to copy a VDisk between a source and a destination MDisk group that have different extent sizes.
- VDisk Mirroring to add a copy of the disk from the destination MDisk group. After the copies are synchronized, you can free up extents by deleting the copy of the data in the source MDisk group.

The choice of extent size affects the total amount of storage that is managed by the cluster. Table 8 shows the maximum amount of storage that can be managed by a cluster for each extent size.

Table 8. 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
1024 MB	4 PB
2048 MB	8 PB

A cluster can manage 4 million extents (4 x 1024 x 1024). For example, with a 16 MB extent size, the cluster can manage up to 16 MB x 4 MB = 64 TB of storage.

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

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, this can cause storage to be wasted at the end of each VDisk.

Virtual disks

A *virtual disk (VDisk)* is a logical disk that the cluster presents to the hosts.

To keep a VDisk accessible even when a managed disk on which it depends has become unavailable, a mirrored copy can be added to a selected VDisk. Each VDisk can have a maximum of two copies. Each VDisk copy is created from a set of extents in an MDisk group.

Application servers on the SAN access VDisks, not managed disks (MDisks).

There are three types of VDisks: striped, sequential, and image.

Types

Each VDisk copy can be one of the following types:

Striped

A VDisk copy that has been striped is at the extent level. One extent is allocated, in turn, from each MDisk that is in the group. For example, an MDisk group that has 10 MDisks takes one extent from each MDisk. The 11th extent is taken from the first MDisk, 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 MDisk group. The round-robin procedure is used across the specified stripe set.

Attention: By default, striped VDisk copies are striped across all MDisks in the group. If some of the MDisks are smaller than others, the extents on the smaller MDisks are used up before the larger MDisks run out of extents. Manually specifying the stripe set in this case might result in the VDisk copy not being created.

If you are unsure if there is sufficient free space to create a striped VDisk copy, 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 copy by not supplying a specific stripe set.

Figure 8 shows an example of an MDisk group that contains three MDisks. This figure also shows a striped VDisk copy that is created from the extents that are available in the group.

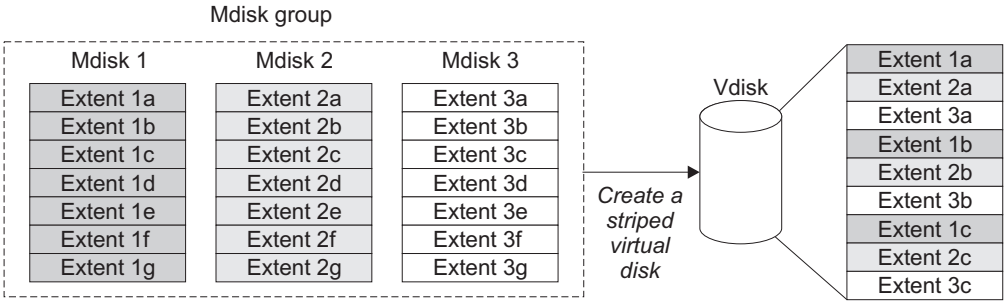

Figure 8. MDisk groups and VDisks

Sequential

When extents are selected, they are allocated sequentially on one MDisk to create the VDisk copy if enough consecutive free extents are available on the chosen MDisk.

Image Image-mode VDisks are special VDisks that have a direct relationship with one MDisk. If you have an MDisk that contains data that you want to merge into the cluster, you can create an image-mode VDisk. When you create an image-mode VDisk, a direct mapping is made between extents

that are on the MDisk and extents that are on the VDisk. The MDisk is not virtualized. The logical block address (LBA) *x* on the MDisk is the same as LBA *x* on the VDisk.

When you create an image-mode VDisk copy, you must assign it to an MDisk group. An image-mode VDisk copy must be at least one extent in size. The minimum size of an image-mode VDisk copy is the extent size of the MDisk group to which it is assigned.

The extents are managed in the same way as other VDisk copies. 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 VDisk copy becomes a virtualized disk, and the mode of the MDisk changes from image to managed.

Attention: If you add a managed mode MDisk to an MDisk group, any data on the MDisk is 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 if it contains partitions or data.

You can use more sophisticated extent allocation policies to create VDisk copies. When you create a striped VDisk, you can specify the same MDisk more than once in the list of MDisks that are used as the stripe set. This is useful if you have an MDisk group in which not all the MDisks are of the same capacity. For example, if you have an MDisk group that has two 18 GB MDisks and two 36 GB MDisks, you can create a striped VDisk copy 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 VDisk, you destroy access to the data that is on the VDisk. The extents that were used in the VDisk are returned to the pool of free extents that is in the MDisk group. The deletion might fail if the VDisk is still mapped to hosts. The deletion might also fail if the VDisk is still part of a FlashCopy, Metro Mirror or Global Mirror mapping. If the deletion fails, you can specify the force-delete flag to delete both the VDisk and the associated mappings to hosts. Forcing the deletion deletes the Copy Services relationship and mappings.

States

A VDisk can be in one of three states: online, offline, and degraded. Table 9 describes the different states of a VDisk.

Table 9. VDisk states

State	Description
Online	At least one synchronized copy of the VDisk is online and available if both nodes in the I/O group can access the VDisk. A single node can only access a VDisk if it can access all the MDisks in the MDisk group that are associated with the VDisk.

Table 9. VDisk states (continued)

State	Description
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 any synchronized copy of the VDisk. The VDisk can also be offline if the VDisk is the secondary of a Metro Mirror or Global Mirror relationship that is not synchronized. A space-efficient VDisk goes offline if a user attempts to write an amount of data that exceeds the available disk space.
Degraded	The status of the VDisk is degraded if one node in the I/O group is online and the other node is either missing or cannot access any synchronized copy of the VDisk. **Note:** If you have a degraded VDisk and all of the associated nodes and MDisks are online, call the IBM Support Center for assistance.

Cache modes

You can select to have read and write operations stored in cache by specifying a cache mode. You must specify the cache mode when you create the VDisk. After the VDisk is created, you cannot change the cache mode.

Table 10 describes the two types of cache modes for a VDisk.

Table 10. VDisk cache modes

Cache mode	Description
readwrite	All read and write I/O operations that are performed by the VDisk are stored in cache. This is the default cache mode for all VDisks.
none	All read and write I/O operations that are performed by the VDisk are not stored in cache.

Virtual disk mirroring:

Virtual disk mirroring allows a VDisk to have two physical copies. Each VDisk copy can belong to a different managed disk (MDisk) group, and each copy has the same virtual capacity as the VDisk.

When a server writes to a mirrored VDisk, the SAN Volume Controller cluster writes the data to both copies. When a server reads a mirrored VDisk, the SAN Volume Controller cluster picks one of the copies to read. If one of the mirrored VDisk copies is temporarily unavailable; for example, because the RAID controller that provides the MDisk group is unavailable, the VDisk remains accessible to servers. The SAN Volume Controller cluster remembers which areas of the VDisk are written and resynchronizes these areas when both copies are available.

You can create a VDisk with one or two copies and convert a non-mirrored VDisk into a mirrored VDisk by adding a copy. When a copy is added in this way, the

SAN Volume Controller cluster synchronizes the new copy so that it is the same as the existing VDisk. Servers can access the VDisk during this synchronization process.

You can convert a mirrored VDisk into a non-mirrored VDisk by deleting one copy or by splitting one copy to create a new non-mirrored VDisk.

The VDisk copy can be any type: image, striped, sequential, and space-efficient or not. The two copies can be of completely different types.

VDisk Mirroring can be used for the following applications:

- Improving availability of VDisks by protecting them from a single storage controller failure.
- Allowing concurrent maintenance of a storage controller that does not natively support concurrent maintenance.
- Providing an alternative method of data migration with better availability characteristics. While a VDisk is being migrated using the data migration feature, it is vulnerable to failures on both the source and target MDisk group. VDisk Mirroring provides an alternative because you can start with a non-mirrored VDisk in the source MDisk group and then add a copy to that VDisk in the destination MDisk group. When the VDisk is synchronized, you can delete the original copy that is in the source MDisk group. During the synchronization process, the VDisk remains available even if there is a problem with the destination MDisk group.
- Maintaining access to data that is stored on SSDs if one of the nodes in an I/O group is being serviced or fails.
- Converting between fully allocated VDisks and space-efficient VDisks.

When you use VDisk mirroring, consider how quorum candidate disks are allocated. VDisk mirroring maintains some state data on the quorum disks. If a quorum disk is not accessible and VDisk mirroring is unable to update the state information, a mirrored VDisk might need to be taken offline to maintain data integrity. To ensure the high availability of the system, ensure that multiple quorum candidate disks, allocated on different controllers, are configured.

Attention: Mirrored VDisks may be taken offline if there is no quorum disk available. This behavior occurs because synchronization status for mirrored VDisks is recorded on the quorum disk. To protect against mirrored VDisks being taken offline, follow the guidelines for setting up quorum disks.

Space-efficient virtual disks:

When you create a virtual disk (VDisk), you can designate it as space-efficient. A space-efficient VDisk has a virtual capacity and a real capacity.

Virtual capacity is the VDisk storage capacity that is available to a host. *Real capacity* is the storage capacity that is allocated to a VDisk copy from a managed disk (MDisk) group. In a fully allocated VDisk, the virtual capacity and real capacity are the same. In a space-efficient VDisk, however, the virtual capacity can be much larger than the real capacity.

The virtual capacity of a space-efficient VDisk is typically significantly larger than its real capacity. A SAN Volume Controller cluster uses the real capacity to store data that is written to the VDisk, and metadata that describes the space-efficient configuration of the VDisk. As more information is written to the VDisk, more of

the real capacity is used. The SAN Volume Controller cluster identifies read operations to unwritten parts of the virtual capacity and returns zeros to the server without using any of the real capacity.

The SAN Volume Controller must maintain extra metadata that describes the contents of space-efficient VDisks. This means the I/O rates that are obtained from space-efficient VDisks are slower than those obtained from fully allocated VDisks that are allocated on the same MDisks.

Space-efficient VDisks can also simplify server administration. Instead of assigning a VDisk with some capacity to an application and increasing that capacity as the application's needs change, you can configure a VDisk with a large virtual capacity for the application and then increase or shrink the real capacity as the application needs change, without disrupting the application or server.

When you configure a space-efficient VDisk, you can use the warning level attribute to generate a warning event when the used real capacity exceeds a specified amount or percentage of the total real capacity. You can also use the warning event to trigger other actions, such as taking low-priority applications offline or migrating data into other MDisk groups.

If a space-efficient VDisk does not have enough real capacity for a write operation, the VDisk is taken offline and an error is logged (error code 1865, event ID 060001). Access to the space-efficient VDisk is restored by either increasing the real capacity of the VDisk or by increasing the size of the MDisk group that it is allocated on.

Note: On a SAN Volume Controller 2145-CF8 node, space is not allocated on a space-efficient VDisk if an incoming host write operation contains all zeros.

When you create a space-efficient VDisk, you can choose the grain size for allocating space in 32 KB, 64 KB, 128 KB, or 256 KB chunks. Generally, smaller grain sizes save space but require more metadata access, which can adversely impact performance. If you are not going to use the space-efficient VDisk as a FlashCopy source or target VDisk, use 256 KB to maximize performance. If you are going to use the space-efficient VDisk as a FlashCopy source or target VDisk, specify the same grain size for the VDisk and for the FlashCopy feature.

When you create a space-efficient VDisk, set the cache mode to `readwrite` to maximize performance. If the cache mode is set to `none`, the SAN Volume Controller cluster cannot cache the space-efficient metadata, which decreases performance.

The autoexpand feature prevents a space-efficient VDisk from using up its capacity and going offline. As a space-efficient VDisk uses capacity, the autoexpand feature maintains a fixed amount of unused real capacity, called the *contingency capacity*. For space-efficient VDisks that are not configured with the autoexpand feature, the contingency capacity can get used up, causing the VDisk to go offline. To determine if an application requires a space-efficient VDisk with the autoexpand feature, create a space-efficient VDisk with the autoexpand feature turned off. If the application causes the VDisk to run out of capacity and go offline, you can then create a space-efficient VDisk with the autoexpand feature turned on.

Image mode space-efficient VDisks:

When you create an image mode virtual disk (VDisk), you can designate it as space-efficient. An image mode space-efficient VDisk has a virtual capacity and a real capacity.

An image mode space-efficient VDisk has a direct relationship with a single MDisk where the contents of the MDisk map to the real capacity that is used by the space-efficient VDisk. Unlike fully-allocated VDisks, the logical block address (LBA) on the MDisk is not necessarily the same as the LBA on the VDisk. You cannot change the real capacity of an image mode space-efficient VDisk, manually or by using the autoexpand feature. To use the autoexpand feature, the VDisk must be in managed mode.

You can use an image mode VDisk to move a space-efficient VDisk between two SAN Volume Controller clusters using the following procedure. The procedure is similar to that used for fully-allocated VDisks, but has an extra step during the import process to specify the existing space-efficient metadata, rather than to create a new, empty VDisk.

1. If the VDisk is not already in image mode, migrate the VDisk to image mode and wait for the migration to complete.
2. Delete the VDisk from the exporting cluster.
3. Disconnect the MDisk from the exporting cluster and connect the MDisk to the importing cluster.
4. Create a new image mode space-efficient VDisk using the MDisk. You must specify the **import** option.
5. Optionally, migrate the VDisk to managed mode.

The **import** option is only valid for SAN Volume Controller space-efficient VDisks. If you use this method to import a space-efficient volume that is created by RAID controllers into a cluster, SAN Volume Controller cannot detect it as a space-efficient VDisk. However, you can use the VDisk mirroring feature to convert an image-mode fully allocated VDisk to a space-efficient VDisk.

Converting space-efficient VDisks:

You can convert space-efficient VDisks into fully allocated VDisks.

You can nondisruptively convert a space-efficient VDisk into a fully allocated VDisk by using the following VDisk mirroring procedure:

1. Start with a single-copy, space-efficient VDisk.
2. Add a fully allocated copy to the VDisk.
3. Wait while the VDisk mirroring feature synchronizes.
4. Remove the space-efficient copy from the VDisk.

Converting fully allocated VDisks:

You can convert fully allocated VDisks to space-efficient VDisks.

You can nondisruptively convert a fully allocated VDisk into a space-efficient VDisk by following this procedure:

1. Start with a single copy, fully allocated VDisk.
2. Add a space-efficient copy to the VDisk. Use a small real capacity and the autoexpand feature.
3. Wait while the VDisk mirroring feature synchronizes the copies.

4. Remove the fully allocated copy from the space-efficient VDisk.

Any grains of the fully allocated VDisk that contain all zeros do not cause any real capacity to be allocated on the space-efficient copy. Before you create the mirrored copy, you can fill the free capacity on the VDisk with a file that contains all zeros.

Host objects

A *host system* is a computer that is connected to the SAN Volume Controller through either a fibre-channel interface or an IP network.

A *host object* is a logical object in the SAN Volume Controller that represents a list of worldwide port names (WWPNs) and a list of iSCSI names that identify the interfaces that the host system uses to communicate with the SAN Volume Controller. iSCSI names can be either iSCSI qualified names (IQNs) or extended unique identifiers (EUIs).

A typical configuration has one host object for each host system that is attached to the SAN Volume Controller. If a cluster of hosts accesses the same storage, you can add HBA ports from several hosts to one host object to make a simpler configuration. A host object can have both WWPNs and iSCSI names.

The cluster does not automatically present virtual disks (VDisks) to the host system. You must map each VDisk to a particular host object to enable the VDisk to be accessed through the WWPNs or iSCSI names that are associated with the host object.

When you create a new host object, the configuration interfaces provide a list of unconfigured WWPNs. These represent the WWPNs that the cluster has detected. Candidate iSCSI names are not available and must be entered manually.

The cluster can detect only WWPNs that have connected to the cluster through the fibre-channel network. Some fibre-channel HBA device drivers do not let the ports remain logged in if no disks are detected on the fabric. This can prevent some WWPNs from appearing in the list of candidate WWPNs. The configuration interface provides a method to manually type the port names.

Note: You must not include a WWPN or an iSCSI name that belongs to a SAN Volume Controller node in a host object.

A WWPN or iSCSI name can be added to only one host object.

Port masks

You can use the port-mask property of the host object to control the fibre-channel ports on each SAN Volume Controller node that a host can access. The port mask applies to logins from the WWPNs that are associated with the host object. The port-mask configuration has no effect on iSCSI connections.

For each login between a host fibre-channel port and node fibre-channel port, the node examines the port mask for the associated host object and determines if access is allowed or denied. If access is denied, the node responds to SCSI commands as if the HBA WWPN is unknown.

The port mask is four binary bits. Valid mask values range from 0000 (no ports enabled) to 1111 (all ports enabled). For example, a mask of 0011 enables port 1 and port 2. The default value is 1111.

Multiple target ports

When you create a VDisk-to-host mapping to a fibre-channel attached host, the host ports that are associated with the host object can view the LUN that represents the VDisk on up to eight fibre-channel ports. Nodes follow the American National Standards Institute (ANSI) Fibre Channel (FC) standards for SCSI LUs that are accessed through multiple node ports. All nodes within a single I/O group present a consistent set of SCSI LUs across all ports on those nodes.

Similarly, all nodes within a single I/O group present a consistent set of SCSI LUs across all iSCSI ports on those nodes.

Node login counts

The number of nodes that can see each WWPN or iSCSI name is reported on a per-node basis and is known as the *node login count*. If the count is less than expected for the current configuration, you might have a connectivity problem.

VDisk-to-host mapping

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

VDisk-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. VDisk-to-host mapping is done at the SAN Volume Controller level.

The act of mapping a VDisk to a host makes the VDisk accessible to the WWPNs or iSCSI names such as iSCSI qualified names (IQNs) or extended-unique identifiers (EUIs) that are configured in the host object.

VDisks and host mappings

Each host mapping associates a VDisk with a host object and allows all WWPNs and iSCSI names in the host object to access the VDisk. You can map a VDisk to multiple host objects. When a mapping is created, multiple paths might exist across the SAN fabric or Ethernet network from the hosts to the SAN Volume Controller nodes that are presenting the VDisk. Without a multipathing device driver, most operating systems present each path to a VDisk as a separate storage device. The multipathing software manages the many paths that are available to the VDisk and presents a single storage device to the operating system. If there are multiple paths, the SAN Volume Controller requires that the multipathing software run on the host.

Note: The iSCSI names and associated IP addresses for the SAN Volume Controller nodes can fail over between nodes in the I/O group, which negates the need for multipathing drivers in some configurations. Multipathing drivers are still recommended, however, to provide the highest availability.

When you map a VDisk to a host, you can optionally specify a SCSI ID for the VDisk. This ID controls the sequence in which the VDisks are presented to the host. Check the host software requirements for SCSI IDs because some require a contiguous set. For example, if you present three VDisks to the host, and those VDisks have SCSI IDs of 0, 1, and 3, the VDisk 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 lowest available SCSI ID if none is specified.

Figure 9 and Figure 10 show two VDisks, and the mappings that exist between the host objects and these VDisks.

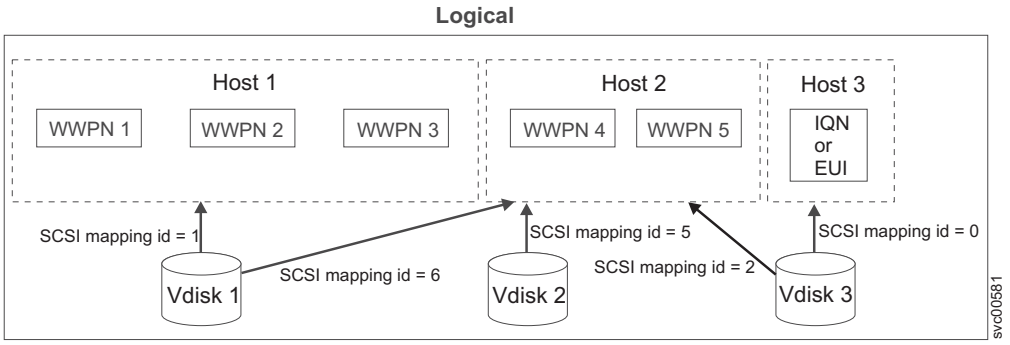

Figure 9. Hosts, WWPNs, IQNs or EUIs, and VDisks

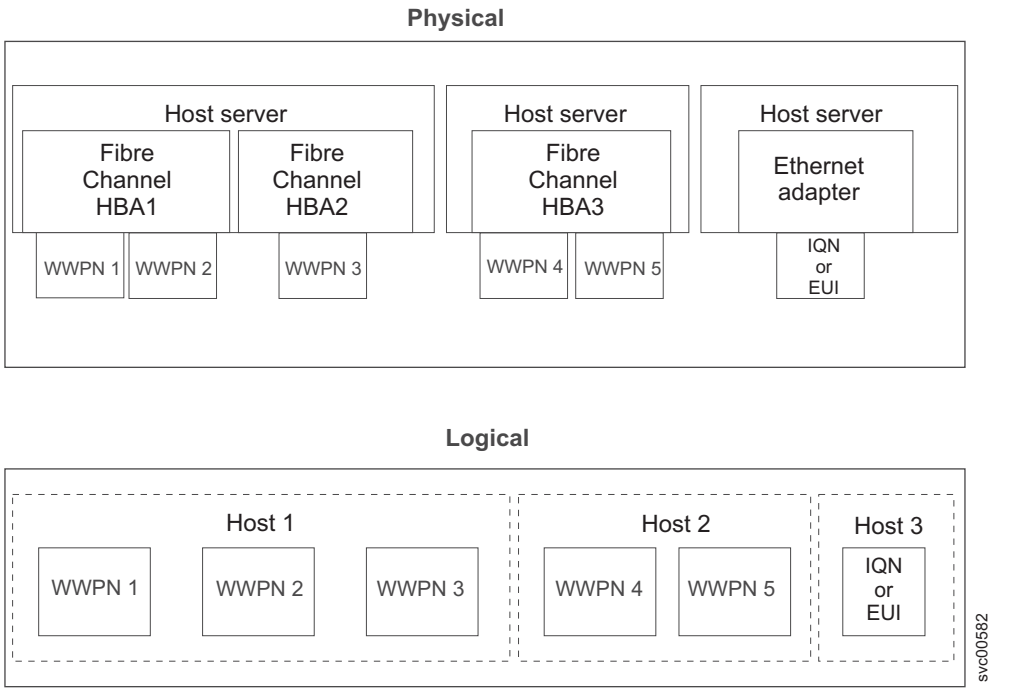

Figure 10. Hosts, WWPNs, IQNs or EUIs, VDisks, and SCSI mappings

LUN masking is usually implemented in the device driver software on each host. The host has visibility of more LUNs than it is intended to use, and device driver software masks the LUNs that are not to be used by this host. After the masking is complete, only some disks are visible to the operating system. The SAN Volume Controller can support this type of configuration by mapping all VDisks to every host object and by using operating system-specific LUN masking technology. The default, and recommended, SAN Volume Controller behavior, however, is to map to the host only those VDisks that the host requires access to.

Standard and persistent reserves

The SCSI Reserve command and the SCSI Persistent Reserve command are specified by the SCSI standards. Servers can use these commands to prevent ports in other servers from accessing the LUN.

This prevents accidental data corruption that is caused when a server overwrites data on another server. The Reserve and Persistent Reserve commands are often used by clustering software to control access to SAN Volume Controller virtual disks (VDisks).

If a server is not shut down or removed from the server cluster in a controlled way, the server reserves and persistent reserves are maintained. This prevents other servers from accessing data that is no longer in use by the server that holds the reservation. In this situation, you might want to release the reservation and allow a new server to access the VDisk.

When possible, you should have the server that holds the reservation explicitly release the reservation to ensure that the server cache is flushed and the server software is aware that access to the VDisk has been lost. In circumstances where this is not possible, you can use operating system specific tools to remove reservations. Consult the operating system documentation for details.

When you use the **svctask rmvdiskhostmap** CLI command or the SAN Volume Controller Console to remove VDisk-to-host mappings, SAN Volume Controller nodes with a software level of 4.1.0 or later can remove the server reservations and persistent reservations that the host has on the VDisk.

SAN Volume Controller maximum configuration

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

See the following Web site for the latest maximum configuration support:

www.ibm.com/storage/support/2145

SAN Volume Controller cluster high availability

A SAN Volume Controller cluster has several features that can be used to deploy a high availability storage system with no single point of failure.

Each I/O group within a cluster consists of a pair of nodes. If a node fails within an I/O group, the other node in the I/O group assumes the I/O responsibilities of the failed node. If the node contains solid-state drives (SSDs), you should create a mirrored virtual disk (VDisk) of any VDisk that uses the SSDs. SSDs can be a single point of failure in the event of an outage to the SSDs or to the node itself.

If a cluster of SAN Volume Controller nodes is split into two partitions (for example due to a SAN fabric fault), the partition with the majority of nodes continues to process I/O operations. If a cluster is split into two equal-sized partitions, a quorum disk is accessed to determine which half of the cluster continues to read and write data.

Each SAN Volume Controller node has four fibre-channel ports, which can be used to attach the node to multiple SAN fabrics. For high availability, attach the nodes in a cluster to at least two fabrics. SAN Volume Controller software incorporates multipathing software that is used for communication among SAN Volume

Controller nodes and for I/O operations among SAN Volume Controller nodes and storage systems. If a SAN fabric fault disrupts communication or I/O operations, the multipathing software recovers and retries the operation through an alternative communication path. Also for high availability, configure your fibre-channel host systems to use multipathing software. If a SAN fabric fault or node failure occurs, I/O operations among fibre-channel host systems and SAN Volume Controller nodes are retried. Subsystem device driver (SDD) multipathing software is available from IBM at no additional charge for use with SAN Volume Controller. For additional information about subsystem device driver (SDD), go to the Support for IBM Systems Web site:

www.ibm.com/systems/support

iSCSI-attached hosts connect to SAN Volume Controller through node Ethernet ports. If a node fails, SAN Volume Controller maintains host availability by failing over the IP addresses of the failed node to the partner node in the I/O group.

The SAN Volume Controller Virtual Disk Mirroring feature can be used to mirror data across storage systems. This feature provides protection against a storage system failure.

The SAN Volume Controller Metro Mirror and Global Mirror features can be used to mirror data between clusters at different physical locations for disaster recovery.

Node management and support tools

The SAN Volume Controller solution offers several management and support tools for you to maintain and manage your nodes.

The following node management tools are available with the SAN Volume Controller solution:

- Master console

  Although it can no longer be purchased, the master console can be upgraded to support clusters running the latest SAN Volume Controller software.
- To manage older levels of the SAN Volume Controller Console, install the corresponding CIM agent on the SSPC server.

Both solutions incorporate the following SAN Volume Controller applications:

- Secure Shell
- Assist On-site

IBM System Storage Productivity Center

The IBM System Storage Productivity Center (SSPC) is an integrated hardware and software solution that provides a single point of entry for managing SAN Volume Controller clusters, IBM System Storage DS8000® systems, and other components of your data storage infrastructure.

SSPC simplifies storage management in the following ways:

- Centralizing the management of storage network resources with IBM storage management software
- Providing greater synergy between storage management software and IBM storage devices

- Reducing the number of servers that are required to manage your software infrastructure
- Providing simple migration from basic device management to storage management applications that provide higher-level functions

SSPC includes the following software components:

- SAN Volume Controller Console
- PuTTY (SSH client software)
- IBM Tivoli® Storage Productivity Center Basic Edition, which can be used to access the IBM System Storage DS8000 Storage Manager and the SAN Volume Controller
- IBM DB2® Enterprise Server Edition

Figure 11 shows an overview of how SSPC and the components of IBM Tivoli Storage Productivity Center, IBM System Storage DS8000, and SAN Volume Controller interrelate with each other.

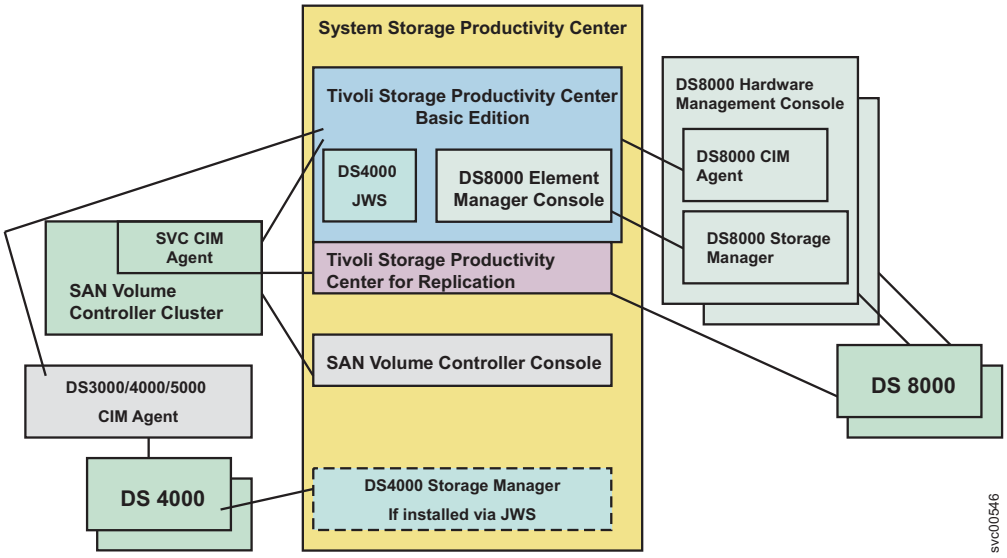

Figure 11. Overview of the IBM System Storage Productivity Center

For more information on SSPC, see the *IBM System Storage Productivity Center Introduction and Planning Guide*.

Secure Shell protocol through PuTTY

Secure Shell (SSH) software is a client-server protocol that can be used from the IBM System Storage Productivity Center or from a host server to enable you to control the SAN Volume Controller through a command-line interface (CLI).

SSH provides a secure communications channel between systems. You can configure SSH to use a key pair (a private key and a public key) to establish the secure connection to a remote system. If you want to create an SSH connection to a server using an SSH key pair (such as to the SAN Volume Controller cluster), you must place the public key on the server.

Assist On-site and remote service

When you contact IBM to help you resolve a problem with your SAN Volume Controller environment, the IBM service representative might suggest using the IBM Assist On-site tool to remotely access the IBM System Storage Productivity Center (SSPC) or master console. This type of remote service can help you reduce service costs and shorten repair times.

The IBM Assist On-site tool is a remote desktop-sharing solution that is offered through the IBM Web site. With it, the IBM service representative can remotely view your system to troubleshoot a problem. You can maintain a chat session with the IBM service representative so that you can monitor the activity and either understand how to fix the problem yourself or allow the representative to fix it for you.

To use the IBM Assist On-site tool, the SSPC or master console must be able to access the Internet. The following Web site provides further information about this tool:

www.ibm.com/support/assistonsite/

When you access the Web site, you sign in and enter a code that the IBM service representative provides to you. This code is unique to each IBM Assist On-site session. A plug-in is downloaded onto your SSPC or master console to connect you and your IBM service representative to the remote service session. The IBM Assist On-site contains several layers of security to protect your applications and your computers. You can also use security features to restrict access by the IBM service representative.

Your IBM service representative can provide you with more detailed instructions for using the tool.

Event notifications

SAN Volume Controller can use Simple Network Management Protocol (SNMP) traps, syslog messages, and Call Home e-mail to notify you and the IBM Support Center when significant events are detected. Any combination of these notification methods can be used simultaneously.

Each event that SAN Volume Controller detects is assigned a notification type of Error, Warning, or Information. You can configure SAN Volume Controller to send each type of notification to specific recipients.

Table 11 describes the types of event notifications.

Table 11. SAN Volume Controller notification types

Notification type	Description
Error	An error notification is sent to indicate a problem that must be corrected as soon as possible. This notification indicates a serious problem with the SAN Volume Controller. For example, the event that is being reported could indicate a loss of redundancy in the system, and it is possible that another failure could result in loss of access to data. The most typical reason that this type of notification is sent is because of a hardware failure, but some configuration errors or fabric errors also are included in this notification type. Error notifications can be configured to be sent as a Call Home e-mail to the IBM Support Center.
Warning	A warning notification is sent to indicate a problem or unexpected condition with the SAN Volume Controller. Always immediately investigate this type of notification to determine the effect that it might have on your operation, and make any necessary corrections. A warning notification does not require any replacement parts and therefore should not require IBM Support Center involvement. However, the event being reported might indicate a condition that could be fatal to your operating environment: such as, for example, a critical FlashCopy operation has failed.
Information	An informational notification is sent to indicate that an expected event has occurred. For example, a FlashCopy operation has completed. No remedial action is required when these notifications are sent.

SNMP traps

SNMP is a standard protocol for managing networks and exchanging messages. SAN Volume Controller can send SNMP messages that notify personnel about an event. You can use an SNMP manager to view the SNMP messages that SAN Volume Controller sends. You can use the SAN Volume Controller Console or the SAN Volume Controller command-line interface to configure and modify your SNMP settings.

You can use the Management Information Base (MIB) file for SNMP to configure a network management program to receive SNMP messages that are sent by the SAN Volume Controller. This file can be used with SNMP messages from all versions of SAN Volume Controller software. More information about the MIB file for SNMP is available at the Support for SAN Volume Controller (2145) Web site:

www.ibm.com/storage/support/2145

Search for **SAN Volume Controller MIB**. Go to the downloads results to find **Management Information Base (MIB) file for SNMP**. Click this link to find download options. The name of this file is `SVC_MIB_<release>.MIB` such as in `SVC_MIB_4.3.1.MIB`.

Syslog messages

The syslog protocol is a standard protocol for forwarding log messages from a sender to a receiver on an IP network. The IP network can be either IPv4 or IPv6. SAN Volume Controller can send syslog messages that notify personnel about an event. SAN Volume Controller can transmit syslog messages in either expanded or concise format. You can use a syslog manager to view the syslog messages that SAN Volume Controller sends. SAN Volume Controller uses the User Datagram Protocol (UDP) to transmit the syslog message. You can use the SAN Volume Controller Console or the SAN Volume Controller command-line interface to configure and modify your syslog settings.

Table 12 shows how SAN Volume Controller notification codes map to syslog security-level codes.

Table 12. SAN Volume Controller notification codes and corresponding syslog level codes

SAN Volume Controller notification code	Syslog level code	Description
SS_EID_UNKNOWN	Not mapped	
SS_EID_ERROR	LOG_ALERT	Error that needs immediate attention
SS_EID_WARNING	LOG_ERROR	Warning that needs attention
SS_EID_INFO	LOG_INFO	Informational messages
SS_EID_TEST	LOG_DEBUG	Test message

Table 13 shows how syslog facility codes map to SAN Volume Controller values of user-defined message origin identifiers.

Table 13. Syslog facility codes and SAN Volume Controller values of user-defined message origin identifiers

Syslog facility code	Syslog value	SAN Volume Controller value
LOG_LOCAL0	16	0
LOG_LOCAL1	17	1
LOG_LOCAL2	18	2
LOG_LOCAL3	19	3
LOG_LOCAL4	20	4
LOG_LOCAL5	21	5
LOG_LOCAL6	22	6
LOG_LOCAL7	23	7

Call Home e-mail

The Call Home feature transmits operational and error-related data to you and IBM through a Simple Mail Transfer Protocol (SMTP) server connection in the form of an event notification e-mail. When configured, this function alerts IBM service personnel about hardware failures and potentially serious configuration or environmental issues.

To send e-mail, you must configure at least one SMTP server. You can specify as many as five additional SMTP servers for backup purposes. The SMTP server must

accept the relaying of e-mail from the SAN Volume Controller cluster IP address. You can then use the SAN Volume Controller Console or the SAN Volume Controller command-line interface to configure the e-mail settings, including contact information and e-mail recipients. Set the reply address to a valid e-mail address. Send a test e-mail to check that all connections and infrastructure are set up correctly. You can disable the Call Home function at any time using the SAN Volume Controller Console or the SAN Volume Controller command-line interface.

Call Home and inventory e-mail information

SAN Volume Controller can use Call Home e-mail and inventory information e-mail to provide data and event notifications to you and to the IBM Support Center.

Call Home e-mail

Call Home support is initiated for the following reasons or types of data:

- Event notifications: Data is sent to the specified e-mail address when events occur. There are three types of events–error, warning, and information. All these types of events can cause an e-mail notification to be sent, depending on how e-mail settings are configured.
- Communication tests: You can test for the successful installation and communication infrastructure.
- Inventory information: A notification is sent to provide the necessary status and hardware information to IBM service personnel.

To send data and notifications to IBM service personnel, use one of the following e-mail addresses:

- For SAN Volume Controller nodes located in North America, Latin America, South America or the Caribbean Islands, use `callhome1@de.ibm.com`
- For SAN Volume Controller nodes located anywhere else in the world, use `callhome0@de.ibm.com`

Call Home e-mail can contain any combination of the following types of information:

- Contact name
- Contact phone number
- Offshift phone number
- Contact e-mail
- Machine location
- Record type
- Machine type
- Machine serial number
- Error ID
- Error code
- Software version
- FRU part number
- Cluster name
- Node ID
- Error sequence number
- Time stamp

- Object type
- Object ID
- Problem data

Inventory information e-mail

Inventory information e-mail is a type of Call Home notification. Inventory information can be sent to IBM to assist IBM service personnel in evaluating your SAN Volume Controller system. Because inventory information is sent using the Call Home e-mail function, you must meet the Call Home function requirements and enable the Call Home e-mail function before you can attempt to send inventory information e-mail. You can adjust the contact information, adjust the frequency of inventory e-mail, or manually send an inventory e-mail using the SAN Volume Controller Console or the SAN Volume Controller command-line interface. Inventory information is automatically reported to IBM when you activate error reporting.

Inventory information that is sent to IBM includes the following information about the cluster on which the Call Home function is enabled:

- The output from the **svcinfo lscluster** command
- The output from the **svcinfo lsnodevpd** command (once for each node).
- The output from the **svcinfo lslicense** command

User roles

Each user of the SAN Volume Controller Console must provide a user name and a password to sign on. Each user also has an associated role such as monitor, copy operator, service, administrator, or security administrator. These roles are defined at the cluster level. For example, a user can perform the administrator role for one cluster and perform the service role for another cluster.

Monitor
Users with the monitor role have access to all viewing actions available with the SAN Volume Controller Console. This user cannot perform any actions that change the state of the cluster or the resources that the cluster manages. The user can access all the information-related panels and commands, back up configuration data, change his or her password, and issue the following commands: finderr, dumperrlog, dumpinternallog, ping, and chcurrentuser.

Copy Operator
Users with the copy operator role can manage all existing FlashCopy, Metro Mirror, and Global Mirror relationships. They can also create and delete FlashCopy mappings, FlashCopy consistency groups, Metro Mirror or Global Mirror relationships, and Metro Mirror and Global Mirror consistency groups. In addition, the user can access all the functions available to the Monitor role.

Service
Users with the service role can view the View Clusters panel, launch the SAN Volume Controller Console, and view the progress of actions on clusters with the View Progress panel, begin disk discovery process, and discover and include disks. The user can access the following commands: applysoftware, setlocale, addnode, rmnode, cherrstate, setevent, writesernum, detectmdisk, and includemdisk. A user with this role can also access all the functions available to the Monitor role.

Administrator
Users with the administrator role can access all functions on the SAN Volume Controller Console and issue any command-line interface (CLI) command, except those that deal with managing users, user groups, and authentication.

Security Administrator
Users with the security administrator role can access all functions on the SAN Volume Controller Console and issue any CLI command. Users with this role can also manage users, user groups, and manage user authentication.

Configuring user authentication

You can configure authentication and authorization for users of the SAN Volume Controller cluster.

You can create two types of users who access the cluster. These types are based on how the users are authenticated to the cluster. Local users must provide either a password, a Secure Shell (SSH) key, or both. Local users are authenticated through the authentication methods that are located on the SAN Volume Controller cluster. If the local user needs access to SAN Volume Controller Console, a password is needed for the user. If the user requires access to the command-line interface, a valid SSH key file is necessary. If a user is working with both interfaces, both a password and SSH key are required. Local users must be part of a user group that is defined on the cluster. User groups define roles that authorize the users within that group to a specific set of operations on the cluster.

A remote user is authenticated on a remote service usually provided by a SAN management application, such as IBM Tivoli Storage Productivity Center. Remote users require no local credentials to access the SAN Volume Controller Console. Remote users have their groups defined by the remote authentication service. If a remote user needs to use the command-line interface, both a password and SSH key are required. If the remote authentication service fails, then remote users cannot access the SAN Volume Controller Console or the command-line interface. In this situation, a local user with the Security Administrator role must change remote users to local users by adding them to the appropriate user group. After logging in to a SAN Volume Controller application, a remote user is granted access to the SAN Volume Controller CLI and Console by default.

Chapter 2. Copy Services features

The SAN Volume Controller provides Copy Services features that enable you to copy virtual disks (VDisks).

The following Copy Services features are available for all supported hosts that are connected to the SAN Volume Controller:

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

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

Global Mirror
Provides a consistent copy of a source VDisk on a target VDisk. Data is written to the target VDisk asynchronously, so that the copy is continuously updated, but the copy might not contain the last few updates in the event that a disaster recovery operation is performed.

FlashCopy

FlashCopy is a Copy Services feature that is available with the SAN Volume Controller.

In its basic mode, the FlashCopy feature copies the contents of a source virtual disk (VDisk) to a target VDisk. Any data that existed on the target VDisk is lost and is replaced by the copied data. After the copy operation has completed, the target VDisks contain the contents of the source VDisks as they existed at a single point in time unless target writes have been performed. The FlashCopy feature 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 to complete, the resulting data on the target VDisk is presented so that the copy appears to have occurred immediately.

Although it is difficult to make a consistent copy of a data set that is constantly updated, point-in-time copy techniques 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, the resulting copy might contain data that 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 copied, the copy contains the referenced object at its new location but the copied reference still points to the old location.

More advanced FlashCopy features allow operations to occur on multiple source and target VDisks. FlashCopy management operations are coordinated to allow a common single point in time for copying target VDisks from their respective source VDisks. This allows a consistent copy of data that spans multiple VDisks. The FlashCopy feature also allows multiple target VDisks to be copied from each source VDisk. This can be used to create images from different points in time for each source VDisk.

FlashCopy associates a source VDisk and a target VDisk in a FlashCopy Mapping. The source VDisks and target VDisks must meet the following requirements:

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

The Cascaded FlashCopy feature allows a FlashCopy target VDisk to be the source VDisk of another FlashCopy mapping.

The incremental FlashCopy feature reduces the amount of time that is required to copy the source VDisk for multiple FlashCopy mappings. The initial FlashCopy mapping copies all of the data from the source VDisk to the target VDisk. Subsequent FlashCopy mappings only copy data that has been modified since the initial FlashCopy mapping. You can define a FlashCopy mapping as incremental only when you create the FlashCopy mapping.

Note: For incremental FlashCopy support information, see the release-specific *IBM System Storage SAN Volume Controller Restrictions* technical note at the Support for SAN Volume Controller (2145) Web site: www.ibm.com/storage/support/2145

The multi-target reverse FlashCopy feature allows you to start a FlashCopy mapping where the target VDisk is the source VDisk in a second FlashCopy mapping. You can use this feature to reverse a FlashCopy mapping direction without having to remove existing mappings, and without losing data from the target VDisk.

A FlashCopy mapping can also be created to mirror an existing mapping; these paired mappings are called *partners*. A mapping can have only one partnership. For example, VDisks A and B can have two mappings; A > B and B > A.

Any VDisk that is part of a FlashCopy operation can be space-efficient. Using a space-efficient VDisk as a FlashCopy target and setting the background FlashCopy rate to 0 (nocopy) can reduce the amount of storage that is required to maintain a point-in-time copy. The source VDisks and target VDisks can also be mirrored to improve availability of the VDisks.

FlashCopy applications

You can use the FlashCopy feature to create consistent backups of dynamic data, test applications, and create copies for auditing purposes and for data mining.

To create consistent backups of dynamic data, use the FlashCopy feature 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 become redundant and can now be discarded. Usually in this backup condition, the target data can be managed as read-only.

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.

You can also use the FlashCopy feature to create restart points for long running batch jobs. This means that if a batch job fails several days into its run, it might be possible to restart the job from a saved copy of its data rather than rerunning the entire multiday job.

Host considerations for FlashCopy integrity

The SAN Volume Controller FlashCopy feature transfers a point-in-time copy of a source virtual disk (VDisk) onto a designated target VDisk. You must create or already have an existing target VDisk before you can transfer the copy. You must also ensure the target VDisk has enough space available to support the amount of data that is being transferred.

After the mapping is started, all of the data that is stored on the source VDisk can be accessed through the target VDisk. This includes any operating system control information, application data, and metadata that was stored on the source VDisk. Because of this, some operating systems do not allow a source VDisk and a target VDisk to be addressable on the same host.

In order to ensure the integrity of the copy that is made, it is necessary to completely flush the host cache of any outstanding reads or writes before you proceed with the FlashCopy operation. You can flush the host cache by unmounting the source VDisks from the source host before you start the FlashCopy operation.

Because the target VDisks are overwritten with a complete image of the source VDisks, it is important that any data held on the host operating system (or application) caches for the target VDisks 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 VDisks prior to starting the FlashCopy operation.

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, they can be used to prepare and start a FlashCopy operation. See your host and application documentation for details.

Some operating systems are unable to use a copy of a VDisk without *synthesis*. Synthesis performs a transformation of the operating system metadata on the target VDisk to allow the operating system to use the disk. See your host documentation on how to detect and mount the copied VDisks.

Flushing data from the host volumes

All outstanding read and write operations must be flushed from the host cache before you use the FlashCopy feature.

Perform the following steps to flush data from your host volumes and start a FlashCopy operation:

1. If you are using UNIX® or Linux® operating systems, perform the following steps:
   a. Quiesce all applications to the source volumes that you want to copy.
   b. Use the **umount** command to unmount the designated drives.
   c. Prepare and start the FlashCopy operation for those unmounted drives.
   d. Remount your volumes with the mount command and resume your applications.

2. If you are using the Microsoft® Windows® operating system using drive letter changes, perform the following steps:
   a. Quiesce all applications to the source volumes that you want to copy.
   b. Go into your disk management window and remove the drive letter on each drive that you want to copy. This unmounts the drive.
   c. Prepare and start the FlashCopy operation for those unmounted drives.
   d. Remount 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 that you want to copy.
   b. Issue the **chkdsk /x** command on each drive you want to copy. The /x option unmounts, scans, and remounts the volume.
   c. Ensure that all applications to the source volumes are still quiesced.
   d. Prepare and start the FlashCopy operation for those unmounted drives.

   Note: If you can ensure that no reads and writes are issued to the source volumes after you unmount the drives, you can immediately remount and then start the FlashCopy operation.

FlashCopy mappings

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

The FlashCopy feature makes an instant copy of a VDisk at the time that it is started. To create an instant copy of a VDisk, 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 VDisks must be of equal size.

A mapping can be created between any two VDisks in a cluster. The VDisks do not have to be in the same I/O group or managed disk (MDisk) group. When a FlashCopy operation starts, 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. Each 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, 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.

During a read operation to the target VDisk, the bitmap is used to determine if 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 also specify a clean rate. The clean rate is used to control the rate that data is copied from the target VDisk of the mapping to the target VDisk of a mapping that is either the latest copy of the target VDisk, or is the next oldest copy of the source VDisk. The clean rate is used in the following situations:

- The mapping is in the stopping state
- The mapping is in the copying state and has a copy rate of zero

- The mapping is in the copying state and the background copy has completed

You can use the clean rate to minimize the amount of time that a mapping is in the stopping state. If the mapping has not completed, the target VDisk is offline while the mapping is stopping. The target VDisk remains offline until the mapping is restarted.

When you create a mapping, you specify a copy rate. When the mapping is in the copying state, the copy rate determines the priority that is given to the background copy process. If you want a copy of the whole source VDisk so that a mapping can be deleted and still be accessed from the target VDisk, you must copy all the data that is on the source VDisk to the target VDisk.

The default values for both the clean rate and the copy rate is 50.

When a mapping is started and the copy rate is greater than zero (or a value other than `NOCOPY`), the unchanged data is copied to the target VDisk, and the bitmap is updated to show that the copy has occurred. After a time, the length of which depends on the priority that was determined by the copy rate and the size of the VDisk, the whole VDisk is copied to the target. The mapping returns to the idle_or_copied state and you can now restart the mapping at any time to create a new copy at the target.

While the mapping is in the copying state, you can set the copy rate to zero and the clean rate to a value other than zero to minimize the amount of time a mapping is in the stopping state.

If you use multiple target mappings, the mapping can stay in the copying state after all of the source data is copied to the target (the progress is 100%). This situation can occur if mappings that were started earlier and use the same source disk are not yet 100% copied.

If the 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 require a temporary copy of the source.

You can stop the mapping at any time after it has been started. Unless the target VDisk already contains a complete copy of the source VDisk, this action makes the target inconsistent and the target VDisk is taken offline. The target VDisk remains offline until the mapping is restarted.

You can also set the autodelete attribute. If this attribute is set to on, the mapping is automatically deleted when the mapping reaches the idle_or_copied state and the progress is 100%.

FlashCopy mapping states

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

Idle or copied

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

If the mapping is incremental and the background copy is complete, the mapping only records the differences between the source and target

VDisks. If the connection to both nodes in the I/O group that the mapping is assigned to is lost, the source and target VDisks will be offline.

Copying
The copy is in progress. Read and write caching is enabled on the source and the target VDisks.

Prepared
The mapping is ready to start. The target VDisk is online, but is not accessible. The target VDisk cannot perform read or write caching. Read and write caching is failed by the SCSI front-end as a hardware error. If the mapping is incremental and a previous mapping has completed, the mapping only records the differences between the source and target VDisks. If the connection to both nodes in the I/O group that the mapping is assigned to is lost, the source and target VDisks will be offline.

Preparing
The target VDisk is online, but not accessible. The target VDisk cannot perform read or write caching. Read and write caching is failed by the SCSI front-end as a hardware error. 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. If the mapping is incremental and a previous mapping has completed, the mapping records only the differences between the source and target VDisks. If the connection to both nodes in the I/O group that the mapping is assigned to is lost, the source and target VDisks will be offline.

Stopped
The mapping is stopped because either you issued a stop command or an I/O error occurred. The target VDisk is offline and its data is lost. To access the target VDisk, you must restart or delete the mapping. The source VDisk is accessible and the read and write cache is enabled. If the mapping is incremental, the mapping is recording write operations to the source VDisk. If the connection to both nodes in the I/O group that the mapping is assigned to is lost, the source and target VDisks will be offline.

Stopping
The mapping is in the process of copying data to another mapping.

- If the background copy process is complete, the target VDisk is online while the stopping copy process completes.
- If the background copy process is not complete, data is discarded from the target VDisk cache. The target VDisk is offline while the stopping copy process runs.

The source VDisk is accessible for I/O operations.

Suspended
The mapping started, but it did not complete. Access to the metadata is lost, which causes both the source and target VDisk to go offline. When access to the metadata is restored, the mapping returns to the copying or stopping state and the source and target VDisks return online. The background copy process resumes. Any data that has not been flushed and has been written to the source or target VDisk before the suspension, is in cache until the mapping leaves the suspended state.

Notes:

1. If a FlashCopy source VDisk goes offline, any FlashCopy target VDisks that depend on that VDisk also go offline.

2. If a FlashCopy target VDisk goes offline, any FlashCopy target VDisks that depend on that VDisk also go offline. The source VDisk remains online.

Before you start the mapping, you must prepare it. Preparing the mapping ensures that the data in the cache is de-staged to disk and 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 Controller nodes; it passes straight through to the MDisks. The prepare operation for the mapping might take some time to complete; 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 and start the mapping.

Note: The svctask startfcmap and svctask startfcconsistgrp commands can take some time to process.

If you do not want to use consistency groups, the SAN Volume Controller allows a mapping to be treated as an independent entity. In this case, the mapping is known as a stand-alone mapping. For mappings that have been configured in this way, use the svctask prestartfcmap and svctask startfcmap commands rather than the svctask prestartfcconsistgrp and svctask startfcconsistgrp commands.

Multiple target FlashCopy mappings

You can copy up to 256 target VDisks from a single source VDisk. Each relationship between a source and target VDisk is managed by a unique mapping such that a single VDisk can be the source VDisk in up to 256 mappings.

Each of the mappings from a single source can be started and stopped independently. If multiple mappings from the same source are active (in the copying or stopping states), a dependency exists between these mappings.

Example 1

Mapping A depends on mapping B if the following is true:

- Mapping A and mapping B both have the same source VDisk
- Mapping A and mapping B are both in the copying or stopping state
- Mapping B was started more recently than mapping A

  Note: If both mappings were in the same consistency group and therefore started at the same time, the order of dependency is decided internally when the consistency group is started.
- Mapping A does not have a complete copy of the source because the copying progress for the mapping is less than 100.
- A mapping does not exist from the same source started more recently than A and later than B which has a complete copy of the source because the copying progress of the mapping is less than 100.

Example 2

Target VDisk A depends on target VDisk B if the mapping that VDisk A belongs to depends on the mapping that target VDisk B belongs to. The target VDisk of the

most recently started mapping from the source VDisk depends on the source VDisk until there is a complete copy of the source (progress is 100%).

Incremental FlashCopy mappings

With incremental FlashCopy mappings, the background copy process copies only the parts of the source or target VDisk that have changed since the last FlashCopy process. This reduces the amount of time that it takes to recreate an independent FlashCopy image.

Note: For incremental FlashCopy support information, see the release-specific *IBM System Storage SAN Volume Controller Restrictions* technical note at the Support for SAN Volume Controller (2145) Web site: www.ibm.com/storage/support/2145

Cascaded FlashCopy mappings

With cascaded FlashCopy mappings, target VDisks can be the source of other mappings.

Up to 256 mappings can exist in a cascade. If cascaded mappings and multiple target mappings are used, a tree of up to 256 mappings can be created.

FlashCopy mapping restoration

You can start a mapping with a target VDisk that is the source VDisk of another active mapping that is in either the idle_copied, stopped, or copying states. If the mapping is in the copying state, the **restore** parameter is required for the svctask startfcmap and svctask prestartfcmap commands. You can restore the contents of a FlashCopy source VDisk by using the target of the same FlashCopy mapping or a different FlashCopy mapping without waiting for the mapping to become idle and without loss of the contents of any other FlashCopy target VDisk.

FlashCopy partner mappings

You can create a mapping to mirror an existing FlashCopy mapping. The resulting pair of mappings are called *partners*. A mapping can have only one partner. For example, if you have VDisk A and VDisk B with two mappings (Mapping 0 from VDisk A to VDisk B and Mapping 1 from VDisk B to VDisk A), Mapping 0 and Mapping 1 are partners.

Incremental FlashCopy mappings share the metadata for recording changes. Therefore, if one mapping in a mirrored pair (partnership) is incremental, the other mapping becomes incremental automatically and remains incremental until it is deleted.

A cluster that is running SAN Volume Controller version 4.3.x can have FlashCopy mappings that mirror each other. If such a cluster is upgraded to SAN Volume Controller version 5.1.0, these mappings become partners. A nonincremental FlashCopy mapping that becomes the partner of an incremental FlashCopy mapping during the upgrade process becomes incremental. If two mirrored FlashCopy mappings are incremental before the upgrade, the resulting partners become incremental and the amount of metadata that is used is reduced. Only one incremental bitmap is required for the pair.

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 provide independent access to both.

FlashCopy mapping events

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

Table 14 provides a description of each FlashCopy mapping event.

Table 14. FlashCopy mapping events

Create	A new FlashCopy mapping is created between the specified source virtual disk (VDisk) and the specified target VDisk. The operation fails if any of the following is true: • For SAN Volume Controller software version 4.1.0 or earlier, the source or target VDisk is already a member of a FlashCopy mapping. • For SAN Volume Controller software version 4.2.0 or later, the source or target VDisk is already a target VDisk of a FlashCopy mapping. • For SAN Volume Controller software version 4.2.0 or later, the source VDisk is already a member of 16 FlashCopy mappings. • For SAN Volume Controller software version 4.3.0 or later, the source VDisk is already a member of 256 FlashCopy mappings. • The node has insufficient bitmap memory. • The source and target VDisks are different sizes.
Prepare	The prepare command is directed to either a consistency group for FlashCopy mappings that are members of a normal consistency group or to the mapping name for FlashCopy mappings that are stand-alone mappings. The prepare command places the FlashCopy mapping into the preparing state. **Attention:** The prepare command can corrupt any data that previously resided on the target VDisk because cached writes are discarded. Even if the FlashCopy mapping is never started, the data from the target might have logically changed during the act of preparing to start the FlashCopy mapping.
Flush done	The FlashCopy mapping automatically moves from the preparing state to the prepared state after all cached data for the source is flushed and all cached data for the target is no longer valid.

Table 14. FlashCopy mapping events (continued)

Start	When all the FlashCopy mappings in a consistency group are in the prepared state, the FlashCopy mappings can be started. To preserve the cross volume consistency group, the start of all of the FlashCopy mappings in the consistency group must be synchronized correctly with respect to I/Os that are directed at the VDisks. This is achieved during the start command. The following occurs during the **start** command: • New reads and writes to all source VDisks in the consistency group are paused in the cache layer until all ongoing reads and writes below the cache layer are completed. • After all FlashCopy mappings in the consistency group are paused, the internal cluster state is set to allow FlashCopy operations. • After the cluster state is set for all FlashCopy mappings in the consistency group, read and write operations are unpaused on the source VDisks. • The target VDisks are brought online. As part of the **start** command, read and write caching is enabled for both the source and target VDisks.
Modify	The following FlashCopy mapping properties can be modified: • FlashCopy mapping name • Clean rate • Consistency group • Copy rate (for background copy or stopping copy priority) • Automatic deletion of the mapping when the background copy is complete
Stop	There are two separate mechanisms by which a FlashCopy mapping can be stopped: • You have issued a command • An I/O error has occurred
Delete	This command requests that the specified FlashCopy mapping is deleted. If the FlashCopy mapping is in the stopped state, the force flag must be used.
Flush failed	If the flush of data from the cache cannot be completed, the FlashCopy mapping enters the stopped state.
Copy complete	After all of the source data has been copied to the target and there are no dependent mappings, the state is set to copied. If the option to automatically delete the mapping after the background copy completes is specified, the FlashCopy mapping is automatically deleted. If this option is not specified, the FlashCopy mapping is not automatically deleted and can be reactivated by preparing and starting again.
Bitmap Online/Offline	The node has failed.

Space-efficient FlashCopy

You can have a mix of space-efficient and fully allocated VDisks in FlashCopy mappings. One common combination is a fully allocated source with a space-efficient target, which allows the target to consume a smaller amount of real storage than the source.

For best performance, the grain size of the space-efficient VDisk must match the grain size of the FlashCopy mapping. However, if the grain sizes are different, the mapping still proceeds.

Consider the following information when you create your FlashCopy mappings:

- If you are using a fully allocated source with a space-efficient target, disable background copy and cleaning mode on the FlashCopy map by setting both the background copy rate and cleaning rate to zero. Otherwise, if these features are enabled, all the source is copied onto the target VDisk. This causes the space-efficient VDisk to either go offline or to grow as large as the source.
- If you are using only space-efficient source, only the space that is used on the source VDisk is copied to the target VDisk. For example, if the source VDisk has a virtual size of 800 GB and a real size of 100 GB of which 50 GB have been used, only the used 50 GB are copied.
- A FlashCopy bitmap contains one bit for every grain on a VDisk. For example, if you have a space-efficient VDisk with 1 TB virtual size (100 MB real capacity), you must have a FlashCopy bitmap to cover the 1 TB virtual size even though only 100 MB of real capacity is allocated.

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 start that group instead of the individual mappings. This ensures that a consistent copy is made of all the source virtual disks (VDisks). Mappings to control at an individual level are known as stand-alone mappings. Do not place stand-alone mappings into a consistency group because they become controlled as part of that consistency group.

When you copy data from one VDisk to another, the data might not include all that you need to enable you to use the copy. Many applications have data that spans multiple VDisks and requires 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 operations 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.

You can set the autodelete attribute for FlashCopy consistency groups. If this attribute is set to on, the consistency group is automatically deleted when the last mapping in the group is deleted or moved out of the consistency group.

Multiple target FlashCopy mappings

Consistency groups aggregate FlashCopy mappings, not the VDisks themselves. Therefore, a source VDisk with multiple FlashCopy mappings can be in different consistency groups. If a VDisk is the source VDisk for several FlashCopy mappings that are in the same consistency group, multiple identical copies of the source VDisk are created when the consistency group is started.

Cascaded FlashCopy mappings

To create a FlashCopy mapping in a consistency group, the source VDisk cannot be the target of a mapping in the same consistency group. In addition, the target VDisk cannot be the source of another FlashCopy mapping in the same consistency group. You cannot move a FlashCopy mapping into a consistency group that contains similar FlashCopy mappings in the cascade.

FlashCopy consistency group states

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

Idle_or_Copied
All FlashCopy Mappings in this consistency group are in the Idle or Copied state.

Preparing
At least one FlashCopy mapping in this consistency group is in the Preparing state.

Prepared
The consistency group is ready to start. While in this state, the target VDisks of all FlashCopy mappings in this consistency group are not accessible.

Copying
At least one FlashCopy mapping in the consistency group is in the Copying state and no FlashCopy mappings are in the Suspended state.

Stopping
At least one FlashCopy mapping in the consistency group is in the Stopping state and no FlashCopy mappings are in the Copying or Suspended state.

Stopped
The consistency group is stopped because either you issued a command or an I/O error occurred.

Suspended
At least one FlashCopy mapping in the consistency group is in the Suspended state.

Empty The consistency group does not have any FlashCopy mappings.

Dependent writes

To preserve the integrity of data that is being written, ensure that dependent writes are run in the intended sequence of the application.

The following list is a typical sequence of write operations for a database update transaction:

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

The database ensures correct ordering of these writes by waiting for each step to complete before starting the next. However, if the database log (updates 1 and 3)

and the database itself (update 2) are on different virtual disks (VDisks) and a FlashCopy mapping is started during this update, the possibility that the database itself is copied slightly before the database log resulting in the target VDisks seeing writes (1) and (3) but not (2) must be excluded. In this case, if the database is restarted from a backup made from the FlashCopy target disks, the database log indicates that the transaction has completed successfully when, in fact, that is not the case. The transaction is lost and the integrity of the database is compromised.

You can perform a FlashCopy operation on multiple VDisks as an atomic operation to create a consistent image of user data. To use FlashCopy this way, the SAN Volume Controller supports the concept of a consistency group. A consistency group can contain an arbitrary number of FlashCopy mappings, up to the maximum number of FlashCopy mappings that are supported by a SAN Volume Controller cluster. You can use the command-line interface (CLI) **svctask startfcconsistgrp** command to start the point-in-time copy for the entire consistency group. All of the FlashCopy mappings in the consistency group are started at the same time, resulting in a point-in-time copy that is consistent across all of the FlashCopy mappings that are contained in the consistency group.

See the following Web site for the latest maximum configuration support:

www.ibm.com/storage/support/2145

Grains and the FlashCopy bitmap

When data is copied between virtual disks (VDisks), it is copied in units of address space known as *grains*.

The grain size is 64 KB or 256 KB. The FlashCopy bitmap contains one bit for each grain. The bit records whether the associated grain has been split by copying the grain from the source to the target.

Write to target VDisk

A write to the newest target VDisk must consider the state of the grain for its own mapping and the grain of the next oldest mapping.

- If the grain of the intermediate mapping or the next oldest mapping has not been copied, it must be copied before the write is allowed to proceed. This is done to preserve the contents of the next oldest mapping. The data written to the next oldest mapping can come from a target or source.
- If the grain of the target that is being written has not been copied, the grain is copied from the oldest already copied grain in the mappings that are newer than the target (or the source if no targets are already copied). After the copy is complete, the write can be applied to the target.

Read to target VDisk

If the grain that is being read has been split, the read returns data from the target that is being read. If the read is to an uncopied grain on an intermediate target VDisk, each of the newer mappings are examined to determine if the grain has been split. The read is surfaced from the first split grain found or from the source VDisk if none of the newer mappings have a split grain.

FlashCopy indirection layer

The FlashCopy feature provides the semantics of a point-in-time copy by using an indirection layer which intercepts I/Os that are targeted at both the source and target virtual disks (VDisks).

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 that are in the consistency group.

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

- The VDisk and LBA to which the I/O is addressed,
- Its direction (read or write)
- The state of an internal data structure, the FlashCopy bitmap.

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

Table 15 provides an overview of the FlashCopy I/O path actions.

Table 15. FlashCopy I/O path actions

VDisk	Grain already copied?	Host I/O operations	
		Read	**Write**
Source	No	Read from source	Copies the grain to the most recently started target VDisk for this source VDisk and then writes to the source VDisk.
	Yes	Read from source	Write to source

Table 15. FlashCopy I/O path actions (continued)

VDisk	Grain already copied?	Host I/O operations	
		Read	**Write**
Target	No	When the grain has not been copied, you can use the following algorithm to determine the VDisk that is being read: 1. If newer target VDisks exist for the source VDisk of the FlashCopy mapping that is the target VDisk being written to and the data has already been copied, the read operation is from the target VDisk. 2. If the source VDisk is not a target of another FlashCopy mapping, the read operation is from the source VDisk. 3. If newer target VDisks exist for the source VDisk of the FlashCopy mapping that is the source VDisk being written to and the data has already been copied, the read operation is from the target VDisk.	When the grain has not been copied or overwritten, you can use the following algorithm: 1. Use the algorithm for the corresponding read to determine the VDisk to read. 2. If there is an older target VDisk and the data has not been copied to this VDisk, the data is written to this VDisk. 3. If there is a target VDisk for this VDisk and the data has not been copied to this VDisk, the data is written to this VDisk. 4. Writes to target.
	Yes	Read from target When the grain has been copied, you can use the following algorithm to determine the VDisk that is being read: 1. If newer target VDisks exist for this source VDisk and the grain has already been copied, the read comes from the oldest target VDisk. 2. If there are no newer target VDisks, the read comes from the source VDisk.	Write to target 1. If newer target VDisks exist for this source VDisk and the grain has already been copied, the read operation comes from the oldest target VDisk. If there are no newer target VDisks, the read operation comes from the source VDisk. 2. If the grain has not already been copied to the next oldest target VDisk for this source VDisk, the same data is also copied to the next oldest target VDisk. 3. Writes to target

Note: For cascaded FlashCopy operations, a VDisk can be both the source and the target. When the VDisk is both the source and target, the I/O path actions are handled as described for a target VDisk.

Source read operations

Source read operations are always passed through to the underlying source VDisk.

Target read operations

To process a read operation from the target VDisk, the FlashCopy mapping must consult the FlashCopy bitmap. If the data has already been copied to the target VDisk, the read operation is sent to the target VDisk. If the data has not already been copied, the target read operation is either sent to the source VDisk, or to another target VDisk if multiple target FlashCopy mappings exist for the source VDisk. While the target read operation is outstanding, no write operations that change the data that is being read are allowed to run.

Background copy and cleaning rates

FlashCopy mapping copy *rate* values can be between 1 and 100 and can be changed when the FlashCopy mapping is in any state.

If NOCOPY is specified, background copy is disabled. You can specify NOCOPY for short-lived FlashCopy mappings that are only used for backups, for example. Because the source data set is not expected to significantly change during the lifetime of the FlashCopy mapping, it is more efficient in terms of managed disk (MDisk) I/Os to not perform a background copy.

Note: For the command-line interface (CLI), the value NOCOPY is equivalent to setting the copy rate to 0 (zero).

Table 16 provides the relationship of the copy and cleaning *rate* values to the attempted number of grains to be split per second. A grain is the unit of data represented by a single bit.

Table 16. Relationship between the rate, data rate and grains per second values

User-specified *rate* **attribute value**	**Data copied/sec**	**256 KB grains/sec**	**64 KB grains/sec**
1 - 10	128 KB	0.5	2
11 - 20	256 KB	1	4
21 - 30	512 KB	2	8
31 - 40	1 MB	4	16
41 - 50	2 MB	8	32
51 - 60	4 MB	16	64
61 - 70	8 MB	32	128
71 - 80	16 MB	64	256
81 - 90	32 MB	128	512
91 - 100	64 MB	256	1024

The data copied/sec and the grains/sec numbers represent standards that the SAN Volume Controller tries to achieve. The SAN Volume Controller is unable to achieve these standards if insufficient bandwidth is available from the nodes to the physical disks that make up the managed disks (MDisks) after taking into account the requirements of foreground I/O. If this situation occurs, background copy I/O contends for resources on an equal basis with I/O that arrives from hosts. Both tend to see an increase in latency and consequential reduction in throughput with respect to the situation had the bandwidth not been limited.

Degradation runs smoothly. Background copy, stopping copy, and foreground I/O continue to make forward progress and do not stop, hang or cause the node to fail.

The background copy is performed by one of the nodes that belong to the I/O group in which the source VDisk resides. This responsibility is moved to the other node in the I/O group in the event of the failure of the node that performs the background and stopping copy.

The background copy starts with the grain that contains the highest logical block numbers (LBAs) and works in reverse towards the grain that contains LBA 0. The

background copy is performed in reverse to avoid any unwanted interactions with sequential write streams from the application.

The stopping copy operation copies every grain that is split on the stopping map to the next map (if one exists) which is dependent on that grain. The operation starts searching with the grain that contains the highest LBAs and works in reverse towards the grain that contains LBA 0. Only those grains that other maps are dependent upon are copied.

Cleaning mode

When you create or modify a FlashCopy mapping, you can specify a cleaning rate for the FlashCopy mapping that is independent of the background copy rate. The cleaning rates shown in Table 16 on page 60 controls the rate at which the cleaning process operates. The cleaning process copies data from the target VDisk of a mapping to the target VDisks of other mappings that are dependent on this data. The cleaning process must complete before the FlashCopy mapping can go to the stopping state.

Cleaning mode allows you to activate the cleaning process when the FlashCopy mapping is in the copying state. This keeps your target VDisk accessible while the cleaning process is running. When operating in this mode, it is possible that host I/O operations can prevent the cleaning process from reaching 100% if the I/O operations continue to copy new data to the target VDisks. However, it is possible to minimize the amount of data that requires cleaning while the mapping is stopping.

Cleaning mode is active if the background copy progress has reached 100% and the mapping is in the copying state, or if the background copy rate is set to 0.

Metro Mirror and Global Mirror

The Metro Mirror and Global Mirror Copy Services features enable you to set up a relationship between two virtual disks (VDisks), so that updates that are made by an application to one VDisk are mirrored on the other VDisk. The VDisks can be in the same cluster or on two different clusters.

Although the application only writes to a single VDisk, the SAN Volume Controller maintains two copies of the data. If the copies are separated by a significant distance, the Metro Mirror and Global Mirror copies can be used as a backup for disaster recovery. A prerequisite for the SAN Volume Controller Metro Mirror and Global Mirror operations between clusters is that the SAN fabric to which they are attached provides adequate bandwidth between the clusters.

For both Metro Mirror and Global Mirror copy types, 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 I/O operations to the secondary VDisk.

The Metro Mirror feature provides a *synchronous*-copy process. When a host writes to the primary VDisk, it does not receive confirmation of I/O completion until the write operation has completed for the copy on both the primary VDisk and the secondary VDisk. This ensures that the secondary VDisk is always up-to-date with the primary VDisk in the event that a failover operation must be performed.

However, the host is limited to the latency and bandwidth limitations of the communication link to the secondary VDisk.

The Global Mirror feature provides an *asynchronous*-copy process. When a host writes to the primary VDisk, confirmation of I/O completion is received before the write operation has completed for the copy on the secondary VDisk. If a failover operation is performed, the application must recover and apply any updates that were not committed to the secondary VDisk. If I/O operations on the primary VDisk are paused for a small length of time, the secondary VDisk can become an exact match of the primary VDisk.

The Metro Mirror and Global Mirror operations support the following functions:

- 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 participate in active Metro Mirror and Global Mirror relationships with itself and up to three other clusters.
- Intercluster and intracluster Metro Mirror and Global Mirror relationships can be used concurrently within a cluster.
- The intercluster link is bidirectional. This means that it can copy data from cluster A to cluster B for one pair of VDisks while copying data from cluster B to cluster A for a different pair of VDisks.
- The copy direction can be reversed for a consistent relationship.
- Consistency groups are supported to manage a group of relationships that must be kept synchronized for the same application. This also simplifies administration, because a single command that is issued to the consistency group is applied to all the relationships in that group.
- SAN Volume Controller supports a maximum of 8192 Metro Mirror and Global Mirror relationships per cluster.

Metro Mirror and Global Mirror relationships

Metro Mirror and Global Mirror relationships define the relationship between two virtual disks (VDisks): a master VDisk and an auxiliary VDisk.

Typically, 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 for disaster recovery.

The master and auxiliary VDisks are defined when the relationship is created, and these attributes never change. However, either VDisk can operate in the primary or secondary role as necessary. The primary VDisk contains a valid copy of the application data and receives updates from the host application, analogous to a source VDisk. The secondary VDisk receives a copy of any updates to the primary VDisk, because these updates are all transmitted across the Mirror link. Therefore, the secondary VDisk is analogous to a continuously updated target VDisk. 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, you can reverse the copy direction from the command-line interface (CLI) or the SAN Volume Controller Console.

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 I/O group.

A relationship can be added to a consistency group, for ease of application management.

Note: Membership of a consistency group is an attribute of the relationship, not the consistency group. Therefore, issue the svctask chrcrelationship command to add or remove a relationship to or from a consistency group.

Copy types

A Metro Mirror copy ensures that updates are committed to both the primary and secondary VDisks before sending confirmation of I/O completion to the host application. This ensures that the secondary VDisk is synchronized with the primary VDisk in the event that a failover operation is performed.

A Global Mirror copy allows the host application to receive confirmation of I/O completion before the updates are committed to the secondary VDisk. If a failover operation is performed, the host application must recover and apply any updates that were not committed to the secondary VDisk.

States

When a Metro Mirror or Global Mirror relationship is created with two VDisks in different clusters, the distinction between the connected and disconnected states is important. These states apply to both clusters, the relationships, and the consistency groups. The following Metro Mirror and Global Mirror relationship states are possible:

InconsistentStopped
The primary VDisk is accessible for read and write I/O operations, but the secondary VDisk is not accessible for either operation. A copy process must be started to make the secondary VDisk consistent.

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

ConsistentStopped
The secondary VDisk contains a consistent image, but it might be out of date with respect to the primary VDisk. This state can occur when a relationship was in the ConsistentSynchronized state and experiences an error that forces a freeze of the consistency group. This state can also occur when a relationship is created with the CreateConsistentFlag parameter set to `TRUE`.

ConsistentSynchronized
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 an auxiliary VDisk operates in the primary role. Consequently the VDisk is accessible for write I/O operations.

IdlingDisconnected
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.

InconsistentDisconnected
The VDisks in this half of the consistency group are all operating in the secondary role and cannot accept read or write I/O operations.

ConsistentDisconnected
The VDisks in this half of the consistency group are all operating in the secondary role and can accept read I/O operations but not write I/O operations.

Metro Mirror and Global Mirror relationships between clusters

Metro Mirror and Global Mirror relationships can exist simultaneously between clusters. In this type of configuration, there can be impacts to performance because write data from both Metro Mirror and Global Mirror relationships is transported over the same intercluster links.

Metro Mirror and Global Mirror relationships manage heavy workload differently. Metro Mirror typically maintains the relationships that are in the copying or synchronized states, which causes the primary host applications to see degraded performance. Global Mirror requires a higher level of write performance to primary host applications. If the link performance is severely degraded, the link tolerance feature automatically stops Global Mirror relationships when the link tolerance threshold is exceeded. As a result, Global Mirror writes can suffer degraded performance if Metro Mirror relationships use most of the capability of the intercluster link.

Metro Mirror and Global Mirror partnerships

Metro Mirror and Global Mirror partnerships define an association between a local cluster and a remote cluster.

Before a Metro Mirror or Global Mirror relationship or consistency group can be created with a remote cluster, a *partnership* between the two clusters must be established. If Global Mirror or Metro Mirror relationships or consistency groups exist between two remote clusters, those clusters must maintain their partnership. Each cluster can maintain up to three partnerships, and each partnership can be with a single remote cluster. As many as four clusters can be directly associated with each other.

SAN Volume Controller clusters also become indirectly associated with each other through partnerships. If two clusters each have a partnership with a third cluster, those two clusters are indirectly associated. A maximum of four clusters can be directly or indirectly associated.

SAN Volume Controller nodes must know not only about the relationship between the two VDisks but also about an association among clusters. A maximum of four clusters can be connected either directly or indirectly.

The following examples show possible partnerships that can be established among SAN Volume Controller clusters.

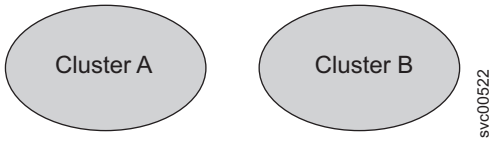

Figure 12. Two clusters with no partnerships

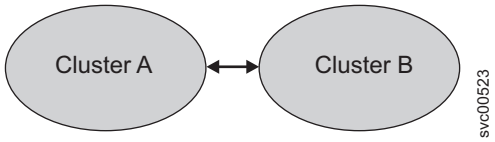

Figure 13. Two clusters with one partnership.

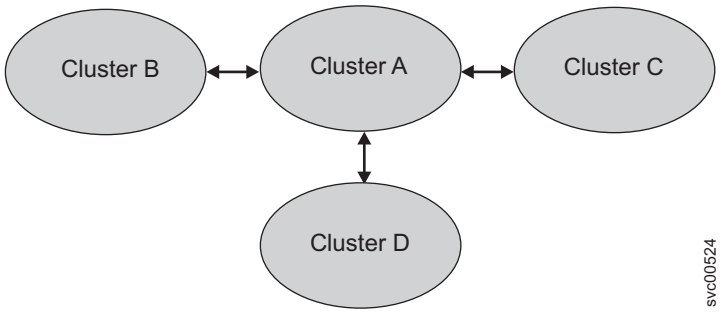

Figure 14. Four clusters in a partnership. Cluster A might be a disaster recovery site.

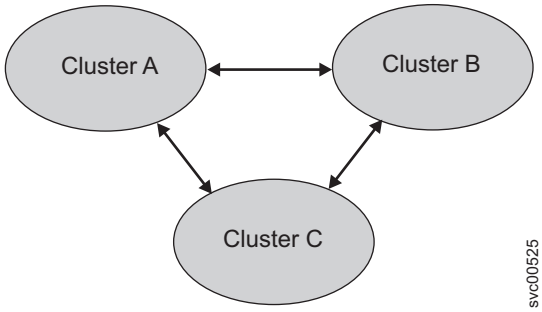

Figure 15. Three clusters in a migration situation. Data Center B is migrating to C. Cluster A is host production, and Cluster B and Cluster C are disaster recovery.

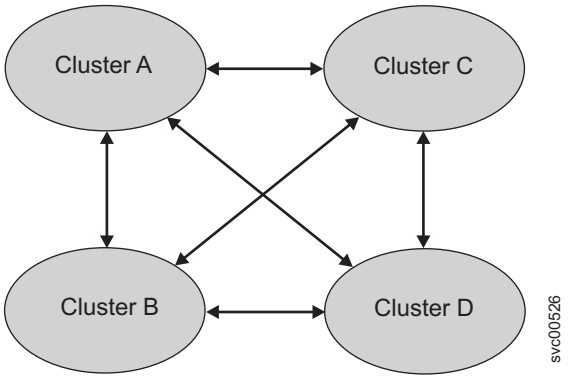

Figure 16. Clusters in a fully connected mesh configuration. Every cluster has a partnership to each of the three other clusters.

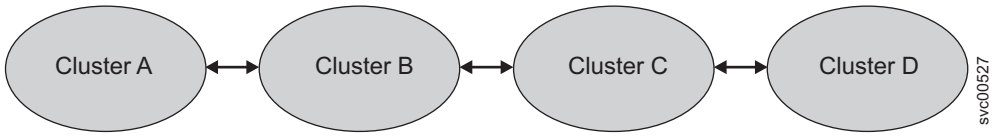

Figure 17. Four clusters in three partnerships.

Figure 18 depicts a cluster configuration that is not supported. Five clusters are in the connected set, even though no individual cluster is in more than two partnerships.

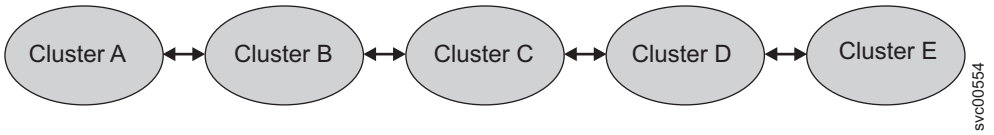

Figure 18. An unsupported cluster configuration.

To establish a Metro Mirror and Global Mirror partnership between two clusters, you must run the svctask mkpartnership command from both clusters. For example, to establish a partnership between clusterA and clusterB, you must run the svctask mkpartnership command from clusterA and specify clusterB as the remote cluster. At this point the partnership is partially configured and is sometimes described as one-way communication. Next, you must run the svctask mkpartnership command from clusterB and specify clusterA as the remote cluster. When this command completes, the partnership is fully configured for two-way communication between the clusters. You can also use the SAN Volume Controller Console to create Metro Mirror and Global Mirror partnerships.

The state of the partnership helps determine whether the partnership operates as expected. In addition to being fully configured, a cluster partnership can have the following states:

Partially Configured
: Indicates that only one cluster partner is defined from a local or remote cluster to the displayed cluster and is started. For the displayed cluster to be configured fully and to complete the partnership, you must define the cluster partnership from the cluster that is displayed to the corresponding local or remote cluster. You can do this by issuing the **mkpartnership**

command on the local and remote cluster that are in the partnership, or by using the SAN Volume Controller Console to create a partnership on both the local and remote clusters.

Fully Configured
Indicates that the partnership is defined on the local and remote clusters and is started.

Remote Not Present
Indicates that the remote cluster is not present to the partnership.

Partially Configured (Local Stopped)
Indicates that the local cluster is only defined to remote cluster and the local cluster is stopped.

Fully Configured (Local Stopped)
Indicates that a partnership is defined on both the local and remote clusters and the remote cluster is present, but the local cluster is stopped.

Fully Configured (Remote Stopped)
Indicates that a partnership is defined on both the local and remote clusters and the remote cluster is present, but the remote cluster is stopped.

Fully Configured (Local Excluded)
Indicates that a partnership is defined between a local and remote cluster; however, the local cluster has been excluded. Usually this state occurs when the fabric link between the two clusters has been compromised by too many fabric errors or slow response times of the cluster partnership. Check the error log for 1720 errors by selecting **Service and Maintenance** → **Analyze Error Log** to resolve these errors.

Fully Configured (Remote Excluded)
Indicates that a partnership is defined between a local and remote cluster; however, the remote cluster has been excluded. Usually this state occurs when the fabric link between the two clusters has been compromised by too many fabric errors or slow response times of the cluster partnership. Check the error log for 1720 errors by selecting **Service and Maintenance** → **Analyze Error Log** to resolve these errors.

Fully Configured (Remote Exceeded)
Indicates that a partnership is defined between a local and remote cluster and the remote is available; however, the remote cluster exceeds the number of allowed clusters within a cluster network. The maximum of four clusters can be defined in a network. If the number of clusters exceeds that limit, SAN Volume Controller determines the inactive cluster or clusters by sorting all the clusters by their unique identifier in numerical order. The inactive cluster partner which is not in the top four of the cluster unique identifiers displays **Fully Configured (Remote Exceeded)**.

To change Metro Mirror and Global Mirror partnerships, use the svctask chpartnership command. To delete Metro Mirror and Global Mirror partnerships, use the svctask rmpartnership command.

Attention: Before you run the svctask rmpartnership command, you must remove all relationships and groups that are defined between the two clusters. To display cluster relationships and groups, run the svcinfo lsrcrelationship and svcinfo lsrcconsistgrp commands. To remove the relationships and groups that are defined between the two clusters, run the svctask rmrcrelationship and svctask rmrcconsistgrp commands.

Background copy management

You can control the rate at which the initial background copy from the local cluster to the remote cluster is performed. The bandwidth parameter specifies this rate in whole megabytes per second.

Global Mirror configuration requirements

To use the Global Mirror feature, all components in the SAN must be capable of sustaining the workload that is generated by application hosts and the Global Mirror background copy process. If all of the components in the SAN cannot sustain the workload, the Global Mirror relationships are automatically stopped to protect your application hosts from increased response times.

When using the Global Mirror feature, follow these best practices:

- Use IBM Tivoli Storage Productivity Center or an equivalent SAN performance analysis tool to monitor your SAN environment. The IBM Tivoli Storage Productivity Center provides an easy way to analyze the SAN Volume Controller performance statistics.
- Analyze the SAN Volume Controller performance statistics to determine the peak application write workload that the link must support. Gather statistics over a typical application I/O workload cycle.
- Set the background copy rate to a value that can be supported by the intercluster link and the backend storage controllers at the remote cluster.
- Do not use cache-disabled VDisks in Global Mirror relationships.
- Set the gmlinktolerance parameter to an appropriate value. The default value is 300 seconds (5 minutes).
- When you perform SAN maintenance tasks, take one of the following actions:
  - Reduce the application I/O workload for the duration of the maintenance task.
  - Disable the gmlinktolerance feature or increase the gmlinktolerance value.

    Note: If the gmlinktolerance value is increased during the maintenance task, do not set it to the normal value until the maintenance task is complete. If the gmlinktolerance feature is disabled for the duration of the maintenance task, enable it after the maintenance task is complete.
  - Stop the Global Mirror relationships.
- Evenly distribute the preferred nodes for the Global Mirror VDisks between the nodes in the clusters. Each VDisk in an I/O group has a preferred node property that can be used to balance the I/O load between nodes in the I/O group. The preferred node property is also used by the Global Mirror feature to route I/O operations between clusters. A node that receives a write for a VDisk is normally the preferred node for that VDisk. If the VDisk is in a Global Mirror relationship, the node is responsible for sending the write to the preferred node of the secondary VDisk. By default, the preferred node of a new VDisk is the node that owns the fewest VDisks of the two nodes in the I/O group. Each node in the remote cluster has a set pool of Global Mirror system resources for each node in the local cluster. To maximize Global Mirror performance, set the preferred nodes for the VDisks of the remote cluster to use every combination of primary nodes and secondary nodes.

Long distance links for Metro Mirror and Global Mirror partnerships

For intracluster partnerships, all clusters can be considered as candidates for Metro Mirror or Global Mirror operations. For intercluster partnerships, cluster pairs must be separated by a number of moderately high bandwidth links.

Figure 19 shows an example of a configuration that uses dual redundant fabrics. Part of each fabric is located at the local cluster and the remote cluster. There is no direct connection between the two fabrics.

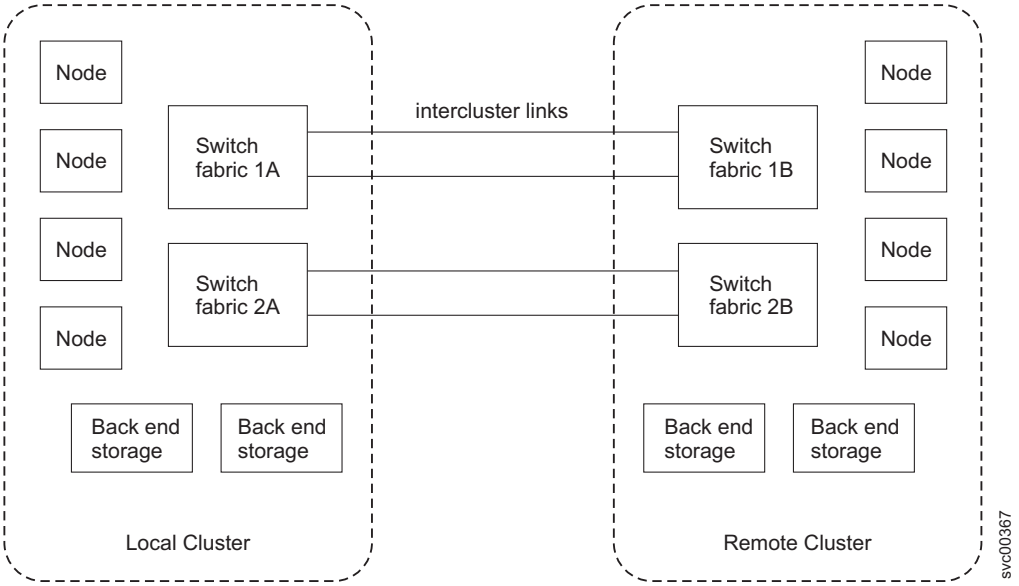

Figure 19. Redundant fabrics

You can use fibre-channel extenders or SAN routers to increase the distance between two clusters. Fibre-channel extenders transmit fibre-channel packets across long links without changing the contents of the packets. SAN routers provide virtual nPorts on two or more SANs to extend the scope of the SAN. The SAN router distributes the traffic from one virtual nPort to the other virtual nPort. The two fibre-channel fabrics are independent of each other. Therefore, nPorts on each of the fabrics cannot directly log into each other. See the following Web site for specific firmware levels and the latest supported hardware:

www.ibm.com/storage/support/2145

If you use fibre-channel extenders or SAN routers, you must meet the following requirements:

- For SAN Volume Controller software level 4.1.0, the round-trip latency between sites cannot exceed 68 ms for fibre-channel extenders or 20 ms for SAN routers.
- For SAN Volume Controller software level 4.1.1 or higher, the round-trip latency between sites cannot exceed 80 ms for either fibre-channel extenders or SAN routers.
- The configuration must be tested with the expected peak workloads.
- Metro Mirror and Global Mirror require a specific amount of bandwidth for intercluster heartbeat traffic. The amount of traffic depends on the number of nodes that are in both the local cluster and the remote cluster. Table 17 on page 70 lists the intercluster heartbeat traffic for the primary cluster and the

secondary cluster. These numbers represent the total traffic between two clusters when there are no I/O operations running on the copied VDisks. Half of the data is sent by the primary cluster and half of the data is sent by the secondary cluster so that traffic is evenly divided between all of the available intercluster links. If you have two redundant links, half of the traffic is sent over each link.

Table 17. Intercluster heartbeat traffic in Mbps

Cluster 1	**Cluster 2**			
	2 nodes	4 nodes	6 nodes	8 nodes
2 nodes	2.6	4.0	5.4	6.7
4 nodes	4.0	5.5	7.1	8.6
6 nodes	5.4	7.1	8.8	10.5
8 nodes	6.7	8.6	10.5	12.4

- The bandwidth between two sites must meet the peak workload requirements and maintain the maximum round-trip latency between the sites. When you evaluate the workload requirement, you must consider the average write workload over a period of one minute or less and the required synchronization copy bandwidth. If there are no active synchronization copies and no write I/O operations for VDisks that are in the Metro Mirror or Global Mirror relationship, the SAN Volume Controller protocols operate with the bandwidth that is indicated in Table 17. However, you can only determine the actual amount of bandwidth that is required for the link by considering the peak write bandwidth to VDisks that are participating in Metro Mirror or Global Mirror relationships and then adding the peak write bandwidth to the peak synchronization bandwidth.
- If the link between two sites is configured with redundancy so that it can tolerate single failures, the link must be sized so that the bandwidth and latency statements are correct during single failure conditions.
- The channel must not be used for links between nodes in a single cluster. Configurations that use long distance links in a single cluster are not supported and can cause I/O errors and loss of access.
- The configuration is tested to confirm that any failover mechanisms in the intercluster links interoperate satisfactorily with the SAN Volume Controller.
- All other SAN Volume Controller configuration requirements are met.

Limitations on host to cluster distances

There is no limit on the fibre-channel optical distance between SAN Volume Controller nodes and host servers. You can attach a server to an edge switch in a core-edge configuration with the SAN Volume Controller cluster at the core. SAN Volume Controller clusters support up to three ISL hops in the fabric. Therefore, the host server and the SAN Volume Controller cluster can be separated by up to five fibre-channel links. If you use longwave SFPs, four of the fibre-channel links can be up to 10 km long.

Using the intercluster link for host traffic

If you use the intercluster link for host traffic, ensure that you have sufficient bandwidth to support all sources of load.

Scenario: The hosts in a local cluster can read and write to the VDisks in a remote cluster

In this scenario, the hosts in the local cluster also exchange heartbeats with the hosts that are in the remote cluster. Because the intercluster link is being used for multiple purposes, you must have sufficient bandwidth to support the following sources of load:

- Global Mirror or Metro Mirror data transfers and the SAN Volume Controller cluster heartbeat traffic.
- Local host to remote VDisk I/O traffic or remote host to local VDisk I/O traffic.
- Local host to remote host heartbeat traffic. If the local host to remote VDisk I/O traffic is allowed to consume a high percentage of intercluster link bandwidth, the latency seen by the hosts that access SAN Volume Controller VDisks that are participating in Metro Mirror or Global Mirror operations can be impacted. The bandwidth congestion can cause the Global Mirror link tolerance threshold to be exceeded. When the Global Mirror link tolerance threshold is exceeded, Global Mirror relationships are stopped.

Metro Mirror and Global Mirror consistency groups

You can group Metro Mirror or Global Mirror relationships into a consistency group so that they can be updated at the same time. A command that is issued to the consistency group is simultaneously applied to all of the relationships in the group.

Metro Mirror or Global Mirror relationships can be based on “loose” or “tight” associations. A more significant use arises when the relationships contain virtual disks (VDisks) with a tight association. A simple example of a tight association is the spread of data for an application 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. In both examples, specific rules exist as to how the relationships can be updated. This ensures that the set of secondary VDisks contain usable data. The key property is that these relationships are consistent.

Metro Mirror or Global Mirror relationships can only belong to one consistency group; however, they do not have to belong to a consistency group. Relationships that are not part of a consistency group are called stand-alone relationships. A consistency group can contain zero or more relationships. All relationships in a consistency group must have matching primary and secondary clusters, which are sometimes referred to as master and auxiliary clusters. All relationships in a consistency group must also have the same copy direction and state.

Metro Mirror and Global Mirror relationships cannot belong to the same consistency group. A copy type is automatically assigned to a consistency group when the first relationship is added to the consistency group. After the consistency group is assigned a copy type, only relationships of that copy type can be added to the consistency group. Each cluster can have a maximum of six different types of consistency groups. The following types of consistency groups are possible:

- Intracluster Metro Mirror
- Intercluster Metro Mirror from the local cluster to remote cluster
- Intercluster Metro Mirror from the remote cluster to local cluster
- Intracluster Global Mirror
- Intercluster Global Mirror from the local cluster to remote cluster

- Intercluster Global Mirror from the remote cluster to local cluster

States

A consistency group can be in one of the following states:

Inconsistent (stopped)
The primary VDisks are accessible for read and write 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 the **svctask startrcconsistgrp** command is issued to a consistency group in the InconsistentStopped state. This state is also entered when the **svctask startrcconsistgrp** 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 occur when a relationship was in the ConsistentSynchronized state and experiences an error that forces a freeze of the consistency group. This state can also occur 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 Both the primary VDisks and the secondary 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 cannot 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 can accept read I/O operations but not write I/O operations.

Empty The consistency group does not contain any relationships.

Background copy bandwidth impact on foreground I/O latency

The background copy bandwidth determines the rate at which the background copy for Metro Mirror or Global Mirror Copy Services are attempted.

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

- If the background copy bandwidth is set too high for the intercluster link capacity, the following results can occur:
  - The background copy I/Os can back up on the intercluster link

  - For Metro Mirror, there is a delay in the synchronous secondary writes of foreground I/Os
  - For Global Mirror, the work is backlogged, which delays the processing of writes and causes the relationship to stop
  - The foreground I/O latency increases as detected 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.
  - For Global Mirror, the work is backlogged and again the relationship is stopped

To set the background copy bandwidth optimally, you must 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. You must also consider concurrent host I/O because if other write operations arrive at the primary cluster for copy to the remote site, these write operations can be delayed by a high level of background copy and the hosts at the primary site receive poor write-operation response times.

This provisioning can be done by the calculation above or by determining 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.

Example

If the bandwidth setting at the primary site for the secondary cluster is set to 200 MBps (megabytes per second) and the relationships are not synchronized, the SAN Volume Controller attempts to resynchronize the relationships at a maximum rate of 200 MBps with a 25 MBps restriction for each individual relationship. The SAN Volume Controller cannot resynchronize the relationship if the throughput is restricted. The following can restrict throughput:

- The read response time of backend storage at the primary cluster
- The write response time of the backend storage at the secondary site
- Intercluster link latency

Migrating a Metro Mirror relationship to a Global Mirror relationship

You can migrate a Metro Mirror relationship to a Global Mirror relationship.

Scenario: I/O operations to the secondary VDisk can be stopped during the migration

In this scenario, you have the ability to stop I/O operations to the secondary VDisk during the migration process.

To stop I/O operations to the secondary VDisk while migrating a Metro Mirror relationship to a Global Mirror relationship, you must specify the synchronized option when you create the Global Mirror relationship.

1. Stop all host I/O operations to the primary VDisk.
2. Verify that the Metro Mirror relationship is consistent.

   Important: If the Metro Mirror relationship is not consistent when it is stopped, or if any host I/O operations run between the Metro Mirror relationship being stopped and the Global Mirror relationship being created, the updates are not copied to the secondary VDisk.
3. Delete the Metro Mirror relationship.
4. Create the Global Mirror relationship between the same two VDisks.

After the Global Mirror relationship is created, you can start the relationship and resume host I/O operations.

Scenario: I/O operations to the secondary VDisk cannot be stopped during the migration

In this scenario, you do not have the ability to stop I/O operations to the secondary VDisk during the migration process.

If I/O operations to the secondary VDisk cannot be stopped, the data on the secondary VDisk becomes out-of-date. When the Global Mirror relationship is started, the secondary VDisk is inconsistent until all of the recent updates are copied to the remote site.

If you do not require a consistent copy of the VDisk at the secondary site, perform the following steps to migrate from a Metro Mirror relationship to a Global Mirror relationship:

Important: The data on the secondary VDisk is not usable until the synchronization process is complete. Depending on your link capabilities and the amount of data that is being copied, this process can take an extended period of time. You must set the background copy bandwidth for the intercluster partnerships to a value that does not overload the intercluster link.

1. Delete the Metro Mirror relationship.
2. Create and start the Global Mirror relationship between the same two VDisks.

If you require a consistent copy of the VDisk at the secondary site, perform the following steps to migrate from a Metro Mirror relationship to a Global Mirror relationship:

1. Delete the Metro Mirror relationship.
2. Create a Global Mirror relationship between VDisks that were not used for the Metro Mirror relationship. This preserves the VDisk so that you can use it if you require a consistent copy at a later time.

Alternatively, you can use the FlashCopy feature to maintain a consistent copy. Perform the following steps to use the FlashCopy feature to maintain a consistent copy:

1. Start a FlashCopy operation for the Metro Mirror VDisk.
2. Wait for the FlashCopy operation to complete.
3. Create and start the Global Mirror relationship between the same two VDisks. The FlashCopy VDisk is now your consistent copy.

Using FlashCopy to create a consistent image before restarting a Global Mirror relationship

For disaster recovery purposes, you can use the FlashCopy feature to create a consistent copy of an image before you restart a Global Mirror relationship.

When a consistent relationship is stopped, the relationship enters the consistent_stopped state. While in this state, I/O operations at the primary site continue to run. However, updates are not copied to the secondary site. When the relationship is restarted, the synchronization process for new data is started. During this process, the relationship is in the inconsistent_copying state. The secondary VDisk for the relationship cannot be used until the copy process completes and the relationship returns to the consistent state. When this occurs, start a FlashCopy operation for the secondary VDisk before you restart the relationship. While the relationship is in the copying state, the FlashCopy feature can provide a consistent copy of the data. If the relationship does not reach the synchronized state, you can use the FlashCopy target VDisk at the secondary site.

The SVCTools package that is available on the IBM alphaWorks® Web site provides an example script that demonstrates how to manage the FlashCopy process. See the copymanager script that is available in the SVCTools package. You can download the SVCTools package from the following Web site:

www.alphaworks.ibm.com/tech/svctools/download

Monitoring Global Mirror performance with the IBM System Storage Productivity Center

You can use the IBM System Storage Productivity Center (SSPC) to monitor key Global Mirror performance measurements.

It is important to use a Storage Area Network (SAN) performance monitoring tool to ensure that all SAN components are performing correctly. This is particularly important when you use an asynchronous copying solution such as the SAN Volume Controller Global Mirror feature. SSPC monitors key performance measures and alerts you when thresholds are exceeded.

Note: If your VDisk or MDisk configuration changes, restart the SSPC performance report to ensure that performance is monitored for the new configuration.

Use SSPC to check the following measurements:

- The *Port to Remote Node Send Response Time* measurement is less than 80 milliseconds. If this measurement is greater than 80 milliseconds during monitoring, the long-distance link has excessive latency. Ensure that the link is operating at its maximum bandwidth.
- The sum of the *Port to Local Node Send Response Time* measurement and the *Port to Local Node Send Queue* measurement is less than 1 millisecond for the primary cluster and the CPU Utilization Percentage is below 50%. A value that exceeds these amounts can indicate that an I/O group is reaching the I/O throughput limit, which can limit performance.
- The sum of the *Backend Write Response Time* measurement and the *Write Queue Time for Global Mirror MDisks* measurement of the secondary cluster is less than 100 milliseconds. A longer response time can indicate that the storage controller is overloaded.
- The sum of the *Backend Write Response Time* measurement and the *Write Queue Time for Global Mirror MDisks* measurement of the primary cluster is less than

100 milliseconds. If the response time is greater than 100 milliseconds, application hosts might see extended response times when the SAN Volume Controller cluster cache is full.

- The *Write Data Rate for Global Mirror MDisk groups* measurement of the secondary cluster indicates the amount of data that is being written by Global Mirror operations. If this value approaches either the intercluster link bandwidth or the storage controller throughput limit, further increases can cause overloading of the system. Monitor for this condition in a way that is appropriate for your network.

The gmlinktolerance feature

You can use the svctask chcluster CLI command or the SAN Volume Controller Console to set the gmlinktolerance feature. The gmlinktolerance feature represents the number of seconds that the primary SAN Volume Controller cluster tolerates slow response times from the secondary cluster.

If the poor response extends past the specified tolerance, a 1920 error is logged and one or more Global Mirror relationships are automatically stopped. This protects the application hosts at the primary site. During normal operation, application hosts see a minimal impact to response times because the Global Mirror feature uses asynchronous replication. However, if Global Mirror operations experience degraded response times from the secondary cluster for an extended period of time, I/O operations begin to queue at the primary cluster. This results in an extended response time to application hosts. In this situation, the gmlinktolerance feature stops Global Mirror relationships and the application hosts response time returns to normal. After a 1920 error has occurred, the Global Mirror auxiliary VDisks are no longer in the consistent_synchronized state until you fix the cause of the error and restart your Global Mirror relationships. For this reason, ensure that you monitor the cluster to track when this occurs.

You can disable the gmlinktolerance feature by setting the gmlinktolerance value to 0 (zero). However, the gmlinktolerance cannot protect applications from extended response times if it is disabled. It might be appropriate to disable the gmlinktolerance feature in the following circumstances:

- During SAN maintenance windows where degraded performance is expected from SAN components and application hosts can withstand extended response times from Global Mirror VDisks.
- During periods when application hosts can tolerate extended response times and it is expected that the gmlinktolerance feature might stop the Global Mirror relationships. For example, if you are testing using an I/O generator which is configured to stress the backend storage, the gmlinktolerance feature might detect the high latency and stop the Global Mirror relationships. Disabling gmlinktolerance prevents this at the risk of exposing the test host to extended response times.

Diagnosing and fixing 1920 errors

A 1920 error indicates that one or more of the SAN components are unable to provide the performance that is required by the application hosts. This can be temporary (for example, a result of maintenance activity) or permanent (for example, a result of a hardware failure or unexpected host I/O workload). If you are experiencing 1920 errors, set up a SAN performance analysis tool, such as the IBM Tivoli Storage Productivity Center, and make sure that it is correctly configured and monitoring statistics when the problem occurs. Set your SAN performance analysis tool to the minimum available statistics collection interval.

For the IBM Tivoli Storage Productivity Center, the minimum interval is five minutes. If several 1920 errors have occurred, diagnose the cause of the earliest error first. The following questions can help you determine the cause of the error:

- Was maintenance occurring at the time of the error? This might include replacing a storage controller's physical disk, upgrading a storage controller's firmware, or performing a code upgrade on one of the SAN Volume Controller clusters. You must wait until the maintenance procedure is complete and then restart the Global Mirror relationships. You must wait until the maintenance procedure is complete to prevent a second 1920 error because the system has not yet returned to a stable state with good performance.
- Were there any unfixed errors on either the source or target system? If yes, analyze them to determine if they might have been the reason for the error. In particular, see if they either relate to the VDisk or MDisks that are being used in the relationship, or if they would have caused a reduction in performance of the target system. Ensure that the error is fixed before you restart the Global Mirror relationship.
- Is the long distance link overloaded? If your link is not capable of sustaining the short-term peak Global Mirror workload, a 1920 error can occur. Perform the following checks to determine if the long distance link is overloaded:
  - Look at the total Global Mirror auxiliary VDisk write throughput before the Global Mirror relationships were stopped. If this is approximately equal to your link bandwidth, your link might be overloaded. This might be due to application host I/O operations or a combination of host I/O and background (synchronization) copy activities.
  - Look at the total Global Mirror source VDisk write throughput before the Global Mirror relationships were stopped. This represents the I/O operations that are being performed by the application hosts. If these operations are approaching the link bandwidth, upgrade the link's bandwidth, reduce the I/O operations that the application is attempting to perform, or use Global Mirror to copy fewer VDisks. If the auxiliary disks show significantly more write I/O operations than the source VDisks, there is a high level of background copy. Decrease the Global Mirror partnership's background copy rate parameter to bring the total application I/O bandwidth and background copy rate within the link's capabilities.
  - Look at the total Global Mirror source VDisk write throughput after the Global Mirror relationships were stopped. If write throughput increases by 30% or more when the relationships are stopped, the application hosts are attempting to perform more I/O operations than the link can sustain. While the Global Mirror relationships are active, the overloaded link causes higher response times to the application host, which decreases the throughput it can achieve. After the Global Mirror relationships have stopped, the application host sees lower response times. In this case, the link bandwidth must be increased, the application host I/O rate must be decreased, or fewer VDisks must be copied using Global Mirror.
- Are the storage controllers at the secondary cluster overloaded? If one or more of the MDisks on a storage controller are providing poor service to the SAN Volume Controller cluster, a 1920 error occurs if this prevents application I/O operations from proceeding at the rate that is required by the application host. If the backend storage controller requirements have been followed, the error might have been caused by a decrease in controller performance. Use IBM Tivoli Storage Productivity Center to obtain the backend write response time for each MDisk at the secondary cluster. If the response time for any individual MDisk

exhibits a sudden increase of 50 ms or more or if the response time is above 100 ms, this indicates a problem. Perform the following checks to determine if the storage controllers are overloaded:

  - Check the storage controller for error conditions such as media errors, a failed physical disk, or associated activity such as RAID array rebuilding. If there is an error, you should fix the problem and then restart the Global Mirror relationships.
  - If there is no error, determine if the secondary controller is capable of processing the required level of application host I/O operations. It might be possible to improve the performance of the controller by adding more physical disks to a RAID array, changing the RAID level of the array, changing the controller's cache settings and checkin the cache battery to ensure it is operational, or changing other controller-specific configuration parameters.

- Are the storage controllers at the primary cluster overloaded? Analyze the performance of the primary backend storage using the same steps as for the secondary backend storage. If performance is bad, limit the amount of I/O operations that can be performed by application hosts. Monitor the backend storage at the primary site even if the Global Mirror relationships have not been affected. If bad performance continues for a prolonged period, a 1920 error occurs and the Global Mirror relationships are stopped.
- Is one of your SAN Volume Controller clusters overloaded? Use IBM Tivoli Storage Productivity Center to obtain the port to local node send response time and the port to local node send queue time. If the total of these two statistics for either cluster is above 1 millisecond, the SAN Volume Controller might be experiencing a very high I/O load. Also check the SAN Volume Controller node CPU utilization. If this figure is above 50%, this can also be contributing to the problem. In either case, contact your IBM service representative for further assistance. If CPU utilization is much higher for one node than for the other node in the same I/O group, this might be caused by having different node hardware types within the same I/O group. For example, a SAN Volume Controller 2145-8F4 in the same I/O group as a SAN Volume Controller 2145-8G4. If this is the case, contact your IBM service representative.
- Do you have FlashCopy operations in the prepared state at the secondary cluster? If the Global Mirror auxiliary VDisks are the sources of a FlashCopy mapping and that mapping is in the prepared state for an extended time, performance to those VDisks can be impacted because the cache is disabled. Start the FlashCopy mapping to enable the cache and improve performance for Global Mirror I/O operations.

Valid combinations of FlashCopy and Metro Mirror or Global Mirror functions

The following table outlines the combinations of FlashCopy and Metro Mirror or Global Mirror functions that are valid for a single virtual disk (VDisk).

FlashCopy	Metro Mirror or Global Mirror Primary	Metro Mirror or Global Mirror Secondary
FlashCopy source	Supported	Supported
FlashCopy target	Not supported	Not supported

Chapter 3. SAN fabric and LAN overview

The *SAN fabric* is an area of the network that contains routers and switches. In fibre-channel environments, *zoning* is the grouping of multiple ports to form a virtual, private, storage network. iSCSI is an IP-based standard for linking data storage devices over a network and transferring data by carrying SCSI commands over IP networks.

SAN fabric and LAN configuration terms

Ensure that you understand the basic terms and definitions when you are configuring the SAN Volume Controller within the SAN fabric or a local area network (LAN).

Table 18 provides terms and definitions that can guide your understanding of the SAN fabric rules and requirements.

Table 18. SAN fabric configuration terms and definitions

Term	Definition
ISL hop	A hop on an interswitch link (ISL). With reference to all pairs of N-ports or end-nodes that are in a fabric, 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 that there is a symmetrical network and a specific workload that is applied equally from all initiators and sent equally to all targets. A symmetrical network is one in which all initiators are connected at the same level and all the controllers are connected at the same level. **Note:** The SAN Volume Controller puts its back-end traffic onto the same symmetrical network. The back-end traffic can vary 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 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. You can create a redundant SAN by splitting the SAN into two independent counterpart SANs.
Counterpart SAN	A nonredundant 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.

Table 18. SAN fabric configuration terms and definitions (continued)

Term	Definition
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 and Global Mirror, significant distances might exist between the components of the local cluster and those of the remote clusters.
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 and Global Mirror, significant distances might exist between the components of the local cluster and those of the remote clusters.
Local/remote fabric interconnect	The SAN components that connect the local fabrics to the remote fabrics. There might be significant distances between the components in the local cluster and those in the remote cluster. These components might be single-mode optical fibers 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 fail over to another port, and the fan in requirements might be exceeded in this degraded mode.
Not valid configuration	The current SAN configuration is not correct. An attempted operation failed and generated an error code that indicates what caused it to become "not valid." The most likely cause is that either a device has failed, or a device has been added to the SAN that has caused the configuration to be marked as not valid.
Unsupported configuration	A configuration that could operate successfully, but that IBM does not guarantee any solutions for if problems occur. Typically, this type of configuration does not create any error log entries.
Valid configuration	A configuration that consists of devices and connections that are identified as valid and supported. Neither of the following two conditions exist with the current configuration: • Not valid • Unsupported configuration
Degraded	A valid configuration that has had a failure, but continues to be valid and supported. Typically, a repair action is required to restore the degraded configuration to a valid configuration.
Fibre-channel extender	A device for long distance communication that connects other SAN fabric components. Generally these components might 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 that are configured to create a larger switched network. With this configuration, four or more switches are connected in a loop with additional direct connections between nonadjacent switches within the loop. An example of this configuration is four switches that are connected in a loop with ISLs for one of the diagonals.

If you plan to use iSCSI to configure a LAN, ensure that you also understand the iSCSI terms and definitions. Table 19 on page 81 highlights the terms and

definitions for iSCSI.

Table 19. iSCSI configuration terms and definitions

Term	Definition
Challenge Handshake Authentication Protocol (CHAP)	An authentication protocol that protects against eavesdropping by encrypting the user name and password.
Clustered Ethernet port	A physical Ethernet port on a node in a cluster that contains configuration settings that are shared by all the ports in a cluster.
Extended-unique identifier (EUI)	A unique iSCSI name that identifies an iSCSI target adapter or an iSCSI initiator adapter as defined by the iSCSI standard (RFC 3722).
Host object	A logical object that represents a list of worldwide port names (WWPNs) and a list of iSCSI names that identify the interfaces that the host system uses to communicate with a device. iSCSI names can be either iSCSI qualified names (IQNs) or extended-unique identifiers (EUIs).
Host system	A computer that is connected to the SAN Volume Controller through either a fibre-channel interface or an IP network.
Initiator	The system component that originates an I/O command over an I/O bus or network. I/O adapters and network interface controllers are typical initiators.
Internet Storage Name Service (iSNS) Protocol	A protocol that is used by a host system to manage iSCSI targets and iSCSI discovery. iSCSI initiators use the ISNS Protocol to locate the appropriate storage resources.
iSCSI alias	An alternative name for the iSCSI-attached host.
iSCSI name	A name that identifies an iSCSI target adapter or an iSCSI initiator adapter. An iSCSI name can be an iSCSI qualified name (IQN) or an extended-unique identifier (EUI). Typically, this identifier has the following format: iqn.datecode.reverse domain.
iSCSI qualified name (IQN)	A specific type of iSCSI name that identifies an iSCSI target adapter or an iSCSI initiator adapter as defined by the iSCSI standard (RFC 3722).
Network interface controller (NIC)	Hardware that provides the interface control between system main storage and external high-speed link (HSL) ports.
Node Ethernet port	A port that represents an iSCSI port on a SAN Volume Controller node. Configuration settings are specific to a single physical Ethernet port.
Subnet	A portion of a network that is divided into smaller independent subgroups, which still are interconnected.
Target	The program or system to which a request for files or processing is sent.

SAN fabric overview

The *SAN fabric* is an area of the network that contains routers and switches. A SAN is configured into a number of zones. A device using the SAN can communicate only with devices that are included in the same zones that it is in. A SAN Volume Controller cluster requires several distinct types of zones: a cluster zone, host zones, and disk zones. The intercluster zone is optional.

In the host zone, the host systems can identify and address the SAN Volume Controller nodes. You can have more than one host zone and more than one disk

zone. The cluster zone contains all ports from all SAN Volume Controller nodes in the cluster, unless you are using a dual-core fabric design. Create one zone for each host fibre-channel port. In a disk zone, the SAN Volume Controller nodes identify the storage systems. Generally, create one zone for each storage system. Host systems cannot operate on the storage systems directly; all data transfer occurs through the SAN Volume Controller nodes. If you are using the Metro Mirror and Global Mirror feature, create a zone with at least one port from each node in each cluster; up to four clusters are supported.

Figure 20 shows an example of a host zone. Figure 21 on page 83 shows an example of a cluster zone. Figure 22 on page 83 shows an example of a disk zone.

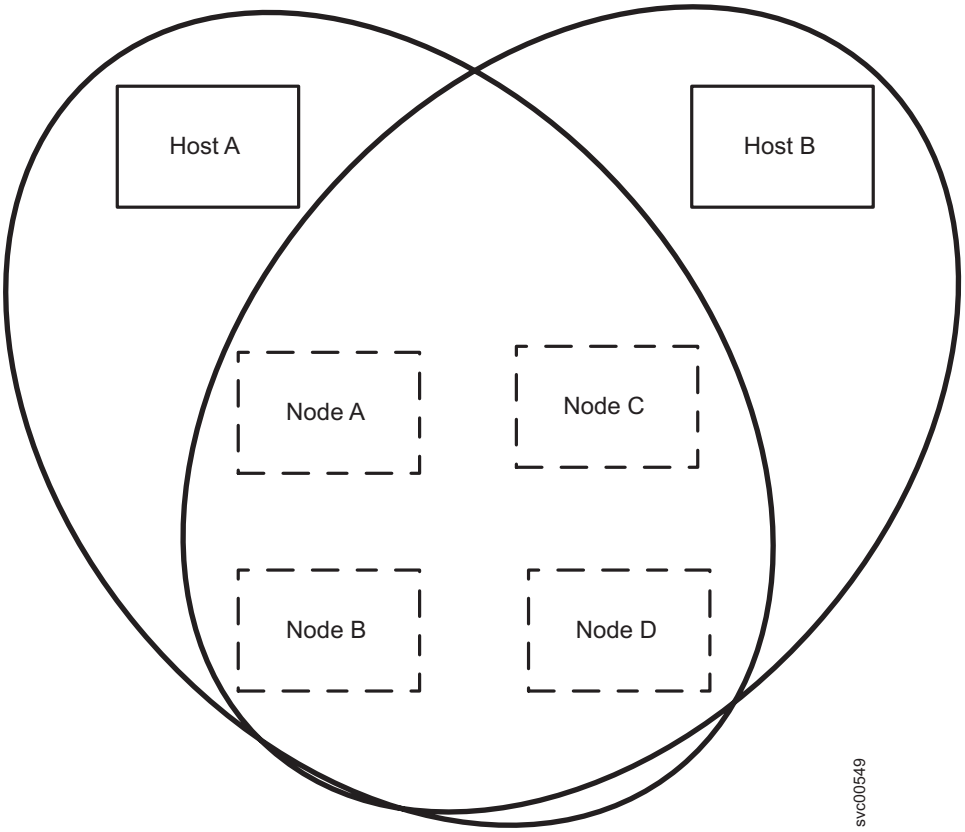

Figure 20. Example of a SAN Volume Controller host zone

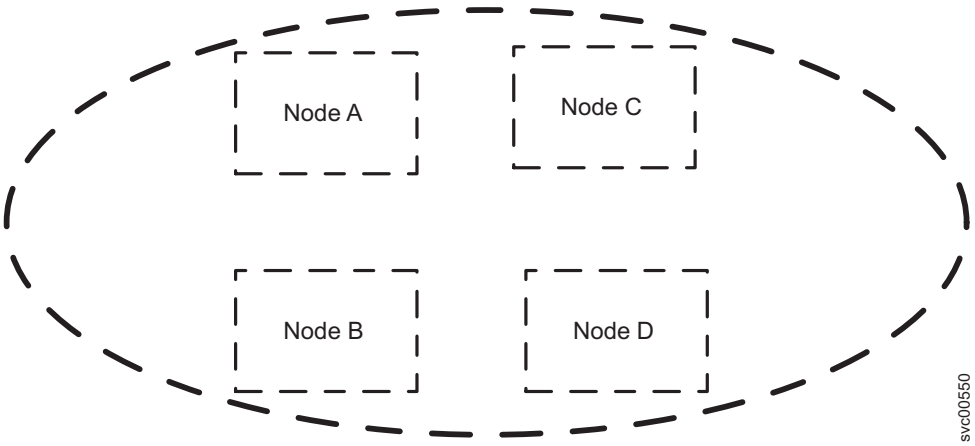

Figure 21. Example of a SAN Volume Controller cluster zone

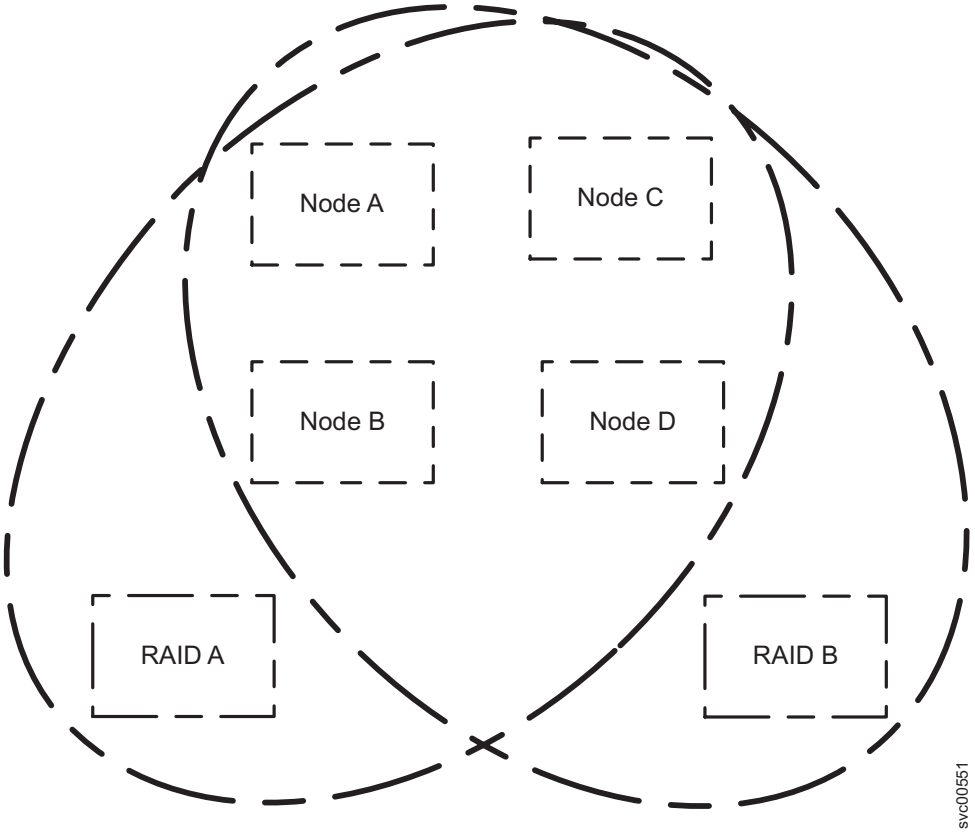

Figure 22. Example of a SAN Volume Controller disk zone

A cluster of SAN Volume Controller nodes is connected to the fibre-channel fabric and presents virtual disks (VDisks) to the host systems. You create these VDisks from units of space within a managed disk (MDisk) group. An MDisk group is a collection of MDisks that are presented by the storage systems (RAID controllers). The MDisk group provides a storage pool. You specify how each group is created, and you can combine MDisks from different manufacturers' controllers in the same MDisk group. However, to optimize the use of resources, ensure that all MDisks in an MDisk group have similar performance characteristics.

Note: Some operating systems cannot tolerate other operating systems in the same host zone, although you might have more than one host type in the SAN fabric. For example, you can have a SAN that contains one host that runs on an IBM AIX® operating system and another host that runs on a Microsoft Windows operating system.

All communication between SAN Volume Controller nodes is performed through the SAN. All SAN Volume Controller node configuration and service commands are sent to the cluster through an Ethernet network.

Each SAN Volume Controller node contains its own vital product data (VPD). Each cluster contains VPD that is common to all the SAN Volume Controller nodes in the cluster, and any system, with the correct access authority, that is connected to the Ethernet network can access this VPD.

iSCSI overview

iSCSI is an IP-based standard for transferring data that supports host access by carrying SCSI commands over IP networks. The iSCSI standard is defined by RFC 3720.

For SAN Volume Controller, connections from iSCSI-attached hosts to nodes are supported. iSCSI connections from SAN Volume Controller nodes to storage systems are not supported.

Table 20 shows that iSCSI and fibre-channel terms have analogous components.

Table 20. Comparison of iSCSI and fibre-channel components

iSCSI components	Fibre-channel components
iSCSI host bus adapter	Fibre-channel host bus adapter
Network interface controller (NIC) and iSCSI software initiator	Fibre-channel host bus adapter
IP switch	Fibre-channel switch
IP router	–
iSCSI name, such as IQN (iSCSI qualified name) or EUI (extended-unique identifier)	WWNN (worldwide node name)

iSCSI initiators and targets

In an iSCSI configuration, the iSCSI host or server sends requests to a node. The host contains one or more initiators that attach to an IP network to initiate requests to and receive responses from an iSCSI target. Each initiator and target are given a unique iSCSI name such as an iSCSI qualified name (IQN) or an extended-unique identifier (EUI). An IQN is a 223-byte ASCII name. An EUI is a 64-bit identifier. An iSCSI name represents a worldwide unique naming scheme that is used to identify each initiator or target in the same way that worldwide node names (WWNNS) are used to identify devices in a fibre-channel fabric.

An iSCSI target is any device that receives iSCSI commands. The device can be an end node such as a storage device, or it can be an intermediate device such as a bridge between IP and fibre-channel devices. Each iSCSI target is identified by a unique iSCSI name. The SAN Volume Controller can be configured as one or more iSCSI targets. Each node that has one or both of its node Ethernet ports configured becomes an iSCSI target.

To transport SCSI commands over the IP network, an iSCSI driver must be installed on the iSCSI host and target. The driver is used to send iSCSI commands and responses through a network interface controller (NIC) or an iSCSI HBA in the host or target hardware.

For maximum performance, use a gigabit Ethernet adapter that transmits 1000 megabits per second (Mbps) for the connection between the iSCSI host and the iSCSI target.

iSCSI host connection options

Figure 23 shows an iSCSI host that connects to SAN Volume Controller over an Ethernet network.

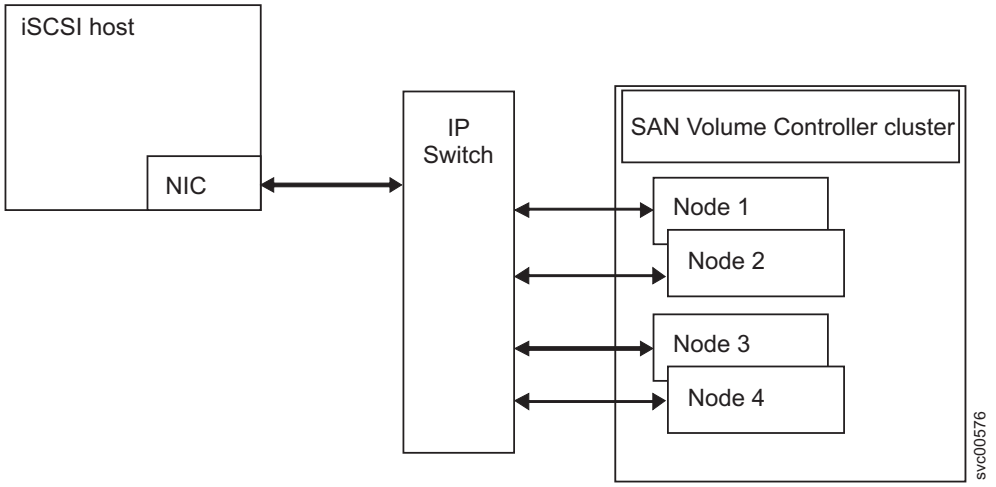

Figure 23. Transmitting SCSI over TCP/IP

Figure 24 on page 86 shows an example where the iSCSI host still connects to an Ethernet network, but a bridge or gateway continues the connection on the fibre-channel network. The bridge or gateway serves to translate between the Ethernet and fibre-channel connections so that the iSCSI host detects the SAN Volume Controller as an iSCSI target.

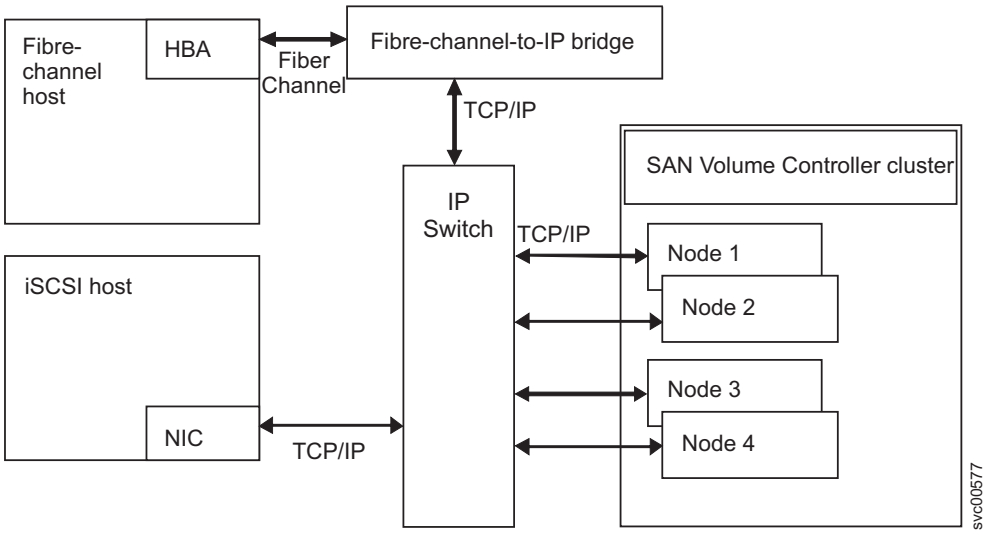

Figure 24. Transmitting SCSI over both TCP/IP and fibre-channel interconnections

Configuration rules

Storage area network (SAN) configurations that contain SAN Volume Controller nodes must be configured correctly.

A SAN configuration that contains SAN Volume Controller nodes must follow configuration rules for the following components:

- Storage systems
- Nodes
- Fibre-channel host bus adapters (HBAs)
- Fibre-channel switches
- iSCSI Ethernet ports
- Fabrics
- Zoning

Storage system configuration rules

Follow these rules when you are planning the configuration of storage systems for use with SAN Volume Controller clusters.

See the following Web site for the latest support information:

www.ibm.com/storage/support/2145

All SAN Volume Controller nodes in a cluster must be able to connect to the same set of storage system ports on each device. A cluster that contains any two nodes that cannot connect to the same set of storage-system ports is considered degraded, and a system error is logged that requires a repair action. This rule can have important effects on a storage system such as an IBM System Storage DS4000® series controller, which has exclusion rules that determine to which host bus adapter (HBA) worldwide node names (WWNNs) a storage partition can be mapped.

A storage-system logical unit (LU) must not be shared between the SAN Volume Controller and a host.

You can configure certain storage controllers to safely share resources between the SAN Volume Controller cluster and direct attached hosts. This type of configuration is described as a split controller. In all cases, it is critical that you configure the controller and SAN so that the SAN Volume Controller cluster cannot access logical units (LUs) that a host or another SAN Volume Controller cluster can also access. This split controller configuration can be arranged by controller logical unit number (LUN) mapping and masking. If the split controller configuration is not guaranteed, data corruption can occur.

Besides a configuration where a controller is split between a SAN Volume Controller cluster and a host, the SAN Volume Controller cluster also supports configurations where a controller is split between two SAN Volume Controller clusters. In all cases, it is critical that you configure the controller and SAN so that the SAN Volume Controller cluster cannot access LUs that a host or another SAN Volume Controller cluster can also access. This can be arranged by controller LUN mapping and masking. If this is not guaranteed, data corruption can occur.

Attention: Avoid configuring a storage system to present the same LU to more than one SAN Volume Controller cluster. This configuration is not supported and is likely to cause undetected data loss or corruption.

Unsupported storage systems

When a storage system is detected on the SAN, the SAN Volume Controller attempts to recognize it using its Inquiry data. If the device is not supported, 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 cluster, especially under failure scenarios. However, the SAN Volume Controller cluster does not regard accessing a generic device as an error condition and does not log an error. Managed disks (MDisks) that are presented by generic devices are not eligible to be used as quorum disks.

Split storage system configuration rules

The SAN Volume Controller cluster is configured to manage LUs that are exported only by RAID storage systems. Non-RAID storage systems are not supported. If you are using SAN Volume Controller to manage solid-state drive (SSD) or other JBOD (just a bunch of disks) LUs that are presented by non-RAID storage systems, the SAN Volume Controller cluster itself does not provide RAID functions; so these LUs are exposed to data loss in the event of a disk failure.

If a single RAID storage system presents multiple LUs, either by having multiple RAID configured or by partitioning one or more RAID into multiple LUs, each LU can be owned by either the SAN Volume Controller cluster or a direct-attach host. LUN masking must also be configured, to ensure that LUs are not shared between SAN Volume Controller nodes and direct-attach hosts.

In a split storage system configuration, a storage system presents some of its LUs to a SAN Volume Controller cluster (which treats the LU as an MDisk) and the remaining LUs to another host. The SAN Volume Controller cluster presents virtual disks (VDisks) that are created from the MDisk to another host. There is no requirement for the multipathing driver for the two hosts to be the same. Figure 25 on page 88 shows that the RAID controller could be an IBM DS4000, for example, 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 cluster and directly by the device.

Note: A connection coming from a host can be either a fibre-channel or an iSCSI connection.

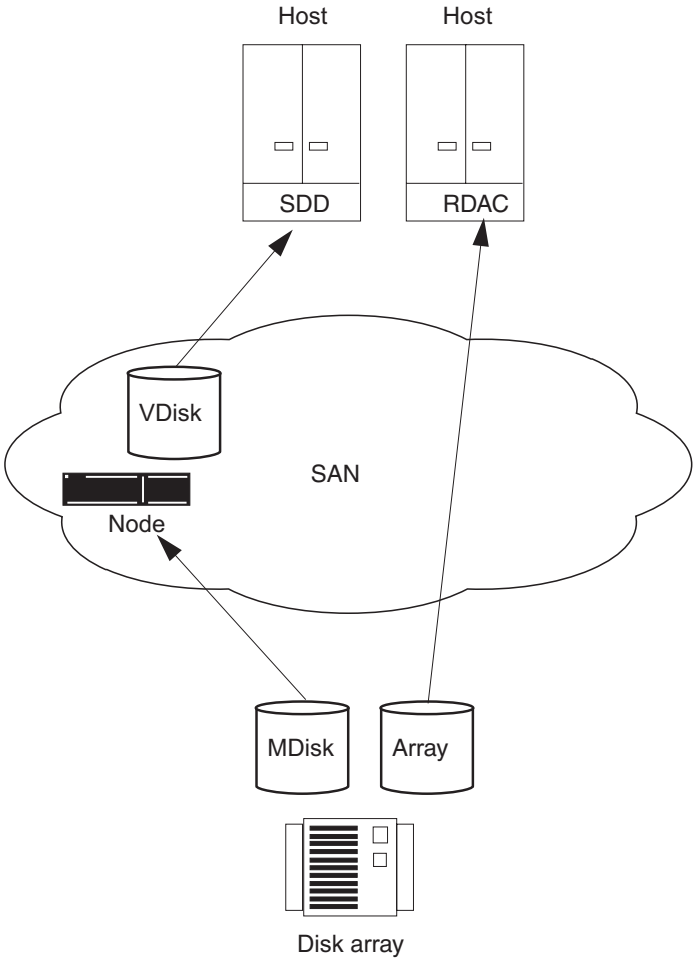

Figure 25. Storage system shared between SAN Volume Controller node and a host

It is also possible to split a host so that it accesses some of its LUNs through the SAN Volume Controller cluster and some directly. In this case, the multipathing software that is used by the storage system must be compatible with the SAN Volume Controller multipathing software. Figure 26 on page 89 is a supported configuration because the same multipathing driver is used for both directly accessed LUNs and VDisks.

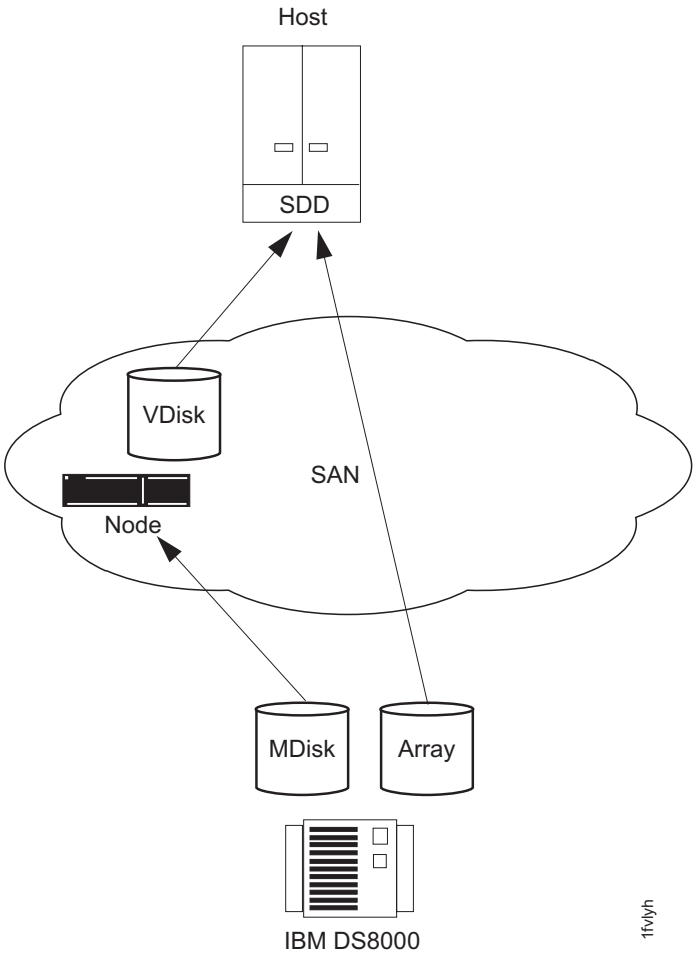

Figure 26. IBM System Storage DS8000 LUs accessed directly with a SAN Volume Controller node

In the case where the RAID storage system uses multipathing software that is compatible with SAN Volume Controller multipathing software (see Figure 27 on page 90), it is possible to configure a system where some LUNs are mapped directly to the host and others are accessed through the SAN Volume Controller. An IBM TotalStorage Enterprise Storage Server (ESS) that uses the same multipathing driver as a SAN Volume Controller node is one example. Another example with IBM DS5000 is shown in Figure 27 on page 90.

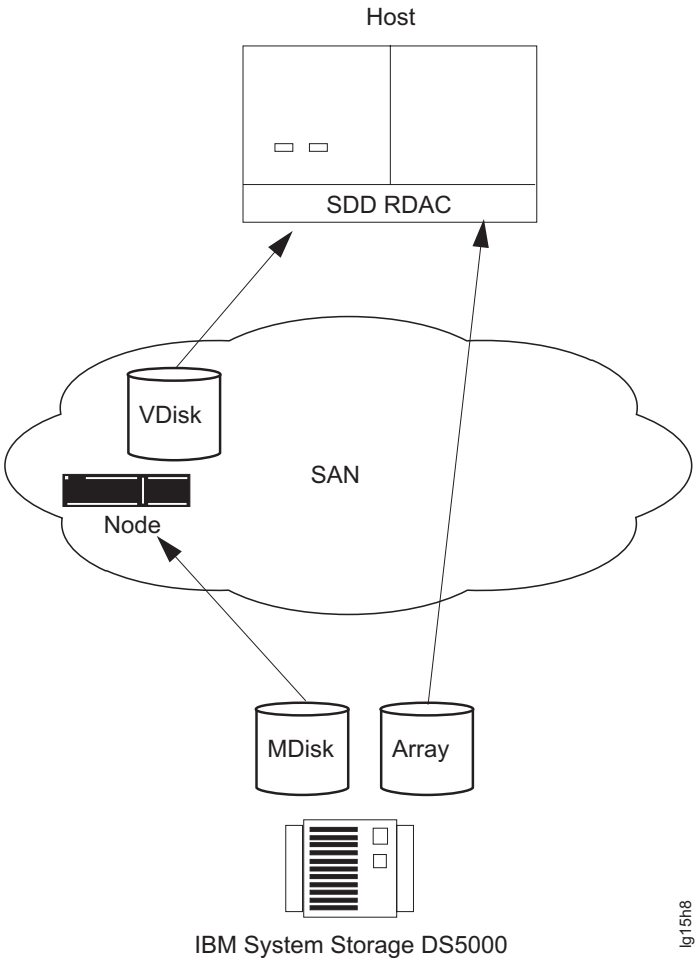

Figure 27. IBM DS5000 direct connection with a SAN Volume Controller node on one host

Fibre-channel host bus adapter configuration rules

You must follow the SAN Volume Controller configuration rules for fibre-channel host bus adapters (HBAs).

The SAN Volume Controller must be configured to export virtual disks (VDisks) only to host fibre-channel ports that are on the list of supported HBAs. See the Support for SAN Volume Controller (2145) Web site for specific firmware levels and the latest supported hardware:

www.ibm.com/storage/support/2145

Operation with other HBAs is not supported.

The SAN Volume Controller does not specify the number of host fibre-channel ports or HBAs that a host or a partition of a host can have. The number of host fibre-channel ports or HBAs are specified by the host multipathing device driver. The SAN Volume Controller supports this number; however it is subject to the configuration rules for the SAN Volume Controller. To obtain optimal performance and to prevent overloading, the workload to each SAN Volume Controller port

must be equal. You can achieve an even workload by zoning approximately the same number of host fibre-channel ports to each SAN Volume Controller fibre-channel port.

The SAN Volume Controller supports configurations that use N-port virtualization in the host bus adapter or SAN switch.

iSCSI configuration rules

You must follow the SAN Volume Controller configuration rules for iSCSI host connections.

You can attach the SAN Volume Controller to Small Computer System Interface Over Internet Protocol (iSCSI) hosts using the SAN Volume Controller's Ethernet ports.

Note: SAN Volume Controller supports SAN devices that bridge iSCSI connections into a fibre-channel network.

iSCSI connections route from hosts to the SAN Volume Controller over the LAN. You must follow the SAN Volume Controller configuration rules for iSCSI host connections:

- SAN Volume Controller supports up to 256 iSCSI sessions per node
- SAN Volume Controller currently supports one iSCSI connection per session
- SAN Volume Controller port limits are now shared between fibre-channel WWPNs and iSCSI names

Each SAN Volume Controller node has two Ethernet ports. For each Ethernet port, a maximum of one IPv4 address and one IPv6 address can be designated for iSCSI I/O.

iSCSI hosts connect to the SAN Volume Controller through the node-port IP address. If the node fails, the address becomes unavailable and the host loses communication with SAN Volume Controller. To allow hosts to maintain access to data, the node-port IP addresses for the failed node are transferred to the partner node in the I/O group. The partner node handles requests for both its own node-port IP addresses and also for node-port IP addresses on the failed node. This process is known as node-port IP failover. In addition to node-port IP addresses, the iSCSI name and iSCSI alias for the failed node are also transferred to the partner node. After the failed node recovers, the node-port IP address and the iSCSI name and alias are returned to the original node.

Multiple configurations are supported. You can have both node Ethernet ports on the same subnet, or you can have each Ethernet port on separate subnets and use different gateways. Before you configure the Ethernet ports on separate subnets, validate that the IP configuration is correct by pinging from the iSCSI host to the nodes, and vice versa.

A SAN Volume Controller VDisk can be mapped the same way either to a fibre-channel host, an iSCSI host, or both.

Each I/O group can map VDisks to the same total maximum number of host objects (256), which could include fibre-channel attachments, iSCSI attachments, or both.

SAN Volume Controller supports the following I/O descriptions:

- I/O from different initiators in the same host to the same I/O group
- I/O from different initiators in different hosts to the same VDisks
- I/O from fibre-channel and iSCSI initiators in different hosts to the same VDisks

I/O from fibre-channel and iSCSI initiators in the same hosts to the same VDisks is not supported.

A clustered Ethernet port consists of one Ethernet port from each node in the cluster that are connected to the same Ethernet switch. Because SAN Volume Controller nodes have two Ethernet ports, two clustered Ethernet ports are possible. Ethernet configuration commands can be used for clustered Ethernet ports or node Ethernet ports. SAN Volume Controller clusters can be configured with redundant Ethernet networks.

Two types of authentication through the Challenge Handshake Authentication Protocol (CHAP) are supported:

1. One-way authentication: iSCSI target (SAN Volume Controller nodes) authenticating iSCSI initiators
2. Two-way (mutual) authentication: iSCSI target (SAN Volume Controller nodes) authenticating iSCSI initiators, and vice versa.

Attention: With the iSCSI initiator, you can set two passwords: one for discovery and another for iSCSI session I/O. However, SAN Volume Controller requires that both passwords be the same.

iSCSI protocol limitations

When using an iSCSI connection, you must consider the iSCSI protocol limitations:

- There is no SLP support for discovery.
- Header and data digest support is provided only if the initiator is configured to negotiate.
- Only one connection per session is supported.
- A maximum of 256 iSCSI sessions is supported.
- Only ErrorRecoveryLevel 0 (session restart) is supported.
- The behavior of a host that supports both fibre-channel and iSCSI connections and accesses a single VDisk can be unpredictable and depends on the multipathing software.
- There can be only one session coming from one iSCSI initiator.

The following iSCSI session parameters are supported:

```
initial_r2t = 1
immediate_data = 0
max_connections = 1
Max_recv_segment_data_length = 32k
max_xmit_data_length = 32k
max_burst_length = 32k
first_burst_length = 32k
default_wait_time = 2
default_retain_time = 20
max_outstanding_r2t = 1
data_pdu_inorder = 1
data_sequence_inorder = 1
error_recovery_level = 0
header_digest = CRC32C,None
data_digest = CRC32C,None
```

```
ofmarker = 0
ifmarker = 0
ofmarkint = 2048
ifmarkint = 2048
```

Node configuration rules

You must follow the configuration rules for SAN Volume Controller nodes to ensure that you have a valid configuration.

Host bus adapters and nodes

SAN Volume Controller 2145-8F2 nodes contain two 2-port host bus adapters (HBAs). If one HBA fails, the node operates in degraded mode. If an HBA is physically removed, the configuration is not supported.

SAN Volume Controller 2145-CF8, SAN Volume Controller 2145-8F4, SAN Volume Controller 2145-8G4, and SAN Volume Controller 2145-8A4 nodes contain one 4-port HBA.

VDisks

Each node presents a virtual disk (VDisk) to the SAN through four ports. Each VDisk is accessible from the two nodes in an I/O group. Each HBA port can recognize up to eight paths to each logical unit (LU) that is presented by the cluster. The hosts must run a multipathing device driver before the multiple paths can resolve to a single device. You can use fabric zoning to reduce the number of paths to a VDisk that are visible by the host.

The number of paths through the network from an I/O group to a host must not exceed eight; configurations that exceed eight paths are not supported. Each node has four ports and each I/O group has two nodes. Therefore, without any zoning, the number of paths to a VDisk is eight multiplied by the number of host ports.

Optical connections

Valid optical connections are based on the fabric rules that the manufacturers impose for the following connection methods:

- Host to a switch
- Back end to a switch
- Interswitch links (ISLs)

Optical fiber connections can be used between a node and its switches.

Clusters that use the intercluster Metro Mirror and Global Mirror feature can use optical fiber connections between the switches, or they can use distance-extender technology that is supported by the switch manufacturer.

Ethernet connection

To ensure cluster failover operations, Ethernet port 1 on all nodes must be connected to the same set of subnets. If used, Ethernet port 2 on all nodes must also be connected to the same set of subnets. However, the subnets for Ethernet port 1 do not have to be the same as Ethernet port 2.

Physical location

The physical distance between SAN Volume Controller nodes in the same cluster is limited to 100 meters due to connectivity requirements and servicing requirements. Several of the SAN Volume Controller service actions in problem situations require that the manipulations be done to both SAN Volume Controller nodes within an I/O group or a cluster within one minute of each other. Set up your cluster environment to enable IBM service personnel to easily perform actions that are almost simultaneous in the required timeframe.

A SAN Volume Controller node must be in the same rack as the uninterruptible power supply from which it is supplied.

The depth of the SAN Volume Controller 2145-8A4 node is less than other components or nodes by approximately 127 mm or 5 inches. SAN Volume Controller 2145-8A4 nodes should not be located in the rack between components or nodes with greater depths; otherwise, it will not be possible to attach cables to a SAN Volume Controller 2145-8A4 node.

Fibre-channel connection

SAN Volume Controller supports shortwave and longwave fibre-channel connections between SAN Volume Controller nodes and the switches that they are connected to.

To avoid communication between nodes that are being routed across interswitch links (ISLs), connect all SAN Volume Controller nodes to the same fibre-channel switches.

No ISL hops are permitted among the SAN Volume Controller nodes within the same I/O group. However, one ISL hop is permitted among SAN Volume Controller nodes that are in the same cluster though different I/O groups. If your configuration requires more than one ISL hop for SAN Volume Controller nodes that are in the same cluster but in different I/O groups, contact your IBM service representative.

To avoid communication between nodes and storage systems that are being routed across ISLs, connect all storage systems to the same fibre-channel switches as the SAN Volume Controller nodes. One ISL hop between the SAN Volume Controller nodes and the storage controllers is permitted. If your configuration requires more than one ISL, contact your IBM service representative.

In larger configurations, it is common to have ISLs between host systems and the SAN Volume Controller nodes.

Port speed

The fibre-channel ports on SAN Volume Controller 2145-CF8 nodes can operate at 2 Gbps, 4 Gbps, or 8 Gbps. The fibre-channel ports on SAN Volume Controller 2145-8F4, SAN Volume Controller 2145-8G4 and SAN Volume Controller 2145-8A4 nodes can operate at 1 Gbps, 2 Gbps, or 4 Gbps. The fibre-channel ports on all these node types autonegotiate the link speed that is used with the FC switch. The ports normally operate at the maximum speed that is supported by both the SAN Volume Controller port and the switch. However, if a large number of link errors occur, the ports might operate at a lower speed than what could be supported.

Fibre-channel ports on SAN Volume Controller 2145-8F2 nodes cannot autonegotiate the speed at which they operate. You must set the required speed manually, and the optical fiber connections between the fibre-channel switches and all SAN Volume Controller 2145-8F2 nodes in a cluster must run at the same speed.

Solid-state drive (SSD) configuration rules

You must follow the SAN Volume Controller configuration rules for solid-state drives (SSDs).

Optional solid-state drives (SSDs) provide high-speed managed disk (MDisk) capability for SAN Volume Controller 2145-CF8 nodes. Each SAN Volume Controller 2145-CF8 node supports up to four SSDs. SSDs are local drives and are not accessible over the SAN fabric.

SSD configuration rules for nodes, I/O groups, and clusters

You must follow the SAN Volume Controller SSD configuration rules for nodes, I/O groups, and clusters:

- Nodes that contain SSDs can coexist in a single SAN Volume Controller cluster with any other supported nodes.
- Do not combine nodes that contain SSDs and nodes that do not contain SSDs in a single I/O group. However, while upgrading an earlier SAN Volume Controller node to a SAN Volume Controller 2145-CF8 node, you can temporarily combine the two node types in a single I/O group.
- Nodes in the same I/O group must share the same number of SSDs.
- Quorum functionality is not supported on SSDs within SAN Volume Controller nodes.

SSD configuration rules for MDisks and MDisk groups

You must follow the SAN Volume Controller SSD configuration rules for MDisks and MDisk groups:

- Each SSD is recognized by the cluster as a single MDisk.
- For each node that contains SSDs, create a single MDisk group that includes only the SSDs that are installed in that node.

SSD configuration rules for VDisks

You must follow the SAN Volume Controller SSD configuration rules for VDisks that use storage from SSDs within SAN Volume Controller nodes. In the following rules, *SAN Volume Controller SSD storage* is a managed disk group that uses SSDs within a SAN Volume Controller node.

Note: SSD storage within SAN-attached storage systems, such as the IBM DS8000, is not subject to these configuration rules.

- VDisks that use SAN Volume Controller SSD storage must be created in the I/O group that the SSDs physically reside in.
- VDisks that use SAN Volume Controller SSD storage must be mirrored to another managed disk group to provide fault tolerance. The following mirroring configurations are supported:

- To maximize performance, create the two VDisk copies in the two MDisk groups that correspond to the SAN Volume Controller SSD storage in two nodes in the same I/O group.
- To maximize utilization of SSD capacity, place the primary VDisk copy on SAN Volume Controller SSD storage, and the secondary copy on Tier 1 storage such as an IBM DS8000.

  Notes on capacity mirroring configuration:

  1. Under certain failure scenarios, VDisk performance degrades to the performance of non-SSD storage.
  2. All read I/O operations are sent to the primary copy of a mirrored VDisk, so read operations match SSD storage performance. Write I/O operations are mirrored to both locations, so write operations match the performance of the slowest copy.

- To balance the read workload, evenly split the primary and secondary VDisk copies on each node that contains SSDs.
- The preferred node for the VDisk must be the node that contains the SSDs that are used by the primary VDisk copy.
- If you shut down a node that contains unmirrored VDisks that use SAN Volume Controller SSD storage, you will lose access to any VDisks that are associated with SSD storage in that node.
- I/O requests to SSDs in other nodes are automatically forwarded, but this produces additional delays. The SSD configuration rules are designed to direct all host I/O operations to the node that contains the relevant SSDs.

SAN switch configuration

You must follow the SAN Volume Controller configuration rules for fibre-channel switches to ensure that you have a valid configuration.

The SAN must contain only supported switches.

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

www.ibm.com/storage/support/2145

Configuring your SAN with at least two independent switches, or networks of switches, ensures a redundant fabric with no single point of failure. If one of the two SAN fabrics fails, the configuration is in a degraded mode, but is still valid. A SAN with only one fabric is a valid configuration but risks loss of access to data if the fabric fails. SANs with only one fabric are exposed to a single point of failure.

Configurations with more than four SANs are not supported.

For fibre-channel connections, the SAN Volume Controller nodes must always be connected to SAN switches only. Each node must be connected to each of the counterpart SANs that are in the redundant fabric. Any fibre-channel configuration that uses a direct physical connection between a host and a SAN Volume Controller node is not supported. When attaching iSCSI hosts to SAN Volume Controller nodes, Ethernet switches must be used.

All backend storage systems must always be connected to SAN switches only. Multiple connections are permitted from redundant storage systems to improve data bandwidth performance. A connection between each redundant storage

system and each counterpart SAN is not required. For example, in an IBM System Storage DS4000 configuration in which the IBM DS4000 contains two redundant storage systems, only two storage system minihubs are usually used. Storage system A is connected to counterpart SAN A, and storage system B is connected to counterpart SAN B. Any configuration that uses a direct physical connection between the SAN Volume Controller node and the storage system is not supported.

When you attach a node to a SAN fabric that contains core directors and edge switches, connect the node ports to the core directors and connect the host ports to the edge switches. In this type of fabric, the next priority for connection to the core directors is the storage systems, leaving the host ports connected to the edge switches.

A SAN Volume Controller SAN must follow all switch manufacturer configuration rules, which might place restrictions on the configuration. Any configuration that does not follow switch manufacturer configuration rules is not supported.

Mixing manufacturer switches in a single SAN fabric

Within an individual SAN fabric, only mix switches from different vendors if the configuration is supported by the switch vendors.

Fibre-channel switches and interswitch links

The SAN Volume Controller supports distance-extender technology, including DWDM (dense wavelength division multiplexing) and FCIP (Fibre Channel over IP) extenders, to increase the overall distance between local and remote clusters. If this extender technology involves a protocol conversion, the local and remote fabrics are regarded as independent fabrics, limited to three ISL hops each.

With ISLs between nodes in the same cluster, the ISLs are considered a single point of failure. This is illustrated in Figure 28.

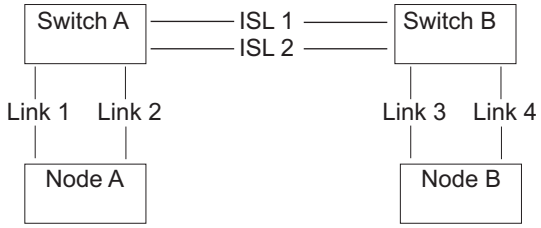

Figure 28. Fabric with ISL between nodes in a cluster

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 fails for a period of time, and the node is not 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 29 on page 98.

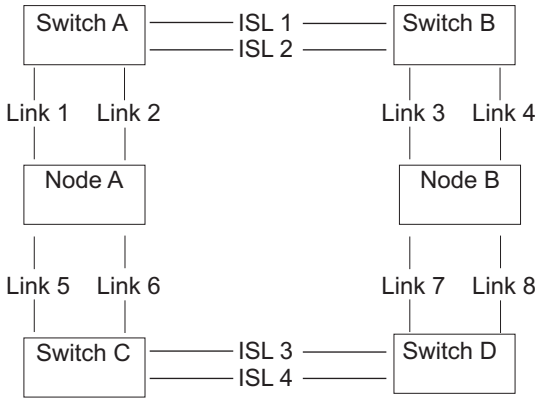

Figure 29. Fabric with ISL in a redundant configuration

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

ISL oversubscription

Perform a thorough SAN design analysis to avoid ISL congestion. Do not configure the SAN to use SAN Volume Controller to SAN Volume Controller traffic or SAN Volume Controller to storage system traffic across ISLs that are oversubscribed. For host to SAN Volume Controller traffic, do not use an ISL oversubscription ratio that is greater than 7 to 1. Congestion on the ISLs can result in severe SAN Volume Controller performance degradation and I/O errors on the host.

When you calculate oversubscription, you must account for the speed of the links. For example, if the ISLs run at 4 Gbps and the host runs at 2 Gbps, calculate the port oversubscription as 7×(4/2). In this example, the oversubscription can be 14 ports for every ISL port.

Note: The SAN Volume Controller port speed is not used in the oversubscription calculation.

SAN Volume Controller in a SAN with director class switches

You can use director class switches within the SAN to connect large numbers of RAID controllers and hosts to a SAN Volume Controller cluster. Because director class switches provide internal redundancy, one director class switch can replace a SAN that uses multiple switches. However, the director class switch provides only network redundancy; it does not protect against physical damage (for example, flood or fire), which might destroy the entire function. A tiered network of smaller switches or a core-edge topology with multiple switches in the core can provide comprehensive redundancy and more protection against physical damage for a network in a wide area. Do not use a single director class switch to provide more than one counterpart SAN because this does not constitute true redundancy.

Example SAN Volume Controller configurations

These examples show typical ways to configure your SAN Volume Controller.

Figure 30 on page 99 illustrates a small SAN configuration. Two fibre-channel switches are use to provide redundancy. Each host system, SAN Volume Controller node, and storage system is connected to both fibre-channel switches.

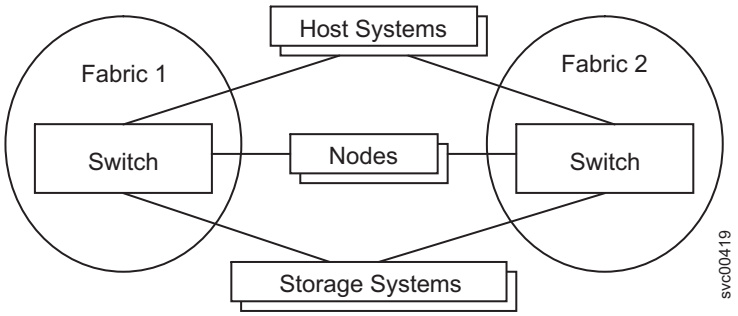

Figure 30. Simple SAN configuration

Figure 31 illustrates a medium-sized fabric that consists of three fibre-channel switches. These switches are interconnected with interswitch links (ISLs). For redundancy, use two fabrics with each host system, SAN Volume Controller node, and storage system that are being connected to both fabrics. The example fabric attaches the SAN Volume Controller nodes and the storage systems to the core switch. There are no ISL hops between SAN Volume Controller nodes or between nodes and the storage systems.

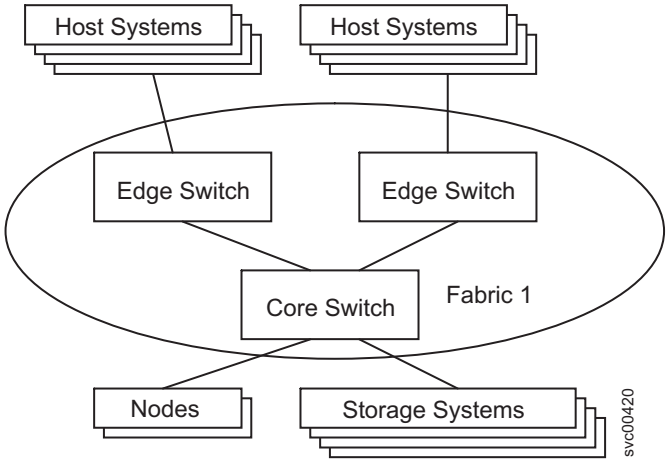

Figure 31. SAN configuration with a medium-sized fabric

Figure 32 on page 100 illustrates a large fabric that consists of two core fibre-channel switches and edge switches that are interconnected with ISLs. For redundancy, use two fabrics with each host system, SAN Volume Controller node, and storage system that is being connected. Both fabrics attach the SAN Volume Controller nodes to both core fabrics and distribute the storage systems between the two core switches. This ensures that no ISL hops exist between SAN Volume Controller nodes or between nodes and the storage systems.

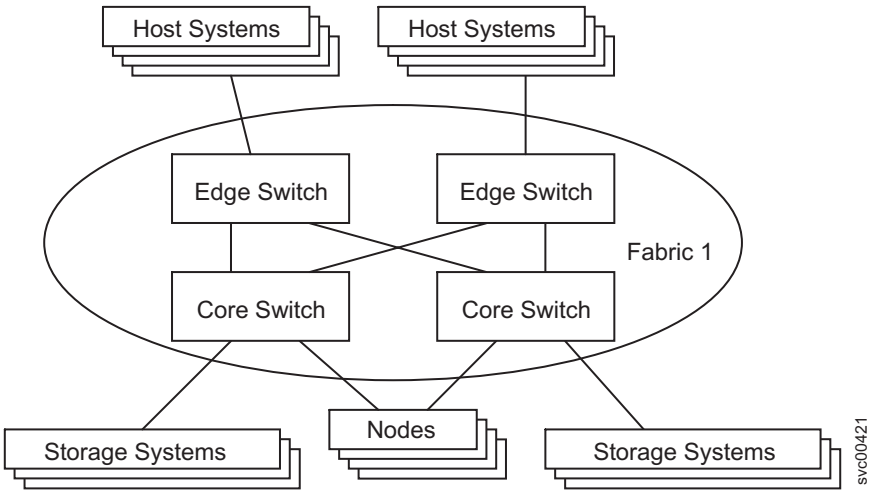

Figure 32. SAN configuration with a large fabric

Figure 33 illustrates a fabric where the host systems are located at two different sites. A long-wave optical link is used to interconnect switches at the different sites. For redundancy, use two fabrics and at least two separate long-distance links. If a large number of host systems are at the remote site, use ISL trunking to increase the available bandwidth between the two sites.

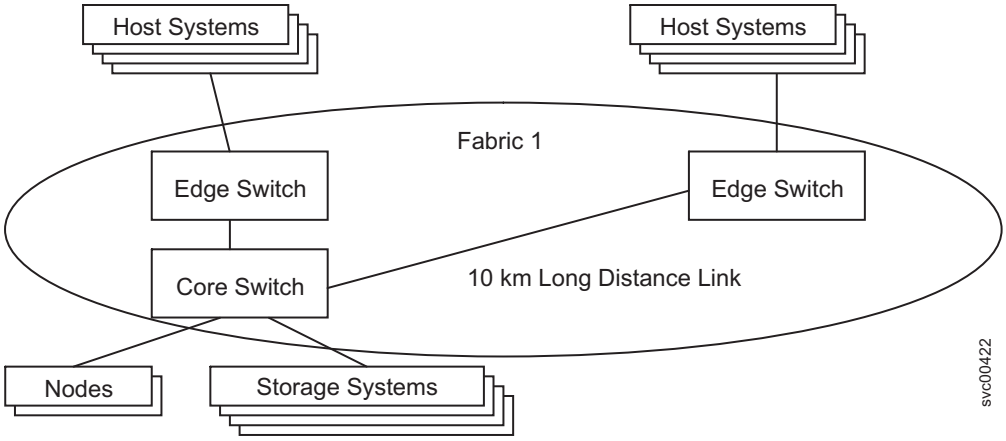

Figure 33. SAN configuration across two sites

Split cluster configuration

For high availability, you can split a SAN Volume Controller cluster across three locations and mirror the data.

To provide protection against failures that affect an entire location, such as a power failure, you can use a configuration that splits a single SAN Volume Controller cluster across three physical locations. However, you must consider that split clusters typically exhibit substantially reduced performance.

Attention: Do not separate nodes in the same I/O group by more than 10 kilometers (6.2 miles).

You must configure a split cluster to meet the following requirements:

- Directly connect each SAN Volume Controller node to one or more SAN fabrics at the primary and secondary sites. Sites are defined as independent power domains that would fail independently. Power domains could be located in the same room or across separate physical locations.
- Use a third site to house a quorum disk.
- The storage system that provides the quorum disk at the third site must support extended quorum disks. Storage systems that provide extended quorum support are listed at the following Web site:

  www.ibm.com/storage/support/2145
- Do not use powered devices to provide distance extension for the SAN Volume Controller to switch connections.
- Place independent storage systems at the primary and secondary sites, and use VDisk mirroring to mirror the host data between storage systems at the two sites.
- SAN Volume Controller nodes that are in the same I/O group and separated by more than 100 meters (109 yards) must use longwave fibre-channel connections. Longwave SFPs can be purchased as an optional SAN Volume Controller component, and must be one of the longwave SFPs listed at the following Web site:

  www.ibm.com/storage/support/2145
- Using inter-switch links (ISLs) in paths between SAN Volume Controller nodes in the same I/O group is not supported.
- Avoid using inter-switch links (ISLs) in paths between SAN Volume Controller nodes and external storage systems. If this is unavoidable, do not oversubscribe the ISLs because of substantial fibre-channel traffic across the ISLs. For most configurations, trunking is required. Because ISL problems are difficult to diagnose, switch-port error statistics must be collected and regularly monitored to detect failures.
- Using a single switch at the third site can lead to the creation of a single fabric rather than two independent and redundant fabrics. A single fabric is an unsupported configuration.
- SAN Volume Controller nodes in the same cluster must be connected to the same Ethernet subnet.
- A SAN Volume Controller node must be located in the same rack as the 2145 UPS or 2145 UPS-1U that supplies its power.
- Some service actions require physical access to all SAN Volume Controller nodes in a cluster. If nodes in a split cluster and separated by more than 100 meters, service actions might require multiple service personnel. Contact your IBM service representative to inquire about multiple site support.

A split cluster configuration locates the active quorum disk at a third site. If communication is lost between the primary and secondary sites, the site with access to the active quorum disk continues to process transactions. If communication is lost to the active quorum disk, an alternative quorum disk at another site can become the active quorum disk.

Although a cluster of SAN Volume Controller nodes can be configured to use up to three quorum disks, only one quorum disk can be elected to resolve a situation where the cluster is partitioned into two sets of nodes of equal size. The purpose of the other quorum disks is to provide redundancy if a quorum disk fails before the cluster is partitioned.

Figure 34 illustrates an example split cluster configuration. When used in conjunction with VDisk mirroring, this configuration provides a high availability solution that is tolerant of a failure at a single site. If either the primary or secondary site fails, the remaining sites can continue performing I/O operations. In this configuration, the connections between SAN Volume Controller nodes in the cluster are greater than 100 meters apart, and therefore must be longwave fibre-channel connections.

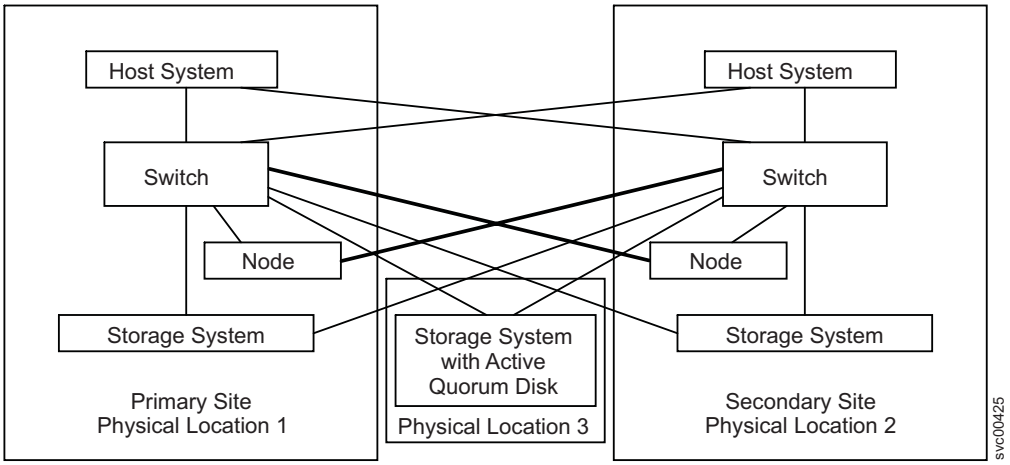

Figure 34. A split cluster with a quorum disk located at a third site

In Figure 34, the storage system that hosts the quorum disks is attached directly to a switch at both the primary and secondary sites using longwave fibre-channel connections. If either the primary site or the secondary site fails, you must ensure that the remaining site has retained direct access to the storage system that hosts the quorum disks.

An alternative configuration can use an additional fibre-channel switch at the third site with connections from that switch to the primary site and to the secondary site. This type of split-site configuration is supported only when the storage system that hosts the quorum disks supports extended quorum. Although SAN Volume Controller can use other types of storage systems for providing quorum disks, access to these quorum disks is always through a single path.

For quorum disk configuration requirements, see the *Guidance for Identifying and Changing Managed Disks Assigned as Quorum Disk Candidates* technote at the following Web site:

http://www.ibm.com/support/docview.wss?rs=591&uid=ssg1S1003311

Zoning guidelines

Ensure that you are familiar with the zoning guidelines for storage system zones and host zones.

Paths to hosts

The number of paths through the network from the SAN Volume Controller nodes to a host must not exceed eight. Configurations in which this number is exceeded are not supported.

- Each node has four ports and each I/O group has two nodes. Therefore, without any zoning, the number of paths to a VDisk would be eight multiplied by the number of host ports.
- This rule exists to limit the number of paths that must be resolved by the multipathing device driver.

If you want to restrict the number of paths to a host, zone the switches so that each host bus adapter (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.

Storage system zones

Switch zones that contain storage system ports must not have more than 40 ports. A configuration that exceeds 40 ports is not supported.

SAN Volume Controller zones

The switch fabric must be zoned so that the SAN Volume Controller nodes can detect the back-end storage systems and the front-end host HBAs. Typically, the front-end host HBAs and the back-end storage systems are not in the same zone. The exception to this is where split host and split storage system configuration is in use.

All nodes in a cluster must be able to detect the same ports on each back-end storage system. Operation in a mode where two nodes detect a different set of ports on the same storage system is degraded, and the system logs errors that request a repair action. This can occur if inappropriate zoning is applied to the fabric or if inappropriate LUN masking is used. This rule has important implications for back-end storage, such as IBM DS4000 storage systems, which impose exclusive rules for mappings between HBA worldwide node names (WWNNs) and storage partitions.

Each SAN Volume Controller port must be zoned so that it can be used for internode communications. When configuring switch zoning, you can zone some SAN Volume Controller node ports to a host or to back-end storage systems.

When configuring zones for communication between nodes in the same cluster, the minimum configuration requires that all fibre-channel ports on a node detect at least one fibre-channel port on each other node in the same cluster. You cannot reduce the configuration in this environment.

It is critical that you configure storage systems and the SAN so that a cluster cannot access logical units (LUs) that a host or another cluster can also access. You can achieve this configuration with storage system logical unit number (LUN) mapping and masking.

If a node can detect a storage system through multiple paths, use zoning to restrict communication to those paths that do not travel over ISLs.

With Metro Mirror and Global Mirror configurations, 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 systems and local nodes or remote nodes, or both, is not valid.

For clusters that are running SAN Volume Controller version 5.1, configure your system so that all fibre-channel node ports detect at least one fibre-channel port on each node in the remote cluster. For best results in Metro Mirror and Global Mirror configurations, zone each node so that it can communicate with at least one fibre-channel port on each node in each remote cluster. This configuration maintains redundancy of the fault tolerance of port and node failures within local and remote clusters. For communications between multiple SAN Volume Controller version 5.1 clusters, this also achieves optimal performance from the nodes and the intercluster links.

However, to accommodate the limitations of some switch vendors on the number of ports or worldwide node names (WWNNs) that are allowed in a zone, you can further reduce the number of ports or WWNNs in a zone. Such a reduction can result in reduced redundancy and additional workload being placed on other cluster nodes and the fibre-channel links between the nodes of a cluster.

The minimum configuration requirement is to zone both nodes in one I/O group to both nodes in one I/O group at the secondary site. The I/O group maintains fault tolerance of a node or port failure at either the local or remote site location. It does not matter which I/O groups at either site are zoned because I/O traffic can be routed through other nodes to get to the destination. However, if an I/O group that is doing the routing contains the nodes that are servicing the host I/O, there is no additional burden or latency for those I/O groups because the I/O group nodes are directly connected to the remote cluster.

For clusters that are running SAN Volume Controller version 4.3.1 or earlier, the minimum configuration requirement is that all nodes must detect at least one fibre-channel port on each node in the remote cluster. You cannot reduce the configuration in this environment.

In configurations with a version 5.1 cluster that is partnered with a cluster that is running a SAN Volume Controller version 4.3.1 or earlier, the minimum configuration requirements of the version 4.3.1 or earlier cluster apply.

If only a subset of the I/O groups within a cluster are using Metro Mirror and Global Mirror, you can restrict the zoning so that only those nodes can communicate with nodes in remote clusters. You can have nodes that are not members of any cluster zoned to detect all of the clusters. You can then add a node to the cluster in the event that you must replace a node.

Host zones

The configuration rules for host zones are different depending upon the number of hosts that will access the cluster. For configurations of less than 64 hosts per cluster, the SAN Volume Controller supports a simple set of zoning rules that enable a small set of host zones to be created for different environments. For configurations of more than 64 hosts per cluster, the SAN Volume Controller supports a more restrictive set of host zoning rules.

Zoning that contains host HBAs must ensure host HBAs in dissimilar hosts or dissimilar HBAs 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.

To obtain the best overall performance of the system and to prevent overloading, the workload to each SAN Volume Controller port must be equal. This can

typically involve zoning approximately the same number of host fibre-channel ports to each SAN Volume Controller fibre-channel port.

Clusters with less than 64 hosts

For clusters with less than 64 hosts attached, zones that contain host HBAs must contain no more than 40 initiators including the SAN Volume Controller ports that act as initiators. A configuration that exceeds 40 initiators is not supported. A valid zone can be 32 host ports plus 8 SAN Volume Controller ports. When it is possible, place each HBA port in a host that connects to a node into a separate zone. Include exactly one port from each node in the I/O groups that are associated with this host. This type of host zoning is not mandatory, but is preferred for smaller configurations.

Note: If the switch vendor recommends fewer ports per zone for a particular SAN, the rules that are imposed by the vendor takes precedence over the SAN Volume Controller rules.

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

Clusters with more than 64 hosts

Each HBA port must be in a separate zone and each zone must contain exactly one port from each SAN Volume Controller node in each I/O group that the host accesses.

Note: A host can be associated with more than one I/O group and therefore access VDisks from different I/O groups in a SAN. However, this reduces the maximum number of hosts that can be used in the SAN. For example, if the same host uses VDisks in two different I/O groups, this consumes one of the 256 hosts in each I/O group. If each host accesses VDisks in every I/O group, there can be only 256 hosts in the configuration.

Zoning examples

These examples describe ways for zoning a switch.

Example 1

Consider the SAN environment in the following example:

- Two nodes (nodes A and B)
- Nodes A and B each have four ports
  - Node A has ports A0, A1, A2, and A3
  - Node B has ports B0, B1, B2, and B3
- Four hosts called P, Q, R, and S
- Each of the four hosts has two ports, as described in Table 21.

Table 21. Four hosts and their ports

P	Q	R	S
P0	Q0	R0	S0
P1	Q1	R1	S1

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

The following is an example configuration:

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.

Create the following host zones on switch X:

5. Create one zone per host port (one port per node) (A0, B0, P0, A0, B0, and Q0)
6. Create one zone per host port (one port per node) (A0, B0, R0, A0, B0, and S0)

Create the following host zones on switch Y:

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.

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.

Example 2

The following example describes a SAN environment that is similar to the previous example except for the addition of two hosts that have two ports each.

- Two nodes called A and B
- Nodes A and B have four ports each
  - Node A has ports A0, A1, A2, and A3
  - 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 the other two hosts have two ports each as described in Table 22.

Table 22. 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

The following is an example configuration:

```
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.

Create the following host zones on switch X:

```
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.
```

Create the following host zones on switch Y:

```
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.
```

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.
```

Zoning considerations for Metro Mirror and Global Mirror

Ensure that you are familiar with the constraints for zoning a switch to support the Metro Mirror and Global Mirror feature.

SAN configurations that use intracluster Metro Mirror and Global Mirror relationships do not require additional switch zones.

SAN configurations that use intercluster Metro Mirror and Global Mirror relationships require the following additional switch zoning considerations:

- A cluster can be configured so that it can detect all nodes in all remote clusters. Alternatively, a cluster can be configured so that it detects only a subset of nodes in remote clusters. For I/O groups that share an intercluster Metro Mirror or Global Mirror relationship, all node ports in each I/O group must be zoned to detect all node ports in the other I/O group.
- Use of interswitch link (ISL) trunking in a switched fabric.
- Use of redundant fabrics.

For intercluster Metro Mirror and Global Mirror relationships, you must perform the following steps to create the additional zones that are required:

1. Configure your SAN so that fibre-channel traffic can be passed between the two clusters. To configure the SAN this way, you can connect the clusters to the same SAN, merge the SANs, or use routing technologies.
2. (Optional) Configure zoning to allow all nodes in the local fabric to communicate with all nodes in the remote fabric.

   Note:

a. If you are using McData Eclipse routers, model 1620, only 64 port pairs are supported, regardless of the number of iFCP links that are used.

3. (Optional) As an alternative to step 2 on page 107, choose a subset of nodes in the local cluster to be zoned to the nodes in the remote cluster. Minimally, you must ensure that one whole I/O group in the local cluster has connectivity to one whole I/O group in the remote cluster. I/O between the nodes in each cluster is then routed to find a path that is permitted by the configured zoning.

   Reducing the number of nodes that are zoned together can reduce the complexity of the intercluster zoning and might reduce the cost of the routing hardware that is required for large installations. Reducing the number of nodes also means that I/O must make extra hops between the nodes in the system, which increases the load on the intermediate nodes and can increase the performance impact; in particular, for Metro Mirror.

4. Optionally, modify the zoning so that the hosts that are visible to the local cluster can recognize the remote cluster. This allows a host to examine data in both the local and remote cluster.

5. Verify that cluster A cannot recognize any of the back-end storage that is owned by cluster B. A cluster cannot access logical units (LUs) that a host or another cluster can also access.

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

The following table provides a description of the ISL trunking and the extended fabric features:

Feature	Description
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 might 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 cannot be used together. Although it is possible to enable extended fabric operation to one link of a trunked pair, this does not offer any performance advantages and adds complexity to the configuration setup. Therefore, do not use mixed mode operations.

Feature	Description
Extended fabric	Extended fabric operation allocates extra buffer credits to a port. This is important over long links that are usually found in intercluster Metro Mirror operation and Global Mirror operations. Because of the time that it takes for a frame to traverse the link, it is possible to have more frames in transmission at any instant in time than is 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. • The Normal option is suitable for short links. • The Extended Normal option provides significantly better performance for the links up to 10 km long. **Note:** The extended fabric license provides two extra options: Medium, 10 - 50 km and Long, 50 - 100 km.

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 run 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 total number of VDisk-to-host mappings is at least 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 run in parallel on a device. It is usually possible to set a limit on the queue depth on the subsystem device driver (SDD) paths (or equivalent) or the host bus adapter (HBA).

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 virtual disks (VDisks).

Note: You might have a number of servers in the configuration that are idle or do not initiate the calculated quantity of I/O operations. If so, you might 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.

The formula for queue depth calculation considers the following factors:

- The maximum number of queued commands is per node and there are two nodes in an input/output (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.

- If a virtual disk (VDisk) is mapped so that it can be seen by more than one server, then each of the servers can send commands to it.
- If a device driver times out of a command, it immediately reissues the command. The SAN Volume Controller will have both commands in its command queue.

Homogeneous queue depth calculation

Ensure you are familiar with the homogeneous queue depth calculation.

The homogeneous queues must meet one of the following statements:

- The queued commands are shared among all paths rather than providing servers with additional resources.
- The virtual disks (VDisks) are distributed evenly among the input/output (I/O) groups in the cluster.

You can set the queue depth for each VDisk on the servers using the following calculation:

q = ((n ×7000) / (v×p×c))

where:

- *q* is the queue depth per device path
- *n* is the number of nodes in the cluster
- *v* is the number of VDisks configured in the cluster
- *p* is the 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* is the 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 is 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.

Nonhomogeneous queue depth calculation

For nonhomogeneous queues, use the following calculation.

Nonhomogeneous 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.

Set the queue depth for each VDisk 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, the server will not experience problems due to a full queue.

Limiting the queue depth

After 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 virtual disk (VDisk) basis.

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

Supported fibre-channel extenders

Fibre-channel extenders extend a fibre-channel link by transmitting fibre-channel packets across long links without changing the contents of those packets.

IBM has tested a number of such fibre-channel extender technologies with SAN Volume Controller and supports fibre-channel extenders of all types for intercluster links, provided they meet the latency requirements for Metro Mirror and Global Mirror.

When you are planning to use fibre-channel extenders, 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 microseconds per kilometer. Bit error rates vary depending on the quality of the circuit that is 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.

See the following Web site for the latest supported hardware:

www.ibm.com/storage/support/2145

Performance of fibre-channel extenders

When you are planning to use fibre-channel extenders, 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 microseconds per kilometer. Bit error rates vary depending on the quality of the circuit that is 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.

Chapter 4. Creating a SAN Volume Controller cluster

After the IBM System Storage Productivity Center is configured, you must complete the two phases that are required to create a SAN Volume Controller cluster before you can configure the cluster.

The first phase to create a cluster is performed from the front panel of the SAN Volume Controller. The second phase is performed from the SAN Volume Controller Console, which is accessible from a Web server that runs on the IBM System Storage Productivity Center (SSPC), or in previous releases, the master console.

Before you can access the SAN Volume Controller Console and command-line interface (CLI), you must configure the IBM System Storage Productivity Center. To access the CLI, this also includes using the PuTTY client to generate Secure Shell (SSH) key pairs that secure data flow between the SAN Volume Controller cluster configuration node and a client. You do not need an SSH key pair if you want to use just the SSPC. For more information, see the *IBM System Storage Productivity Center User's Guide*.

Initiating cluster creation from the front panel

After you have installed a pair of nodes, you can use the front panel of one of the SAN Volume Controller nodes to initiate the creation of the cluster. To create a cluster, you do not need to repeat these instructions on more than one node. After you complete the steps for initiating cluster creation from the front panel, you can use the SAN Volume Controller Console to create the cluster and add additional nodes to complete cluster configuration.

Before you create a cluster, ensure that you verify the SAN Volume Controller installation after the installation has completed.

When you create the cluster, you must specify either an IPv4 or an IPv6 cluster address for port 1. After the cluster is created, you can specify additional IP addresses for port 1 and port 2 until both ports have an IPv4 address and an IPv6 address.

If you choose to have the IBM service representative or IBM Business Partner initially create the cluster, you must provide the following information prior to configuring the cluster:

- For a cluster with an IPv4 address:
  - Cluster IPv4 address
  - Subnet mask
  - Gateway IPv4 address
- For a cluster with an IPv6 address
  - Cluster IPv6 address
  - IPv6 prefix
  - Gateway IPv6 address

These addresses should be defined on the Configuration Data Table planning chart which is used when installing a cluster.

Attention: The Cluster IPv4 address and the IPv6 address must not be the same as any other device accessible on the network.

The IBM service representative or IBM Business Partner uses the front panel of the node to enter the information that you have provided. The node generates a random password on the display panel. The IBM service representative or IBM Business Partner gives you this password. You must record the password and the IPv4 address or the IPv6 address. The password and the IP address are used to connect to the node and to create the cluster.

Using the front panel, follow these steps to create and configure the cluster:

1. Choose a node that you want to make a member of the cluster that you are creating.

   Note: You can add additional nodes after you have successfully created and initialized the cluster.
2. Press and release the up or down button until `Node:` is displayed on the node service panel.
3. Press and release the left or right button until `Create Cluster?` is displayed.

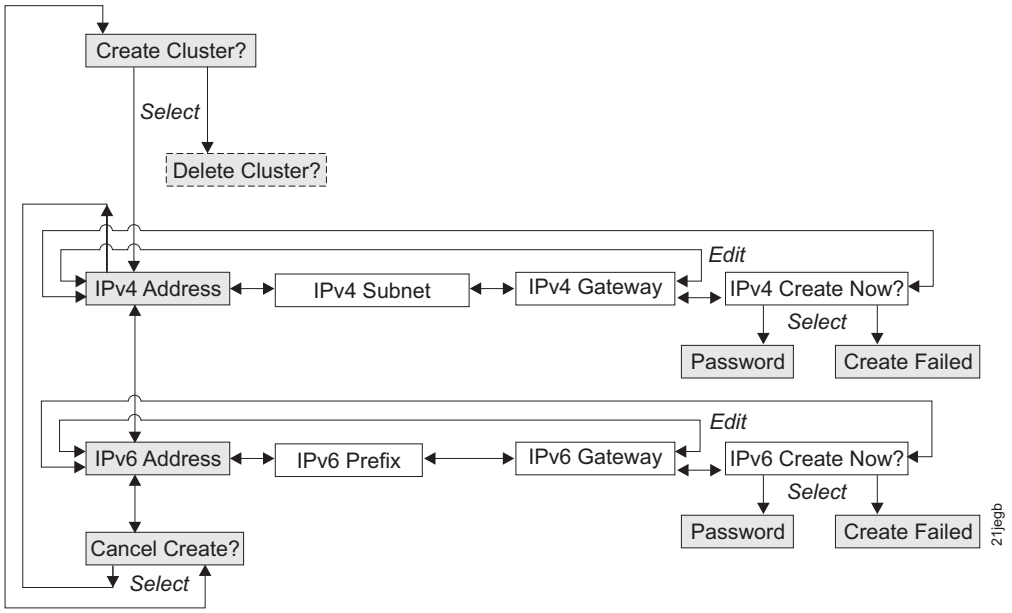

Figure 35. Create Cluster? navigation

4. Press and release the select button.
   - If `Delete Cluster?` is displayed on the first line of the service display panel, this node is already a member of a cluster. Press and release the Up button until `Cluster:` is displayed on the service display panel. The name of the cluster the node belongs to is displayed on line 2 of the service display panel. If you want to delete the node from this cluster, refer to the instructions in "Deleting a node from a cluster using the SAN Volume Controller Console" on page 143. If you do not want to delete this node from the cluster, review the situation and determine the correct nodes to include in the new cluster. Then go to step 1 and begin again.
   - If you are creating a cluster with an IPv4 address and `IPv4 Address:` is displayed on line 1 of the panel, go to "Creating a cluster with an IPv4 address" on page 115.

- If you are creating a cluster with an IPv6 address, press and release the down button to see `IPv6 Address:` on line 1 of the panel. Go to "Creating a cluster with an IPv6 address" on page 116.

Creating a cluster with an IPv4 address

Your cluster IP address can be an IPv4 or IPv6 address.

To create a cluster with an IPv4 address, complete the following steps:

1. Put the panel into edit mode by pressing and releasing the select button. The first IPv4 address number is highlighted.
2. Press the up button if you want to increase the value that is displayed; press the down button if you want to decrease that value. If you want to quickly increase the highlighted value, hold the up button. If you want to quickly decrease the highlighted value, hold the down button.

   Note: To change the address scrolling speed, see the note at the end of this section.
3. Press the right or left buttons to move to the number field that you want to update.
4. Use the right button to move to the next field and use the up or down button to change the value of this field.
5. Repeat step 4 for each of the remaining fields of the IPv4 Address.
6. After you have changed the last field of the IPv4 Address, press and release the select button to put the data in view rather than edit mode.
7. Press the right button. `IPv4 Subnet:` is displayed.
8. Press and release the select button.
9. Use the up or down button to quickly increase or decrease the value of the first field of the IPv4 Subnet to the value that you have chosen.

   Note: To change the address scrolling speed, see the note at the end of this section.
10. Use the right button to move to the next field and use the up or down buttons to change the value of this field.
11. Repeat step 10 for each of the remaining fields of the IPv4 Subnet.
12. After you have changed the last field of IPv4 Subnet, press the select button to put the data in view mode.
13. Press the right button. `IPv4 Gateway:` is displayed.
14. Press and release the select button.
15. Press the up button if you want to increase the value that is displayed; press the down button if you want to decrease that value. If you want to quickly increase the highlighted value, hold the up button. If you want to quickly decrease the highlighted value, hold the down button.

    Note: To change the address scrolling speed, see the note at the end of this section.
16. Use the right button to move to the next field and use the up or down button to change the value of this field.
17. Repeat step 16 for each of the remaining fields of the IPv4 Gateway.
18. After you have changed the last field of IPv4 Gateway, press the select button to put the data in view mode.
19. Press and release the right button until `IPv4 Create Now?` is displayed.

20. Ensure that you have a pen and paper to record the cluster superuser password before you create the cluster.

    Attention: The password is displayed for 60 seconds, or until the up, down, left or right button is pressed, which deletes it. You must be ready to record the password before you select one of the following actions:

    - 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 `IPv4 Create Now?`, and then press the select button.
    - If you are satisfied with your settings, press the select button. If the cluster is created successfully, `Password:` is displayed on line 1 of the service display screen. Line 2 contains a password that you can use to access the cluster. Record this password now.

      Important: If you do not record the password, you must restart the cluster configuration procedure. After you have recorded the password, press the up, down, left, or right button to clear the password from the screen.

After you complete this task, the following information is displayed on the service display screen:

- `Cluster:` is displayed on line 1.
- A temporary, system-assigned cluster name that is based on the IP address is displayed on line 2.

Note: To disable the fast increase and decrease address scrolling speed function using the front panel, press and hold the down arrow button, press and release the select button, and then release the down arrow button. The disabling of the fast increase and decrease function lasts until cluster creation is completed or until the feature is enabled again. If you press and hold the up or down arrow button while the function is disabled, the value increases or decreases once every two seconds. To enable the fast increase and decrease function again, press and hold the up arrow button, press and release the select button, and then release the up arrow button.

After you have created the cluster on the front panel with the correct IP address format, you can finish the cluster configuration by accessing the SAN Volume Controller Console, completing the Create Cluster wizard, and adding additional nodes to the cluster.

Creating a cluster with an IPv6 address

Your cluster IP address can be an IPv4 or IPv6 address.

To create the cluster with an IPv6 address, complete the following steps:

1. From the IPv4 Address panel, press the down button. The IPv6 Address option is displayed.
2. Press the select button again to put the panel into edit mode. The IPv6 address and the IPv6 gateway address consist of eight 4-digit hexadecimal values. Enter the full address by working across a series of four panels to update each of the 4-digit hexadecimal values that make up the IPv6 addresses. The panels consists of eight fields, where each field is a 4-digit hexadecimal value.
3. Press the right button or left button to move to the number field that you want to set.

4. Use the right button to move to the next field and use the up or down button to change the value of this field.
5. Repeat step 4 for each of the remaining fields of the IPv6 Address.
6. After you have changed the last field of the IPv6 Address, press and release the select button to put the data in view mode.
7. Press and release the right button until `IPv6 Prefix:` is displayed.
8. Press and release the select button.
9. Use the up or down button to quickly increase or decrease the value of the first field of the IPv6 Prefix to the value that you have chosen.

   Note: To change the address scrolling speed, see the note at the end of this section.
10. Press the select button to put the data in view mode.
11. Press the right button. `IPv6 Gateway:` is displayed.
12. Press and release the select button.
13. Use the right button to move to the next field and use the up or down button to change the value of this field.
14. Repeat step 13 for each of the remaining fields of the IPv6 Gateway.
15. After you have changed the last field of IPv6 Gateway, press the select button to put the data in view mode.
16. Press and release the right button until `IPv6 Create Now?` is displayed.
17. Ensure that you have a pen and paper to record the cluster superuser password before you create the cluster.

    Attention: The password is displayed for 60 seconds, or until the up, down, left or right button is pressed, which deletes it. You must be ready to record the password before you select one of the following actions:

    - 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 `IPv6 Create Now?`, and then press the select button.
    - If you are satisfied with your settings, press the select button. If the cluster is created successfully, `Password:` is displayed on line 1 of the service display panel. Line 2 contains a password that you must use when you first access the cluster. Record this password now.

      Important: If you do not record the password, you must restart the cluster configuration procedure. After you have recorded the password, press the up, down, left, or right button to clear the password from the screen.

After you complete this task, the following information is displayed on the service display screen:

- `Cluster:` is displayed on line 1.
- A temporary, system-assigned cluster name that is based on the IP address is displayed on line 2.

Note: To disable the fast increase and decrease address scrolling speed function using the front panel, press and hold the down arrow button, press and release the select button, and then release the down arrow button. The disabling of the fast increase and decrease function lasts until cluster creation is completed or until the feature is enabled again. If you press and hold the up or down arrow button while the function is disabled, the value

increases or decreases once every two seconds. To enable the fast increase and decrease function again, press and hold the up arrow button, press and release the select button, and then release the up arrow button.

After you have created the cluster on the front panel with the correct IP address format, you can finish the cluster configuration by accessing the SAN Volume Controller Console, completing the Create Cluster wizard, and adding additional nodes to the cluster.

Creating a cluster using the SAN Volume Controller Console

After you have created the cluster using the SAN Volume Controller front panel, you can use the Add SAN Volume Controller Cluster function from the SAN Volume Controller Console to identify the cluster to the IBM System Storage Productivity Center (SSPC) or the master console.

Complete the following steps to create a cluster:

1. Start the SAN Volume Controller Console by selecting **Start** → **All Programs** → **IBM System Storage SAN Volume Controller** → **Launch SVC Console** on your IBM System Storage Productivity Center (SSPC) . The IBM System Storage SAN Volume Controller Welcome panel is displayed.
2. If this is the first time that you have accessed the SAN Volume Controller Console, go to step 3. Otherwise, go to step 4.
3. Click **Add SAN Volume Controller Cluster** from the Welcome panel. The Adding a Cluster panel displays. Go to step 6 and proceed.
4. Select **Clusters** from the portfolio. The Viewing Clusters panel is displayed.
5. From the task list, select **Add a Cluster** and click **Go**.
6. The Adding a Cluster panel is displayed. Type the IP address of the cluster. This address should be the same IP address that you entered on the front panel. The SAN Volume Controller Console supports both IPv4 and IPv6 address formats. For IPv4, SAN Volume Controller Console supports the standard format for these addresses; 208.77.188.166 is an example of an IPv4 address. For IPv6 addresses, the following formats are supported:
   - Eight colon-separated groups of four hexadecimal digits; for example, 1234:1234:abcd:0123:0000:0000:7689:6576
   - Eight colon-separated groups of hexadecimal digits with the leading zeros omitted; for example, 1234:1234:abcd:123:0:0:7689:6576
   - Zero suppression format; for example, 1234:1234:abcd:123::7689:6576

   Note: You can only suppress one set of zeros in an address.
7. Select the **Create (Initialize) Cluster** check box to create the new cluster. If the cluster is already in use and you are adding this cluster to the list of managed clusters for this installation of the SAN Volume Controller Console, *do not* select the Create (Initialize) Cluster check box.

   Click **OK**. Depending on the browser type and version, a security alert might display. Complete the instructions that are presented on the browser's security panels.
8. The Connecting to *ipaddress* panel is displayed, where *ipaddress* is the IP address of the system that you are connecting to. Type the cluster user name `superuser` and the password that was generated when you created the cluster from the front panel.
9. Click **OK**.

10. On the Welcome panel of the Create Cluster wizard, click **Continue**. The Enter Cluster Settings panel is displayed.
11. Complete the Enter Cluster Settings panel.
    a. Type a name for your cluster. A valid cluster name is 1 to 15 ASCII characters. The following characters can be used: a - z, A - Z, 0 - 9, -, or _. The cluster name cannot begin with a number or the dash (-) character. You must choose a name that is different from any other cluster name in your storage area network, even remotely attached clusters. The cluster name forms part of the iSCSI qualified name (IQN) that might be used when connecting hosts to SAN Volume Controller using iSCSI. The cluster name should therefore be finalized before configuring any iSCSI hosts.
    b. Type the cluster superuser password that you want to use to access all cluster maintenance functions. This password must be different from the temporary cluster superuser password that was initially provided on the front panel of the node. The temporary cluster superuser password is used to create the cluster and to initially sign on to the cluster. Passwords must meet the following requirements:
       - Passwords are a maximum of 6 to 64 printable ASCII characters.
       - Passwords are case-sensitive.
       - The first or last character cannot be a blank character.

       Important: Record this password because you will need it to access the cluster after it is created.
    c. If you want the ability to reset the cluster superuser password from the front panel, accept the default setting for the **Allow cluster superuser password reset from the front panel** check box. With this option, you can reset the cluster superuser password from the front panel if the password is forgotten. You must ensure adequate physical security to the cluster hardware. If you do not have adequate physical security for the cluster hardware, deselect this option.
    d. Type the service password. The service password is used for routine cluster service activities. With a service password, the service user can access a limited subset of the maintenance functions available that is defined for the superuser. This can be useful if you want to give different users different levels of maintenance access. Re-enter the service password to verify the password.

       Passwords must meet the following requirements:
       - Passwords are a maximum of 15 alphanumeric characters.
       - Passwords are case-sensitive.
       - Valid characters are uppercase letters (A - Z), lowercase letters (a - z), digits (0 - 9), dash (-), and underscore (_).
       - The first character cannot be a dash (-).

       Important: Record this password because you will need it if you cannot access the cluster using the cluster superuser ID and password.
    e. Type the service IP address for the cluster. The system uses the service IP address when an individual node is in service mode. The service mode IP address can be either a fixed IP address or Dynamic Host Configuration Protocol (DHCP) IP address. When using a fixed IP address, you cannot mix IPv4 and IPv6 addresses in the basic cluster configuration. Therefore, if the cluster IP address, previously configured on the front panel, is an IPv6 address, then the service IP address must also be an IPv6 address and

similarly for IPv4. When using DHCP, an IP address is dynamically allocated when the node enters service mode. With DHCP, you can have multiple nodes in service mode simultaneously. Service mode IP addresses, including DHCP-assigned addresses, display on the front panel when a node is in service mode.

f. Select the fabric speed. If your cluster contains nodes that automatically negotiate the fabric speed, this setting has no effect on those nodes. Node model SAN Volume Controller 2145-8F2 does not automatically negotiate its fabric speed and only operates at 1 or 2 Gbps; therefore, the fabric speed for these node models can be set at 1 or 2 Gbps. If your cluster contains nodes that automatically negotiate the fabric speed, set this value to 2 Gbps, even if the fibre channel operates at 4 Gbps.

12. Click **Create New Cluster** when you have completed this panel. After a few seconds, the cluster is created and the series of status panels are displayed. Click **Continue** on each of these panels. The **License Setting** panel is displayed.
13. Click **Continue** to complete the wizard. A message displays indicating that the cluster has been created successfully. Click the **X** that is located in the right corner of the window to close the wizard. You have successfully connected and configured the cluster. The cluster should be listed on the Viewing Clusters panel.

You have successfully connected and configured the cluster. The cluster should be listed on the Viewing Clusters panel.

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

After you have verified that the cluster has been created successfully, you can sign on to the cluster and complete the following tasks to continue setting up your cluster environment:

1. Add additional nodes to the cluster
2. Configure user authentication and authorization
3. Set up Call home options
4. Set up event notifications and inventory reporting
5. Configure Secure Shell (SSH) keys for command-line interface users
6. Create managed disk (MDisk) groups
7. Add MDisks to MDisk groups
8. Identify and create virtual disks (VDisks)
9. Create and map host objects
10. Identify and configure FlashCopy mappings, Metro Mirror relationships, or Global Mirror relationships
11. Back up cluster configuration

Chapter 5. Using the SAN Volume Controller Console

The SAN Volume Controller Console is a Web-browser based graphical user interface. The SAN Volume Controller Console can be used to create and maintain the configuration of storage that is associated with SAN Volume Controller clusters. A single SAN Volume Controller Console can manage multiple clusters.

See the SAN Volume Controller Console support information on the following Web site for the supported operating systems and Web browsers:

www.ibm.com/storage/support/2145

Key functions

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

- Initial set up of the cluster, its nodes, and the I/O groups (or node pairs).
- Set up and maintain managed disks and managed disk groups.
- Set up and maintain Secure Shell keys.
- Set up and maintain virtual disks.
- Set up 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 mappings and FlashCopy consistency groups.
  - Metro Mirror and Global Mirror relationships and consistency groups.
- Perform service and maintenance tasks.

SAN Volume Controller Console port requirements

Using the SAN Volume Controller console requires that specific ports be enabled.

To use the SAN Volume Controller console for certain functions, specific ports must be enabled for firewall access.

1. To use the SAN Volume Controller console online help system, ensure that the following port is enabled:

   ```
   Console online help port: 9001
   ```

   Note: This port value is fixed and cannot be changed.
2. To access IBM WebSphere® through the SAN Volume Controller console, ensure that the following ports are enabled:

   ```
   SOAP port: 8884
   RMI port: 2809
   HTTP port: 9080
   HTTPS port: 9443
   ```

   Note: During installation, you can modify these ports to other values.

SAN Volume Controller Console layout

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

Figure 36 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.

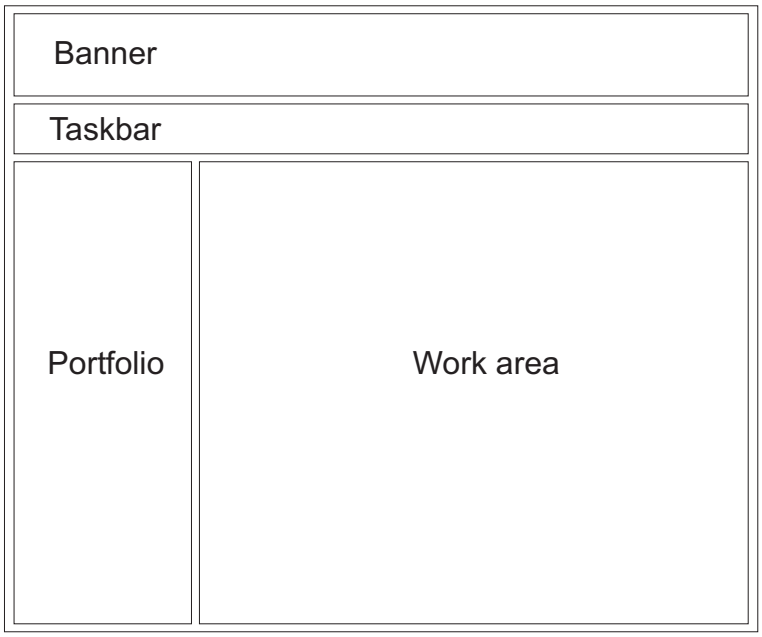

Figure 36. Basic frame layout

SAN Volume Controller Console banner

The banner of the SAN Volume Controller Console provides product or customer identification.

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.

Figure 37 shows the task bar. You can click the **question mark (?)** icon on the right side to display the information center in a separate browser window. You can click the (i) icon to display a help topic for the panel that is currently displayed in the work area.

Figure 37. Task bar

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.

The following task-based links are available from the Welcome panel of the SAN Volume Controller Console:

- Welcome
- Clusters

The following task-based links are available after you have launched the SAN Volume Controller Console:

- Welcome
- Manage Cluster
  - View Cluster Properties
  - Modify IP Addresses
  - Configure iSNS Server
  - Remove iSNS Server
  - Configure iSCSI Authentication
  - Set Cluster Time
  - Start Statistics Collection
  - Stop Statistics Collection
  - Shut Down Cluster
- Manage Authentication
  - Modify Current User
  - Users
  - User Groups
  - Cluster Passwords
  - Remote Authentication
- Work with Nodes
  - I/O Groups
  - Nodes
  - Node Ethernet Ports
- Manage Progress
  - View Progress
- Work with Managed Disks
  - Disk Controller Systems
  - Discovery Status
  - Managed Disks
  - Quorum Disks
  - Managed Disk Groups
- Work with Hosts
  - Hosts
- Work with Virtual Disks
  - Virtual Disks
  - Virtual Disk-to-Host Mappings

- Manage Copy Services
  - FlashCopy Mappings
  - FlashCopy Consistency Groups
  - Metro and Global Mirror Relationships
  - Metro and Global Mirror Consistency Groups
  - Metro and Global Mirror Cluster Partnership
- Service and Maintenance
  - Upgrade Software
  - Run Maintenance Procedures
  - Set SNMP Event Notification
  - Set Syslog Event Notification
  - Set E-mail Features
  - Analyze Error Log
  - License Settings
  - View License Settings Log
  - Dump Configuration
  - List Dumps
  - Backup Configuration
  - Delete Backup
  - Fabrics
  - CIMOM Log Configuration

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. For each panel that displays a table, you can optionally set filters to sort the information and arrange the data that is displayed. Table filter views are persistent for the duration of the login session.

Checking your Web browser and settings before accessing the SAN Volume Controller Console

To access the SAN Volume Controller Console, you must ensure that your Web browser is supported and set to allow pop-up windows.

See the SAN Volume Controller Console support information on the following Web site for the supported operating systems and Web browsers:

www.ibm.com/storage/support/2145

Complete the following steps to configure your Web browser:

1. Ensure that the Web browser is not set to block or suppress pop-up windows.

   Note: If you are using Internet Explorer 7.0 and receive a message that a pop-up window has been blocked, click the Information Bar at the top of the browser and select **Always allow popups** from this site. If you receive a message that content was blocked because it was not signed by

a valid security certificate, click the Information Bar at the top of the screen and select **Show blocked** content.

Note: If you are using Mozilla Firefox 3.x, you must manually add a pop-up exception for the SAN Volume Controller Console by completing the following steps:

a. Copy the SAN Volume Controller Console Web address from the browser's navigation toolbar.
b. Select **Tools** → **Options** → **Content**.
c. Click **Exceptions** to the right of **Block pop-up windows**.
d. Paste the SAN Volume Controller Console Web address in the **Address of Web site** field and click **Allow**.
e. Click **Close**.
f. Click **OK**.

2. Ensure that you have not installed any applications on the Web browser that block or suppress pop-up windows. If such an application is installed with the Web browser, uninstall it or turn it off.
3. Disable the proxy setting by completing the following steps:

   For Netscape:

   a. Open your Netscape browser and click **Edit** → **Preferences**. The Preferences window displays.
   b. From the left side category, click **Advanced** to expand the secondary options. The suboption Proxies displays.
   c. Click **Proxies**. The Proxies window displays.
   d. Select **Direct connection to Internet**.

   For Internet Explorer:

   a. Click **Tools** → **Internet Options** → **Connections** → **LAN Settings**.
   b. Click to clear the **Use a proxy server** box.

   For Mozilla Firefox:

   a. Open your Firefox browser and click **Tools** → **Options** → **Advanced**. The Advanced window displays.
   b. Select the Network tab and click **Settings** under **Connections** heading. The Connection Setting panel displays
   c. Under Configure Proxies to Access the Internet, ensure that **No Proxies** is selected.
   d. Click **OK**.

4. (Optional) Complete the following steps to add password protection so that your password does not display when you type it in:

   For Netscape:

   a. Start a Netscape session.
   b. Click **Edit** → **Preferences** from the menu bar.
   c. Click **Privacy and Security**.
   d. Click **Web Passwords**.
   e. Ensure that the **Remember passwords for sites that require me to log in** box is unchecked.
   f. Click **OK**.

   For Internet Explorer:

a. Start an Internet Explorer session.
b. Click **Tools** → **Internet Options** from the menu bar. The Internet Options panel is displayed.
c. Click the **Content** tab.
d. Click **AutoComplete**. The AutoComplete Settings panel is displayed.
e. Ensure that the **User names and passwords on forms** box is unchecked.
f. Click **OK**.

For Mozilla Firefox

a. Open your Firefox browser and click **Tools** → **Options** → **Security**. The Security window displays.
b. Under the password heading, ensure that the **Remember passwords for sites** box is unchecked.
c. Click **OK**.

Accessing the SAN Volume Controller Console

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 pop-up windows because this can prevent the windows in the SAN Volume Controller Console from opening. If you are using Internet Explorer 7.0 and receive a message that a pop-up window has been blocked, click the Information Bar at the top of the browser and select **Always allow popups from this site**. If you receive a message that content was blocked because it was not signed by a valid security certificate, click the Information Bar at the top of the screen and select **Show blocked content**.

If you are using Mozilla Firefox 3.x, you must manually add a pop-up exception for the SAN Volume Controller Console by completing the following steps:

1. Copy the SAN Volume Controller Console Web address from the browser's navigation toolbar.
2. Select **Tools** → **Options** → **Content**.
3. Click **Exceptions** to the right of **Block pop-up windows**.
4. Paste the SAN Volume Controller Console Web address in the **Address of Web site** field and click **Allow**.
5. Click **Close**.
6. Click **OK**.

You have two options for accessing your the SAN Volume Controller Console.

If you are accessing the SAN Volume Controller Console from the server running the IBM System Storage Productivity Center (SSPC), you can access the SAN Volume Controller Console by selecting **Start** → **All Programs** → **IBM System Storage SAN Volume Controller** → **Launch SVC Console** Alternatively, you can access the SAN Volume Controller Console from the workstation where SSPC is installed by pointing your browser to the following Web address:

http://*localhost*:9080/ica

where *localhost* is the address of the machine where your SAN Volume Controller Console is installed.

As an alternative, you can use the SAN Volume Controller Console from any workstation that can access the SSPC. Start a supported Web browser on your workstation and point to the following Web address:

http://*svcconsoleip*:9080/ica

where *svcconsoleip* is the IP address of the SSPC server on which the SAN Volume Controller Console is running.

1. If you are accessing SAN Volume Controller Console for the first time, select **Add SAN Volume Controller Cluster** on the Welcome panel.
2. On the Adding a Cluster panel, enter the IP address of the cluster that you are adding. If you are completing a cluster-creation process that was started from the front panel of the node, select **Create (Initialize) Cluster**. However, if the cluster is already in use and you are only adding the cluster to the list of managed clusters for this installation of the SAN Volume Controller Console, do not select the **Create (Initialize) Cluster** check box.
3. After the cluster has been added to the SAN Volume Controller Console, specify the cluster and select **Launch the SAN Volume Controller Console**. If the availability status for the cluster is `Unauthenticated`, you are prompted to sign on to the cluster with a valid user name and password.

Launching the SAN Volume Controller Console to manage a cluster

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

The SAN Volume Controller Console is the centralized Web application that is used to manage your clusters.

Perform the following steps to launch the SAN Volume Controller Console for a specific cluster:

1. Start the SAN Volume Controller Console by selecting **Start** → **All Programs** → **IBM System Storage SAN Volume Controller** → **Launch SVC Console** or by pointing your Web browser to http://*svcconsoleip*:9080/ica, where *svcconsoleip* is the IP address of the IBM System Storage Productivity Center or the master console. Either an IPv4 or IPv6 connection is allowed. For example, the appropriate Web browser address could be of the form `http://9.134.5.6:9080/ica` or `http://[2020:1234::1234]:9080/ica`.

   Note: If you are using Internet Explorer 7.0 and receive a message that a popup has been blocked, click the Information Bar at the top of the browser and select Always allow popups from this site. If you receive a message that content was blocked because it was not signed by a valid security certificate, click the Information Bar at the top of the screen and select Show blocked content.

   Note: If you are using Mozilla Firefox 3.x, you must manually add a pop-up exception for the SAN Volume Controller Console by completing the following steps:

   a. Copy the SAN Volume Controller Console Web address from the browser's navigation toolbar.

b. Select **Tools** → **Options** → **Content**.
c. Click **Exceptions** to the right of **Block pop-up windows**.
d. Paste the SAN Volume Controller Console Web address in the **Address of Web site** field and click **Allow**.
e. Click **Close**.
f. Click **OK**.

2. Click **Clusters** in the portfolio. The Viewing Clusters panel is displayed.
3. Select the cluster you wish to manage. If the cluster that you want to manage is not displayed, select **Add Cluster** from the task list and click **Go** On the Add Cluster panel, specify the IP address for the cluster you want to add and click **OK**. When adding a fully functional cluster, do not select the **Create (Initialize) Cluster** check box. This option is only necessary after you create the cluster from the front panel for the first time. Additional clusters that are added after this process do not need to be initialized.
4. Select **Launch the SAN Volume Controller Console** from the task list.

   Note: If the Availability Status for the cluster is Unauthenticated, you are prompted to sign on with your cluster user name and password.
5. Click **Go**. A secondary browser window opens.

Setting cluster date and time

You can set the date and time for a SAN Volume Controller cluster from the Cluster Date and Time Settings panel.

This task assumes that you have already launched the SAN Volume Controller Console.

You can set the cluster date and time manually, or by specifying an NTP server:

1. Click **Manage Clusters** → **Set Cluster Time** in the portfolio. The Cluster Date and Time Settings panel is displayed.
2. To use NTP to manage the cluster date and time, enter an IPv4 address and click **Set NTP Server**.

   Note: If you are using a remote authentication service to authenticate users to the SAN Volume Controller cluster, then both the cluster and the remote service should use the same NTP server. Consistent time settings between the two systems ensure interactive performance of the SAN Volume Controller Console and correct assignments for user roles.
3. To set the cluster date and time manually, continue with the following steps.
4. Type your changes into the **Date**, **Month**, **Year**, **Hours** and **Minutes** fields and select a new time zone from the **Time Zone** list
5. Select **Update cluster time and date**, **Update cluster time zone**, or both.
6. Click **Update** to submit the update request to the cluster.

Modifying the cluster IP addresses

You can display and change the IP addresses that are associated with a cluster from the Modify IP Addresses panel.

This task assumes that you have already launched the SAN Volume Controller Console.

If you change the cluster IP address, the cluster stops serving Web pages that use the old IP address. You must use the new IP address to reconnect your Web browser to the cluster. When you reconnect to the cluster, accept the new site certificate.

Perform the following steps to change the IP addresses:

1. Click **Manage Cluster** → **Modify IP Address** in the portfolio. The Modify IP Addresses panel is displayed. Both IPv4 and IPv6 addresses can be defined on this panel. The Modify IP Addresses panel displays the cluster Ethernet ports and the IP addresses that are associated with them. The first Ethernet port IP address is configured during the initial configuration of the cluster from the SAN Volume Controller node. The second port is optional and can be configured to provide redundancy to the cluster.
2. Select the cluster Ethernet port that you want to change IP address settings for and select **Modify Port Settings** from the task list and click **Go**. The Modify Port Settings panel displays.
3. On the Modify Port Settings panel, enter the new IP address values for the selected cluster Ethernet port. If you are changing IP setting that you are currently using to connect to the cluster, you may need to re-establish the connection to the cluster.
   - If you are updating IPv4 addresses, provide the following information for the cluster:

     IP address
     Enter a valid IPv4 address for the selected cluster Ethernet port.

     Service IP Address
     Select the service IP address which connects to a node if the node has been removed from the cluster and is being serviced. You can select one of the following options:

     Assign automatically (DHCP)
     Select this if you want the service IP address to be assigned automatically through a DHCP server.

     Static IP:
     Enter an IP address for the service IP address.

     Subnet Mask
     Enter the subnet mask for the cluster.

     Gateway
     Enter the gateway IP address used for the cluster.
   - If you are updating IPv6 addresses, provide the following information for the cluster:

     IP address
     Enter a valid IPv6 address for the selected cluster Ethernet port.

     Service IP Address
     Select the service IP address which connects to a node if the node has been removed from the cluster and is being serviced. You can select one of the following options:

     Assign automatically (DHCP)
     Select this if you want the service IP address to be assigned automatically through a DHCP server.

     Static IP:
     Enter an IP address for the service IP address.

The SAN Volume Controller Console supports the following IPv6 formats:

- Eight colon-separated groups of four hexadecimal digits; for example, 1234:1234:abcd:0123:0000:0000:7689:6576
- Eight colon-separated groups of hexadecimal digits with leading zeros omitted; for example, 1234:1234:abcd:123:0:0:7689:6576
- Zero suppression format; for example, 1234:1234:abcd:123::7689:6576

  Note: You can only suppress one set of zeros in an address.

IPv6 Network Prefix
Displays the network prefix that is associated with the IPv6 cluster and service IP addresses. Valid values are 0 - 127.

Gateway
Enter the gateway IP address used for the cluster.

4. Click **OK**.

A new SSL certificate is generated by the cluster to display the new IP address. This new certificate displays when the Web browser first connects to the cluster.

Changing from an IPv4 to an IPv6 address

Your cluster IP address can be an IPv4 or IPv6 address.

Perform the following steps to change the cluster IP address from IPv4 to IPv6:

1. To change the cluster to accept both IPv4 and IPv6 addresses, complete the following steps:
   a. Click **Manage Cluster → Modify IP Addresses** in the portfolio. The Modify IP Addresses panel is displayed. On this panel there are two tables display for each of the supported address structures.
   b. Select the Cluster Ethernet Port 1 for the IPv4 cluster, and select **Modify Port Settings** from the task list for the Cluster Ethernet Port 1 and click **Go**.
   c. On the Modify Port Settings panel, update the IPv4 addresses to the new IPv6 address and click **OK**.
   d. Repeat Steps 1 and 2 for the Cluster Ethernet Port 2, if it is configured.
2. Remove the managed cluster with the IPv4 address, by completing the following tasks:
   a. Click **View Cluster Properties** in the portfolio. The View Cluster Properties panel is displayed.
   b. Select the cluster you want to remove and select **Remove a Cluster** from the list. Click **Go**. The Confirming the Removal of Cluster panel is displayed.
   c. Click **Yes** to remove the cluster.
   d. Return to the View Cluster Properties panel.
3. Verify the cluster is available on the new address by issuing the ping command for the new IP address. A successful ping indicates that the cluster is available at the new IP address.
4. Add the cluster with the new IPv6 address, by completing these steps:
   a. On the Viewing Clusters panel, select **Add a Cluster** from the list and click **Go**. The Adding a Cluster panel is displayed.
   b. Type the IPv6 address for the cluster.
   c. Ensure that the **Create (Initialize) Cluster** check box is not selected.

d. Click **OK**.
e. Click **Yes** to confirm adding the cluster.

5. Return to the Modify IP Addresses panel, and select the cluster Ethernet port with the IPv4 setting and select **Clear Port Settings** from the task list and click **Go**. The Delete Cluster IPv4 Settings panel displays
6. Click **Delete** to delete the IPv4 settings.

The cluster is now only available from the IPv6 address.

Changing from an IPv6 to an IPv4 address

Your cluster IP address can be an IPv4 or IPv6 address.

Perform the following steps to change the cluster from IPv6 to IPv4:

1. Change the cluster to accept both IPv4 and IPv6 addresses, by completing these steps:
   a. Click **Manage Cluster** → **Modify IP Addresses** in the portfolio. The Modify IP Addresses panel is displayed. On this panel there are two tables display for each of the supported address structures.
   b. Select the Cluster Ethernet Port 1 for the IPv6 cluster, and select **Modify Port Settings** from the task list for the Cluster Ethernet Port 1 and click **Go**.
   c. On the Modify Port Settings panel, update the IPv6 addresses to the new IPv4 address and click **OK**.
   d. Repeat Steps 1 and 2 for the Cluster Ethernet Port 2, if it is configured.
2. Remove the managed cluster with the IPv6 address, by completing the following tasks:
   a. Click **View Cluster Properties** in the portfolio. The View Cluster Properties panel is displayed.
   b. Select the cluster with the IPv6 address that you want to remove and select **Remove a Cluster** from the list. Click **Go**. The Confirming the Removal of Cluster panel is displayed.
   c. Click **Yes** to remove the cluster.
   d. Return to the View Cluster Properties panel.
3. Verify the cluster is available on the new settings by issuing the ping command with the new IP address. A successful ping indicates that the cluster is available at the new IP address.
4. Add the cluster with the new IPv4 address, by completing these steps:
   a. On the Viewing Clusters panel, select **Add a Cluster** from the list and click **Go**. The Adding a Cluster panel is displayed.
   b. Type the IPv4 address for the cluster.
   c. Ensure that the **Create (Initialize) Cluster** check box is not selected.
   d. Click **OK**.
   e. Click **Yes** to confirm adding the cluster.
5. Return to the Modify IP Addresses panel, and select the cluster Ethernet port with the IPv4 setting and select **Clear Port Settings** from the task list and click **Go**. The Delete Cluster IPv4 Settings panel displays
6. Click **Delete** to delete the IPv4 settings.

The cluster is now only available from the IPv4 address.

Modifying service password

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to maintain cluster passwords:

1. Click **Manage Authentication** → **Cluster Passwords** in the portfolio. The Modify Service Passwords panel is displayed.
2. Type the new service password in the appropriate fields and click **OK** to change the password.

   Passwords must meet the following requirements:

   - Passwords are a maximum of 15 alphanumeric characters.
   - Passwords are case-sensitive.
   - Valid characters are uppercase letters (A - Z), lowercase letters (a - z), digits (0 - 9), dash (-), and underscore (_).
   - The first character cannot be a dash (-).

   Note: Passwords must be typed twice to allow verification.

Viewing cluster properties

You can use the SAN Volume Controller Console to view the properties for a cluster.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the properties of a cluster:

1. Click **Manage Cluster** → **View Cluster Properties** in the portfolio. The Viewing General Properties panel is displayed.
2. Click the following tabs:
   a. **General** to display the general properties
   b. **IP Addresses** to view the IP addresses that are used by the cluster
   c. **Remote Authentication** to view attributes for the remote authentication service, which is used by remote users to access the cluster
   d. **Space** to view the space and capacity for managed disks (MDisks), MDisk groups, and virtual disks (VDisks)
   e. **Statistics** to view the cluster statistics details
   f. **Metro Mirror and Global Mirror** to view the Metro Mirror or Global Mirror properties of the cluster
   g. **iSCSI** to view the iSCSI properties for the cluster
   h. **SNMP** to view the SNMP properties for the cluster
   i. **Syslog** to view the syslog properties for the cluster
   j. **E-Mail Servers** to view the e-mail server properties for the cluster
   k. **E-Mail Users** to view the e-mail user properties for the cluster
3. Click **Close** to close the panel.

Viewing remote cluster properties

You can display attributes for a remote cluster in a cluster partnership on the Viewing Remote Cluster Properties panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to view the properties of a remote cluster in the selected partnership:

1. Click **Manage Copy Services** → **Metro & Global Mirror Cluster Partnerships** in the portfolio. The Viewing Metro & Global Mirror Cluster Partnership panel displays.
2. Click the name of the partnership that you want to display the remote cluster properties. The Viewing Remote Cluster Properties panel is displayed.
3. Click **Close** to return to the Viewing Metro & Global Mirror Cluster Partnership panel.

Adding nodes to a cluster

Before you add a node to a cluster, you must 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.

If you are replacing a node and the switch is zoned by worldwide port name (WWPN) rather than by switch port, make sure that the switch is configured such that the node being added is in the same VSAN/zone.

Special procedures when adding a node to a cluster

If you are adding a node that has been used previously, either within a different I/O group within this cluster or within a different cluster, consider the following situations before adding the node. If you add a node to the cluster without changing its worldwide node name (WWNN), hosts might detect the node and use it as if it were in its old location. This could cause the hosts to access the wrong virtual disks (VDisks).

- If you are re-adding a node back to the same I/O group after a service action required the node to be deleted from the cluster and the physical node has not changed, then no special procedures are required and the node can be added back to the cluster.
- In a service situation, a node should normally be added back into a cluster using the original node name. As long as the partner node in the I/O group has not been deleted too, this is the default name used if **-name** is not specified.
- If you are replacing a node in a cluster, either because of a node failure or an upgrade, you must change the WWNN of the new node to match that of the original node before you connect the node to the fibre channel network and add the node to the cluster.
- If you are creating a new I/O group in the cluster and are adding new node, then there are no special procedures since this node has never been added to a cluster before and its WWNN that has not existed before.
- If you are creating a new I/O group in the cluster and are adding new node, but this node has been added to a cluster before, then host system might still be configured to the node WWPNs and the node might still be zoned in the fabric. Since you cannot change the WWNN for the node, you must ensure other

components in your fabric are configured correctly. Verify that any host that was previously configured to use the node has been correctly updated. Also verify the fabric zoning does not currently include this node in any zones.

- If the node you are adding was previously replaced, either for a node repair or upgrade, you might have used the node's WWNN for the replacement node. Ensure that the WWNN of this node was updated so that you do not have two nodes with the same WWNN attached to your fabric. Also ensure that the WWNN of the node that you are adding is not 00000. If it is 00000, contact your support representative.

Note: Consider the following information when using multipathing device drivers:

- Applications on the host systems direct I/O operations to file systems or logical volumes that are mapped by the operating system to virtual paths (vpaths), which are pseudo disk objects that are supported by the multipathing device drivers. Multipathing device drivers maintain an association between a vpath and a SAN Volume Controller virtual disk (VDisk). This association uses an identifier (UID) which is unique to the VDisk and is never reused. The UID allows multipathing device drivers to directly associate vpaths with VDisks.
- Multipathing device drivers operate within a protocol stack that 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 that is provided by these SCSI and fibre-channel device drivers uses a combination of a SCSI logical unit number (LUN) and the worldwide node name (WWNN) for the fibre-channel node and ports.
- If an error occurs, the 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 WWNN and LUN numbers that were previously used.
- Multipathing device drivers do not check the association of the VDisk with the vpath on every I/O operation that it performs.

Adding nodes to a cluster using the SAN Volume Controller Console

Attention:

1. If you are re-adding a node to the SAN, ensure that you are adding the node to the same I/O group from which it was removed. Failure to do this can result in data corruption. You must 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 the IBM Support Center to add the node back into the cluster without corrupting the data.
2. In a service situation, a node should normally be added back into a cluster using the original node name. As long as the partner node in the I/O group has not been deleted too, this is the default name used if **-name** is not specified.
3. The LUNs that are presented to the ports on the new node must be the same as the LUNs that are presented to the nodes that currently exist in the cluster. You must ensure that the LUNs are the same before you add the new node to the cluster.
4. LUN masking for each LUN must be identical on all nodes in a cluster. You must ensure that the LUN masking for each LUN is identical before you add the new node to the cluster.
5. You must ensure that the model type of the new node is supported by the SAN Volume Controller software level that is currently installed on the cluster. If the model type is not supported by the SAN Volume Controller software level, upgrade the cluster to a software level that supports the model type of the new node. See the following Web site for the latest supported software levels:

   www.ibm.com/storage/support/2145

Each node in an I/O group must:

- Be connected to a different uninterruptible power supply.
- Have a unique name. If you do not provide a name, the cluster assigns a default name to the object.

Note: Whenever possible you must provide a meaningful name for objects to make identifying that object easier in the future.

This task assumes that you have already launched the SAN Volume Controller Console.

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

1. Click **Work with Nodes** → **Nodes** in the portfolio. The Viewing Nodes panel is displayed.
2. Select **Add a node** from the task list and click **Go**. The Adding a Node to a Cluster panel is displayed.
3. Select the node that you want to add to the cluster from the **Available Candidate Nodes** list.
4. Select the I/O group from the **I/O Groups** list.
5. In the **Node Name** field, type the name that you want to assign to the node.
6. Click **OK**.
7. If you are adding a node into the cluster for the first time, record the following information:
   - Node serial number

- All WWPNs
- The I/O group that the node belongs to

Important: You need this information to avoid possible data corruption if you have to remove and re-add the node to the cluster.

Viewing the node status

You can view the properties for a node from the Viewing General Details panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the node properties:

1. Click **Work with Nodes** → **Nodes** in the portfolio. The Viewing Nodes panel is displayed.
2. Click the name of the node for which you want to view detailed information. The Viewing General Details panel is displayed.
3. In the General navigation area, click **Fibre Channel Ports** to view the worldwide port name (WWPN) details. The Viewing Fibre Channel Port Details panel is displayed. You are shown the attributes and values for the node, which includes the state, WWNN, and I/O group to which it belongs.
4. Click **Vital Product Data** to view the node hardware details. The Viewing Vital Product Data panel is displayed.
5. Click **Ethernet Ports** to view the iSCSI port details. The Viewing Node Ethernet Port panel is displayed.
6. Click **Ethernet Ports–IP** to view the iSCSI port IP address details. The Viewing Node Ethernet Port–IP panel is displayed.
7. Click **Close** to close the panel.

Increasing the size of a cluster

You can use the SAN Volume Controller Console to increase the size of a cluster.

Increasing the size of cluster by adding the number of node in the cluster increases throughput. The nodes must be added in pairs and assigned to a new I/O group.

Complete the following steps to increase the size of your cluster:

1. Add a node to your cluster and repeat this step for the second node.
2. If you want to balance the load between the existing I/O groups and the new I/O groups, you can migrate your virtual disks (VDisks) to new I/O groups; see "Migrating VDisks" on page 159 for details. Repeat this step for all VDisks that you want to assign to the new I/O group.

Adding two nodes to increase the size of a cluster

You can use the SAN Volume Controller Console to increase the size of a cluster by adding two nodes to create a new I/O group.

Attention: If you are adding a node that was previously removed from a cluster, ensure that either the following two conditions have been met:

- The WWPN for the removed node is swapped with the node that replaces it.
- All hosts that accessed the removed node through its WWPNs have been reconfigured to use the WWPN for the new node.

Failure to do either of these actions can result in data corruption.

Before completing these steps, install the new nodes and connect them to the fibre-channel.

This task assumes that you have already launched the SAN Volume Controller Console.

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

1. Click **Work with Nodes → I/O groups** to determine the I/O group where the node will be added. The Viewing Input/Output Groups panel is displayed.
2. Record the name or ID of the first I/O group that has a node count of zero (0). You only need to do this step for the first node that is added. The second node of the pair uses the same I/O group number.
3. Click **Work with Nodes → Nodes**. The Viewing Nodes panel is displayed.
4. Select the node that you want to add from the list of available candidate nodes.
5. From the task list, select **Add a Node** and click **Go**. The Adding a Node to a Cluster panel is displayed.
6. Select the node that you want to add to the cluster from the **Available Candidate Nodes** list.
7. Select the I/O group from the **I/O Groups** list.

   Important: If you are adding a node that was previously removed from the cluster, you must select the name of the I/O group from which the node was previously removed. If you are adding a node that has never been in a cluster, select the name of the I/O group that you recorded in step 2.
8. Click **OK**.
9. Verify that the node is online by refreshing the Viewing Nodes panel. You might have to close the panel and reopen it to refresh the panel.
10. Click the name of the node that you have added to the cluster. The Viewing General Details panel is displayed.
11. Click the **General, Ports** and **Vital Product Data** tabs and record the following information:
    - 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
12. Click **Close** to close the panel.

You may need to reconfigure your storage systems to allow the new I/O group nodes to access them. If the disk controller uses mapping to present RAID arrays or partitions to the cluster and the WWNNs or the worldwide port names have changed, you must modify the port groups that belong to the cluster.

Replacing a faulty node with a spare node

You can use the SAN Volume Controller Console and the SAN Volume Controller front panel to replace a faulty node in a cluster.

Before you attempt to replace a faulty node with a spare node you must ensure that you meet the following requirements:

- SAN Volume Controller version 3.1.0 or higher is installed on the cluster and on the spare node.
- 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.
- Record of the last five characters of the original worldwide node name (WWNN) of the spare node.

  Note: A repaired faulty node, which has been successfully replaced in the cluster with a spare node using the original WWPN of the faulty node, must be assigned a new unique WWNN. You can use the original WWNN of the spare node as the new WWNN of the repaired node.

Attention: Never connect a node with a WWNN of `00000` to the cluster. If this node is no longer required as a spare and is to be used for normal attachment to a cluster, you must change the WWNN to the number you recorded when a spare was created. Using any other number might cause data corruption.

If a node fails, the cluster continues to operate with degraded performance, until the faulty node is repaired. If the repair operation takes an unacceptable amount of time, it is useful to replace the faulty node with a spare node. However, the appropriate procedures must be followed and precautions must be taken so you do *not* interrupt I/O operations and compromise the integrity of your data.

The following table describes the changes that are made to your configuration when you replace a faulty node in the cluster:

Node attributes	Description
Front panel ID	This is the number that is printed on the front of the node and is used to select the node that is added to a cluster.
Node ID	This is the ID that is assigned to the node. 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. This ID will change during this procedure.

Node attributes	Description
Node name	This is the name that is assigned to the node. If you do not specify a name, the SAN Volume Controller assigns a default name. The SAN Volume Controller creates a new default name each time a node is added to a cluster. If you choose to assign your own names, you must type the node name on the Adding a node to a cluster panel. You cannot manually assign a name that matches the naming convention used for names assigned automatically by SAN Volume Controller. If you are using scripts to perform management tasks on the cluster and those scripts use the node name, you can avoid the need to make changes to the scripts by assigning the original name of the node to a spare node. This name might change during this procedure.
Worldwide node name	This is the WWNN that is assigned to the node. The WWNN is used to uniquely identify the node and the fibre-channel ports. During this procedure, the WWNN of the spare node is changed to that of the faulty node. The node replacement procedures must be followed exactly to avoid any duplication of WWNNs. This name does not change during this procedure.
Worldwide port names	These are the WWPNs that are assigned to the node. WWPNs are derived from the WWNN that is written to the spare node as part of this procedure. For example, if the WWNN for a node is 50050768010000F6, the four WWPNs for this node are 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 These names do not change during this procedure.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to replace a faulty node in the cluster:

1. Verify the name and ID of the node that you want to replace.

   Complete 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 → Nodes** in the portfolio. The Viewing Nodes panel is displayed. If the node is faulty, it is shown as offline.
   c. Ensure the partner node in the I/O group is online.

   - If the other node in the I/O group is offline, start the Directed Maintenance Procedures (DMPs) to determine the fault.
   - If you have been directed here by the DMPs, and subsequently the partner node in the I/O group has failed, recover the offline VDisks.
   - If you are replacing the node for other reasons, determine the node that you want to replace and ensure that the partner node in the I/O group is online.
   - If the partner node is offline, you will lose access to the VDisks that belong to this I/O group. Start the DMPs and fix the other node before proceeding to the next step.

2. Click the name of the faulty (offline) node. The Viewing General Details panel is displayed.
3. Click **General** and record the following attributes for the faulty node:
   - ID
   - WWNN
   - I/O Group
   - UPS Serial Number
   - Uninterruptible power supply serial number
4. Click **Close**. Click Fibre Channel Port and record the following attribute for the faulty node:
   - WWPNs
5. Click **Close**. Click Vital Product Data and record the following attribute for the faulty node:
   - System Serial Number
6. Ensure that the faulty node has been powered off.
7. Use the SAN Volume Controller Console to delete the faulty node from the cluster.

   Remember: You must record the following information to avoid data corruption when this node is re-added to the cluster:
   - Node serial number
   - WWNN
   - All WWPNs
   - I/O group that contains the node
8. Disconnect all four fibre-channel cables from the node.

   Important: Do not plug the fibre-channel cables into the spare node until the spare node is configured with the WWNN of the faulty node.
9. Connect the power and signal cables from the spare node to the uninterruptible power supply that has the serial number you recorded in step 3.

   Note: For 2145 UPS-1U units, you must disconnect the cables from the faulty node.
10. Power on the spare node.
11. You must change the WWNN of the spare node to that of the faulty node. The procedure for doing this depends on the SAN Volume Controller version that is installed on the spare node. Press and release the down button until the Node: panel displays. Then press and release the right button until the WWNN: panel displays. If repeated pressing of the right button returns you to the Node: panel, without displaying a WWNN: panel, go to step 13 on page 141; otherwise, continue with step 12.
12. Change the WWNN of the spare node (with SAN Volume Controller V4.3 and above installed) to match the WWNN of the faulty node by performing the following steps:
    a. With the Node WWNN: panel displayed, press and hold the down button, press and release the select button, and then release the down button. The display switches into edit mode. `Edit WWNN` is displayed on line 1. Line 2 of the display contains the last five numbers of the WWNN.

b. Change the WWNN that is displayed to match the last five numbers of the WWNN that you recorded in step 3 on page 140. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.

c. When the five numbers match the last five numbers of the WWNN that you recorded in step 3 on page 140, press the select button to accept the numbers.

13. Change the WWNN of the spare node (with SAN Volume Controller versions prior to V4.3 installed) to match the WWNN of the faulty node by performing the following steps:

    a. Press and release the right button until the Status: panel is displayed.

    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 numbers of the WWNN.

    c. With the WWNN displayed on the front panel; press and hold the down button; press and release the select button; release the down button. The display switches into edit mode.

    d. Change the WWNN that is displayed to match the last five numbers of the WWNN that you recorded in step 3 on page 140. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.

    e. When the five numbers match the last five numbers of the WWNN that you recorded in step 3 on page 140, press the select button to accept the numbers.

    f. Press the select button to retain the numbers that you have updated and return to the WWNN panel.

14. Connect the four fibre-channel cables that you disconnected from the faulty node and connect them to the spare node.

    If the spare node has less Ethernet cables connected than the faulty node, move the Ethernet cables from the faulty node to the spare node. Ensure you connect the cable into the same port on the spare node as it was in on the faulty node.

15. Use the SAN Volume Controller Console to add the spare node to the cluster. If possible, use the same node name that was used for the faulty node. If necessary, the spare node is updated to the same SAN Volume Controller version as the cluster. This update can take up to 20 minutes.

16. Use the tools that are provided with your multipathing device driver on the host systems to verify that all paths are now online. See the documentation that is provided with your multipathing device driver for more information. For example, if you are using the subsystem device driver (SDD), see the *IBM System Storage Multipath Subsystem Device Driver User's Guide* for instructions on how to use the SDD management tool on host systems. It might take up to 30 minutes for the paths to come online.

17. Repair the faulty node.

    Attention: When the faulty node is repaired, do not connect the fibre-channel cables to it. Connecting the cables might cause data corruption because the spare node is using the same WWNN as the faulty node.

    If you want to use the repaired node as a spare node, perform the following steps.

For SAN Volume Controller V4.3 and above:

a. With the Node WWNN: panel displayed, press and hold the down button, press and release the select button, and then release the down button. The display switches into edit mode. `Edit WWNN` is displayed on line 1. Line 2 of the display contains the last five numbers of the WWNN.

b. Change the displayed number to `00000`. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.

c. Press the select button to accept the numbers.

   This node can now be used as a spare node.

For SAN Volume Controller versions prior to V4.3:

a. Press and release the right button until the Status: panel is displayed.

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 numbers of the WWNN.

c. With the WWNN displayed on the front panel; press and hold the down button; press and release the select button; release the down button. The display switches into edit mode.

d. Change the displayed number to `00000`. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.

e. Press the select button to accept the numbers.

f. Press the select button to retain the numbers that you have updated and return to the WWNN panel.

   This node can now be used as a spare node.

Renaming a node

You can rename a node from the Renaming Node panel. If you are renaming a node that is associated with an iSCSI host, renaming the node alters the iSCSI qualified name (IQN), which can require reconfiguration of the iSCSI-attached host.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to rename a node:

1. Click **Work with Nodes** → **Nodes** in the portfolio. The Viewing Nodes panel is displayed.
2. Select the node you want to rename and select **Rename a Node** from the list. Click **Go**. The Renaming Node panel is displayed.
3. Type the new name of the node and click **OK** or click **Force Rename** if you want the node to be renamed even though renaming the node alters the IQN. This might require a reconfiguration of the iSCSI-attached host.

Deleting a node from a cluster using the SAN Volume Controller Console

You might have to remove a node 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.

The cache on the selected node is flushed before the node is taken offline. In some circumstances, such as when the system is already degraded (for example, when both nodes in the I/O group are online and the virtual disks within the I/O group are degraded), the system ensures that data loss does not occur as a result of deleting the only node with the cache data. The cache is flushed before the node is deleted to prevent data loss if a failure occurs on the other node in the I/O group.

Before deleting a node from the cluster, record the node serial number, worldwide node name (WWNN), all worldwide port names (WWPNs), and the I/O group that the node is currently part of. Recording this node information can avoid data corruption if the node is re-added to the cluster at a later time.

Attention:

- If you are removing a single node and the remaining node in the I/O group is online, the data on the remaining node goes into write-through mode. This data can be exposed to a single point of failure if the remaining node fails.
- If virtual disks (VDisks) are already degraded before you delete a node, redundancy to the VDisks is degraded. Removing a node might result in a loss of access to data and data loss.
- Removing the last node in the cluster destroys the cluster. Before you delete the last node in the cluster, ensure that you want to destroy the cluster.
- When you delete a node, you remove all redundancy from the I/O group. As a result, new or existing failures can cause I/O errors on the hosts. The following failures can occur:
  - Host configuration errors
  - Zoning errors
  - Multipathing software configuration errors
- If you are deleting the last node in an I/O group and there are VDisks assigned to the I/O group, you cannot delete the node from the cluster if the node is online. You must back up or migrate all data that you want to save before you delete the node. If the node is offline, you can delete the node.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to delete a node from a cluster:

1. Unless this is the last node in the cluster, turn the power off to the node that you are removing using the Shut Down a Node option on the SAN Volume Controller Console. This step ensures that the multipathing device driver does not rediscover the paths that are manually removed before you issue the delete node request.

Attention:

- When you remove the configuration node, the configuration function moves to a different node within the cluster. This process can take a short time, typically less than a minute. The SAN Volume Controller Console reattaches to the new configuration node transparently.
- If you turn the power on to the node that has been removed and it is still connected to the same fabric or zone, it attempts to rejoin the cluster. At this point, the cluster tells the node to remove itself from the cluster and the node becomes a candidate for addition to this cluster or another cluster.
- If you are adding this node into the cluster, ensure that you add it to the same I/O group that it was previously a member of. Failure to do so can result in data corruption.

2. In the portfolio, click **Work with Nodes** → **Nodes**. The Viewing Nodes panel is displayed.
3. Find the node that you want to delete.

   If the node that you want to delete is shown as `Offline`, then the node is not participating in the cluster.

   If the node that you want to delete is shown as `Online`, deleting the node can result in the dependent VDisks to also go offline. Verify whether or not the node has any dependent VDisks.
4. To check for dependent VDisks before attempting to delete the node, select the node and click **Show Dependent VDisks** from the drop-down menu.

   If any VDisks are listed, you should determine why and if access to the VDisks is required while the node is deleted from the cluster. If the VDisks are assigned from MDisk groups that contain solid-state drives (SSDs) that are located in the node, you should check why the VDisk mirror, if it is configured, is not synchronized. There can also be dependent VDisks because the partner node in the I/O group is offline. Fabric issues can also prevent the VDisk from communicating with storage systems. You should resolve these problems before continuing with the node deletion.
5. Select the node that you want to delete and select **Delete a Node** from the task list. Click **Go**. The Deleting Node from Cluster panel is displayed.
6. Click **OK** to delete the node. Before a node is delete the SAN Volume Controller checks to determine if there are any virtual disks (VDisks) that depend on that node. If the node that you selected contains VDisks within the following situations, VDisks go offline and become unavailable if the node is deleted:

   - The node contains solid-state drives (SSD) and also contains the only synchronized copy of a mirrored VDisk
   - The other node in the I/O group is offline

   If you select a node to delete that has these dependencies, another panel displays confirming the deletion. To delete the node in this case, click **Force Delete** message panel that displays.

Renaming an I/O group

You can rename an I/O group from the Viewing I/O Groups panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to rename an I/O group:

1. Click **Work with Nodes → I/O Groups** in the portfolio. The Viewing Input/Output Groups panel is displayed.
2. Select the I/O group that you want to rename and select **Rename an I/O Group** from the list. Click **Go**. The Renaming I/O Group panel is displayed.
3. Type the new name of the I/O Group in the **New Name** field.
4. Click **OK**.

Modifying a cluster

You can rename a cluster and change the fabric speed from the Modifying Cluster panel.

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

Perform the following steps to modify a cluster:

1. Click **Clusters** in the portfolio. The Viewing Clusters panel is displayed.
2. Select the cluster to modify and select **Modify a Cluster** from the task list. Click **Go**. The Modifying Cluster panel is displayed. You can perform the following from this panel:
   - Type a new name for the cluster.

     Note: Renaming a cluster changes the iSCSI qualified name (IQN) of all the nodes in the cluster and might require reconfiguration of all iSCSI-attached hosts.
   - Select a fabric speed from the **Fabric Speed** list.
3. Click **OK** to modify the cluster.

Shutting down a cluster

You can shut down a SAN Volume Controller cluster from the Shutting Down cluster panel.

If you want to remove all input power to a cluster (for example, the machine room power must be shutdown for maintenance), we you should use the Shutdown Down Cluster operation prior to removing power. Note: If you do not shut down the cluster before turning off input power to the uninterruptible power supply units, the cluster completes an emergency shutdown, which is powered from the uninterruptible power supply battery. This method drains power from the uninterruptible power supply needlessly, and the restart of the cluster is delayed while the uninterruptible power supply charges.

When input power is restored to the uninterruptible power supply units, they start to recharge. However, the SAN Volume Controller do not permit any I/O activity to be performed to the virtual disks (VDisks) 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 process might take as long as two hours. Therefore shutting down the cluster prior to removing input power to the uninterruptible power supply units prevents the battery power from being drained and makes it possible for I/O activity to resume as soon as input power is restored.

Before shutting down a cluster, quiesce all I/O operations that are destined for this cluster. Failure to do so can result in failed I/O operations being reported to your host operating systems.

Attention: If you are shutting down the entire cluster, you lose access to all VDisks that are provided by this cluster. Shutting down the cluster also shuts down all SAN Volume Controller nodes. This shutdown causes the hardened data to be dumped to the internal hard drive.

Begin the following process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDisks that are provided by the cluster.

1. Determine which hosts are using the VDisks that are provided by the cluster.
2. Repeat the previous step for all VDisks.

When input power is restored, you must press the power button on the uninterruptible power supply units before you press the power buttons on the SAN Volume Controller nodes.

Perform the following steps to shut down a cluster:

1. Click **Manage Clusters** → **Shut down Cluster** in the portfolio. The Shutting Down cluster panel is displayed.
2. Click **Yes**.

Shutting down a node

You can shut down a SAN Volume Controller node from the Shutting Down Node panel.

If you are shutting down the last SAN Volume Controller node in an I/O group, quiesce all I/O operations that are destined for this SAN Volume Controller node. Failure to do so can result in failed I/O operations being reported to your host operating systems.

This task assumes that you have already launched the SAN Volume Controller Console.

Do not disrupt the access of the hosts to their data when shutting down a node.

Follow MAP 5350 to shut down a node by using either the CLI or the SAN Volume Controller Console.

Configuring the cluster for iSCSI

You need to complete several tasks to configure the cluster to work with iSCSI-attached hosts.

Before completing any iSCSI-configuration tasks on the cluster, it is important that you complete all the necessary iSCSI-related configuration on the host machine. Because the SAN Volume Controller supports a variety of host machines, consult the documentation for specific instructions and requirements for a particular host. For a list of supported hosts, see the Support for SAN Volume Controller (2145) Web site:

www.ibm.com/storage/support/2145

To configure a cluster for iSCSI, follow these general tasks on the host system:

1. Select a software-based iSCSI initiator, such as Microsoft Windows iSCSI Software Initiator and verify the iSCSI driver installation.
2. If required, install and configure a multipathing driver for the host system.

In addition, determine a naming convention for iSCSI names such as iSCSI qualified names (IQNs) for your cluster. Hosts use iSCSI names to connect to the SAN Volume Controller node. Each node, for example, has a unique IQN, and the cluster name and node name are used as part of that IQN. Each node has a unique IQN, and the cluster name and node name are used as part of that IQN.

Configuring node Ethernet ports

You can use the SAN Volume Controller Console to configure node Ethernet ports, which represent iSCSI ports on a SAN Volume Controller node. Select the type of IP connection and enter the appropriate IP configuration.

This task assumes that you have already launched the SAN Volume Controller Console. To configure a node Ethernet port, complete these steps:

1. In the portfolio, click **Work with Nodes** → **Node Ethernet Ports**. The Viewing Node Ethernet Ports panel is displayed.
2. On the Viewing Node Ethernet Ports panel, select **Configure Ethernet Node Port** from the task list and click **Go**. The Configure Ethernet Node Port panel displays.
3. On the Configure Ethernet Node Port panel, select either IPv4 or IPv6 for the type of IP connection that the port uses. Both IPv4 and IPv6 can be concurrently specified; however, they must be configured separately. For example, you can configure the IPv4 address for the node Ethernet port first and then use this panel to configure the IPv6 address for the port. Depending on your selection, enter the following information for the port:

   IPv4 Enter the following information for IPv4 connections:
   - IP address
   - Subnet mask
   - Gateway

   IPv6 Enter the following information for IPv6 connections:
   - IP address
   - Prefix
   - Gateway

4. Click **OK**.

After you have configured the node Ethernet port on the cluster to represent the iSCSI-attached host, create a node Ethernet port for the partner node in the I/O group. This step is optional. However, when the node Ethernet port is configured for the partner node, it can handle requests to both its own node-port IP addresses and node-port IP address for the other node in the I/O group if that node goes offline.

Configuring partner node Ethernet port

You can use the SAN Volume Controller Console to set the node Ethernet port addresses for a partner node that is currently offline.

This task sets the IP configuration for a partner node, which is offline or not added to the I/O group. The failover configuration is activated on the local node until the partner node comes back online and activates the settings. To configure the other node Ethernet port, select the other local Ethernet node port and run this task again.

This task assumes that you have already launched the SAN Volume Controller Console. To configure a partner node Ethernet port, complete these steps:

1. In the portfolio, click **Work with Nodes** → **Node Ethernet Ports**. The Viewing Node Ethernet Ports panel is displayed.
2. On the Viewing Node Ethernet Ports panel, select **Configure Partner Node Ethernet Port** from the task list and click **Go**. The Configure Partner Node Ethernet Port panel displays.
3. On the Configure Partner Node Ethernet Port panel, select either IPv4 or IPv6 for the type of IP connection that the port uses. Depending on your selection, enter the following information for the port:

   IPv4 Enter the following information for IPv4 connections:
   - Node name (optional)
   - IP address
   - Subnet mask
   - Gateway

   IPv6 Enter the following information for IPv6 connections:
   - Node name (optional)
   - IP address
   - Prefix
   - Gateway
4. Click **OK**.

After you set the node Ethernet port addresses for a partner node, you can create iSCSI aliases.

Configuring or modifying an iSCSI alias

You can use the SAN Volume Controller Console to create or change the iSCSI alias for the selected node. An iSCSI alias is a user-assigned name that identifies the SAN Volume Controller node to the iSCSI-attached host. It is possible to change the iSCSI alias of a node even if the node is offline.

This task assumes that you have already launched the SAN Volume Controller Console. To change the iSCSI alias for a node, complete these steps:

1. In the portfolio, click **Work with Nodes** → **Node**. The Viewing Nodes panel is displayed.
2. On the Viewing Nodes panel, select the node that you want to work with. Select **Modify iSCSI Alias** from the task list and click **Go**. The Modify iSCSI Alias panel displays.
3. Enter the new iSCSI alias for the selected node and click **OK**.

After you have created iSCSI aliases, you can optionally configure the address for the Internet Storage Name Service (iSNS) server for the cluster.

Configuring the iSNS server address

If you are using iSCSI-attached hosts with the SAN Volume Controller cluster, you can optionally configure the address for the Internet Storage Name Service (iSNS) server for the cluster. Host systems use the iSNS server to manage iSCSI targets and for iSCSI discovery.

This task assumes that you have already launched the SAN Volume Controller Console. To configure the address for an iSNS server, follow these steps:

1. In the portfolio, click **Configure iSNS Server**. The Configure iSNS Server panel is displayed.
2. On the Configure iSNS Server panel, select either IPv4 or IPv6 for the protocol that is used by the iSNS server and enter the IP address for the server.
3. Click **OK**.

After you configure the iSNS server address for the cluster, you can configure cluster iSCSI authentication.

Note: To help in problem determination, this step can be delayed until after the first one or two hosts have been configured and their connectivity has been tested without authentication configured.

Configuring cluster iSCSI authentication

You can use the SAN Volume Controller Console to configure the Challenge-Handshake Authentication Protocol (CHAP) to authenticate the SAN Volume Controller cluster to the iSCSI-attached hosts. After the CHAP is set for the cluster, all attached hosts must be configured to authenticate this way. To help in problem determination, this step can be delayed until after the first one or two hosts have been configured and their connectivity has been tested without authentication configured.

This task assumes that you have already launched the SAN Volume Controller Console. To configure authentication between the SAN Volume Controller cluster to the iSCSI-attached hosts, follow these steps:

1. In the portfolio, click **Configure iSCSI Authentication**. The Configure Cluster iSCSI Authentication panel is displayed.
2. On the Configure Cluster iSCSI Authentication panel, enter the shared passphrase that is used to authenticate the SAN Volume Controller to the host in the CHAP Authentication Secret field.
3. Click **OK**.

After you configure the CHAP secret for the SAN Volume Controller cluster, ensure that the cluster CHAP secret is added to each iSCSI-attached host. On all iSCSI-attached hosts, specify a CHAP secret that the hosts use to authenticate to the SAN Volume Controller cluster.

Configuring host objects on a cluster

You can configure both fibre-channel or iSCSI-attached hosts to the cluster.

If you are configuring a host object on a fibre-channel attached host, ensure that you have completed all zone and switch configuration. Also test the configuration to ensure that zoning was created correctly.

If you are configuring a host object on the cluster that uses iSCSI connections, ensure that you have completed the necessary host-system configurations and have configured the cluster for iSCSI connections.

Discovering MDisks

You can have the cluster rescan the fibre-channel network. The rescan discovers any new managed disks (MDisks) that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to discover MDisks:

1. Click **Work with Managed Disks** → **Managed Disks** in the portfolio. The Viewing Managed Disks panel is displayed.
2. Select **Discover MDisks** from the task list and click **Go**. The Discovering Managed Disks panel is displayed. The newly discovered MDisks are displayed in a table on the Discovering Managed Disks panel.
3. Click **Close** to return to the Viewing Managed Disks panel.

Viewing discovery status

You can view the status of a managed disk (MDisk) discovery from the Viewing Discovery Status panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view status of an MDisk discovery:

1. Click **Work with Managed Disks** → **Discovery Status**. The Viewing Discovery Status panel is displayed. The following status values are possible:

   Active This status indicates an MDisk discovery is currently in progress.

   Inactive
   This status indicates an MDisk discovery is not currently in progress.
2. Click **Close** to close this panel.

Renaming MDisks

You can rename a managed disk (MDisk) from the Renaming Managed Disk panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to rename an MDisk:

1. Click **Work with Managed Disks** → **Managed Disks** in the portfolio. The Viewing MDisks panel is displayed.
2. Select the MDisk you want to rename and select **Rename an MDisk** from the list. Click **Go**. The Renaming Managed Disk panel is displayed.
3. Type a new name for the MDisk.
4. Click **OK**.

Adding excluded MDisks to a cluster

You can add managed disks (MDisks) that have been excluded from the cluster back into the cluster from the Including Managed Disk panel.

You must fix the fabric-related problem that caused the MDisk to become excluded from the cluster before you can add the MDisk to the cluster.

This task assumes that you have already launched the SAN Volume Controller Console.

The MDisk might have been excluded from the cluster because of multiple I/O failures that are caused by noisy links.

Perform the following steps to add an excluded MDisk to a cluster:

1. Click **Work with Managed Disks** → **Managed Disks** in the portfolio. The Viewing MDisks panel is displayed.
2. Select the excluded MDisk to add to the cluster and select **Include an MDisk** from the list. Click **Go**. The Including Managed Disk panel is displayed.
3. Follow the instructions that are displayed on the Including Managed Disk panel.

Setting quorum disks

You can set a managed disk (MDisk) as a quorum disk from the Setting a Quorum Disk panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Quorum disks are used when there is a problem in the SAN fabric or when nodes are shut down, leaving half of the nodes remaining in the cluster. This type of problem causes a loss of communication between the nodes that remain in the cluster and those that do not. The nodes are split into groups where the nodes in each group can communicate with each other, but not with the other group of nodes that were formerly part of the cluster.

In this situation, some nodes must stop operating and processing I/O requests from hosts to preserve data integrity while maintaining data access. If a group contains less than half the nodes that were active in the cluster, the nodes in that group stop operating and processing I/O requests from hosts.

It is possible for a cluster to split into two groups with each group containing half the original number of nodes in the cluster. A quorum disk determines which group of nodes stops operating and processing I/O requests. In this tie-break situation, the first group of nodes that accesses the quorum disk marks their ownership of the quorum disk and as a result continues to operate as the cluster, handling all I/O requests. If the other group of nodes cannot access the quorum disk or finds it owned by another group of nodes, it stops operating as the cluster and does not handle I/O requests.

A quorum disk is formed by taking a small amount of space from a managed disk (MDisk). There is only one active quorum disk that is used to determine which nodes operate as the cluster. However, up to three candidate quorum disks are maintained. The purpose of the other quorum disks is to provide redundancy if a quorum disk fails prior to the cluster being split into two equal groups of nodes.

The SAN Volume Controller automatically chooses up to three MDisk as quorum disk candidates and then automatically selects one of these as the active quorum disk. You can also manually set quorum disk candidates and the active quorum disk.

You must set quorum disks on multiple controllers to avoid the possibility of losing all of the quorum disks with a single failure. Mirrored virtual disks (VDisks) also are taken offline when all the quorum disks are on a single controller that experiences a failure. Mirrored VDisks are taken offline because synchronization data is stored on the quorum disks. To protect against mirrored VDisks being taken offline, follow these guidelines for changing the MDisk that are assigned as quorum disk candidates:

- When possible, distribute the quorum candidate disks so that each MDisk is provided by a different storage system. For a list of storage systems that support quorum disks, refer to the Supported Hardware List located at the SAN Volume Controller (2145) Web site:

  www.ibm.com/storage/support/2145
- Ensure that the status of the MDisk that is being set as a quorum disk is online before setting it as the quorum disk.
- Wait at least two minutes between setting quorum disks. This delay allows time for the update to the quorum disk to complete before setting another quorum disk.

Attention: MDisks that are associated with solid-state drives (SSDs) on SAN Volume Controller 2145-CF8 nodes cannot be used as quorum disks.

Complete the following steps to set an MDisk as a quorum disk:

1. In the portfolio, click **Work with Managed Disks** → **Quorum Disks**. The Viewing Quorum Disks panel is displayed.
2. Select the MDisk to set as a quorum disk and select **Set a Quorum Disk** from the list. Click **Go**. The Setting a Quorum Disk panel is displayed.
3. Select a quorum index number from the **Quorum Index** list. and click **OK**.
4. If you want this disk to be used as the active quorum disk, select **Set as Active Quorum Disk**. The active quorum disk is the disk which is used first in the event of a cluster partition. You should set an active quorum disk when nodes in an I/O group are split between different physical locations to ensure that the most highly-available quorum disk is used.

Setting the active quorum disk

You can assign one of the quorum disks to be the active quorum disk in the SAN Volume Controller Console. In a cluster partition, the active quorum disk is used first to establish a quorum. Selecting a disk as the active quorum disk is useful in split-site configurations. Setting the active quorum disk ensures that the most highly available quorum disk is used in normal processing.

The cluster automatically chooses three managed disks (MDisks) as quorum disk candidates. Each disk is assigned a number (either 0, 1, or 2) that is called a *quorum index*. The number identifies the MDisk as a quorum disk to the cluster. You can also assign one of the quorum disks to be the active quorum disk. If a tie-break situation occurs, the active quorum disk is used first to establish a quorum. If SAN connectivity is lost, the cluster can continue to operate. One half

of the cluster obtains the active quorum disk, and continues to process transactions. The other half ceases to process transactions as part of the cluster, until SAN connectivity is restored.

Attention: MDisks that are associated with solid-state drives (SSDs) on SAN Volume Controller 2145-CF8 nodes cannot be used as quorum disks.

This task assumes that you have already launched the SAN Volume Controller Console. To set the active quorum disk, follow these steps:

1. In the portfolio, click **Work with Managed Disks** → **Quorum Disks**. The Viewing Quorum Disks panel is displayed.
2. Select the quorum disk that you want to set as the active quorum disk. From the task list, select **Set a Quorum Disk**. Click **Go**. The Setting Active Quorum Disk panel is displayed.
3. Confirm that the selected disk is the one that you want to make the active quorum disk by clicking **Set Active Quorum Disk**.

Determining the relationship between MDisks and VDisks

You can use the SAN Volume Controller Console to determine the relationship between managed disks (MDisks) and virtual disks (VDisks).

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to determine the relationship between MDisks and VDisks:

1. Click **Work with Managed Disks** → **Managed Disks** in the portfolio. The Viewing Managed Disks panel is displayed.
2. Select the MDisk that you want to view.
3. Select **Show VDisks** from the task list and click **Go**. The Viewing Virtual Disks panel is displayed. This panel lists the VDisks that use this MDisk.

Determining the relationship between MDisks and RAID arrays or LUNs

Each managed disk (MDisk) corresponds with a single RAID array, or a single partition on a given RAID array. Each RAID controller defines a LUN number for this disk. The LUN number and controller name or ID are needed to determine the relationship between MDisks and RAID arrays or partitions.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to determine the relationship between MDisks and RAID arrays:

1. Click **Work with Managed Disks** → **Managed Disks** in the portfolio. The Viewing Managed Disks panel is displayed.
2. Click the name of the MDisk that you want to view. The Viewing Managed Disk (MDisk) Details panel is displayed.
3. Record the controller name and controller LUN number.
4. Click **Work with Managed Disks** → **Disk Controller Systems** in the portfolio.
5. Click the name of the controller that you recorded in step 3 to show the detailed view of the controller. The Viewing General Details panel is displayed.

6. Record the vendor ID, the product ID and worldwide node name (WWNN).
7. Use the vendor ID, the product ID and WWNN to determine which controller presents this MDisk.
8. From the native user interface for the controller that presents this MDisk, list the LUNs that the controller presents and match the LUN number with that noted in step 2 on page 153. This is the exact RAID array and partition that corresponds with the MDisk.

Creating MDisk 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 (VDisk) allocation within one disk controller system, ensure that the MDisk group that corresponds with a single disk controller system is presented by that disk controller system. This also enables nondisruptive migration of data from one disk controller system to another disk controller system and simplifies the decommissioning process if you want to decommission a disk controller system at a later time.

Ensure all MDisks that are allocated to a single MDisk group are of the same RAID-type. Using the same RAID-type 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 have three RAID-5 arrays in one group and add a non-RAID disk to this group, you lose access to all the data that is striped across the group if the non-RAID disk fails. Similarly, for performance reasons, you should not mix RAID types.

If you are using a SAN Volume Controller solid-state drive (SSD) managed disk, ensure that you are familiar with the SSD configuration rules.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to create a new MDisk group:

1. Click **Work with Managed Disks → Managed Disk Groups** in the portfolio. The Viewing Managed Disk Groups panel is displayed.
2. Select **Create an MDisk Group** from the task list and click **Go**. The Create a Managed Disk Group wizard begins.
3. Complete the Create a Managed Disk Group wizard.

Adding MDisks to MDisk groups

You can add managed disks (MDisks) to an MDisk group from the Adding Managed Disks to Managed Disk Group panel.

If you are using a SAN Volume Controller solid-state drive (SSD) managed disk, ensure that you are familiar with the SSD configuration rules.

This task assumes that you have already launched the SAN Volume Controller Console.

Solid-state drives (SSDs) that are located in SAN Volume Controller 2145-CF8 nodes are presented to the cluster as MDisks. To determine whether the selected MDisk is an SSD, click the link on the MDisk name to display the Viewing MDisk

Details panel. If the selected MDisk is an SSD that is located on a SAN Volume Controller 2145-CF8 node, the Viewing MDisk Details panel displays values for the Node ID, Node Name, and Node Location attributes. Alternatively, you can select **Work with Managed Disks → Disk Controller Systems** from the portfolio. On the Viewing Disk Controller panel, you can match the MDisk to the disk controller system that has the following values for these attributes.

Table 23. Disk controller attributes for SSDs

Attribute	SSD value
Vendor ID	IBM
Product ID Low	2145
Product ID High	Internal

Note: The first time that you add a new solid-state drive (SSD) to an MDisk group, the SSD is automatically formatted and set to a block size of 512 bytes.

Perform the following steps to add MDisks to an MDisk group:

1. Click **Work with Managed Disks → Managed Disk Groups** in the portfolio. The Viewing Managed Disk Groups panel is displayed.
2. Select the MDisk group to add MDisks to and select **Add MDisks** from the list. Click **Go**. The Adding Managed Disks to Managed Disk Group panel is displayed.
3. Select the MDisks to add and click **OK**.

Displaying MDisk groups

You can display the managed disk (MDisk) group that an MDisk is a part of from the Viewing Managed Disk Groups panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to display the MDisk group:

1. Click **Work with Managed Disks → Managed Disks** in the portfolio. The Viewing Managed Disks panel is displayed.
2. Select the MDisk and select **Show MDisk Group** from the list. Click **Go**. The Viewing Managed Disk Groups panel is displayed. The MDisk group is displayed in a table on the Viewing Managed Disk Groups panel.

Removing MDisks from an MDisk group

You can remove managed disks (MDisks) from an MDisk group.

This task assumes that you have already launched the SAN Volume Controller Console.

Solid-state drives (SSDs) that are located in SAN Volume Controller 2145-CF8 nodes are presented to the cluster as MDisks. To determine whether the selected MDisk is an SSD, click the link on the MDisk name to display the Viewing MDisk Details panel. If the selected MDisk is an SSD that is located on a SAN Volume Controller 2145-CF8 node, the Viewing MDisk Details panel displays values for the Node ID, Node Name, and Node Location attributes. Alternatively, you can select **Work with Managed Disks → Disk Controller Systems** from the portfolio. On the

Viewing Disk Controller panel, you can match the MDisk to the disk controller system that has the following values for these attributes.

Table 24. Disk controller attributes for SSDs

Attribute	SSD value
Vendor ID	IBM
Product ID Low	2145
Product ID High	Internal

Complete the following steps to remove an MDisk from an MDisk group:

1. Click **Work with Managed Disks → Managed Disk Groups** in the portfolio. The Viewing Managed Disk Groups panel is displayed.
2. Select the MDisk Group that you want to delete MDisks from and select **Remove MDisks** from the list. Click **Go**. The Deleting Managed Disks from Managed Disk Group panel is displayed.
3. Select the MDisk that you want to remove.
4. Click **OK**.

Viewing the progress of an MDisk removal

You can view the progress of a managed disk (MDisk) removal from the Viewing MDisk Removal Progress panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the progress of an MDisk removal:

1. Click **Manage Progress → View Progress**. The View Progress panel is displayed.
2. Click the **MDisk Removal** link. The Viewing MDisk Removal Progress panel is displayed.

Renaming MDisk groups

You can rename a managed disk (MDisk) group from the Modify Managed Disk Group panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to rename an MDisk group:

1. Click **Work with Managed Disks → Managed Disk Groups** in the portfolio. The Viewing Managed Disk Groups panel is displayed.
2. Select the MDisk group that you want to rename and select **Modify an MDisk Group** from the list. Click **Go**. The Modify Managed Disk Group panel is displayed.

Deleting MDisk groups

You can delete a managed disk (MDisk) group using the Deleting a Managed Disk Group panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to delete an MDisk group:

1. Click **Work with Managed Disks → Managed Disk Groups** in the portfolio. The Viewing Managed Disk Groups panel is displayed.
2. Select the MDisk group to delete and select **Delete an MDisk Group** from the list. Click **Go**. The Deleting a Managed Disk Group panel is displayed.

Creating VDisks

You can create virtual disks (VDisks) using the Create Virtual Disks wizard.

If the VDisk that you are planning to create use disk extents on a solid-state drive (SSD) that is located on a SAN Volume Controller 2145-CF8 node, the data stored on the VDisk is not protected against SSD failures or node failures. To avoid data loss, add a VDisk copy that maps to the SSD on another SAN Volume Controller 2145-CF8 node.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to create VDisks:

1. Click **Work with Virtual Disks → Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.
2. Select **Create VDisks** from the task list and click **Go**. The Create Virtual Disks wizard begins. This wizard allows you to create mirrored VDisks and space-efficient VDisks.
3. Complete the Create Virtual Disks wizard.

Creating VDisks for FlashCopy targets

You can select an existing source virtual disk (VDisk) to use to create target VDisks for FlashCopy mappings. The Creating VDisk for use as FlashCopy Targets wizard in the SAN Volume Controller Console lets you create one or more source VDisks for mappings.

The FlashCopy feature makes an instant copy of a VDisk at the time that it is started. To create an instant copy of a VDisk, 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 wizard ensures that the source and target VDisks are equal size.

The wizard takes you through the following steps to create a target VDisk for FlashCopy mappings:

- Specifying the I/O group and the preferred node for the target VDisk
- Setting attributes for the source VDisk, such as the number of VDisks that you are creating and space-efficient attributes
- Selecting the managed disk (MDisk) groups and MDisk for all the target VDisks
- Naming the VDisk

Note: The target name is added to the source VDisk name and is incremented if multiple target VDisks are being created. Full names for the target VDisks cannot exceed 15 characters. Only VDisks with unique names are created. If you have previously created target VDisks, they are not created. To check whether the specified name for the VDisk exists on the cluster, select **Check Name Availability** on the Name the VDisk panel in the

wizard. This action displays names that are currently defined on the cluster and, as a result, are not created as a target VDisk.

- Verifying the target VDisk attributes and settings

This task assumes that you have already launched the SAN Volume Controller Console.

To create target VDisks for FlashCopy mappings, complete the following steps:

1. In the portfolio, click **Work with Virtual Disks** → **Virtual Disks**. The Viewing Virtual Disks panel is displayed.
2. Select the VDisk to use as the source VDisk for the FlashCopy mapping. From the task list, select **Create VDisks For Use as FlashCopy Targets**. Click **Go**.
3. Select the host that you want to modify. From the task list, select **Modify a Host**. Click **Go**. The Create VDisks For Use as FlashCopy Targets wizard begins.
4. Complete the **Create VDisks For Use as FlashCopy Targets** wizard. For detailed information regarding the fields and options that are displayed on the wizard panels, use the online help in the SAN Volume Controller Console.

Displaying VDisks

You can display the virtual disks (VDisks) that use a managed disk (MDisk) group from the Viewing Virtual Disks panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to display the VDisks that use an MDisk group:

1. Click **Work with Managed Disks** → **Managed Disk Groups** in the portfolio. The Viewing Managed Disk Groups panel is displayed.
2. Select the MDisk group to display VDisks for and select **Show VDisks Using This Group** from the list. Click **Go**. The Viewing Virtual Disks panel is displayed.

Moving a VDisk to a new I/O group

You can move a virtual disk (VDisk) to a new I/O group to manually balance the workload across the nodes in the cluster.

This task assumes that you have already launched the SAN Volume Controller Console.

Attention: This is a disruptive procedure. Access to the VDisk is lost while you follow this procedure. Under no circumstances should VDisks be moved to an offline I/O group. You must ensure that the I/O group is online before moving the VDisks to avoid data loss.

Complete the following steps to move a single VDisk:

1. Quiesce all I/O operations for the VDisk. You might have to determine the hosts that are using this VDisk.
2. Update the multipathing device driver configuration to remove all device identifiers that are presented by the VDisk you intend to move. If you are using the system device driver (SDD), the device identifiers are referred to as virtual paths (vpaths).

Attention: Failure to perform this step can result in data corruption.

3. Stop and delete all FlashCopy mappings, Metro Mirror, or Global Mirror relationships that use this VDisk. To check if the VDisk is part of a mapping or relationship, perform the following steps:
   a. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.
   b. Click on the name of the VDisk that you want to migrate. The View VDisk General Details panel is displayed.
   c. Look for the **FlashCopy Map Count** and **Relationship ID** fields. If these fields are not blank or contain zeros, the VDisk is part of a mapping or relationship.
   d. Click **Close** to close the panel.
4. Move the VDisk by selecting the VDisk from the Viewing Virtual Disks panel and select **Modify a VDisk** from the task list and click **Go**. The Modifying Virtual Disk panel is displayed.
5. Select the new I/O group from the **I/O Group** list and click **OK**.
6. Follow your multipathing device drivers instructions for discovering new device identifiers. For example, if you are using SDD, see the *IBM System Storage Multipath Subsystem Device Driver User's Guide* and follow the instructions for discovering vpaths.

Viewing the progress of VDisk formatting

You can view the progress of virtual disk (VDisk) formatting from the Viewing VDisk Formatting Progress panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the progress of VDisk formatting:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click the **VDisk Formatting** link. The Viewing VDisk Formatting Progress panel is displayed.

Migrating VDisks

You can migrate a virtual disk (VDisk) from one managed disk (MDisk) group to another from the Migrating VDisks panel.

This task assumes that you have already launched the SAN Volume Controller Console.

The SAN Volume Controller provides various data migration features. You can use these features to move the placement of data both within MDisk groups and between MDisk groups. These features can be used concurrently with I/O operations. There are two ways that you can migrate data:

1. Migrate data (extents) from one MDisk to another MDisk within the same MDisk group. This can be used to remove active or overutilized MDisks. This can only be performed using the command-line interface (CLI).
2. Migrate VDisks from one MDisk group to another. This can be used to remove active MDisk groups; for example, you can reduce the utilization of a group of MDisks.

You can determine the usage of MDisks by gathering I/O statistics about nodes, MDisks, and VDisks. After you have gathered this data, you can analyze it to determine which VDisks or MDisks are active.

When a migration command is issued, a check ensures that the migration target has sufficient free extents available. If there are sufficient free extents, the command proceeds.

Notes:

- 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 start with a non-mirrored VDisk in one MDisk group and then add a mirrored copy to that VDisk in another MDisk group. You can also create a FlashCopy mapping to create an instant copy of a VDisk that is in another MDisk group.
- Migration commands fail if the target or source VDisk is offline, or if there is insufficient quorum disk space to store the metadata. Correct the offline or quorum disk condition and retry the command.

While the migration proceeds, 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 situation, 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 (by marking the error as fixed).
2. Migrate one or more VDisks that are already created from the MDisk group to another group. This frees up extents in the group and allows the original migrations to be restarted.

Perform the following steps to migrate VDisks between MDisk groups:

1. Perform the following steps to determine if VDisks are overused:
   a. Click **Manage Cluster → Start statistics collection** in the portfolio. The Starting the Collection of Statistics panel is displayed.
   b. Enter 15 minutes for the interval and click **OK**. This generates a new I/O statistics dump file approximately every 15 minutes.
   c. Wait at least 15 minutes before you proceed to the next step.
2. View the I/O statistics log.
   a. Click **Service and Maintenance → List dumps** in the portfolio. The List Dumps panel is displayed.
   b. Click **I/O Statistics Logs**. This lists the I/O statistics files that have been generated on a per-node basis. These are prefixed with `Nm` for MDisk statistics, `Nv` for VDisk statistics, and `Nn` for node statistics.
   c. Click a filename to view the contents of the log.
   d. Analyze the dumps to determine which VDisks are active. It might be helpful to also determine which MDisks are heavily utilized so you can spread the data that they contain more evenly across all the MDisks in the group. Either create a new MDisk group or determine an existing group that is not yet over used. You can do this by checking the I/O statistics files that were previously generated and ensuring that the MDisks or VDisks in the target MDisk group are less utilized than the source group.

3. Stop the statistics collection by clicking **Manage Cluster** → **Stop statistics collection** in the portfolio.
4. Migrate the VDisk.
   a. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.
   b. Select the VDisk to migrate and select **Migrate a VDisk** from the task list. Click **Go**. The Migrating Virtual Disks panel is displayed.
   c. Select a target MDisk group from the **Target MDisk Group** list.
   d. Click **OK**.

Viewing the progress of VDisk migration

You can view the progress of virtual disk (VDisk) migration from the Viewing VDisk Migration Progress panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the progress of VDisk migration:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click the **VDisk Migration** link. The Viewing VDisk Migration Progress panel is displayed.

Shrinking VDisks

You can use the SAN Volume Controller Console to make a target or auxiliary virtual disk (VDisk) the same size as the source or master VDisk when you create FlashCopy® mappings, Metro Mirror relationships, or Global Mirror relationships.

You can shrink a VDisk from the Shrinking Virtual Disks panel in the SAN Volume Controller Console. Shrinking a VDisk reduces its total capacity. Use this function when you want a FlashCopy target disk to have the same capacity as its source. However, if the VDisk contains data, do not shrink the size of the disk.

Attention:

1. The SAN Volume Controller arbitrarily reduces the capacity of the VDisk by removing one or more extents from those that are allocated to the VDisk. You cannot control which extents are removed so you cannot guarantee that it is unused space that is removed.
2. If the VDisk contains data that is being used, but you still want to reduce its size, ensure that you back up your data before you attempt this operation.
3. If the VDisk is being used by hosts, ensure that the target VDisk is not mapped to any hosts before you shrink the VDisk.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to shrink a VDisk:

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.
2. Select the VDisk to shrink and select **Shrink a VDisk** from the task list. Click **Go**. The Shrinking Virtual Disks panel is displayed.
3. Enter the capacity to reduce the size of the VDisk by in the **Reduce By Capacity** field and click **OK**.

Shrinking or expanding space-efficient VDisks

You can use the SAN Volume Controller Console to increase or decrease the real capacity of a space-efficient virtual disk (VDisk).

You can use the Shrink/Expand Space-Efficient Disks panel in the SAN Volume Controller Console to change the real capacity of a space-efficient VDisk, unless the VDisk is in image mode.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to shrink or expand a space-efficient VDisk:

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.
2. Select the VDisk and select **Expand/Shrink Space-efficient VDisk** from the task list. Click **Go**. The Shrink/Expand Space-Efficient Disks panel is displayed.
3. Choose which copies of the VDisk that you want to shrink or expand.
4. Enter the amount to either decrease or increase the real capacity of the selected space-efficient VDisk in the **Amount to shrink/expand** field.

   Note: You cannot shrink the real capacity of a space-efficient VDisk below its used capacity. Further shrinking is restricted to prevent data loss.
5. Select the **Shrink** or **Expand** option and then click **OK**.
6. If you are expanding a striped space-efficient VDisk, you can select an MDisk candidate to be used to allocate new extents to the space-efficient VDisk.

   If you are working with VDisk copies, perform the following steps:

   a. Select the MDisk to use to allocate new extents to the space-efficient VDisk.
   b. Select **MDisk for each copy**.
   c. Click **Add** to add the selected MDisk to the **Managed Disks Striped in This Order** list.
   d. After you have added all the MDisks to use to expand the capacity of the VDisk, you can use the arrows to determine the order in which they are used.

Configuring bitmap space for Copy Services or VDisk mirroring

You can use the Modify Copy Service Space panel in the SAN Volume Controller Console to modify the amount of memory that is available for the FlashCopy, Metro Mirror, or Global Mirror Copy Services features or virtual disk (VDisk) mirroring.

The total bitmap space that is available for Copy Services features and VDisk mirroring in an I/O group is 512 MB. The following table provides an example of the amount of memory that is required for VDisk Mirroring and each Copy Service feature:

Feature	Grain size	1 MB of memory provides the following VDisk capacity for the specified I/O group
Metro Mirror or Global Mirror	256 KB	2 TB of total Metro Mirror and Global Mirror VDisk capacity

Feature	Grain size	1 MB of memory provides the following VDisk capacity for the specified I/O group
FlashCopy	256 KB	2 TB of total FlashCopy source VDisk capacity
FlashCopy	64 KB	512 GB of total FlashCopy source VDisk capacity
Incremental FlashCopy	256 KB	1 TB of total incremental FlashCopy source VDisk capacity
Incremental FlashCopy	64 KB	256 GB of total incremental FlashCopy source VDisk capacity
VDisk Mirroring	256 KB	2TB of mirrored VDisk capacity

Notes:

1. For multiple FlashCopy targets, you must consider the number of mappings. For example, for a mapping with a grain size of 256 KB, 8 KB of memory allows one mapping between a 16 GB source VDisk and a 16 GB target VDisk. Alternatively, for a mapping with a 256 KB grain size, 8 KB of memory allows two mappings between one 8 GB source VDisk and two 8 GB target VDisks.
2. When creating a FlashCopy mapping, if you specify an I/O group other than the I/O group of the source VDisk, the memory accounting goes towards the specified I/O group, not towards the I/O group of the source VDisk.
3. For VDisk mirroring, the full 512 MB of memory space enables 1 PB of total VDisk Mirroring capacity.

Before you specify the memory settings on the Modify Copy Service Space panel, consider the following factors.

- For FlashCopy relationships, only the source VDisk allocates space in the bitmap table.
- For Metro Mirror or Global Mirror relationships, two bitmaps exist. One is used for the master cluster and one is used for the auxiliary cluster, because the direction of the relationship can be reversed.
- The smallest possible bitmap is 4 KB; therefore, a 512 byte VDisk requires 4 KB of bitmap space.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to change the amount of memory that is available for Copy Services features or virtual disk (VDisk) mirroring:

1. Click **Work with Nodes → I/O Groups** in the portfolio. The Viewing Input/Output Groups panel is displayed.
2. Select the I/O Group and select **Modify Copy Services Space** from the task list. Click **Go**. The Modify Copy Service Space panel is displayed.

   The current setting of total memory that is allocated for Global / Metro Mirror, FlashCopy, and VDisk Mirroring is displayed. The total free memory that can be allocated for Global / Metro Mirror, FlashCopy, and VDisk Mirroring is also displayed.
3. Enter a value between 0 MB and 512 MB for the new total amount of memory to allocate for Global / Metro Mirror, FlashCopy, or virtual disk (VDisk) mirroring.

4. Click **OK** to change the total bitmap space allocated for the selected Copy Services feature or VDisk mirroring in an I/O group.

Adding a mirrored copy to a VDisk

Use the Add Copy to VDisk panel in the SAN Volume Controller Console to add a mirrored copy to the selected virtual disk (VDisk). Each VDisk can have a maximum of two copies.

Creating mirrored copies of a VDisk allows the VDisk to remain accessible even when a managed disk (MDisk) that the VDisk depends on becomes unavailable. You can create copies of a VDisk either from different MDisk groups or by creating an image mode copy of the VDisk. Copies allow for availability of data; however, they are not separate objects. You can only create or change mirrored copies from the VDisk.

In addition, you can use VDisk mirroring as an alternative method of migrating VDisks between MDisk groups. For example, if you have a nonmirrored VDisk in one MDisk group and want to migrate that VDisk to a second MDisk group, you can add a new copy of the VDisk by specifying the second MDisk group for that VDisk copy. After the copies have synchronized, you can delete the copy in the first MDisk group. The VDisk is migrated to the second MDisk group while remaining online during the migration.

This alternative method of migrating VDisks has the following advantages:

- Access to the VDisk data is not lost if the second MDisk group goes offline during the migration.
- The speed of the migration can be adjusted, using the VDisk synchronization rate, and the migration can be paused.
- The migration can be ended by deleting the VDisk copy in the second MDisk group before migration completes.
- The MDisk groups can have different extent sizes.

This alternative method of migrating VDisks has the following limitations:

- You cannot use this method for VDisks that are already mirrored.
- There are more manual steps that are associated with this method.
- Write I/O performance is slightly affected during the migration, because the mirrored copies must be kept synchronized.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to add a mirrored copy to a VDisk:

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.
2. Select the VDisk to copy and select **Add Mirrored VDisk Copy** from the task list. Click **Go**. The Add Copy to VDisk panel is displayed.
3. Select the type of VDisk copy and then select the managed disk (MDisk) group to create the VDisk copy from. View the panel Help for more information about the options that are available on this panel.
4. Click **OK** to add a mirrored copy to the selected VDisk.

Splitting a VDisk copy

You can create a separate virtual disk (VDisk) by splitting a synchronized VDisk copy. You can select a copy to split off from the VDisk and set its attributes.

1. Click **Work with Virtual Disks** → **Virtual Disks**. The Viewing Virtual Disk panel is displayed.
2. Select a VDisk that contains copies and select **Split a VDisk Copy** from the task list and click **Go**. The Split a Copy from VDisk panel displays.
3. Select the VDisk copy that is displayed in the table to create a new VDisk.
4. Enter a name for the new VDisk.
5. Select the **Force Split** option to force the split to proceed even though the copy you are trying to split is not synchronized.

   Note: If you select this option, the split copy might not be point-in-time consistent.
6. Select the I/O group for the new VDisk that you are creating from the VDisk copy. By default, the VDisk is created in the same I/O group as the VDisk that the copy is split from.
7. Select the preferred node for the new VDisk that you are creating from the VDisk copy. When you let the system choose the preferred node, the system performs workload balancing by managing the I/O traffic for the VDisks across multiple nodes.
8. Select the cache mode for the new VDisk that you are creating from the VDisk copy.
9. Optional: Enter a unit device identifier for the new VDisks in the **Unit Device Identifier** field. This field is used only by hosts that are using the OpenVMS operating system. For other operating systems, setting a unit device identifier is not required.
10. Click **OK**.

Deleting a copy from a VDisk

Use the Deleting a Copy from VDisk panel in the SAN Volume Controller Console to delete a mirrored copy from the selected virtual disk (VDisk).

If the VDisk that you are deleting a copy from maps to a solid-state drive (SSD), the data that is stored on the VDisk is not protected against SSD failures or node failures. To avoid data loss, ensure that a VDisk copy exists that maps to an SSD on another node.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to delete a mirrored copy from a VDisk:

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.
2. Select the VDisk with a copy to delete and select **Delete a Mirrored VDisk Copy** from the task list. Click **Go**. The Deleting a Copy from VDisk panel is displayed.
3. If necessary, click **Force delete** to force the deletion of the VDisk copy in the following situations:
   - Migration to an image mode VDisk is in progress for the selected VDisk copy.

- If the selected VDisk copy is an image mode with virtual medium errors.
- The cache is not empty for an image mode VDisk copy.
- The image mode VDisk copy is not synchronized.

If the copy being deleted is the last synchronized VDisk copy, the VDisk and all its copies are deleted.

4. Select the VDisk copy to delete and click **OK**.

Viewing virtual disk-to-host mappings

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

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view your virtual disk-to-host mappings:

1. Click **Work with Virtual Disks** → **Virtual Disk-to-Host Mappings** in the portfolio. The Virtual Disk-to-Host Mappings panel is displayed.
2. Click **Close** to close the panel.

Creating a VDisk-to-host mapping

You can create a new mapping between a virtual disk (VDisk) and a host from the Creating Virtual Disk-to-Host Mappings panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to create a new mapping:

1. In the portfolio, click **Work with Virtual Disks** → **Virtual Disks**. The Viewing Virtual Disk panel is displayed.
2. Select the VDisk to map to your host and select **Map VDisks to a Host** from the list. Click **Go**. The Creating Virtual Disk-to-Host Mappings panel is displayed.
3. Select the host that you want to map the VDisk to and click **OK**.

After you have mapped VDisks to hosts, discover the disks on the host machine. This step requires accessing the host system and using the host system utilities to discover the new disks that are made available by the SAN Volume Controller. You also have the option of creating a file system for those new disks. Consult your host system documentation for more information on completing this task.

Deleting a virtual disk-to-host mapping

You can delete a mapping between a virtual disk (VDisk) and a host object from the Deleting a Virtual Disk-to-Host Mapping panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to delete a mapping between a VDisk and a host object:

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.

2. Select the VDisk that you no longer want mapped to your host and select **Delete a VDisk-to-Host-Mapping** from the list and click **Go**. The Deleting a VDisk-to-Host Mapping panel is displayed.
3. Select the VDisk from which you want to remove the VDisk mapping and click **OK**.

Determining the relationship between VDisks and MDisks

You can use the SAN Volume Controller Console to determine the relationship between virtual disks (VDisks) and managed disks (MDisks).

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to determine the relationship between VDisks and MDisks:

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.
2. Select the VDisk that you want to view.
3. Select **Show MDisks This VDisk is Using** from the task list and click **Go**. The Viewing Managed Disks panel is displayed. This panel lists the MDisks that the selected VDisk uses.

Verifying and repairing mirrored VDisk copies

The virtual disk (VDisk) copy verification process checks if data on mirrored VDisk copies match. You can choose repair options if differences are found during the verification process.

Attention: Proceed with this task only if all VDisk copies are synchronized.

Use the Verifying VDisk Copies panel to start the VDisk copy verification process for a selected VDisk. If differences are found during verification, you can choose one of the following actions.

- Stop the process when the first difference is found. Select this option if you only want to verify that the mirrored VDisk copies are identical. You can run this option, starting at a different logical block address (LBA) each time to count the number of differences on a VDisk.
- Automatically repair the copy by overwriting sectors with data from the primary VDisk copy. Select the resync option if you are sure that either the primary VDisk copy data is correct or that your host applications can handle incorrect data.
- Create a virtual medium error at the VDisk level. Select this option if you are unsure what the correct data is and you do not want an incorrect version of the data to be used.

If a medium error is encountered on one of the copies, the VDisk copy is automatically repaired if the data can be read from another copy.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to verify mirrored VDisk copies:

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.

2. Select the VDisk to verify and then select **Verify VDisk Copies** from the task list. Click **Go**. The Verifying VDisk Copies panel is displayed.
3. Select the repair action if errors are found and click **OK**. You can also specify an LBA from which to start the verification. Start at different LBAs to count the number of differences on a VDisk.

Viewing the progress of VDisk copy verification

You can view the progress of verification of one or more mirror copies for a virtual disk (VDisk) from the Viewing Mirror Copy Verification Progress panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the progress of mirror copy verification:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click the **VDisk Copy Verification** link. The Viewing Mirror Copy Verification Progress panel is displayed.
3. Click **Close** to close the panel.

Repairing offline space-efficient VDisks

When a space-efficient virtual disk (VDisk) is taken offline because its metadata is corrupted, you can use the Repairing Space-Efficient VDisk panel to repair the metadata. The repair operation automatically detects corrupted metadata and performs any necessary repair actions.

This task assumes that you have already launched the SAN Volume Controller Console.

Use the Repairing Space-Efficient VDisk panel when directed through maintenance procedures. When the repair operation completes successfully, the error is automatically marked as fixed and the volume is brought back online. If the repair operation fails, an error is logged (error ID 060003) and the volume remains offline.

Once started, the VDisk remains offline for the duration of the repair, but you can move the VDisk to another I/O group.

Attention: You can only use this panel to repair a space-efficient VDisk that has reported corrupt metadata.

Perform the following steps to repair the offline space-efficient VDisk:

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.
2. Select the VDisk to repair and then select **Repair Space-efficient VDisk** from the task list. Click **Go**. The Repairing Space-Efficient VDisks panel is displayed.
3. Select the VDisk copy to repair and click **OK**.

Viewing the progress of space-efficient VDisk copy repair

You can view the progress of space-efficient virtual disk (VDisk) copy repair from the Viewing Space-Efficient Copy Repair Progress panel.

The time that is needed to complete a space-efficient VDisk copy repair depends on the amount of data that is currently on the copy. The repair process might complete very quickly.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the progress of space-efficient VDisk copy repair:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click the **Space-Efficient Copy Repair** link. The Viewing Space-Efficient Copy Repair Progress panel is displayed.
3. Click **Close** to close the panel.

Recovering from offline VDisks

You can use the SAN Volume Controller Console to recover from an offline virtual disk (VDisk) after a node or an I/O group has 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, you must perform one of the following procedures to regain access to your VDisks. Depending on the failure type, you might have lost data that was cached for these VDisks and the VDisks are now offline.

Data loss scenario 1

One node in an I/O group has failed and failover has started on the second node. During the failover process, the second node in the I/O group fails before the data in the write cache is written to hard disk. The first node is successfully repaired but its hardened data is not the most recent version that is committed to the data store; therefore, it cannot be used. The second node is repaired or replaced and has lost its hardened 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 when one node has earlier-level hardened data and the other node has lost hardened data:

1. Recover the node and add it back into the cluster.
2. Delete all FlashCopy, Metro Mirror, and Global Mirror mappings and relationships that use the offline VDisks.
3. Run the **recovervdisk**, **recovervdiskbyiogrp** or **recovervdiskbycluster** command.
4. Recreate all FlashCopy, Metro Mirror, and Global Mirror mappings and relationships that use the VDisks.

Data loss scenario 2

Both nodes in the I/O group have failed and have been repaired. The nodes have lost their hardened data, therefore, the nodes have no way of recognizing that they are part of the cluster.

Perform the following steps to recover from an offline VDisk when both nodes have lost their hardened data and cannot be recognized by the cluster:

1. Delete all FlashCopy, Metro Mirror, and Global Mirror mappings and relationships that use the offline VDisks.
2. Run the **recovervdisk**, **recovervdiskbyiogrp** or **recovervdiskbycluster** command.

3. Recreate all FlashCopy, Metro Mirror, and Global Mirror mappings and relationships that use the VDisks.

Recovering VDisks

Virtual disks (VDisks) or VDisk copies are corrupted if they have lost cached data or space-efficient metadata, usually as a result of hardware failure. A Fast Write State of Corrupt indicates this data loss.

You can recover one or more corrupted VDisks and VDisk copies. This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to recover VDisks and VDisk copies:

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.
2. Select the offline VDisks and select **Recover VDisk** from the task list and click **Go**. The Recovering VDisks panel is displayed.
3. Verify that the VDisks and VDisk copies have completed recovery by monitoring the VDisk Recovery Results panel.

Recovering VDisks by I/O group

Virtual disks (VDisks) or VDisk copies are corrupted if they have lost cached data or space-efficient metadata, usually as a result of hardware failure. A Fast Write State of Corrupt indicates this data loss.

You can recover all corrupted VDisks and VDisk copies in an I/O group. This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to recover VDisks and VDisk copies by I/O group:

1. Click **Work with Nodes** → **I/O Groups** in the portfolio. The Viewing I/O Groups panel is displayed.
2. Select the I/O Group and select **Recover VDisks by I/O Group** from the task list and click **Go**. The Recovering VDisks by I/O Group panel is displayed. Click **OK**.
3. Verify that the VDisks and VDisk copies have completed recovery by monitoring the VDisk Recovery Results panel.

Recovering VDisks by cluster

Virtual disks (VDisks) or VDisk copies are corrupted if they have lost cached data or space-efficient metadata, usually as a result of hardware failure. A Fast Write State of Corrupt indicates this data loss.

You can recover all corrupted VDisks and VDisk copies in a cluster. This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to recover VDisks and VDisk copies by cluster:

1. Click **Work with Virtual Disks** → **Virtual Disks** → **Viewing VDisks** in the portfolio. The Viewing Virtual Disks panel is displayed.
2. Select **Recover All VDisks** from the task list and click **Go**. The Recovering All VDisks panel is displayed. Click **OK**.
3. Verify that the VDisks and VDisk copies have completed recovery by monitoring the VDisk Recovery Results panel.

Moving offline VDisks to their original I/O group

After a node or an I/O group fails, you can move offline virtual disks (VDisks) to their original I/O group.

This task assumes that you have already launched the SAN Volume Controller Console.

Attention: Under no circumstances should VDisks be moved to an offline I/O group. Ensure that the I/O group is online before moving back the VDisks to avoid any further data loss.

Perform the following steps to move offline VDisks to their original I/O group:

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.
2. Select the offline VDisk and select **Modify a VDisk** from the task list and click **Go**. The Modifying Virtual Disk panel is displayed.
3. From the **I/O Group** list, select the name of the VDisk's original I/O group. You might be asked to confirm and force the move; select to force the move. Click **OK**. The Viewing Virtual Disks panel is displayed.
4. Verify that the VDisk is online.
5. Repeat these steps for each offline VDisk.

Deleting VDisks

You can delete a virtual disk (VDisk) from the Deleting Virtual Disk panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to delete a VDisk:

1. Click **Work with Virtual Disks** → **Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.
2. Select the VDisk you want to delete and select **Delete a VDisk** from the list. Click **Go**. The Deleting Virtual Disk panel is displayed.
3. Click **OK**.

Using image mode VDisks

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

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

Image mode MDisks are members of an MDisk group but, they do not contribute to free extents. Image mode VDisks are not affected by the state of the MDisk group because the MDisk group controls image mode VDisks through the VDisks association to an MDisk. Therefore, if an MDisk that is associated with an image mode VDisk is online and the MDisk group of which they are members goes offline, the image mode VDisk remains online. Conversely, the state of an MDisk group is not affected by the state of the image mode VDisks in the group.

An image mode VDisk behaves just as a managed mode VDisk in terms of the Metro Mirror, Global Mirror, and FlashCopy Copy Services. 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.

Creating an image mode VDisk

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

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

- 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 MDisks. The newly detected disk is displayed.
- Do *not* manually add an unmanaged-mode MDisk that contains existing data to an MDisk group. If you do, the data is lost. When you use the command to convert an image mode VDisk from an unmanaged-mode disk, select the MDisk group where you want to add the VDisk.

See the following Web site for more information:

www.ibm.com/storage/support/2145

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to create an image mode VDisk:

1. Stop all I/O operations from the hosts.
2. Unmap the logical disks that contain the data from the hosts.
3. Perform the following steps to create one or more MDisk groups:
   a. Click **Work with Managed Disks → Managed Disk Groups** in the portfolio. The Viewing Managed Disk Groups panel is displayed.
   b. Select **Create an MDisk Group** from the task list and click **Go**. The Create Managed Disk Group wizard begins.
   c. Use the wizard to create the MDisk group.
4. Perform the following steps to convert the unmanaged-mode MDisk to an image mode VDisk:
   a. Click **Work with Managed Disks → Managed Disks** in the portfolio. The Viewing Managed Disks panel is displayed.

      If the new unmanaged-mode MDisk is not listed, you can perform a fabric-level discovery. Select **Discover MDisks** from the task list and click **Go**. When this process is complete, refresh the list of MDisks, and the unmanaged-mode MDisk should appear in the list.

b. Select the unmanaged-mode MDisk and select **Create VDisk in Image Mode** from the task list. Click **Go**. The Create Image mode Virtual Disk wizard begins.

   Note: If the VDisk that you are creating maps to a solid-state drive (SSD), the data that is stored on the VDisk is not protected against SSD failures or node failures. To avoid data loss, add a VDisk copy that maps to an SSD on another node.

c. Use the wizard to select the MDisk group where the image mode VDisk should be added and the I/O group that will provide the data path for the VDisk.

5. Perform the following steps to map the new VDisk to the hosts that were previously using the data that the MDisk now contains:
   a. Click **Work with Virtual Disks → Virtual Disks** in the portfolio. The Viewing Virtual Disks panel is displayed.
   b. Select the VDisks and select **Map VDisks to a host** from the task list. Click **Go**. The Creating Virtual Disk-to-Host Mappings panel is displayed.
   c. Select the host that you want to map the VDisk to and click **OK**.

After the image mode VDisk is mapped to a host object, it 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 VDisk. Migrate the data on the image mode VDisk to managed-mode disks in another MDisk group.

Migration methods

Several methods can be used to migrate image mode virtual disks (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. The VDisk is automatically converted into a managed mode disk whenever any kind of migration activity is attempted. After the image mode to managed mode migration operation has occurred, the VDisk becomes a managed mode VDisk and is treated the same way as any other managed mode VDisk.

If the image mode disk has a partial last extent, this last extent in the image mode VDisk must be the first to be migrated. This migration is processed as a special case. After 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 changed into a managed mode VDisk and is treated the same way as any other managed mode VDisk.

An image mode disk can also be migrated to another image mode disk. The image mode disk becomes managed while the migration is ongoing, but returns to image mode when the migration is complete.

You can perform the following types of migrations:

- Migrate extents
- Migrate a VDisk
- Migrate to image mode

Note: Migration commands fail if the target or source VDisk is offline, or if there is insufficient quorum disk space to store the metadata. Correct the offline or quorum disk condition and reissue the command.

Perform the following steps to migrate VDisks:

1. Dedicate one MDisk group to image mode VDisks.
2. Dedicate one MDisk group to managed mode VDisks.
3. Use the migrate VDisk function to move the VDisks.

Viewing the progress of image mode migration

You can view the progress of image mode migration from the Viewing Image Mode Migration Progress panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the progress of image mode migration:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click the **Image Mode Migration** link. The Viewing Image Mode Migration Progress panel is displayed.

Viewing the progress of extent migration

You can view the progress of image mode migration from the Viewing Extent Migration Progress panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the progress of extent migration:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click the **Extent Migration** link. The Viewing Extent Migration Progress panel is displayed.

Creating hosts

You can create a new host object from the Creating Hosts panel.

Before you create a host using the SAN Volume Controller Console, ensure that you know the connection type you are using to attach to the host. Depending on whether the host attaches to the SAN Volume Controller cluster with fibre channel connections or with iSCSI connections, the information that you need differs.

This task assumes that you have already launched the SAN Volume Controller Console.

To create a new host object, complete the following steps:

1. Click **Work with Hosts** → **Hosts** in the portfolio. The Viewing Hosts panel is displayed.
2. Select **Create a Host** from the task list and click **Go**. The Creating Hosts panel is displayed.
3. Type the name for the host in the **Host Name** field. If you do not specify a name, a default name is assigned.

4. Select the type of host from the **Type** list.
5. Select the I/O groups to map to this host from the **I/O Groups** list.
6. Select either Fibre Channel or iSCSI for the connection type.

   If you select Fibre Channel as the connection type, complete the following steps:

   a. In the Fibre Channel Port Mask field, enter a valid port login mask for the host. For each login between a host HBA port and node port, the node examines the port mask that is associated with the host object to determine if access is allowed or denied. If access is denied, the node responds to SCSI commands as if the HBA port is unknown. Valid mask values range from 0000 (no ports enabled) to 1111 (all ports enabled). For example, a mask of 0011 enables port 1 and port 2. The default value is 1111 (all ports enabled). You can use a port mask to control the node target ports that a host can access.
   b. Assign a worldwide port name (WWPN). A WWPN consists of 16 hexadecimal digits (for example, 210100e08b251dd4). You can either 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 WWPNs to a single logical host object. Click Add to include the WWPN in the list of Selected Ports.

   If you select iSCSI as the connection type, complete the following steps:

   a. In the iSCSI Names field, enter the IQN name or EUI name for the iSCSI-attached host.
   b. In the CHAP Authentication Secret, enter the shared passphrase which is used to authenticate the host to the cluster.
7. Click **OK**.
8. Repeat steps 2 on page 174 through 7 for each host object to create.

After you have created the host object on the cluster, you can map virtual disks (VDisks) to host.

If you are unable to discover the disk on the host system or if there are fewer paths available for each disk than expected, test the connectivity between your host system and the cluster. Depending on the connection type to the host, these steps might be different. For iSCSI-attached hosts, test your connectivity between the host and SAN Volume Controller ports by pinging SAN Volume Controller from the host. Ensure that the firewall and router settings are configured correctly and validate that the values for the subnet mask and gateway are specified correctly for the SAN Volume Controller host configuration.

For fibre-channel attached hosts, ensure that the active switch configuration includes the host zone and check the host port link status. To verify end-to-end connectivity, you can use the svcinfo lsfabric CLI command or the View Fabric panel under Service and Maintenance container in the SAN Volume Controller Console.

Viewing host details

You can view details about a host object from the Viewing General Details panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view details for a host object:

1. Click **Work with Hosts → Hosts** in the portfolio. The Viewing Hosts panel is displayed.
2. Click the name of the host for which you want to view details. The Viewing General Details panel is displayed.
3. Click **Close** to return to the Viewing Hosts panel.

Viewing port details for hosts

You can view the ports that are attached to a host object from the Viewing Port Details panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the ports for a host object:

1. Click **Work with Hosts → Hosts** in the portfolio. The Viewing Hosts panel is displayed.
2. Click the name of the host for which you want to view port details. The Viewing General Details panel is displayed.
3. Click **Ports** to view the ports that are attached to the host object. The Viewing Port Details panel is displayed.
4. Click **Close** to return to the Viewing Hosts panel.

Viewing mapped I/O groups

You can view the I/O groups that are mapped to a host object from the Viewing Mapped I/O Groups panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the I/O groups that are mapped to a host object:

1. Click **Work with Hosts → Hosts** in the portfolio. The Viewing Hosts panel is displayed.
2. Click the name of the host for which you want to view the mapped I/O groups. The Viewing General Details panel is displayed.
3. Click **Mapped I/O Groups** to view the I/O groups that are mapped to the host object. The Viewing Mapped I/O Groups panel is displayed.
4. Click **Close** to return to the Viewing Hosts panel.

Displaying VDisks that are mapped to a host

You can display the virtual disks (VDisks) that are mapped to a host by using the Viewing Virtual Disks panel.

This task assumes that you have already launched the SAN Volume Controller Console.

If a large number of new VDisks are mapped to a host and a large number of devices are already running I/O operations, a significant number of errors might be logged. When the new VDisk is mapped, multiple recoverable errors can be logged in the event log. The event log displays the errors that are caused by a check condition. The errors state that there has been a change to the device information since the last logical unit number (LUN) operation.

Perform the following steps to show the VDisks that are mapped to a host:

1. Click **Work with Hosts** → **Hosts** in the portfolio. The Viewing Hosts panel is displayed.
2. Select the host and select **Show the VDisks Mapped to this Host** from the task list. Click **Go**.

Modifying a host

You can modify a host from the Modifying Host panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to modify a host:

1. Click **Work with Hosts** → **Hosts** in the portfolio. The Viewing Hosts panel is displayed.
2. Select the host that you want to modify and select **Modify a Host** from the task list. Click **Go**. The Modifying Host panel is displayed.

   You can modify the following attributes for a host:

   - Name
   - Type
   - I/O group
3. Select either Fibre Channel or iSCSI for the connection type.

   If you select Fibre Channel as the connection type, complete the following steps:

   a. In the Fibre Channel Port Mask field, enter a valid port login mask for the host. For each login between a host HBA port and node port, the node examines the port mask that is associated with the host object to determines if access is allowed or denied. If access is denied, the node responds to SCSI commands as if the HBA port is unknown. Valid mask values range from 0000 (no ports enabled) to 1111 (all ports enabled). For example, a mask of 0011 enables port 1 and port 2. The default value is 1111 (all ports enabled). You can use a port mask to control the node target ports that a host can access.

   If you select iSCSI as the connection type, complete the following steps:

   a. In the CHAP Authentication Secret, enter the shared passphrase which is used to authenticate the cluster to the host.
4. Click **OK** after you have selected the new attributes. If you are modifying a host-to-I/O group mapping that results in the loss of a VDisk-to-host mapping, the Forcing the Deletion of a Host to I/O Group Mappings panel is displayed. Perform one of the following steps:

   - Click **Force Remove** to remove the host-to-I/O group mapping.
   - Click **Cancel** to preserve the host-to-I/O group mapping.

Adding ports to a host

You can add ports to a host from the Adding Ports panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to add ports to a host:

1. Click **Work with Hosts → Hosts** in the portfolio. The Viewing Hosts panel is displayed.
2. Select the host that you want to add ports to and select **Add Ports** from the task list. Click **Go**. The Adding Ports panel is displayed.
3. Select either Fibre Channel or iSCSI for the connection type.

   If you select Fibre Channel as the connection type, complete the following steps:

   a. Assign a worldwide port name (WWPN). A WWPN consists of 16 hexadecimal digits (for example, 210100e08b251dd4). You can either 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 WWPNs to a single logical host object. Click Add to include the WWPN in the list of Selected Ports.

   If you select iSCSI as the connection type, complete the following steps:

   a. In the iSCSI Names field, enter the IQN name or EUI name for the iSCSI-attached host.
4. Click **OK**.

Deleting ports from a host

You can delete ports from the Deleting Ports panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to delete ports from a host:

1. Click **Work with Hosts → Hosts** in the portfolio. The Viewing Hosts panel is displayed.
2. Select the host that you want to delete ports from and select **Delete Ports** from the task list. Click **Go**. The Deleting Ports panel is displayed.
3. Select either Fibre Channel or iSCSI for the connection type.

   If you select Fibre Channel as the connection type, complete the following steps:

   a. Select a worldwide port name (WWPN) from the **Available Ports** list that your want to delete and click **Add** to include the WWPN in the **Selected Ports** list.

   If you select iSCSI as the connection type, complete the following steps:

   a. Select the iSCSI name from the **Available iSCSI Names** list that you want to delete, and click **Add** to include the iSCSI name in the **Selected iSCSI Initiator** list.
4. Click **OK**. The Deleting Ports--Confirmation panel displays. Note: When you delete the port, all virtual disks (VDisks) no longer use the port. If this is the last port for a host, the host is also automatically deleted. If this is the case, select Force Delete on the Confirmation panel to delete the port and its associated host.

Replacing an HBA in a host

It is sometimes necessary to replace the host bus adapter (HBA) that connects the host to the SAN. You must notify the SAN Volume Controller cluster of the new worldwide port name (WWPN) that this HBA contains.

Before you begin this task, you must ensure that the switch is zoned correctly.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to notify the SAN Volume Controller cluster 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.
2. Click **Work with Hosts** → **Hosts** in the portfolio. The Viewing Hosts panel is displayed.
3. Select the host object and then select **Add Ports** from the task list. Click **Go**. The Adding ports panel is displayed.
4. Select the candidate WWPNs from the **Available Ports** list and click **Add**. Click **OK**. The Viewing Hosts panel is displayed.
5. Select the host object and select **Delete Ports** from the task list. Click **Go**. The Deleting Ports panel is displayed.
6. Select the WWPNs that you want to remove (the ones that correspond with the old HBA that was replaced) and click **Add**. Click **OK**.

Any mappings that exist between the host object and VDisks are automatically applied to the new WWPNs. Therefore, the host sees the VDisks as the same SCSI LUNs as before. See the *IBM System Storage Multipath Subsystem Device Driver User's Guide* or your multipathing device driver user's guide for adding device identifiers (virtual paths if you are using SDD) to existing device identifiers.

Deleting hosts

You can delete a host object from the Deleting Hosts panel.

A deletion fails if there are any virtual disk (VDisk)-to-host mappings for the host. If you attempt to delete the host and it fails due to the existence of VDisk mappings, you are presented with the opportunity to perform a forced deletion, which deletes the VDisk mappings before the host is deleted.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to delete a host object:

1. Click **Work with Hosts** → **Hosts** in the portfolio. The Viewing Hosts panel is displayed.
2. Select the host to delete and select **Delete a host** from the task list. Click **Go**. The Deleting Hosts panel is displayed.
3. Verify that you are deleting the correct host and click **OK**.

When you delete a host object, all active ports are added to the **Available Ports** list.

Viewing fabrics

You can view the fabrics that are associated with a cluster from the Viewing Fabrics panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the fabrics:

1. Click **Service and Maintenance** → **Fabrics**. The Viewing Fabrics panel is displayed.
2. Click **Close** to close the panel.

Creating FlashCopy mappings

You can create a FlashCopy mapping using the Create a FlashCopy Mapping wizard.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to create FlashCopy mappings:

1. Click **Manage Copy Services** → **FlashCopy mappings** in the portfolio. The Viewing FlashCopy Mappings panel is displayed.
2. Select **Create a Mapping** from the task list and click **Go**. The Create a FlashCopy Mapping wizard begins.
3. Complete the Create a FlashCopy Mapping wizard.

Starting FlashCopy mappings

You can use the SAN Volume Controller Console to start FlashCopy mappings.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to start a FlashCopy mapping:

1. Click **Manage Copy Services** → **FlashCopy Mappings** in the portfolio. The Viewing FlashCopy Mappings panel is displayed.
2. Select the appropriate mapping's row from the table.
3. Select **Start a Mapping** from the task list and click **Go**. The Starting FlashCopy Mapping panel is displayed.
4. Click **Start**. The Viewing FlashCopy Mappings panel is displayed.

Before a FlashCopy mapping is started it must be in the Prepared state. Preparing the FlashCopy mapping flushes the cache of any data that is destined for the source VDisk and forces the cache into the write-through mode until the mapping is started. If you have selected a FlashCopy mapping to start and it is not yet in the prepared state, the SAN Volume Controller Console displays the Starting FlashCopy Mapping–Prepare and Start panel.

In addition, the SAN Volume Controller cluster detects if the selected FlashCopy mapping is for restoring data only. This panel indicates that the selected FlashCopy mapping has a target virtual disk (VDisk) that is a source VDisk in another active FlashCopy mapping.

Starting a FlashCopy mapping–Prepare and start

You can use the Starting a FlashCopy Mapping–Prepare and Start panel in the SAN Volume Controller Console to prepare and start the selected FlashCopy mapping.

You need to do this first.

Before a FlashCopy mapping is started it must be in the Prepared state. Preparing the FlashCopy mapping flushes the cache of any data that is destined for the source VDisk and forces the cache into the write-through mode until the mapping is started. If you have selected a FlashCopy mapping to start and it is not yet in the prepared state, the SAN Volume Controller Console displays this panel with the option to prepare and start the FlashCopy mapping.

Complete the following steps to start a FlashCopy mapping:

1. Click **Manage Copy Services** → **FlashCopy Mappings** in the portfolio. The Viewing FlashCopy Mappings panel is displayed.
2. Select the FlashCopy mapping that you want to start and select **Start a Mapping** from the task list, and click **Go**. If the selected FlashCopy mapping has not been prepared, the Starting FlashCopy Mapping–Prepare and Start panel displays.
3. Click **Prepare and Start**. The Viewing FlashCopy Mappings panel is displayed.

In addition, the SAN Volume Controller cluster detects if the selected FlashCopy mapping is for restoring data only. If the FlashCopy mapping is for restoring data only, then the Start FlashCopy Mapping–Restore panel displays. This panel indicates that the selected FlashCopy mapping has a target virtual disk (VDisk) that is a source VDisk in another active FlashCopy mapping.

Starting a FlashCopy mapping–Restore

You can use the Starting FlashCopy Mapping–Restore panel in the SAN Volume Controller Console to start a selected FlashCopy mapping for restoring data. This panel indicates that the selected FlashCopy mapping has a target virtual disk (VDisk) that is a source VDisk in another active FlashCopy mapping.

This task assumes that you have already launched the SAN Volume Controller Console.

When the FlashCopy mapping is being started with restore option, the start operation can fail if dependencies exist between the current FlashCopy mapping being started and other FlashCopy mappings on the cluster that have also been started with the restore option. To determine if dependencies exist between the selected FlashCopy mappings and other mappings, use the Viewing FlashCopy Mappings panel to determine whether other FlashCopy mappings are linked to the selected mapping through shared source or target VDisks. If the restore option is set on any of these linked mappings, then you cannot start the selected FlashCopy mapping with the restore option.

Complete the following steps to start a FlashCopy mapping:

.

1. Click **Manage Copy Services** → **FlashCopy Mappings** in the portfolio. The Viewing FlashCopy Mappings panel is displayed.
2. Select the FlashCopy mapping that you want to start and select **Start a Mapping** from the task list, and click **Go**. If the selected FlashCopy mapping has been flagged for restoring, the Starting FlashCopy Mapping–Restore panel displays.
3. Click **Continue for Restoring**. The Viewing FlashCopy Mappings panel is displayed.

Preparing a FlashCopy mapping–Restore

You can use the Preparing FlashCopy Mapping–Restore panel in the SAN Volume Controller Console to prepare the FlashCopy mapping for a restore operation.

During the prepare operation for a FlashCopy mapping, the cluster verifies that the target virtual disk (VDisk) is also not a source VDisk in another active FlashCopy mapping. If no dependencies for the selected mapping exist, then the mapping is prepared and can be started. If a dependency exists, you can prepare the mapping to restore data only. The restore option allows you to recover from application failures without changing existing backup configurations or without losing target data. To determine if dependencies exist between the selected FlashCopy mappings and other mappings, use the Viewing FlashCopy Mappings panel to determine whether other FlashCopy mappings are linked to the selected mapping through shared source or target VDisks. If the restore option is set on any of these linked mappings, then you cannot prepare the selected FlashCopy mapping with the restore option.

Complete the following steps to prepare a FlashCopy mapping:

1. Click **Manage Copy Services** → **FlashCopy Mappings** in the portfolio. The Viewing FlashCopy Mappings panel is displayed.
2. Select the FlashCopy mapping that you want to start and select **Prepare a Mapping** from the task list, and click **Go**. If the selected FlashCopy mapping has target VDisk that are also source VDisk in another active FlashCopy mapping, then the Prepare FlashCopy Mapping–Restore panel displays.
3. Click **Prepare for Restoring**. The Viewing FlashCopy Mappings panel is displayed.

Viewing the progress of a FlashCopy

You can view the progress of a FlashCopy from the Viewing FlashCopy Progress panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to view the progress of a FlashCopy:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click the **FlashCopy** link. The Viewing FlashCopy Progress panel is displayed.

Stopping FlashCopy mappings

You can use the SAN Volume Controller Console to stop FlashCopy mappings.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to stop a FlashCopy mapping:

1. Click **Manage Copy Services** → **FlashCopy Mappings** in the portfolio. The Viewing FlashCopy Mappings panel is displayed.
2. Select the appropriate mapping's row from the table.
3. Select **Stop a Mapping** from the task list and click **Go**. The Stopping FlashCopy Mapping panel is displayed.
4. Click **Stop**, **Split Stop, Forced Stop**, or **Cancel**. The Viewing FlashCopy Mappings panel is displayed.

Modifying FlashCopy mappings

You can use the SAN Volume Controller Console to change the attributes for a FlashCopy mapping.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to change the attributes for a FlashCopy mapping:

1. Click **Manage Copy Services** → **FlashCopy Mappings** in the portfolio. The Viewing FlashCopy Mapping panel is displayed.
2. Select **Modify a Mapping** from the task list and click **Go**. The Modifying FlashCopy Mapping panel is displayed.
3. Specify the modifications and click **OK**or **Cancel**. The Viewing FlashCopy Mappings panel is displayed.

Deleting FlashCopy mappings

You can delete a FlashCopy mapping from the Deleting FlashCopy Mappings panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to delete a FlashCopy mapping:

1. Click **Manage Copy Services** → **FlashCopy mappings** in the portfolio. The Viewing FlashCopy mappings panel is displayed.
2. Select the appropriate mapping's row from the table.
3. Select **Delete a Mapping** from the task list and click **Go**. The Deleting FlashCopy Mapping panel is displayed.

   Note: If the FlashCopy mapping is in active state, the Forcing the Deletion of a FlashCopy Mapping panel is displayed. Follow the instructions that are displayed on the Forcing the Deletion of a FlashCopy Mapping panel.

Creating FlashCopy consistency groups

You can use the SAN Volume Controller Console to create FlashCopy consistency groups.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to create a FlashCopy consistency group:

1. Click **Manage Copy Services** → **FlashCopy Consistency Groups** in the portfolio. The Viewing FlashCopy Consistency Groups panel is displayed.
2. Select **Create a Consistency Group** from the task list and click **Go**. The Creating FlashCopy Consistency Groups panel is displayed.
3. Type the name of the FlashCopy consistency group in the **FlashCopy Consistency Group Name** field. If you do not specify a name, a default name is assigned to the FlashCopy consistency group.

4. Optionally, check the **Automatically delete consistency group when empty** box. Checking this box automatically deletes the consistency group when the last FlashCopy mapping in the group is deleted or removed from the consistency group.
5. Select the mappings in the consistency group from the **FlashCopy Mappings** list and click **OK**.

   Note: You can create the FlashCopy consistency group before you create the mappings and then add the FlashCopy mappings to the consistency group. To add FlashCopy mappings this way, you must use the Modifying FlashCopy Mapping panel.

Starting FlashCopy consistency groups

You can start a FlashCopy consistency group from the Starting FlashCopy Consistency Group panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to start a FlashCopy consistency group:

1. Click **Manage Copy Services** → **FlashCopy Consistency Groups** in the portfolio. The Viewing FlashCopy Consistency Groups panel is displayed.
2. Select the appropriate group's row from the table.
3. Select **Start a Consistency Group** from the task list and click **Go**. The Starting FlashCopy Consistency Groups panel is displayed.
4. Click **Start** or **Cancel**. The Viewing FlashCopy Consistency Groups panel is displayed.

If you have selected a FlashCopy consistency group to start and it contains mappings that are not yet in the prepared state, the SAN Volume Controller Console displays the Starting FlashCopy Mapping–Prepare and Start panel. Before a FlashCopy mapping is started it must be in the Prepared state. Preparing the FlashCopy mapping flushes the cache of any data that is destined for the source VDisk and forces the cache into the write-through mode until the mapping is started.

If you have selected a FlashCopy consistency group to start and it has mappings where target virtual disks (VDisks) are also source VDisks in other active FlashCopy mappings, then the Starting FlashCopy Consistency Group–Restore panel displays.

Starting a FlashCopy consistency group–Prepare and start

You can use the Starting FlashCopy Consistency Group panel in the SAN Volume Controller Console to prepare and start the FlashCopy consistency group. Before FlashCopy mappings in the selected FlashCopy consistency can start, all the mappings must first be prepared. When the FlashCopy mappings are prepared, any data that resides in the cache for the source virtual disks (VDisks) is transferred to disks thus ensuring that the copy operation includes all data currently on the source VDisks. After the preparation action completes for the FlashCopy mappings in the consistency group, the system starts the FlashCopy consistency group.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to start a FlashCopy consistency group:

1. Click **Manage Copy Services** → **FlashCopy Consistency Groups** in the portfolio. The Viewing FlashCopy Consistency Groups panel is displayed.
2. Select the appropriate group's row from the table.
3. Select **Start a Consistency Group** from the task list and click **Go**. If the selected FlashCopy consistency group contains mappings that are not in the prepared state then the Starting FlashCopy Consistency Groups–Prepare and Start panel displays.
4. Click **Prepare and Start**.

Starting a FlashCopy consistency group–Restore

You can use the Starting FlashCopy Consistency Group–Restore panel in the SAN Volume Controller Console to start a selected FlashCopy consistency group for restoring data.

This panel indicates that the selected FlashCopy consistency group has mappings where target virtual disks (VDisks) are also source VDisks in other active FlashCopy mappings. When the FlashCopy consistency group is being started with restore option, the start operation can fail if dependencies exist between the any of the FlashCopy mapping in the consistency group and other FlashCopy mappings on the cluster that have also been started with the restore option.

To determine if dependencies exist between the selected FlashCopy mappings and other mappings, use the Viewing FlashCopy Mappings panel to determine whether other FlashCopy mappings are linked to the selected mapping through shared source or target VDisks. If the restore option is set on any of these linked mappings, then you cannot start the selected FlashCopy mapping with the restore option.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to start a FlashCopy consistency group:

1. Click **Manage Copy Services** → **FlashCopy Consistency Groups** in the portfolio. The Viewing FlashCopy Consistency Groups panel is displayed.
2. Select the appropriate group's row from the table.
3. Select **Start a Consistency Group** from the task list and click **Go**. If the selected FlashCopy consistency group contains has mappings where target VDisks are also source VDisks in other active FlashCopy mappings, then the Starting FlashCopy Consistency Groups–Restore panel displays.
4. Click **Continue for Restoring**.

Preparing a FlashCopy consistency group–Restore

Use the Preparing a FlashCopy Consistency Group–Restore in the SAN Volume Controller Console to prepare all the FlashCopy mappings within a consistency group.

This panel indicates that the target virtual disks (VDisks) of the mappings in this group are the source VDisks of some other active FlashCopy mappings. When the FlashCopy consistency group is being prepared with restore option, the prepare operation can fail if dependencies exist between the any of the FlashCopy mapping

in the consistency group and other FlashCopy mappings on the cluster that have also been prepared or started with the restore option.

To determine if dependencies exist between the selected FlashCopy mappings and other mappings, use the Viewing FlashCopy Mappings panel to determine whether other FlashCopy mappings are linked to the selected mapping through shared source or target VDisks. If the restore option is set on any of these linked mappings, then you cannot prepare the selected FlashCopy mapping with the restore option.

Complete the following steps to start a FlashCopy consistency group:

1. Click **Manage Copy Services** → **FlashCopy Consistency Groups** in the portfolio. The Viewing FlashCopy Consistency Groups panel is displayed.
2. Select the appropriate group's row from the table.
3. Select **Prepare a Consistency Group** from the task list and click **Go**. If the selected FlashCopy consistency group contains has mappings where target VDisks are also source VDisks in other active FlashCopy mappings, then the Prepare FlashCopy Consistency Groups–Restore panel displays.
4. Click **Prepare for Restoring**.

Stopping FlashCopy consistency groups

You can stop a FlashCopy consistency group from the Stopping FlashCopy Consistency Group panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to stop a FlashCopy consistency group:

1. Click **Manage Copy Services** → **FlashCopy Consistency Groups** in the portfolio. The Viewing FlashCopy Consistency Groups panel is displayed.
2. Select the appropriate consistency group's row from the table.
3. Select **Stop a Consistency Group** from the task list and click **Go**. The Stopping FlashCopy Consistency Group panel is displayed.
4. Click **Stop**, **Split Stop,Forced Stop**, or **Cancel**. The Viewing FlashCopy Consistency Groups panel is displayed.

Renaming FlashCopy consistency groups

You can rename a FlashCopy consistency group from the Modifying FlashCopy Consistency Group panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to rename a consistency group:

1. Click **Manage Copy Services** → **FlashCopy Consistency Groups** in the portfolio. The Viewing FlashCopy Consistency Groups panel is displayed.
2. Select the appropriate group row from the table.
3. Select **Rename a Consistency Group** from the task list and click **Go**. The Modifying FlashCopy Consistency Group panel is displayed.
4. Enter the new name and click **OK**. The Viewing FlashCopy Consistency Groups panel is displayed.

Deleting FlashCopy consistency groups

You can delete a FlashCopy consistency group from the Deleting FlashCopy consistency groups panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to delete a FlashCopy consistency groups:

1. Click **Manage Copy Services** → **FlashCopy Consistency Groups** in the portfolio. The Viewing FlashCopy Consistency Groups panel is displayed.
2. Select the appropriate consistency group's row from the table.
3. Select **Delete a Consistency Group** from the task list and click **Go**. The Deleting FlashCopy Consistency Group panel is displayed.
4. Click **OK** to delete. The Viewing FlashCopy Consistency Groups panel is displayed.

Creating Metro Mirror and Global Mirror relationships

You can use the SAN Volume Controller Console to create Metro Mirror and Global Mirror relationships. The maximum number of Metro Mirror and Global Mirror relationship is 8192 per cluster.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to create a Metro Mirror or Global Mirror relationship:

1. Click **Manage Copy Services** → **Metro & Global Mirror Relationships** in the portfolio. The Viewing Metro & Global Mirror Relationships panel is displayed.
2. Select **Create a Relationship** from the list and click **Go**. The Create a Metro or Global Mirror Relationship wizard begins.
3. Complete the Create a Metro or Global Mirror Relationship wizard.

Starting a Metro Mirror or Global Mirror copy process

You can use the SAN Volume Controller Console to start a Metro Mirror or Global Mirror copy process.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to start a Metro Mirror or Global Mirror copy process:

1. Click **Manage Copy Services** → **Metro & Global Mirror Relationships** in the portfolio. The Viewing Metro & Global Mirror Relationships panel is displayed.
2. Select the relationship for which you want to start the copy process.
3. Select **Start Copy Process** and click **Go**. The Starting Copy Process panel is displayed.

Viewing the progress of Metro Mirror and Global Mirror copy processes

You can use the SAN Volume Controller Console to view the progress of Metro Mirror and Global Mirror copy processes.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to view the progress of Metro Mirror and Global Mirror copy processes:

1. Click **Manage Progress** → **View Progress**. The View Progress panel is displayed.
2. Click **Metro & Global Mirror**. The Viewing Metro & Global Mirror Progress panel is displayed.

Stopping a Metro Mirror or Global Mirror copy process

You can use the SAN Volume Controller Console to stop a Metro Mirror or Global Mirror copy process.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to stop a Metro Mirror or Global Mirror copy process:

1. Click **Manage Copy Services** → **Metro & Global Mirror Relationships** in the portfolio. The Viewing Metro & Global Mirror Relationships panel is displayed.
2. Select the relationship for which you want to stop the copy process.
3. Select **Stop Copy Process** and click **Go**. The Stopping Copy Process panel is displayed.
4. Click **OK** to stop the copy process.

Modifying Metro Mirror and Global Mirror relationships

You can use the SAN Volume Controller Console to modify the attributes for Metro Mirror and Global Mirror relationships.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to modify the attributes for Metro Mirror and Global Mirror relationships:

1. Click **Manage Copy Services** → **Metro & Global Mirror Relationships** in the portfolio. The Viewing Metro & Global Mirror Relationships panel is displayed.
2. Select the relationship to modify.
3. Select **Modify a Relationship** from the task list and click **Go**. The Modifying Metro & Global Mirror Relationship panel is displayed.

   You can change the following attributes from this panel:

   - The relationship name
   - The consistency group that contains this relationship

Switching the copy direction of a Metro Mirror or Global Mirror relationship

You can use the SAN Volume Controller Console to reverse the roles of the primary and secondary virtual disks (VDisks) in a Metro Mirror or Global Mirror relationship.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to reverse the roles of the primary and secondary VDisks:

1. Click **Manage Copy Services** → **Metro & Global Mirror Relationships** in the portfolio. The Viewing Metro & Global Mirror Relationships panel is displayed.
2. Select **Switch Copy Direction** from the task list and click **Go**. The Switching the Direction of Mirror Relationship panel is displayed.

Deleting Metro Mirror or Global Mirror relationships

You can use SAN Volume Controller Console the to delete a Metro Mirror or Global Mirror relationship

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to delete a Metro Mirror or Global Mirror relationship:

1. Click **Manage Copy Services** → **Metro & Global Mirror Relationships** in the portfolio. The Viewing Metro & Global Mirror Relationships panel is displayed.
2. Select the relationship to delete by clicking on the appropriate line in the **Select** column.
3. Select **Delete a Relationship** from the task list and click **Go**. The Deleting Mirror Relationship panel is displayed.
4. Click **OK** to delete the relationship.

Creating Metro Mirror or Global Mirror consistency groups

You can create Metro Mirror or Global Mirror consistency groups using the wizard.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to create a Metro Mirror or Global Mirror consistency group:

1. Click **Manage Copy Services** → **Metro & Global Mirror Consistency Groups** in the portfolio.
2. Select **Create a Consistency Group** from the task list and click **Go**. The wizard begins.
3. Complete the wizard.

Renaming a Metro Mirror or Global Mirror consistency group

You can use the SAN Volume Controller Console to rename a Metro Mirror or Global Mirror consistency group.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to rename a Metro Mirror or Global Mirror consistency group:

1. Click **Manage Copy Services** → **Metro & Global Mirror Consistency Groups** in the portfolio.
2. Select the consistency group that you want to change.
3. Select **Rename a Consistency Group** from the task list and click **Go**. The Renaming Mirror Consistency Group panel is displayed.

4. Type a new name for the consistency group in the **New Name** field.
5. Click **OK**.

Starting a Metro Mirror or Global Mirror consistency group copy

You can use the SAN Volume Controller Console to start a Metro Mirror or Global Mirror copy process.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to start a Metro Mirror or Global Mirror copy process:

1. Click **Manage Copy Services** → **Metro & Global Mirror Consistency Groups** in the portfolio
2. Select the relationship for which you want to start the copy process.
3. Select **Start Copy Process** and click **Go**. The Starting Copy Process panel is displayed.

Stopping a Metro Mirror or Global Mirror consistency group copy process

You can use the SAN Volume Controller Console to stop a Metro Mirror or Global Mirror copy process.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to stop a Metro Mirror or Global Mirror copy process:

1. Click **Manage Copy Services** → **Metro & Global Mirror Consistency Groups** in the portfolio.
2. Select the group for which you want to stop the copy process.
3. Select **Stop Copy Process** and click **Go**. The Stopping Copy Process panel is displayed.
4. Follow the directions that are displayed on this panel.

Deleting Metro Mirror and Global Mirror consistency groups

You can use the SAN Volume Controller Console to delete Metro Mirror and Global Mirror consistency groups.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to delete a Metro Mirror or Global Mirror consistency group:

1. Click **Manage Copy Services** → **Metro & Global Mirror Consistency Groups** in the portfolio.
2. Select the group to delete.
3. Select **Delete a Consistency Group** from the task list and click **Go**.
4. Click **OK** to delete the consistency group.

Creating Metro Mirror and Global Mirror partnerships

You can use the SAN Volume Controller Console to create Metro Mirror and Global Mirror partnerships.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to create a Metro Mirror or Global Mirror partnership:

1. Click **Manage Copy Services** → **Metro & Global Mirror Cluster Partnership** in the portfolio.
2. Click **Create**. The Create Cluster Partnerships panel is displayed.
3. Follow the instructions that are displayed on this panel to create the cluster partnership.

Viewing Metro Mirror and Global Mirror cluster partnerships

Use the Viewing Metro and Global Mirror Cluster Partnerships panel in the SAN Volume Controller Console to view or create all Metro Mirror or Global Mirror cluster partnerships.

Introduction

Metro Mirror and Global Mirror partnerships define the relationship between a local cluster and a remote cluster. Each cluster can have a maximum of three partner clusters, with a maximum of four clusters connected directly or indirectly. By creating multiple partnerships between clusters, administrators can design topologies for disaster recovery and data migration.

For example, administrators can create multiple partnerships between four clusters: Cluster A to Cluster B partnership, Cluster A to Cluster C partnership, and Cluster A to Cluster D partnership. In this example, Cluster A might be the central disaster recovery site for the other three clusters. In the event of a disruption of service to Cluster B, C, or D, Cluster A can be used to recover important applications and to resume normal operations from a remote site. Another example topology for multiple cluster partnerships provides data migration and disaster recovery using three clusters. In this example, Cluster A and Cluster B are in a partnership; Cluster A and Cluster C are partners; and Cluster B and Cluster C are in a partnership. An administrator can use this topology to migrate a data center from Cluster B to Cluster C. If Cluster A hosts production, then Cluster B and C can also provide disaster recovery.

On this panel you can either view all the cluster partnerships that are defined on this cluster, or if cluster partnerships are not defined, you can click **Create** to open the Create Cluster Partnership panel.

This panel displays two sets of properties that are related to cluster partnerships. The first table contains properties that pertain to Global Mirror partnerships that are defined on the selected cluster. The second table displays properties for each of the clusters that are defined in partnerships with the selected cluster.

Attributes

The following properties pertain to all the Global Mirror partnership in the cluster:

Link Tolerance
Displays the length of time that the primary response time remains impacted before relationships are suspended. Valid values are 0 - 86400 seconds. A value of zero disables this feature. The default value is 300 seconds.

Inter-Cluster Delay Simulation
Displays the amount of time to delay I/O writes to the intercluster secondary VDisk. Valid values are 0 - 100 milliseconds. A value of zero disables this feature. The default value is zero.

Intra-Cluster Delay Simulation
Displays the amount of time to delay I/O writes to the intracluster secondary VDisk. Valid values are 0 - 100 milliseconds. A value of zero disables this feature. The default value is zero.

Attributes

The following attributes are displayed in the table if a cluster partnership exists:

Select Select the object to perform an action from the task list.

Note: Only one row can be selected at a time.

Cluster Name
Displays the name of the selected cluster.

IP address (IPv4)
Displays the IPv4 IP address of the local or remote partner cluster.

IP address (IPv6)
Displays the IPv6 IP address of the local or remote partner cluster.

State Displays the current state of the partnership. The following values are possible:

Partially Configured
Indicates that only one cluster partner is defined from a local or remote cluster to the displayed cluster and is started. For the displayed cluster to be configured fully and to complete the partnership, you must define the cluster partnership from the cluster that is displayed to the corresponding local or remote cluster. You can do this by issuing the **mkpartnership** command on the local and remote cluster that are in the partnership, or by using the SAN Volume Controller Console to create a partnership on both the local and remote clusters.

Fully Configured
Indicates that the partnership is defined on the local and remote clusters and is started.

Remote Not Present
Indicates that the remote cluster is not present to the partnership.

Partially Configured (Local Stopped)
Indicates that the local cluster is only defined to remote cluster and the local cluster is stopped.

Fully Configured (Local Stopped)
Indicates that a partnership is defined on both the local and remote clusters and the remote cluster is present, but the local cluster is stopped.

Fully Configured (Remote Stopped)
Indicates that a partnership is defined on both the local and remote clusters and the remote cluster is present, but the remote cluster is stopped.

Fully Configured (Local Excluded)
Indicates that a partnership is defined between a local and remote cluster; however, the local cluster has been excluded. Usually this state occurs when the fabric link between the two clusters has been compromised by too many fabric errors or slow response times of the cluster partnership. Check the error log for 1720 errors by selecting **Service and Maintenance → Analyze Error Log** to resolve these errors.

Fully Configured (Remote Excluded)
Indicates that a partnership is defined between a local and remote cluster; however, the remote cluster has been excluded. Usually this state occurs when the fabric link between the two clusters has been compromised by too many fabric errors or slow response times of the cluster partnership. Check the error log for 1720 errors by selecting **Service and Maintenance → Analyze Error Log** to resolve these errors.

Fully Configured (Remote Exceeded)
Indicates that a partnership is defined between a local and remote cluster and the remote is available; however, the remote cluster exceeds the number of allowed clusters within a cluster network. The maximum of four clusters can be defined in a network. If the number of clusters exceeds that limit, SAN Volume Controller determines the inactive cluster or clusters by sorting all the clusters by their unique identifier in numerical order. The inactive cluster partner which is not in the top four of the cluster unique identifiers displays **Fully Configured (Remote Exceeded)**.

Bandwidth (MBps)
Displays the bandwidth that was specified, shown in megabytes per second (MBps).

Actions

You can select one of the following tasks from the task list:

- **Create a Partnership**
- **Modify a Partnership**
- **Delete a Partnership**
- **Start a Partnership**
- **Stop a Partnership**

Note: With the exception of the **Create a Partnership** task, you must first select a partner cluster before you can select a task.

Go Click this button after selecting a task from the task list. Clicking **Go** launches the panel for the task that you selected.

Modify
Click this button to change the properties for Global Mirror partnerships.

Refresh
Click this button to refresh the table with changed values.

Modifying Global Mirror partnerships

You can change properties, such as link tolerance, for Global Mirror cluster partnerships

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to change a Global Mirror partnership:

1. Click **Manage Copy Services → Metro & Global Mirror Cluster Partnerships** in the portfolio. The Viewing Metro & Global Mirror Cluster Partnership panel displays.
2. Select the cluster partnership that you want to modify from the partnership list and select **Modify a Partnership** from the task list. Click **Go**. The Modifying Global Mirror Properties panel is displayed.
3. You can update the following properties for a Global Mirror cluster partnership:
   - Link Tolerance
   - Inter-Cluster Delay Simulation (ms)
   - Intra-Cluster Delay Simulation (ms)
4. Click **OK**.

Modifying Metro Mirror and Global Mirror partnership bandwidth

You can change the current setting for bandwidth, which is also known as background copy. The partnership bandwidth controls the rate at which data is sent from the local cluster to the remote cluster. The partnership bandwidth can be changed to help manage the use of inter-cluster links. It is measured in megabytes per second (MBps).

This task assumes that you have already launched the SAN Volume Controller Console.

To change the bandwidth setting for a selected Metro or Global Mirror partnership, complete the following steps:

1. In the portfolio, click **Manage Copy Services → Metro & Global Mirror Cluster Partnerships**. The Viewing Metro & Global Mirror Cluster Partnership panel displays.
2. From the partnership list, select the cluster partnership that you want to modify. Select **Modify a Partnership Bandwidth** from the task list. Click **Go**. The Modifying Metro & Global Mirror Bandwidth panel is displayed.
3. Type the new rate for the background copy.

   Note: You can set the bandwidth attribute for the path from cluster A to cluster B to a different setting from the setting that is used for the path from cluster B to cluster A.
4. Click **OK**.

Starting and stopping Metro Mirror or Global Mirror partnerships

You can start and stop Metro Mirror or Global Mirror partnerships using the Modify Cluster Partnership panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to start or stop a Metro Mirror or Global Mirror partnership:

1. Click **Manage Copy Services** → **Metro & Global Mirror Cluster Partnerships** in the portfolio.
2. Click **Modify**. The Modify Cluster Partnership panel is displayed.
3. To stop a partnership, click **Stop**.
4. To start a partnership, click **Start**.
5. Click **OK**.

Deleting Metro Mirror or Global Mirror partnerships

You can use the SAN Volume Controller Console to delete a Metro Mirror or Global Mirror partnership on the local cluster.

This task assumes that you have already launched the SAN Volume Controller Console.

The Metro Mirror or Global Mirror partnership must be deleted on both the local and remote cluster for the partnership to be completely removed.

Perform the following steps to delete a Metro Mirror or Global Mirror partnership on the local cluster.

Note: If a Metro Mirror or Global Mirror partnership has configured relationships or groups, you must delete the relationships and groups before you can delete the partnership.

1. Click **Manage Copy Services** → **Metro & Global Mirror Cluster Partnerships** in the portfolio.
2. Click **Delete**. The Delete Cluster Partnership panel is displayed.
3. Click **Delete** to delete the Partnership on the local cluster or click **Cancel** to return to the previous panel.

Viewing the license settings log

You can view the license settings log for the cluster from the License Settings panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following step to view the license settings log for the cluster:

Click **Service and Maintenance** → **View License Settings Log** in the portfolio. The License Settings panel is displayed.

Updating license settings

You can use the SAN Volume Controller Console to update your license settings.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to update the license settings:

1. In the portfolio, click **Service and Maintenance** → **License Settings**. The License Settings panel is displayed.
2. Choose Capacity Licensing or Physical Disk Licensing and click **Go**.
3. Enter your license settings and click **Update License Settings**.
4. The updated license information is displayed. To confirm that the settings match your license agreement, click **I Agree**.

Running the cluster maintenance procedure

You can use the SAN Volume Controller Console to run the cluster maintenance procedure.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to run the cluster maintenance procedure:

1. Click **Service and Maintenance** → **Run Maintenance Procedures** in the portfolio. The Maintenance Procedures panel is displayed.
2. Click **Start Analysis** to analyze the cluster error log. The Maintenance panel is displayed.

   If you click the error code of a error log entry, you are guided through a series of actions that help you estimate the state of the cluster and determine if the error was an isolated event or a component failure. 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 an unfixed error to a fixed error.

Configuring remote authentication

You can configure SAN Volume Controller to use a remote authentication service. Remote authentication allows users of SAN management applications, such as IBM Tivoli Storage Productivity Center, to authenticate to the cluster using the authentication service provided by the SAN management application.

Ensure that the remote authentication service is configured for the SAN management application. To complete this task, you should have the following information regarding the remote authentication service:

- Web Address for the remote authentication service.
- User name and password for HTTP basic authentication. These credentials are created by and obtained from the administrator of the remote authentication service.
- SSL certificate.

Note: An SSL certificate is required only if you are using a secure Web address for the remote authentication service.

This task assumes that you have already launched the SAN Volume Controller Console. To enable and configure remote authentication service for the cluster, follow these steps:

1. In the portfolio, click **Manage Authentication** → **Remote Authentication**. The Configuring Remote Authentication panel is displayed.
2. To enable remote authentication service, select **Enable**.

   Note: You can also disable remote authentication by deselecting **Enable**.
3. Enter the following attributes for the remote authentication service:

   Service Web Address (IPv4 or IPv6)
   Enter the Web address of the remote authentication service. SAN Volume Controller supports both IPv4 or IPv6 network addresses for the remote authentication service. You can use the following characters: a - z, A - Z, 0 - 9, _, ~, :, [,], %, or /. The maximum length of the Web address is 100 characters. The Web address can have either of the following formats:
   - `http://`*`network_address`*`:`*`http remote authentication service port number/path_to_service`*
   - `https://`*`network_address`*`:https `*`remote authentication service port number/path_to_service`*

     For example, if the system network IPv4 address is 9.71.45.108, you could enter either of the following corresponding addresses:

     ```
     http://9.71.45.108:16310/TokenService/services/Trust
     https://9.71.45.108:16311/TokenService/services/Trust
     ```

     Note:
     a. To obtain the correct remote authentication service port numbers and service path, consult the documentation for your remote authentication service software.
     b. To use a secure Web address, an SSL certificate in privacy enhanced mail (PEM) format is required.

   User Name
   Enter the HTTP basic authentication user name that is required to obtain service from the remote authentication server. The user name cannot start or end with a blank. The user name can consist of a string of 1 - 64 ASCII characters with the exception of the following characters: %:",*' .

   Password
   Enter the HTTP basic authentication password that is required to obtain service from the remote authentication server. The password cannot start or end with a blank. The password can consist of a string of 6 - 64 printable ASCII characters.

   Re-enter Password
   Re-enter the HTTP basic authentication password.

   SSL Certificate
   Enter the fully qualified name of the file that contains the SSL certificate in PEM format for the remote authentication service. The maximum file length for the SSL certificate is 2048 bytes. An SSL

certificate is required to authenticate to the remote authentication service when a secure Web address is configured.

4. Click **OK**.

Viewing remote authentication properties

You can view attributes for remote authentication service which is used to authenticate remote users to the cluster. Remote authentication allows users of SAN management applications, such as IBM Tivoli Storage Productivity Center, to authenticate to the cluster using the authentication service provided by the SAN management application.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following tasks to view remote authentication properties:

1. In the portfolio, click **Manage Cluster**.
2. In the portfolio, click **View Cluster Properties**. The Cluster Properties notebook is displayed.
3. Click the **Remote Authentication** tab in the notebook to display the remote authentication properties.
4. Click **Close** to return to the Cluster Properties notebook.

Creating a user group

You can create user groups to organize users of the SAN Volume Controller cluster by role. Roles define access to different cluster functions. Administrators can create role-based user groups where any users added to the group adopts the role that is assigned to that group. User groups simplify management of user access to the cluster.

You must have the Security Administrator role to create, delete, or change a user group.

Roles apply to both local and remote users on the cluster and are based on the user group to which the user belongs. A local user can only belong to a single group; therefore, the role of a local user is defined by the single group that the user belongs to. Remote users can belong to one or more groups; therefore, the roles of remote users are assigned according to the groups that the remote user belongs to.

This task assumes that you have already launched the SAN Volume Controller Console. To create a user group, complete the following steps:

1. In the portfolio, click **Manage Authentication** → **User Groups**. The Viewing User Groups panel is displayed.
2. Select **Create a User Group** from the task list and click **Go**. The Creating a User Group panel is displayed.
3. Enter a name for the user group.
4. Select the role that all users adopt when they are added to this user group. The following roles can be selected:

   Monitor
   Select this role if you want the user to access all viewing actions available with the SAN Volume Controller Console. This user cannot

perform any actions that change the state of the cluster or the resources that the cluster manages. The user can access all the information-related panels and commands, back up configuration data, change his or her password, and issue the following commands: finderr, dumperrlog, dumpinterallog, and chcurrentuser.

Copy Operator
Select this role if you want the user to manage all existing FlashCopy, Metro Mirror, and Global Mirror relationships. They can also create and delete FlashCopy mappings, FlashCopy consistency groups, Metro Mirror or Global Mirror relationships, and Metro Mirror and Global Mirror consistency groups. In addition, the user can access all the functions available to the Monitor role.

Service
Select this role if you want the user to view the View Clusters panel, launch the SAN Volume Controller Console, and view the progress of actions on clusters with the View Progress panel, begin disk discovery process, and discover and include disks. The user can access the following commands: applysoftware, setlocale, addnode, rmnode, cherrstate, setevent, writesernum, detectmdisk, and includemdisk. A user with this role can also access all the functions available to the Monitor role.

Administrator
Select this role if you want the user to access all functions on the SAN Volume Controller Console and issue any command-line interface (CLI) command, except those that deal with managing users, user groups, and authentication.

Security Administrator
Select this role if you want the user to access all functions on the SAN Volume Controller Console and issue any CLI command. Users with this role can also manage users, user groups, and manage user authentication.

5. Select **Enable this user group to be visible to a remote authentication service** if you want the user group to match the access that is defined in user groups on a remote authentication service. For each group of users on the remote authentication service, there must be an SAN Volume Controller user group with the same name and the user group must be visible to the remote authentication service. Security administrators can control what user groups can match the access of user groups on the remote authentication service. When SAN Volume Controller Console authenticates a remote user, it requests a list of groups that the user belongs to from the remote authentication service. The system then assigns a role to the remote user based on whether there is an existing user group on the SAN Volume Controller with the same name and if that user group allows remote visibility. When these criteria are met, the SAN Volume Controller assigns the role based on the user group role specification. If the user is a member of multiple groups that match multiple roles, the user is given the most powerful role. In the case where a user has a combination of Copy Operator and Service roles, the SAN Volume Controller assigns both roles to the user.
6. Click **OK**.

Viewing user groups

You can display all user groups that are currently defined on the cluster. From the Viewing User Groups panel in the SAN Volume Controller Console, you can work with the displayed user groups by selecting an action from the task list.

Users with the Security Administrator role can organize users of the SAN Volume Controller cluster by role through user groups. Administrators can create role-based user groups where any users added to the group adopt the role that is assigned to that group. Roles apply to both local and remote users on the cluster and are based on the user group to which the user belongs. A local user can only belong to a single group; therefore, the role of a local user is defined by the single group that the user belongs to. Remote users can belong to one or more groups; therefore, the roles of remote users are assigned according to the groups that the remote user belongs to.

This task assumes that you have already launched the SAN Volume Controller Console. To view user groups, complete the following steps:

1. Click **Manage Authentication** → **User Groups**. The Viewing User Groups panel is displayed.
2. Click **Close**.

Viewing user-group details

You can display detailed information on user groups that are currently defined on the cluster.

This task assumes that you have already launched the SAN Volume Controller Console. Perform the following steps to display detailed information on user of the cluster.

1. In the portfolio, click **Manage Authentication** → **User Groups**. The Viewing User Groups panel is displayed.
2. On the Viewing User Groups panel, select the name of the user group that you want to display details for. The Viewing User Group Details panel displays.
3. Click **Close**.

Modifying user groups

You can modify existing user groups that are used to organize users of the SAN Volume Controller cluster by role. Administrators can change properties of user groups by using the Modifying User Groups panel in the SAN Volume Controller Console

You must have the Security Administrator role to create, delete, or change a user group.

This task assumes that you have already launched the SAN Volume Controller Console. To change the properties of a user group, complete the following steps:

1. In the portfolio, click **Manage Authentication** → **User Groups** . The Viewing User Groups panel is displayed.
2. Select the user group that you want to change and select **Modify a Group** from the task list. Click **Go**. The Modifying User Group panel is displayed.
3. Select the role that all users adopt when they are added to this user group. The following roles can be selected:

Monitor
Select this role if you want the user to access all viewing actions available with the SAN Volume Controller Console. This user cannot perform any actions that change the state of the cluster or the resources that the cluster manages. The user can access all the information-related panels and commands, back up configuration data, change his or her password, and issue the following commands: finderr, dumperrlog, dumpinterallog, and chcurrentuser.

Copy Operator
Select this role if you want the user to manage all existing FlashCopy, Metro Mirror, and Global Mirror relationships. They can also create and delete FlashCopy mappings, FlashCopy consistency groups, Metro Mirror or Global Mirror relationships, and Metro Mirror and Global Mirror consistency groups. In addition, the user can access all the functions available to the Monitor role.

Service
Select this role if you want the user to view the View Clusters panel, launch the SAN Volume Controller Console, and view the progress of actions on clusters with the View Progress panel, begin disk discovery process, and discover and include disks. The user can access the following commands: applysoftware, setlocale, addnode, rmnode, cherrstate, setevent, writesernum, detectmdisk, and includemdisk. A user with this role can also access all the functions available to the Monitor role.

Administrator
Select this role if you want the user to access all functions on the SAN Volume Controller Console and issue any command-line interface (CLI) command, except those that deal with managing users, user groups, and authentication.

Security Administrator
Select this role if you want the user to access all functions on the SAN Volume Controller Console and issue any CLI command. Users with this role can also manage users, user groups, and manage user authentication.

4. Select **Enable this user group to be visible to a remote authentication service** if you want the user group to match the access that is defined in user groups on a remote authentication service. Security administrators can control what user groups can match the access of user groups on the remote authentication service. When SAN Volume Controller Console authenticates a remote user, it requests a list of groups that the user belongs to from the remote authentication service. The system then assigns a role to the remote user based on whether there is an existing user group on the SAN Volume Controller with the same name and if that user group allows remote visibility. When these criteria are met, the SAN Volume Controller assigns the role based on the user group role specification. If the user is a member of multiple groups that match multiple roles, the user is given the most powerful role. In the case where a user has a combination of Copy Operator and Service roles, the SAN Volume Controller assigns both roles to the user.
5. Click **OK**.

Deleting a user group

You can delete a user group by using the Deleting User Groups panel in the SAN Volume Controller Console.

You must have the Security Administrator role to create, delete, or change a user group.

This task assumes that you have already launched the SAN Volume Controller Console. To delete a user group, complete the following steps:

1. In the portfolio, click **Manage Authentication** → **User Groups**. The Viewing User Groups panel is displayed.
2. Select the user group that you want to delete and select **Delete a Group** from the task list. Click **Go**. The Deleting User Group panel is displayed.
3. Click **Delete**. If users currently exist in the user group a confirmation panel displays asking if you want to force the deletion of the user group and move the users in that group to the Monitor user group. Click **Force Delete** to move the user and delete the user group.

Creating users

You can create either a local or a remote user to access a SAN Volume Controller cluster.

You can create two categories of users that access the cluster. These types are based on how the users are authenticated to the cluster. Local users must provide either a password, a Secure Shell (SSH) key, or both. Local user are authenticated through the authentication methods that are located on the SAN Volume Controller cluster. If the local user needs access to SAN Volume Controller Console, a password is needed for the user. If the user requires access to the command-line interface (CLI) then a valid SSH key file is necessary. If a user is working with both interfaces, then both a password and SSH key are required. Local users must be part of a user group that is defined on the cluster. User groups define roles that authorize the users within that group to a specific set of operations on the cluster.

A remote user is authenticated on a remote service usually provided by a SAN management application, such as IBM Tivoli Storage Productivity Center, and does not need local authentication methods. For a remote user, both a password and SSH key are required to use the command-line interface. Remote users only need local credentials to access to the SAN Volume Controller Console if the remote service is down. Remote users have their groups defined by the remote authentication service.

This task assumes that you have already launched the SAN Volume Controller Console. Complete the following steps to create either a local or remote user:

1. Click **Manage Authentication** → **Users** in the portfolio. The Viewing Users panel is displayed.
2. Select **Create a User** from the task list and click **Go**. The Creating a User panel is displayed.
3. Enter a name for the user.
4. Enter a password for the user. The password cannot start or end with a blank character. The password can consist of a string of 6 - 64 printable ASCII characters.
5. Enter the SSH key file that is associated with the user. Click **Browse** to select the file. An SSH key is needed if this user plans to use the command-line interface to manage the cluster. Any SAN Volume Controller users that use the remote authentication service and require SSH keys to access the command-line interface must have the same password on the cluster and the remote authentication service. In addition the user group that the user belongs to must

be visible to the remote authentication service. The remote visibility setting instructs SAN Volume Controller to check the remote authentication service for that user's group information to determine the user's role on the cluster.

6. Select the appropriate authentication type for the user. Select **Remote** if the user is authenticate to the cluster by a remote authentication service. Select **Local** if the user is authenticated to the cluster using cluster authentication methods.

   Note: Local is the default setting for the authentication type.

7. If you selected to create a local user, you must also specify the user group that the user belongs to. The user group defines roles that provide the user with access to specific operations on the cluster.
8. Click **OK**.

Viewing user details

You can display detailed information on a user of the SAN Volume Controller cluster.

This task assumes that you have already launched the SAN Volume Controller Console. To display detailed information on user of the cluster, follow these steps:

1. In the portfolio, click **Manage Authentication** → **Users**. The Viewing Users panel is displayed.
2. On the Viewing Users panel, select the name of the user that you want to display details for. The Viewing User Details panel displays.
3. Click **Close**.

Modifying a user

You can change properties and settings for either a local or remote user of a SAN Volume Controller cluster.

You must have security administrator role to create, delete, or modify local and remote users.

This task assumes that you have already launched the SAN Volume Controller Console. To change the properties of either a local or remote user, complete the following steps:

1. Click **Manage Authentication** → **Users** in the portfolio. The Viewing Users panel is displayed.
2. Select the user that you want to change and select **Modify a User** from the task list. Click **Go**. The Modifying a User panel is displayed.
3. You can change the following credentials for the selected user:

   New Name
   Enter the new name of the selected user. The user name cannot start or end with a blank character. The user name can consist of a string of 1 - 256 ASCII characters with the exception of the following characters: %:",* '. This field is required.

   Password
   Enter a password for the user. The password cannot start or end with a blank character. The password can consist of a string of 6 - 64 ASCII characters. The following characters are not allowed: %:",* '. To confirm the new password, re-enter the password. You can also select **Remove Password** to remove the password for the current user.

SSH Key Public File
Enter the SSH key file that is associated with the user. Click **Browse** to select the file. An SSH key is needed if this user plans to use the command-line interface to manage the cluster. You can select **Remove SSH Public Key** to remove the SSH public key from the user.

Authentication Type
Select the authentication type for the selected user. Select **Remote** if the user is authenticated to the cluster by anremote authentication service. Select **Local** if the user is authenticated by the cluster.

User Group
Select the user group that you want the user to belong to.User groups allow you to manage authenticated users into groups based on their access level or role. The role determines the access to cluster functions for the users in the group.

4. Click **OK**.

Modifying current user

Local and remote users can change their own password and Secure Shell (SSH) key using the Modify Current® User panel in the SAN Volume Controller Console.

Local users must provide either a password, a SSH key, or both. Local users are authenticated through the authentication methods that are located on the SAN Volume Controller cluster. If the local user needs access to SAN Volume Controller Console, a password is needed for the user. If the user requires access to the command-line interface, a valid SSH key file is necessary. If a user is working with both interfaces, both a password and SSH key are required. Local users must be part of a user group that is defined on the cluster. User groups define roles that authorize the users within that group to a specific set of operations on the cluster.

A remote user is authenticated on a remote service usually provided by a SAN management application, such as IBM Tivoli Storage Productivity Center, and does not need local authentication methods. For a remote user, both a password and SSH key are required to use the command-line interface. Remote users only need local credentials to access to the SAN Volume Controller Console if the remote service is down. Remote users have their groups defined by the remote authentication service.

This task assumes that you have already launched the SAN Volume Controller Console. To change the password, SSH key, or both for the current user, complete the following steps:

1. Click **Manage Authentication** → **Modify Current User** in the portfolio. The Modify Current User panel is displayed.
2. You can change the following credentials for the current user:

Password
Enter a new password for the user. The password cannot start or end with a blank character. The password can consist of any string of 6 - 64 printable ASCII characters. To confirm the new password, re-enter the password. You can also select **Remove Password** to remove the password for the current user.

Note: This password is only changed on the SAN Volume Controller cluster. If you use remote authentication, then you must also update the password on the remote authentication service.

SSH Key Public File

Enter the SSH key file that is associated with the user. Click **Browse...** to select the file. An SSH key is needed if this user plans to use the command-line interface to manage the cluster. You can select **Remove SSH Public Key** to remove the SSH public key from the user.

3. Click **OK**.

Deleting users

You can delete a user of the SAN Volume Controller Console from the Viewing Users panel.

You must have security administrator role to create, delete, or modify local and remote users.

This task assumes that you have already launched the SAN Volume Controller Console. To delete local or remote user, complete the following steps:

1. Click **Manage Authentication** → **Users** in the portfolio. The Viewing Users panel is displayed.
2. Select the user to delete and select **Deleting a User** from the taks list. Click **Go**. The Deleting User panel is displayed.
3. To delete the user, click **OK**, or to exit this panel without deleting the user, click **Cancel**.

Adding an SNMP server

You can use the SAN Volume Controller Console to specify a Simple Network Management Protocol (SNMP) server to receive event notifications from the cluster. You can specify up to six SNMP servers to receive notifications.

This task assumes that you have already launched the SAN Volume Controller Console.

SNMP is the standard protocol for managing networks and exchanging messages. SNMP enables the SAN Volume Controller to send external messages that notify personnel about an event. You can use an SNMP manager to view the messages that the SNMP agent sends.

Complete the following steps to add a new SNMP server to receive event notifications from the cluster:

1. Click **Service and Maintenance** → **Set SNMP Event Notification** in the portfolio. The Viewing SNMP Error Notification Settings panel displays.
2. Select **Add a Server** from the task list and click **Go**. The Add a New SNMP Server panel displays.
3. On the Adding a New SNMP Server panel, provide the following information for the new SNMP server:

   Server Name

   Enter a new name for the SNMP server. A valid name is 1 to 15 ASCII characters. The following characters can be used: a - z, A - Z, 0 - 9, -, _. The name cannot begin with a number or the hyphen (-) character. Names starting with snmp are reserved for SAN Volume Controller default names.

IP Address
Select the appropriate address format and specify a valid IP address for the SNMP server.

Port Enter the port value that corresponds with the IP address for the SNMP server.

Note: The usual value for the SNMP port is 162, but this must be verified.

Community
Enter a name for the SNMP community. A valid name is 1 to 60 ASCII characters.

Send Error Notifications
Select this option to send notifications with the type of *error* to the SNMP server.

Send Warning Notifications
Select this option to send notifications with the type of *warning* to the SNMP server.

Send Informational Notifications
Select this option to send notifications with the type of *informational* to the SNMP server.

4. Click **OK** to add the SNMP server.

Modifying SNMP server settings

You can use the SAN Volume Controller Console to change Simple Network Management Protocol (SNMP) server settings for the cluster.

This task assumes that you have already launched the SAN Volume Controller Console.

SNMP is the standard protocol for managing networks and exchanging messages. SNMP enables the SAN Volume Controller to send external messages that notify personnel about an event. You can use an SNMP manager to view the messages that the SNMP agent sends.

Complete the following steps to change settings for an SNMP server:

1. Click **Service and Maintenance → Set SNMP Event Notification** in the portfolio. The Viewing SNMP Error Notification Settings panel displays.
2. Select the SNMP server to change and select **Modify Server** from the task list and click **Go**. The Modifying SNMP Server Settings panel displays.
3. On the Modifying SNMP Server Settings panel, change one or more of the following SNMP server properties:

Server Name
Enter a new name for the SNMP server. The current value is displayed on the panel.

IP Address
Select the appropriate address format and specify the new IP address for the SNMP server. The current value is displayed on the panel.

Port Enter the port value for the SNMP IP address.

Note: The usual value for the SNMP port is 162, but this must be verified.

Community
Enter a name for the SNMP community. A valid name is 1 to 60 ASCII characters.

Send Error Notifications
Select this option to send events with the notification type of *error* to the SNMP server.

Send Warning Notifications
Select this option to send events with the notification type of *warning* to the SNMP server.

Send Informational Notifications
Select this option to send events with the notification type of *informational* to the SNMP server.

4. Click **OK** to modify the SNMP server with these new settings.

Deleting an SNMP server

You can use the SAN Volume Controller Console to delete a Simple Network Management Protocol (SNMP) server or multiple SNMP servers.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to delete an SNMP server or servers:

1. Click **Service and Maintenance** → **Set SNMP Event Notification** in the portfolio. The **Viewing SNMP Error Notification Settings** panel displays.
2. Select the SNMP server or servers to delete and select **Delete Server** from the task list and click **Go**. The Delete SNMP Server Setting panel displays.
3. On the Delete SNMP Server Settings panel, confirm that the selected SNMP servers can be deleted.
4. Click **OK** to delete the selected SNMP server or servers.

Adding a syslog server

You can use the SAN Volume Controller Console to specify a syslog server to receive event notifications from the cluster.

This task assumes that you have already launched the SAN Volume Controller Console.

The syslog protocol is a client-server standard for forwarding log messages from a sender to a receiver on an IP network. The IP network can be either IPv4 or IPv6. The User Datagram Protocol (UDP) is used to transmit the syslog message. The syslog protocol can be used across multiple platforms, which enables you to integrate log data from different types of storage systems into a central repository.

Complete the following steps to add a syslog server to receive messages from the cluster:

1. Click **Service and Maintenance** → **Set Syslog Event Notification** in the portfolio. The **Viewing Syslog Notification Settings** panel displays.

2. Select **Add a Server** from the task list and click **Go**. The Add a New Syslog Server panel displays.
3. On the Adding a New Syslog Server panel, provide the following information for the new syslog server:

   Server Name
   Enter a new name for the syslog server. A valid name is 1 to 15 ASCII characters. The following characters can be used: a - z, A - Z, 0 - 9, -, _. The name cannot begin with a number or the hyphen (-) character. Names starting with syslog are reserved for SAN Volume Controller default names.

   IP Address
   Select the appropriate address format and specify a valid IP address for the syslog server.

   SVC Facility
   Select the SAN Volume Controller facility for the syslog server. This value is translated to the corresponding syslog code and value. The SAN Volume Controller facility of 0 - 3 results in the syslog server receiving messages in concise format, which provides standard detail on the event. A SAN Volume Controller facility of 4 - 7 results in the syslog server receiving message in fully expanded format, where more details are provided regarding the event.

   Send Error Notifications
   Select this option to send events with the notification type of *error* to the syslog server.

   Send Warning Notifications
   Select this option to send events with the notification type of *warning* to the syslog server.

   Send Informational Notifications
   Select this option to send events with the notification type of *informational* to the syslog server.
4. Click **OK** to add the syslog server.

Modifying syslog server settings

You can use the SAN Volume Controller Console to change properties for a syslog server that is configured to receive event notifications from the cluster.

This task assumes that you have already launched the SAN Volume Controller Console.

The syslog protocol is a client-server standard for forwarding log messages from a sender to a receiver on an IP network. The IP network can be either IPv4 or IPv6. The User Datagram Protocol (UDP) is used to transmit the syslog message. The syslog protocol can be used across multiple platforms, which enables you to integrate log data from different types of storage systems into a central repository.

Complete the following steps to modify a syslog server:

1. Click **Service and Maintenance** → **Set Syslog Event Notification** in the portfolio. The **Viewing Syslog Notification Settings** panel displays.
2. Select **Modify Server** from the task list and click **Go**. The Modify Syslog Server Settings panel displays.

3. On the Modify Syslog Server Settings panel, update any of the following properties for the selected syslog server:

   Server Name
   Enter a new name for the syslog server. The current value is displayed on the panel.

   IP Address
   Select the appropriate address format and specify a valid IP address for the syslog server. The current value is displayed on the panel.

   SVC Facility
   Select the SAN Volume Controller facility for the syslog server. This value is translated to the corresponding syslog code and value. The SAN Volume Controller facility of 0 - 3 results in the syslog server receiving messages in concise format, which provides standard detail on the event. A SAN Volume Controller facility of 4 - 7 results in the syslog server receiving message in fully expanded format, where more details are provided regarding the event.

   Send Error Notifications
   Select this option to send events with the notification type of *error* to the syslog server.

   Send Warning Notifications
   Select this option to send events with the notification type of *warning* to the syslog server.

   Send Informational Notifications
   Select this option to send events with the notification type of *informational* to the syslog server.
4. Click **OK** to add the syslog server.

Deleting syslog server settings

You can use the SAN Volume Controller Console to delete syslog server settings.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to delete an syslog server or servers:

1. Click **Service and Maintenance** → **Set Syslog Event Notification** in the portfolio. The **Viewing Syslog Notification Settings** panel displays.
2. Select the syslog server or servers to delete and select **Delete Server** from the task list and click **Go**. The Delete SyslogServer Setting panel displays.
3. On the Delete Syslog Server Settings panel, confirm that the selected syslog servers can be deleted.
4. Click **OK** to delete the selected syslog server or servers.

Creating e-mail event notifications and inventory reports

You can use the Create E-mail Event Notification wizard to enable the SAN Volume Controller Console e-mail service to send event notification and inventory reports to the IBM Support Center.

The wizard guides you through all the actions that are necessary to set up e-mail notifications and inventory reports for the first time. Before starting the wizard, ensure that you have the following information:

- Valid IP address for the e-mail server
- Contact information for the user who receive event notifications
- If you do not use the IBM Support Center, you need an alternative e-mail address of an alternate e-mail user.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to create e-mail event notifications and inventory reports:

1. Click **Service and Maintenance** → **Set Email Features**. The Create E-mail Event Notification–Introduction panel is displayed. Click **Continue**.

   Note: If you have not configured e-mail event notifications, the Create E-mail Event Notification wizard launches; however, if e-mail event notifications have been previously configured, then the Manage E-mail Event Notifications panel displays. The wizard also displays if e-mail service is disabled for any reason, such as there are no users or servers configured. If the e-mail service has been disabled, you can bypass the wizard.
2. On the Set Contact Details panel, specify the contact details for e-mail event notifications, such e-mail reply address and machine location of the SAN Volume Controller cluster. Click **Continue**. The Set Initial E-mail Server panel displays.
3. On the Set Initial E-mail Server panel, enter the IP address and other information on the e-mail server that receives event notifications from the SAN Volume Controller cluster. Initially one server can be set up to receive event notifications; however, up to 6 can be added at a later time. Click **Continue**. The Set Initial E-mail User panel displays.
4. On the Set Initial E-mail User panel, you can select the initial e-mail user to receive notifications. It is recommended that you use one of the IBM Support Centers for your initial e-mail user. The IBM Support Center initial user receives event notifications with the type *error*, which indicates a serious problem with the SAN Volume Controller cluster and should be fixed immediately. Inventory reports are also automatically sent to this user. You can add additional users to receive event notifications and inventory reports at a later time. Select user type and press **Continue**. If you selected an IBM Support Center for your user type, press **Continue** to confirm your choice. If you selected a custom user, you must enter your settings and then press **Continue**. If you choose to define a custom support user or a custom local user, you must ensure that the alternate e-mail address is valid. The Start E-mail Service panel displays. For support users, specify event notifications with the type *error* to ensure critical problems with the SAN Volume Controller cluster are handled quickly and enable inventory reporting. For a custom local user, you can enable inventory reporting and choose the following event notification types:

   Error Notifications
   When an event notification with the type *error* is received by the specified user, the problem must be fixed immediately. This notification indicates a serious problem with the SAN Volume Controller cluster. For example, the event that is being reported could indicate a loss of redundancy in the system, and it is possible that another failure could result in loss of access to data. The most typical reason this type of notification is sent is because of a hardware failure, but some configuration errors or fabric errors also are included in this notification type.

Warning Notifications
A warning notification is sent to indicate a problem or unexpected condition with the SAN Volume Controller. Always immediately investigate this type of notification to determine the effect that it might have on your operation and correct if necessary.

Informational Notifications
For these types of notifications, no action is necessary to resolve them. This notification indicates an event that was expected has occurred. For example, a FlashCopy mapping event has completed successfully. No remedial action is required when these notification types are sent.

5. On the next panel, click Start E-mail Service to start the e-mail service. The e-mail notifications and inventory reports are now sent to the specified user.
6. Optional: After the e-mail service successfully starts, the Confirmation panel displays. You can optionally test these settings by clicking the Send Test E-mail button.

Adding e-mail users

You can use the Add E-mail User panel to add additional e-mail users who receive the event notifications and inventory reports. From this panel you can add either a support user or a local user.

This task assumes that you have already launched the SAN Volume Controller Console.

1. Click **Service and Maintenance** → **Set Email Features**. The Manage E-mail Event Notifications panel displays.

   Note: If you have not configured e-mail event notifications, the Create E-mail Event Notification wizard launches; however, if e-mail event notifications have been previously configured, then the Manage E-mail Event Notifications panel displays. The wizard also displays if e-mail service is disabled for any reason, such as there are no users or servers configured. If the e-mail service has been disabled, you can bypass the wizard.
2. On the Manage E-mail Event Notifications, click **Manage E-mail Users**. The Manage E-mail Users panel displays. This panel displays two tables with details on support users and local users that are defined on the cluster. From the task list that corresponds with the user type, select **Add User** from the task list and click **Go**.
3. If you select a support user to add, then the Add Support E-mail User panel displays. Complete the following steps:
   a. Select one of the IBM Support Center based on the geography where your cluster is located or select to enter a custom support user. A support user is a permanent e-mail address that event notifications are sent to. Support user addresses are normally used for dedicated support organizations. If you add a custom support user, you must ensure that the support center and user address are valid prior to setting these values.
   a. Select **Send Error Notifications** for the type of event notification that the support user receives.
   b. Select **Enabled** to have inventory reports sent to the specified user.
   c. Click **OK**.
4. If you select a local user to add, then the Add Local E-mail User panel displays. Complete the following steps:
   a. Specify the user name and e-mail address for the new local user.

b. Select the any of the following event notification settings for the new local user:

Send Error Notification
When an event notification with the type *error* is received by the specified user, the problem must be fixed immediately. This notification indicates a serious problem with the SAN Volume Controller cluster. For example, the event that is being reported could indicate a loss of redundancy in the system, and it is possible that another failure could result in loss of access to data. The most typical reason this type of notification is sent is because of a hardware failure, but some configuration errors or fabric errors also are included in this notification type.

Send Warning Notification
A warning notification is sent to indicate a problem or unexpected condition with the SAN Volume Controller. Always immediately investigate this type of notification to determine the effect that it might have on your operation and correct if necessary. A warning notification does not require any replacement parts and therefore should not require IBM Support Center involvement. However, the event being reported might indicate a condition that could be fatal to your operating environment; such as, for example, a critical FlashCopy operation has failed.

Send Informational Notification
For these types of notifications, no action is necessary to resolve them. This notification indicates an event that was expected has occurred. For example, a FlashCopy mapping event has completed successfully.

a. Select **Enabled** to have inventory reports sent to the specified user.
b. Click **OK**.

Modifying e-mail users

You can use the Modify E-Mail User panel in the SAN Volume Controller Console to change support for e-mail users and local users that you currently have defined to receive event notifications.

This task assumes that you have already launched the SAN Volume Controller Console.

1. Click **Service and Maintenance** → **Set Email Features**. The Manage E-mail Event Notifications panel displays.

   Note: If you have not configured e-mail event notifications, the Create E-mail Event Notification wizard launches; however, if e-mail event notifications have been previously configured, then the Manage E-mail Event Notifications panel displays. The wizard also displays if e-mail service is disabled for any reason, such as there are no users or servers configured. If the e-mail service has been disabled, you can bypass the wizard.
2. On the Manage E-mail Event Notifications panel, click **Manage E-mail Users**. The Manage E-mail Users panel displays. This panel displays two tables with details on support users and local users that are defined on the cluster.
3. Select the user that you want to change from either the Support User table or the Local User table and select **Modify User** from the task list. Click **Go**.

4. If you selected a support user to change, the Modify Support E-mail User panel displays.

   Note: A support e-mail user is recommended. Changing these settings could cause unexpected results. Change these settings only if the new values are correct and another support center is present. If this is the last user that is specified to receive event notifications, deleting this user disables event notifications. Event notification can be restarted by adding another user and restarting the e-mail service. To change the support user settings, follow these steps:

   a. Enter a new support user description and e-mail address. A support user is an e-mail address that event notifications are sent to. Support user addresses are normally used for dedicated support organizations. If you add a custom support user, you must ensure that the support center and user address are valid prior to setting these values.
   a. Select **Send Error Notifications** for the type of event notification that the support user receives.
   b. Select **Enabled** to have inventory reports sent to the specified user.
   c. Click **OK**.
5. If you select a local user to change, the Modify Local E-mail User panel displays. Complete the following steps:
   a. Specify the new user name and e-mail address for the selected local user.
   b. Select the any of the following event notification settings for the selected local user:

      Send Error Notification
      When an event notification with the type *error* is received by the specified user, the problem must be fixed immediately. This notification indicates a serious problem with the SAN Volume Controller cluster. For example, the event that is being reported could indicate a loss of redundancy in the system, and it is possible that another failure could result in loss of access to data. The most typical reason this type of notification is sent is because of a hardware failure, but some configuration errors or fabric errors also are included in this notification type.

      Send Warning Notification
      A warning notification is sent to indicate a problem or unexpected condition with the SAN Volume Controller. Always immediately investigate this type of notification to determine the effect that it might have on your operation and correct if necessary. A warning notification does not require any replacement parts and therefore should not require IBM Support Center involvement. However, the event being reported might indicate a condition that could be fatal to your operating environment; such as, for example, a critical FlashCopy operation has failed.

      Send Informational Notification
      For these types of notifications, no action is necessary to resolve them. This notification indicates an event that was expected has occurred. For example, a FlashCopy mapping event has completed successfully.

   a. Select the appropriate setting for inventory reporting for the the specified user.
   b. Click **OK**.

Deleting an e-mail user

You can use the Delete E-Mail User panel in the SAN Volume Controller Console to delete support e-mail users and local users that you currently have defined to receive event notifications. If this is the last user that is specified to receive event notifications, then deleting this user disables event notifications. Event notification can be restarted by adding another user and restarting the e-mail service.

This task assumes that you have already launched the SAN Volume Controller Console.

1. Click **Service and Maintenance** → **Set Email Features**. The Manage E-mail Event Notifications panel displays.

   Note: If you have not configured e-mail event notifications, the Create E-mail Event Notification wizard launches; however, if e-mail event notifications have been previously configured, then the Manage E-mail Event Notifications panel displays. The wizard also displays if e-mail service is disabled for any reason, such as there are no users or servers configured. If the e-mail service has been disabled, you can bypass the wizard.
2. On the Manage E-mail Event Notifications panel, click **Manage E-mail Users**. The Manage E-mail Users panel displays. This panel displays two tables with details on support users and local users that are defined on the cluster.
3. Select the user that you want to delete from the either the Support User table or the Local User table and select **Delete User** from the task list. Click **Go**.
4. If you selected a support user to change, the Delete Support E-mail User panel displays.

   Note: A support e-mail user is recommended and deleting a support e-mail user stops it from receiving event notifications at the specified e-mail address. Support addresses are usually permanent, and if deleted, the specified support center is not notified of potentially severe events, which could negatively impact the system. To delete the support user settings, follow these steps:

   a. Verify that the selected support user can be deleted. Ensure that the support user is no longer necessary before you delete.
   b. Click **Delete**.
5. If you select a local user to delete, the Delete Local E-mail User panel displays. Complete the following steps:
   a. Verify that the selected local user can be deleted.
   b. Click **Delete**.

Adding an e-mail server

You can use the Add E-mail Server panel in the SAN Volume Controller Console to add a new e-mail server object. You can add up to 6 e-mail servers for the cluster.

This task assumes that you have already launched the SAN Volume Controller Console.

1. Click **Service and Maintenance** → **Set Email Features**. The Manage E-mail Event Notifications panel displays.

   Note: If you have not configured e-mail event notifications, the Create E-mail Event Notification wizard launches; however, if e-mail event notifications

have been previously configured, then the Manage E-mail Event Notifications panel displays. The wizard also displays if e-mail service is disabled for any reason, such as there are no users or servers configured. If the e-mail service has been disabled, you can bypass the wizard.

2. On the Manage E-mail Event Notifications panel, click **Manage E-mail Service**. The Manage E-mail Service panel displays. This panel displays two tables that display details about the e-mail server and contact for the notifications. On this panel, you can also start or stop e-mail service for a selected server.
3. Select **Add E-mail Server** from the task list and click **Go**. The Add E-mail Server panel displays.
4. Enter the name, IP address, and port for the e-mail server that you are adding. Click **OK**.

Modifying an e-mail server

You can use the Modify E-mail Server panel in the SAN Volume Controller Console to change attributes, such as the IP address and port number, for a selected e-mail server.

This task assumes that you have already launched the SAN Volume Controller Console.

1. Click **Service and Maintenance** → **Set Email Features**. The Manage E-mail Event Notifications panel displays.

   Note: If you have not configured e-mail event notifications, the Create E-mail Event Notification wizard launches. However, if e-mail event notifications have been previously configured, the Manage E-mail Event Notifications panel displays.
2. On the Manage E-mail Event Notifications panel, click **Manage E-mail Service**. The Manage E-mail Service panel displays. This panel displays two tables that display details about the e-mail server and contact for the notifications. On this panel, you can also start or stop e-mail service for a selected server.
3. Select the e-mail server from the E-mail Server Details table and select Modify E-mail Server from the task list. Click **Go**. The Modify E-mail Server panel displays.
4. Enter the new name, IP address, or port for the e-mail server that you are changing. Click **OK**.

Deleting an e-mail server

You can use the Delete E-mail Server panel in the SAN Volume Controller Console to delete e-mail server objects from the cluster. If this is the last e-mail server, no event notifications are sent until a new e-mail server object is added and the e-mail service is restarted.

This task assumes that you have already launched the SAN Volume Controller Console.

1. Click **Service and Maintenance** → **Set Email Features**. The Manage E-mail Event Notifications panel displays.

   Note: If you have not configured e-mail event notifications, the Create E-mail Event Notification wizard launches. However, if e-mail event notifications have been previously configured, the Manage E-mail Event Notifications panel displays.

2. On the Manage E-mail Event Notifications panel, click **Manage E-mail Service**. The Manage E-mail Service panel displays. This panel displays two tables that display details about the e-mail server and contact for the notifications. On this panel, you can also start or stop e-mail service for a selected server.
3. Select the e-mail server from the E-mail Server Details table and select Delete E-mail Server from the task list. Click **Go**. The Delete E-mail Server panel displays.
4. Verify that this is the correct server to be deleted and click **Delete**.

Starting the e-mail service

You can use the Start E-mail Service panel in the SAN Volume Controller Console to start the e-mail service.

This task assumes that you have already launched the SAN Volume Controller Console.

1. Click **Service and Maintenance** → **Set Email Features**. The Manage E-mail Event Notifications panel displays.

   Note: If you have not configured e-mail event notifications, the Create E-mail Event Notification wizard launches; however, if e-mail event notifications have been previously configured, then the Manage E-mail Event Notifications panel displays. The wizard also displays if e-mail service is disabled for any reason, such as there are no users or servers configured. If the e-mail service has been disabled, you can bypass the wizard.
2. On the Manage E-mail Event Notifications panel, click **Manage E-mail Service**. The Manage E-mail Service panel displays.
3. Click **Start E-mail Service**. The Start E-mail Service panel displays with a list of users that are currently configured to receive notifications of events. You can change the users who receive event notifications by clicking **Manage E-mail Users**.
4. To start the e-mail service and begin event notification, click **Start E-mail Service**.

Displaying and saving log and dump files

You can save the log and dump files for nodes.

You can save dump data for any node in the cluster. When you use this procedure to display dump data only, the dump files for the configuration node are displayed. An option on the dumps menu allows you to display data from other nodes. If you choose to display or save data from another node, that data is first copied to the configuration node.

The software dump files contain dumps of the SAN Volume Controller memory. Your IBM 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.

The **List dumps** option supports the following file types:

- Error logs
- I/O statistic logs
- I/O trace logs

- Feature logs
- Software dumps
- Audit logs
- CIMOM logs
- Managed Disks (MDisks) logs

Complete the following steps to display log and dump files:

This task assumes that you have already launched the SAN Volume Controller Console.

1. Click **Service and Maintenance** → **List Dumps** in the portfolio. The List Dumps panel is displayed.

   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, 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 are displayed. Dumps and logs on all nodes in the cluster can be deleted on or copied to the configuration 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.
2. Copy the files to your local workstation by right-clicking on the filename and using the **Save Link As...** (Netscape) or **Save Target As...** (Internet Explorer) option from the Web browser.

Analyzing the error log

You can analyze the error log from the Analyze Error Log panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Note: Log files that are copied to the configuration node are *not* automatically deleted by the SAN Volume Controller.

Perform the following steps to analyze the error log:

1. Click **Service and Maintenance** → **Analyze Error Log** in the portfolio. The Analyze Error Log panel is displayed.

   The Analyze Error Log panel lets you 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 that the table is sorted by either error priority or time. For error priority, the most serious errors are the lowest-numbered errors. Therefore, they are displayed first in the table.

   Either the oldest or the latest entry can be displayed first in the table. You can also select how many error log entries are 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 99.
2. 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 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 as unfixed.

3. Click **Clear log** to erase the entire cluster error log.

   Note: Clicking **Clear log** does *not* fix the existing errors.

Recovering a node and returning it to the cluster

After a node or an I/O group fails, you can use the SAN Volume Controller to recover a node and return it to the cluster.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to recover a node and return it to the cluster:

1. Click **Work with Nodes** → **Nodes** in the portfolio. The Viewing Nodes panel is displayed.
2. Verify that the node is offline.
3. Select the offline node.
4. Select **Delete a Node** from the task list and click **Go**. The Deleting Node from Cluster panel is displayed.
5. Click **Yes**.
6. Verify that the node can be seen on the fabric.
7. If the nodes are repaired by replacing the front panel module or a node is repaired by replacing it with another node, the worldwide node name (WWNN) for the node changes. In this case, you must follow these additional steps:
   a. At the end of the recovery process, you must follow your multipathing device driver's procedure to discover the new paths and to check that each device identifier is now presenting the correct number of paths. If you are using the subsystem device driver (SDD), the device identifiers are referred to as virtual paths (vpaths). See the *IBM System Storage Multipath Subsystem Device Driver User's Guide* or the documentation that is provided with your multipathing device driver for more information.
   b. You might also have 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 must modify the port groups that belong to the cluster because the WWNN or worldwide port names (WWPNs) 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 from which it was removed. 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 the IBM Support Center 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, you must record the following information:

- Node serial number
- WWNN
- All WWPNs
- I/O group that the node belongs to

8. Add the node back into the cluster.
   a. From the Viewing Nodes panel, select **Add a Node** from the task list and click **Go**. The Adding a Node to a Cluster panel is displayed.
   b. 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.
   c. Click **OK**.
9. Verify that the node is online by refreshing the Viewing Nodes panel.

   Note: If the panel does not refresh, close the panel and reopen it.

Chapter 6. Using the CLI

The SAN Volume Controller cluster command-line interface (CLI) is a collection of commands that you can use to manage the SAN Volume Controller.

Overview

The CLI commands use 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 you can use the CLI, you must have already created a cluster.

You must perform the following actions to use the CLI from a client system:

- Install and set up SSH client software on each system that you plan to use to access the CLI.
- Generate an SSH key pair on each SSH client.
- Store the SSH public key for each SSH client on the SAN Volume Controller.

  Note: After the first SSH public key is stored, you can add additional SSH public keys using either the SAN Volume Controller Console or the CLI.

You can use the CLI to perform the following functions:

- Set up of the cluster, its nodes, and the I/O groups
- Analyze error logs
- Set up and maintenance of managed disks (MDisk) and MDisk groups
- Set up and maintenance of client public SSH keys on the cluster
- Set up and maintenance of virtual disks (VDisks)
- Set up of logical host objects
- Map VDisks to hosts
- Navigate from managed hosts to VDisks and to MDisks, 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
  - Asynchronous Global Mirror and Global Mirror consistency groups

Configuring a PuTTY session for the CLI

You must configure a PuTTY session using the Secure Shell (SSH) key pair that you have generated before you can use the command-line interface (CLI).

Attention: Do not run scripts that create child processes that run in the background and call SAN Volume Controller commands. This can cause the system to lose access to data and cause data to be lost.

Perform the following steps to configure a PuTTY session for the CLI:

1. Select **Start** → **Programs** → **PuTTY** → **PuTTY**. The PuTTY Configuration window opens.

2. Click **Session** in the Category navigation tree. The Basic options for your PuTTY session are displayed.
3. Click **SSH** as the Protocol option.
4. Click **Only on clean exit** as the Close window on exit option. This ensures that connection errors are displayed.
5. Click **Connection** → **SSH** in the Category navigation tree. The options controlling SSH connections are displayed.
6. Click **2** as the Preferred SSH protocol version.
7. Click **Connection** → **SSH** → **Auth** in the Category navigation tree. The Options controller SSH authentication are displayed.
8. Click **Browse** or type the fully qualified file name and location of the SSH client and private key in the **Private key file for authentication** field.
9. Click **Connection** → **Data** in the Category navigation tree.
10. Type `admin` in the **Auto-login username** field.
11. Click **Session** in the Category navigation tree. The Basic options for your PuTTY session are displayed.
12. Click **Default Settings** and then click **Save**.
13. In the **Host Name (or IP Address)** field, type the name or IP address of one of the SAN Volume Controller cluster IP addresses or host names.
14. Type `22` in the **Port** field. The SAN Volume Controller cluster uses the standard SSH port.
15. Type the name that you want to use to associate with this session in the **Saved Sessions** field. For example, you can name the session SAN Volume Controller Cluster 1.
16. Click **Save**.

You have now configured a PuTTY session for the CLI.

Note: If you configured more than one IP address for the SAN Volume Controller cluster, repeat the previous steps to create another saved session for the second IP address. This can then be used if the first IP address is unavailable.

Preparing the SSH client system for the CLI

Before you can issue command-line interface (CLI) commands from the host to the cluster, you must prepare the Secure Shell (SSH) client system.

Microsoft Windows operating systems

The IBM System Storage Productivity Center (SSPC) and the master console for the SAN Volume Controller include the PuTTY client program, which is a Microsoft Windows SSH client program. The PuTTY client program can be installed on your SSPC or master console server in one of the following ways:

- If you purchased the SSPC or the master console hardware option from IBM, the PuTTY client program has been preinstalled on the hardware.
- You can use the master console software installation CD to install the PuTTY client program. The SSPC, master console hardware option, and the software-only master console each provide this CD.
- You can use the separate PuTTY client program-installation wizard, **putty-*version*-installer.exe**. You can download the PuTTY client program from the following Web site:

www.chiark.greenend.org.uk/~sgtatham/putty/

Note: Before you install the PuTTY client program, ensure that your Windows system meets the system requirements. See the *IBM System Storage Productivity Center Introduction and Planning Guide* for system requirements.

If you want to use an SSH client other than the PuTTY client, the following Web site offers SSH client alternatives for Windows:

www.openssh.org/windows.html

IBM AIX operating systems

For IBM AIX 5L™ for POWER®, versions 5.1, 5.2, 5.3, and AIX version 6.1 for IBM POWER6™ architecture, you can obtain the OpenSSH client from the bonus packs, but you also must obtain its prerequisite, OpenSSL, from the IBM AIX toolbox for Linux applications for IBM Power Systems™. For AIX 4.3.3, you can obtain the software from the AIX toolbox for Linux applications.

You can also obtain the AIX installation images from IBM developerWorks® at the following Web site:

oss.software.ibm.com/developerworks/projects/openssh

Linux operating systems

The OpenSSH client is installed by default on most Linux distributions. If it is not installed on your system, consult your Linux installation documentation or visit the following Web site:

www.openssh.org/portable.html

The OpenSSH client can run on a variety of additional operating systems. For more information about the openSSH client, visit the following Web site:

www.openssh.org/portable.html

Preparing the SSH client system to issue CLI commands

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 is accepted by the SSH server on the cluster.

To use a host that requires a different type of SSH client, OpenSSH for example, follow the instructions for that software.

Perform the following steps to enable your host to issue CLI commands:

1. For the IBM System Storage Productivity Center or master console and Windows hosts:
   a. Generate an SSH key pair using the PuTTY key generator.
   b. Store the SSH clients public key on the cluster (using a browser that points to the SAN Volume Controller Console).
   c. Configure the PuTTY session for the CLI.
2. For other types of hosts:

a. Follow the instructions that are specific to the SSH client to generate an SSH key pair.
b. Store the SSH clients public key on the cluster (using a Web browser to point to the SAN Volume Controller Console or the CLI from an already established host).
c. Follow the instructions that are specific to the SSH client to establish an SSH connection to the SAN Volume Controller cluster.

Preparing the SSH client on an AIX host

When you use AIX hosts, Secure Shell (SSH) logins are authenticated on the SAN Volume Controller cluster using the RSA-based authentication that is supported in the OpenSSH client available for AIX.

RSA-based authentication uses public-key cryptography to allow the encryption and decryption to use separate keys. Therefore, 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 (the SAN Volume Controller cluster in this case) knows the public key, and only the user (the AIX host) knows the private key. Because possession of the private key allows access to the cluster, the private key must be kept in a protected place. You can store the private key in the /.ssh directory on the AIX host with restricted access permissions.

When you use the AIX host to log into the SAN Volume Controller cluster, the SSH program on the SAN Volume Controller cluster sends the AIX host the key pair that it wants to use for authentication. The AIX server checks if this key is permitted, and if so, sends the SSH program that is running on behalf of the user a challenge. The challenge is a random number that is encrypted by the user's public key. The challenge can only be decrypted using the correct private key. The user's client (the AIX host) uses the private key to decrypt the challenge and prove that the user has the private key. The private key is not shown to the server (the SAN Volume Controller cluster) or to anyone who might be intercepting the transmissions between the AIX host and the SAN Volume Controller cluster.

Perform the following steps to set up an RSA key pair on the AIX host and the SAN Volume Controller cluster:

1. Create an RSA key pair by issuing a command on the AIX host that is similar to the following command:

   ```
   ssh-keygen -t rsa
   ```

   Tip: Issue the command from the $HOME/.ssh directory.
   This process generates two user named files. If you select the name *key*, the files are named *key* and *key.pub*. Where *key* is the name of the private key and *key.pub* is the name of the public key.
2. Store the private key from this key pair on the AIX host, in the $HOME/.ssh directory, in the $HOME.ssh/identity file. If you are using multiple keys, all of the keys must appear in the identity file.
3. Store the public key on the IBM System Storage Productivity Center or the master console of the SAN Volume Controller cluster. Typically this can be done with ftp; however, the IBM System Storage Productivity Center or the master console might have ftp disabled for security reasons, in which case an alternative method, such as secure copy is required. You can then use the SAN Volume Controller Console to transfer the public key to the cluster. Select an access level of either administrator or service.

You can now access the cluster from the AIX host using an SSH command similar to the following:

```
ssh admin@my_cluster
```

Where *my_cluster* is the name of the cluster IP. Always use *admin* as the SSH user name. The SAN Volume Controller software determines which user is logging in from the key they are using.

Refer to your client's documentation for SSH on your host system for more host specific details regarding this task.

Issuing CLI commands from a PuTTY SSH client system

You can issue command-line interface (CLI) commands from a PuTTY SSH client system.

Perform the following steps to issue CLI commands:

1. Open a command prompt.
2. Issue the following command to set the path environment variable to include the PuTTY directory:

   ```
   set path=C:\Program Files\putty;%path%
   ```

   Where *Program Files* is the directory where PuTTY is installed.
3. Use the PuTTY plink utility to connect to the SSH server on the cluster.

Starting a PuTTY session for the CLI

You must start a PuTTY session to connect to the command-line interface (CLI).

This task assumes that you have already configured and saved a PuTTY session using the Secure Shell (SSH) key pair that you created for the CLI.

Perform the following steps to start a PuTTY session:

1. Select **Start** → **Programs** → **PuTTY** → **PuTTY**. The PuTTY Configuration window opens.
2. Select the name of your saved PuTTY session and click **Load**.
3. Click **Open**.

   Note: If this is the first time that the PuTTY application is being used since you generated and uploaded the SSH key pair, a PuTTY Security Alert window is displayed. Click **Yes** to accept the change and trust the new key.
4. Type `admin` in the **login as:** field and press Enter.

Setting the cluster time using the CLI

You can use the command-line interface (CLI) to set the cluster time.

Perform the following steps to set the cluster time:

1. Issue the svcinfo showtimezone CLI command to display the current time-zone settings for the cluster. The time zone and the associated time-zone ID are displayed.

2. Issue the svcinfo lstimezones CLI command to list the time zones that are available on the cluster. A list of valid time-zone settings are displayed. Each time zone is assigned an ID. The time zone and the associated ID are indicated in the list.
3. Issue the following CLI command to set the time zone for the cluster.

   ```
   svctask settimezone -timezone time_zone_setting
   ```

   where *time_zone_setting* is the new time zone ID that you have chosen from the list of time zones that are available on the cluster.
4. Issue the following CLI command to set the time for the cluster:

   ```
   svctask setclustertime -time 031809142005
   ```

   where *031809142005* is the new time that you want to set for the cluster. You must use the *MMDDHHmmYYYY* format to set the time for the cluster.

Viewing and updating license settings using the CLI

You can use the command-line interface (CLI) to view and update your license settings.

SAN Volume Controller provides two license options: Physical Disk Licensing and Capacity Licensing. Perform the following steps to view and update your SAN Volume Controller license settings:

1. Issue the **svcinfo lslicense** CLI command to view the current license settings for the cluster.
2. Issue the **svctask chlicense** CLI command to change the licensed settings of the cluster.

   Attention:

   - License settings are entered when the cluster is first created; do not update the settings unless you have changed your license.
   - To select Physical Disk Licensing, run the **svctask chlicense** command with one or more of the **physical_disks**, **physical_flash**, and **physical_remote** parameters.
   - To select Capacity Licensing, run the **svctask chlicense** command with one or more of the **-flash**, **-remote**, and **-virtualization** parameters.

   For detailed license command usage information, see the *IBM System Storage SAN Volume Controller Command-Line Interface User's Guide*.

Displaying cluster properties using the CLI

You can use the command-line interface (CLI) to display the properties for a cluster.

Perform the following step to display cluster properties:

Issue the **svcinfo lscluster** command to display the properties for a cluster. The following is an example of the command you can issue:

```
svcinfo lscluster -delim : ldcluster-19
```

where *ldcluster-19* is the name of the cluster.

```
IBM_2145:ldcluster-19:admin>svcinfo lscluster -delim : ldcluster-19
id:00000200602052F0
name:ldcluster-19
location:local
partnership:
bandwidth:
total_mdisk_capacity:4205812.3GB
space_in_mdisk_grps:4162044.4GB
space_allocated_to_vdisks:1737236.64GB
total_free_space:2468575.7GB
statistics_status:on
statistics_frequency:10
required_memory:8192
cluster_locale:en_US
time_zone:357 Europe/Athens
code_level:5.1.0.0 (build 16.3.0906260000)
FC_port_speed:2Gb
console_IP:x.xx.xx.xx:xxxx
id_alias:00000200602052F0
gm_link_tolerance:300
gm_inter_cluster_delay_simulation:0
gm_intra_cluster_delay_simulation:0
email_reply:manager@mycompany.com
email_contact:manager
email_contact_primary:01202 123456
email_contact_alternate:
email_contact_location:city
email_state:running
inventory_mail_interval:8
total_vdiskcopy_capacity:2009500.80GB
total_used_capacity:1737157.41GB
total_overallocation:47
total_vdisk_capacity:1322910.92GB
cluster_ntp_IP_address:x.xx.xx.xxx
cluster_isns_IP_address:
iscsi_auth_method:none
iscsi_chap_secret:
auth_service_configured:no
auth_service_enabled:no
auth_service_url:
auth_service_user_name:
auth_service_pwd_set:no
auth_service_cert_set:no
relationship_bandwidth_limit:25
```

Maintaining passwords for the front panel using the CLI

You can use the command-line interface (CLI) to view and change the status of the password reset feature for the SAN Volume Controller front panel.

The cluster superuser password can be reset using the front panel of the configuration node. To meet varying security requirements, this functionality can be enabled or disabled using the CLI.

Complete the following steps to view and change the status of the password reset feature:

1. Issue the svctask setpwdreset CLI command to view and change the status of the password reset feature for the SAN Volume Controller front panel.
2. Record the cluster superuser password because you cannot access the cluster without it.

Re-adding a repaired node to a cluster using the CLI

You can use the command-line interface (CLI) to re-add a failed node back into a cluster after it was repaired.

Before you add a node to a cluster, you must 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. If you are replacing a node and the switch is zoned by worldwide port name (WWPN) rather than by switch port, make sure that the switch is configured such that the node being added is in the same VSAN/zone.

Attention:

1. If you are re-adding a node to the SAN, ensure that you are adding the node to the same I/O group from which it was removed. Failure to do this can result in data corruption. You must 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 the IBM Support Center to add the node back into the cluster without corrupting the data.
2. The LUNs that are presented to the ports on the new node must be the same as the LUNs that are presented to the nodes that currently exist in the cluster. You must ensure that the LUNs are the same before you add the new node to the cluster.
3. LUN masking for each LUN must be identical on all nodes in a cluster. You must ensure that the LUN masking for each LUN is identical before you add the new node to the cluster.
4. You must ensure that the model type of the new node is supported by the SAN Volume Controller software level that is currently installed on the cluster. If the model type is not supported by the SAN Volume Controller software level, upgrade the cluster to a software level that supports the model type of the new node. See the following Web site for the latest supported software levels:

   www.ibm.com/storage/support/2145

Special procedures when adding a node to a cluster

Applications on the host systems direct I/O operations to file systems or logical volumes that are mapped by the operating system to virtual paths (vpaths), which are pseudo disk objects supported by the Subsystem Device Driver (SDD). SDD maintains an association between a VPath and a SAN Volume Controller virtual disk (VDisk). This association uses an identifier (UID) which is unique to the VDisk and is never reused. The UID allows SDD to directly associate vpaths with VDisks.

SDD operates within a protocol stack that 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 node name (WWNN) for the fibre channel node and ports.

If an error occurs, the 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 WWNN and LUN numbers that were previously used.

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

Before you add a node to the cluster, you must 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 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 back-end storage.

If any of the previous conditions are true, the following special procedures apply:

- The node must be added to the same I/O group that it was previously in. You can use the command-line interface (CLI) command svcinfo lsnode or the SAN Volume Controller Console to determine the WWN of the cluster nodes.
- Before you add the node back into the cluster, you must shut down all of the hosts using the cluster. The node must then be added before the hosts are rebooted. If the I/O group information is unavailable or it is inconvenient to shut down and reboot all of the hosts using the cluster, then do the following:
  - On all of the hosts connected to the cluster, unconfigure the fibre-channel adapter device driver, the disk device driver and multipathing driver before you add the node to the cluster.
  - Add the node to the cluster and then reconfigure the fibre-channel adapter device driver, the disk device driver, and multipathing driver.

Scenarios where the special procedures can apply

The following two scenarios describe situations where the special procedures can apply:

- Four nodes of an eight-node cluster have been lost because of the failure of a pair of 2145 UPS or four 2145 UPS-1U. In this case, the four nodes must be added back into the cluster using the CLI command svctask addnode 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 CLI command svctask addnode or the SAN Volume Controller Console.

For 5.1.0 nodes, the SAN Volume Controller automatically re-adds nodes that have failed back to the cluster. If the cluster reports an error for a node missing (error code 1195) and that node has been repaired and restarted, the cluster automatically re-adds the node back into the cluster. This process can take up to 20 minutes, so you can manually re-add the node by completing the following steps:

1. Issue the svcinfo lsnode CLI command to list the nodes that are currently part of the cluster and determine the I/O group for which to add the node.

   The following is an example of the output that is displayed:

```
svcinfo lsnode -delim :

id:name:UPS_serial_number:WWNN:status:IO_group_id:IO_group_name
:config_node:UPS_unique_id:hardware:iscsi_name:iscsi_alias
1:node1:10L3ASH:0000000000000000:offline:0:io_grp0:no:1000000000003206:
8A4:iqn.1986-03.com.ibm:2145.ndihill.node1:
2:node2:10L3ASH:50050768010050B0:online:0:io_grp0:yes:10000000000050B0:
8A4:iqn.1986-03.com.ibm:2145.ndihill.node2:
```

2. Issue the **svcinfo lsnodecandidate** CLI command to list nodes that are not assigned to a cluster and to verify that a second node is added to an I/O group.

   The following is an example of the output that is displayed:

```
svcinfo lsnodecandidate -delim :

id:panel_name:UPS_serial_number:UPS_unique_id:hardware
5005076801000001:000341:10L3ASH:202378101C0D18D8:8A4
5005076801000009:000237:10L3ANF:202378101C0D1796:8A4
50050768010000F4:001245:10L3ANF:202378101C0D1796:8A4
....
```

3. Issue the **svctask addnode** CLI command to add a node to the cluster.

   Important: Each node in an I/O group must be attached to a different uninterruptible power supply.

   The following is an example of the CLI command you can issue to add a node to the cluster using the panel name parameter:

```
svctask addnode -panelname 000237
 -iogrp io_grp0
```

   Where *000237* is the panel name of the node, *io_grp0* is the name of the I/O group that you are adding the node to.

   The following is an example of the CLI command you can issue to add a node to the cluster using the WWNN parameter:

```
svctask addnode -wwnodename 5005076801000001
 -iogrp io_grp1
```

   Where *5005076801000001* is the WWNN of the node, *io_grp1* is the name of the I/O group that you are adding the node to.

4. Issue the svcinfo lsnode CLI command to verify the final configuration.

   The following example shows output that is displayed:

```
svcinfo lsnode -delim :

id:name:UPS_serial_number:WWNN:status:IO_group_id:IO_group_name:config_node:UPS_unique_id:
hardware:iscsi_name:iscsi_alias
1:node1:10L3ASH:0000000000000000:offline:0:io_grp0:no:1000000000003206:
8A4:iqn.1986-03.com.ibm:2145.ndihill.node1:
```

   Record the following information for the new node:

   - Node name
   - Node serial number
   - WWNN
   - IQNs (if using hosts attached using iSCSI connections)
   - All WWPNs
   - I/O group that contains the node

Note: If this command is issued quickly after you have added nodes to the cluster, the status of the nodes might be adding. The status is shown as adding if the process of adding the nodes to the cluster is still in progress. You do not have to wait for the status of all the nodes to be online before you continue with the configuration process.

The nodes have been added to the cluster.

Displaying node properties using the CLI

You can use the command-line interface (CLI) to display node properties.

Perform the following steps to display the node properties:

1. Issue the svcinfo lsnode CLI command to display a concise list of nodes in the cluster.

   The following is an example of the CLI command you can issue to list the nodes in the cluster:

   ```
   svcinfo lsnode -delim :
   ```

   The following is an example of the output that is displayed:

```
id:name:UPS_serial_number:WWNN:status:IO_group_id:IO_group_name:config_node:UPS_unique_id:hardware:iscsi_name:iscsi_alias
1:group1node1:10L3ASH:500507680100002C:online:0:io_grp0:yes:202378101C0D18D8:8A4:iqn.1986-03.com.ibm:2145.cluster1:group1node1:
2:group1node2:10L3ANF:5005076801000009:online:0:io_grp0:no:202378101C0D1796:8A4:iqn.1986-03.com.ibm:2145.cluster1:group1node2:
3:group2node1:10L3ASH:5005076801000001:online:1:io_grp1:no:202378101C0D18D8:8A4:iqn.1986-03.com.ibm:2145.cluster1:group2node1:
4:group2node2:10L3ANF:50050768010000F4:online:1:io_grp1:no:202378101C0D1796:8A4:iqn.1986-03.com.ibm:2145.cluster1:group2node2:
```

2. Issue the svcinfo lsnode CLI command and specify the node ID or name of the node that you want to receive detailed output.

   The following is an example of the CLI command you can issue to list detailed output for a node in the cluster:

   ```
   svcinfo lsnode -delim : group1node1
   ```

   Where *group1node1* is the name of the node for which you want to view detailed output.

   The following is an example of the output that is displayed:

```
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_speed:2GB
port_id:500507680120002C
port_status:active
port_speed:2GB
port_id:500507680130002C
port_status:active
port_speed:2GB
port_id:500507680140003C
port_status:active
port_speed:2GB
hardware:8A4
iscsi_name:iqn.1986-03.com.ibm:2145.ndihill.node2
iscsi_alias
failover_active:no
failover_name:node1
failover_iscsi_name:iqn.1986-03.com.ibm:2145.ndihill.node1
failover_iscsi_alias
```

Discovering MDisks using the CLI

You can use the command-line interface (CLI) to discover managed disks (MDisks).

When back-end controllers are added to the fibre-channel SAN and are included in the same switch zone as a SAN Volume Controller cluster, the cluster automatically discovers the back-end controller and integrates the controller to determine the storage that is presented to the SAN Volume Controller nodes. The SCSI logical units (LUs) that are presented by the back-end controller are displayed as unmanaged MDisks. However, if the configuration of the back-end controller is modified after this has occurred, the SAN Volume Controller cluster might be unaware of these configuration changes. You can request that the SAN Volume Controller cluster rescans the fibre-channel SAN to update the list of unmanaged MDisks.

Note: The automatic discovery that is performed by SAN Volume Controller cluster does not write anything to an unmanaged MDisk. You must instruct the SAN Volume Controller cluster to add an MDisk to an MDisk group or use an MDisk to create an image mode virtual disk (VDisk).

Perform the following steps to discover and then view a list of MDisks:

1. Issue the **svctask detectmdisk** CLI command to manually scan the fibre-channel network. The scan discovers any new MDisks that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

   Notes:

   a. Only issue the **svctask detectmdisk** command when you are sure that all of the disk controller ports are working and correctly

configured in the controller and the SAN zoning. Failure to do this can result in errors that are not reported.

b. Although it might appear that the **detectmdisk** command has completed, extra time might be required for it to run. The **detectmdisk** is asynchronous and returns a prompt while the command continues to run in the background. You can use the **lsdiscoverystatus** command to view the discovery status.

2. When the detection is complete, issue the **svcinfo lsmdiskcandidate** CLI command to show the unmanaged MDisks. These MDisks have not been assigned to an MDisk group.
3. Issue the **svcinfo lsmdisk** CLI command to view all of the MDisks.

You have now seen that the back-end controllers and switches have been set up correctly and that the SAN Volume Controller cluster recognizes the storage that is presented by the back-end controller.

The following example describes a scenario where a single back-end controller is presenting eight SCSI LUs to the SAN Volume Controller cluster:

1. Issue `svctask detectmdisk`.
2. Issue `svcinfo lsmdiskcandidate`.

   The following output is displayed:

```
id
0
1
2
3
4
5
6
7
```

3. Issue `svcinfo lsmdisk -delim : -filtervalue mode=unmanaged`

   The following output is displayed:

```
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
```

Creating MDisk groups using the CLI

You can use the command-line interface (CLI) to create a managed disk (MDisk) group.

Attention: If you add an MDisk to an MDisk group as an MDisk, any data on the MDisk is lost. If you want to keep the data on an MDisk (for example, because you want to import storage that was previously not managed by SAN Volume Controller), you must create image mode virtual disks (VDisks) instead.

Assume that the cluster has been set up and that a back-end controller has been configured to present new storage to the SAN Volume Controller.

If you are using a SAN Volume Controller solid-state drive (SSD) managed disk, ensure that you are familiar with the SSD configuration rules.

This task assumes that you have already launched the SAN Volume Controller Console.

Consider the following factors as you decide how many MDisk groups to create:

- A VDisk can only be created using the storage from one MDisk group. Therefore, if you create small MDisk groups, you might lose the benefits that are provided by virtualization, namely more efficient management of free space and a more evenly distributed workload for better performance.
- If any MDisk in an MDisk group goes offline, all the VDisks in the MDisk group go offline. Therefore you might want to consider using different MDisk groups for different back-end controllers or for different applications.
- If you anticipate regularly adding and removing back-end controllers or storage, this task is made simpler by grouping all the MDisks that are presented by a back-end controller into one MDisk group.
- All the MDisks in an MDisk group should have similar levels of performance or reliability, or both. If an MDisk group contains MDisks with different levels of performance, the performance of the VDisks in this group is limited by the performance of the slowest MDisk. If an MDisk group contains MDisks with different levels of reliability, the reliability of the VDisks in this group is that of the least reliable MDisk in the group.

Note: When you create an MDisk group with a new solid-state drive (SSD), the new SSD is automatically formatted and set to a block size of 512 bytes.

Even with the best planning, circumstances can change and you must reconfigure your MDisk groups after they have been created. The data migration facilities that are provided by the SAN Volume Controller enable you to move data without disrupting I/O.

Choosing an MDisk group extent size

Consider the following factors as you decide the extent size of each new MDisk group:

- You must specify the extent size when you create a new MDisk group.
- You cannot change the extent size later; it must remain constant throughout the lifetime of the MDisk group.
- MDisk groups can have different extent sizes; however, this places restrictions on the use of data migration.
- The choice of extent size affects the maximum size of a VDisk in the MDisk group.

Table 25 on page 235 compares the maximum VDisk capacity for each extent size. The maximum is different for space-efficient VDisks. Because the SAN Volume Controller allocates a whole number of extents to each VDisk that is created, using a larger extent size might increase the amount of storage that is wasted at the end of each VDisk. Larger extent sizes also reduces the ability of the SAN Volume Controller to distribute sequential I/O workloads across many MDisks and therefore can reduce the performance benefits of virtualization.

Table 25. Maximum VDisk capacity by extent size

Extent size (MB)	Maximum VDisk capacity in GB (not space-efficient VDisks)	Maximum VDisk capacity in GB (space-efficient VDisks)
16	2048 (2 TB)	2000
32	4096 (4 TB)	4000
64	8192 (8 TB)	8000
128	16,384 (16 TB)	16,000
256	32,768 (32 TB)	32,000
512	65,536 (64 TB)	65,000
1024	131,072 (128 TB)	130,000
2048	262,144 (256 TB)	260,000

Important: You can specify different extent sizes for different MDisk groups; however, you cannot migrate VDisks between MDisk groups with different extent sizes. If possible, create all your MDisk groups with the same extent size.

Perform the following steps to create an MDisk group:

Issue the svctask mkmdiskgrp CLI command to create an MDisk group. The following is an example of the CLI command you can issue to create an MDisk group:

```
svctask mkmdiskgrp -name maindiskgroup -ext 32
 -mdisk mdsk0:mdsk1:mdsk2:mdsk3
```

where *maindiskgroup* is the name of the MDisk group that you want to create, *32* MB is the size of the extent you want to use, and *mdsk0, mdsk1, mdsk2, mdsk3* are the names of the four MDisks that you want to add to the group.

You created and added MDisks to an MDisk group.

The following example provides a scenario where you want to create an MDisk group, but you do not have any MDisks available to add to the group. You plan to add the MDisks at a later time. You use the svctask mkmdiskgrp CLI command to create the MDisk group *bkpmdiskgroup* and later used the svctask addmdisk CLI command to add *mdsk4, mdsk5, mdsk6, mdsk7* to the MDisk group.

1. Issue `svctask mkmdiskgrp -name bkpmdiskgroup -ext 32`
   where *bkpmdiskgroup* is the name of the MDisk group that you want to create and *32* MB is the size of the extent that you want to use.
2. You find four MDisks that you want to add to the MDisk group.
3. Issue `svctask addmdisk -mdisk mdsk4:mdsk5:mdsk6:mdsk7 bkpdiskgroup`
   where *mdsk4, mdsk5, mdsk6, mdsk7* are the names of the MDisks that you want to add to the MDisk group and *bkpdiskgroup* is the name of the MDisk group for which you want to add MDisks.

Adding MDisks to MDisk groups using the CLI

You can use the command-line interface (CLI) to add managed disks (MDisks) to MDisk groups.

The MDisks must be in unmanaged mode. Disks that already belong to an MDisk group cannot be added to another MDisk group until they have been deleted from their current MDisk group. You can delete an MDisk from an MDisk group under the following circumstances:

- If the MDisk does not contain any extents in use by a virtual disk (VDisk)
- If you can first migrate the extents in use onto other free extents within the group

Important: Do not add an MDisk using this procedure if you are mapping the MDisk to an image mode VDisk. Adding an MDisk to an MDisk group allows the SAN Volume Controller to write new data to the MDisk; therefore, any existing data on the MDisk is lost. If you want to create an image mode VDisk, use the svctask mkvdisk command instead of svctask addmdisk.

If you are using a SAN Volume Controller solid-state drive (SSD) managed disk, ensure that you are familiar with the SSD configuration rules.

When you are adding MDisks to an MDisk group using the svctask addmdisk command or when you are creating an MDisk group using the svctask mkmdiskgrp -mdisk command, the SAN Volume Controller performs tests on the MDisks in the list before the MDisks are allowed to become part of an MDisk group. These tests include checks of the MDisk identity, capacity, status and the ability to perform both read and write operations. If these tests fail or exceed the time allowed, the MDisks are not added to the group. However, with the svctask mkmdiskgrp -mdisk command, the MDisk group is still created even if the tests fail, but it does not contain any MDisks. If tests fail, confirm that the MDisks are in the correct state and that they have been correctly discovered.

The following events contribute to an MDisk test failure:

- The MDisk is not visible to all SAN Volume Controller nodes in the cluster.
- The MDisk identity has changed from a previous discovery operation.
- The MDisk cannot perform read or write operations.
- The status of the MDisk can be either degraded paths, degraded ports, excluded, or offline.
- The MDisk does not exist.

The following events contribute to an MDisk test timeout:

- The disk controller system on which the MDisk resides is failing.
- A SAN fabric or cable fault condition exists that is preventing reliable communication with the MDisk.

Note: The first time that you add a new solid-state drive (SSD) to an MDisk group, the SSD is automatically formatted and set to a block size of 512 bytes.

Perform the following steps to add MDisks to MDisk groups:

1. Issue the **svcinfo lsmdiskgrp** CLI command to list the existing MDisk groups.

   The following is an example of the CLI command you can issue to list the existing MDisk groups:

   ```
   svcinfo lsmdiskgrp -delim :
   ```

   The following is an example of the output that is displayed:

```
id:name:status:mdisk_count:vdisk_count:
capacity:extent_size:free_capacity:virtual_capacity:
used_capacity:real_capacity:overallocation:warning
0:mdiskgrp0:online:3:4:33.3GB:16:32.8GB:64.00MB:64.00MB:64.00MB:0:0
1:mdiskgrp1:online:2:1:26.5GB:16:26.2GB:16.00MB:16.00MB:16.00MB:0:0
2:mdiskgrp2:online:2:0:33.4GB:16:33.4GB:0.00MB:0.00MB:0.00MB:0:0
```

2. Issue the **svctask addmdisk** CLI command to add MDisks to the MDisk group.

   The following is an example of the CLI command you can issue to add MDisks to an MDisk group:

   ```
   svctask addmdisk -mdisk mdisk4:mdisk5:mdisk6:mdisk7 bkpmdiskgroup
   ```

   Where *mdisk4:mdisk5:mdisk6:mdisk7* are the names of the MDisks that you want to add to the MDisk group and *bkpmdiskgroup* is the name of the MDisk group for which you want to add the MDisks.

Locating a solid-state drive (SSD) using the CLI

You can identify the node slot that a solid-state drive (SSD) is located in using the SAN Volume Controller CLI.

To identify the node slot that an SSD is located in, run the lsmdisk command. To identify additional information about an SSD, run the lsnodevpd command.

1. To locate an SSD, run the following command :

   ```
   svcinfo lsmdisk mdisk_name | mdisk_id
   ```

   Where *mdisk_name | mdisk_id* is the name or ID of the MDisk that is an SSD.

   Output is similar to the following example. The `node_id` and `node_name` fields identify the node that the SSD is installed in. The `location` field indicates the drive slot that the SSD is located in. If the field is blank, the MDisk is not an SSD in a SAN Volume Controller node.

   ```
   id 0
   name mdisk0
   status online
   mode managed
   mdisk_grp_id 1
   mdisk_grp_name ssd_n1
   capacity 136.7GB
   quorum_index
   block_size 512
   controller_name controller0
   ctrl_type 6
   ctrl_WWNN 5005076801E00047
   controller_id 0
   path_count 1
   max_path_count 1
   ctrl_LUN_# 0000000000000000
   UID 5000a72000008391000000000000000000\
   00000000000000000000000000000000
   preferred_WWPN 5000A72A00008391
   active_WWPN 5000A72A00008391
   node_id 1
   node_name node1
   location 2
   ```

2. Run the following command to list additional details about the SSD, including the serial number and firmware level.

   ```
   svcinfo lsnodevpd node_name | node_id
   ```

   Where *node_name | node_id* is the name or ID of the node that contains the SSD.

Output is similar to the following example. The `drive_location` identifies the specific SSD that is listed:

```
drive_location 2
manufacturer IBM
model Z16IZD2B-73
capacity 00014337400
serial_number S092901FG008
label_serial_number S092901F
supplier_serial_number S092901F
part_number 41Y8476
firmware_level G008
FPGA_revision F5
type SSD
speed
enclosure
connection_type SAS
```

Collecting SSD dump files using the CLI

You can use the command-line interface (CLI) to collect dump files from solid-state drives (SSDs).

To collect internal log files from solid-state drive (SSD) MDisks, run the triggermdiskdump command. Subsequently, you can list, delete or copy the dump files.

The triggermdiskdump command generates a dump file and saves it in the /dumps/mdisk directory on the node that contains the SSD.

1. Issue the **svctask triggermdiskdump** CLI command.

   The following example shows the CLI format for generating a dump file for the specified SSD MDisk:

   ```
   svctask triggermdiskdump mdisk_id | mdisk_name
   ```

2. Issue the **svcinfo lsmdiskdumps** command to list files in the /dumps/drive directory on the specified node.

   The following example shows the CLI format for listing the dump files for the specified node:

   ```
   svcinfo lsmdiskdumps node_id | node_name
   ```

3. Issue the **svctask cleardumps** command to delete all files from the /dumps directory and all subdirectories on the specified node. To delete files from a subdirectory of /dumps only, specify the -prefix parameter.

   The following example shows the CLI format for deleting all dump files from the specified node:

   ```
   svctask cleardumps node_id | node_name
   ```

   The following example shows the CLI format for deleting only the dump files in the specified /elogs/ directory:

   ```
   svctask cleardumps -prefix "/dumps/elogs/*"
   ```

4. Issue the **svctask cpdumps** command to copy dump files to the configuration node. If the /dumps directory on the configuration node becomes full before the copy completes, no message is returned. To avoid this scenario, clear the /dumps directory after migrating data from the configuration node.

   The following example shows the CLI format for copying all dump files from the specified node to the configuration node:

   ```
   svctask cpdumps -prefix /dumps node_id | node_name
   ```

Setting a quorum disk using the CLI

You can set a managed disk (MDisk) as a quorum disk by using the command-line interface (CLI).

To set the MDisk to a specified quorum index or to set the MDisk as the active quorum disk, use the setquorum command.

- When possible, distribute the quorum candidate disks so that each MDisk is provided by a different storage system. For a list of storage systems that support quorum disks, refer to the Supported Hardware List located at the SAN Volume Controller (2145) Web site:

  www.ibm.com/storage/support/2145
- Ensure that the status of the MDisk that is being set as a quorum disk is online before setting it as the quorum disk. To check the status of the MDisk that is being set as the quorum disk, use the svcinfo lsquorum command.
- Wait at least 20 seconds between issuing a consecutive setquorum command. This delay allows time for the first update to the quorum disk to complete before setting another quorum disk.

Note: Quorum functionality is not supported for solid-state drive (SSD) MDisks.

To set an MDisk as a quorum disk, follow these steps:

1. Issue the following CLI command to set the MDisk that is currently assigned the quorum index number to a nonquorum disk. The cluster automatically assigns quorum indexes.

   ```
   svctask setquorum -quorum quorum_index mdisk_id | mdisk_name
   ```

   where *quorum_index* specifies the particular set of managed disks that you want to continue as a quorum disk to the cluster and *mdisk_id | mdisk_name* is the ID or the name of the MDisk to assign as a quorum candidate disk.
2. Issue the following command to make the specified MDisk the active quorum disk. To identify the quorum disk that is active, use the svcinfo lsquorum command.

   ```
   svctask setquorum -quorum quorum_index -active mdisk_id | mdisk_name
   ```

   where *quorum_index* specifies the quorum index with a value of 0, 1, or 2, and *mdisk_id | mdisk_name* is the ID or the name of the MDisk to assign as a quorum candidate disk. The **active** parameter makes the MDisk that was specified in the previous example the active quorum disk.

Modifying the amount of available memory for Copy Service and VDisk Mirroring features using the CLI

You can use the command-line interface (CLI) to modify the amount of memory that is available for the VDisk Mirroring feature and the FlashCopy, Metro Mirror, or Global Mirror Copy Services features.

Table 26 on page 240 provides an example of the amount of memory that is required for VDisk Mirroring and each Copy Service feature.

Table 26. Memory required for VDisk Mirroring and Copy Services

Feature	Grain size	1 MB of memory provides the following VDisk capacity for the specified I/O group
Metro Mirror or Global Mirror	256 KB	2 TB of total Metro Mirror and Global Mirror VDisk capacity
FlashCopy	256 KB	2 TB of total FlashCopy source VDisk capacity
FlashCopy	64 KB	512 GB of total FlashCopy source VDisk capacity
Incremental FlashCopy	256 KB	1 TB of total incremental FlashCopy source VDisk capacity
Incremental FlashCopy	64 KB	256 GB of total incremental FlashCopy source VDisk capacity
VDisk Mirroring	256 KB	2 TB of mirrored VDisk capacity

Notes:

1. For multiple FlashCopy targets, you must consider the number of mappings. For example, for a mapping with a grain size of 256 KB, 8 KB of memory allows one mapping between a 16 GB source VDisk and a 16 GB target VDisk. Alternatively, for a mapping with a 256 KB grain size, 8 KB of memory allows two mappings between one 8 GB source VDisk and two 8 GB target VDisks.
2. When creating a FlashCopy mapping, if you specify an I/O group other than the I/O group of the source VDisk, the memory accounting goes towards the specified I/O group, not towards the I/O group of the source VDisk.
3. For VDisk Mirroring, the full 512 MB of memory space provides 1 PB of total VDisk Mirroring capacity.
4. In this table, *capacity* refers to the virtual capacity of the VDisk. For space-efficient VDisks with different virtual capacities and real capacities, the virtual capacity is used for memory accounting.

To modify and verify the amount of memory that is available, perform the following steps:

1. Issue the following command to modify the amount of memory that is available for VDisk Mirroring or a Copy Service feature:
   `svctask chiogrp -feature flash|remote|mirror -size memory_size io_group_id | io_group_name`

   where *flash | remote | mirror* is the feature that you want to modify, *memory_size* is the amount of memory that you want to be available, and *io_group_id | io_group_name* is the ID or name of the I/O group for which you want to modify the amount of available memory.
2. Issue the following command to verify that the amount of memory has been modified:
   `svcinfo lsiogrp object_id | object_name`

   where *object_id | object_name* is the ID or name of the I/O group for which you have modified the amount of available memory.

   The following information is an example of the output that is displayed.

```
id 0
name io_grp 0
node_count 2
vdisk_count 28
host_count 2
flash_copy_total_memory 20.0MB
flash_copy_free_memory 20.0MB
remote_copy_total_memory 20.0MB
remote_copy_free_memory 20.0MB
mirroring_total_memory 10.0MB
mirroring_free_memory 10.0MB
```

Creating VDisks using the CLI

You can use the command-line interface (CLI) to create a virtual disk (VDisk).

If the VDisk that you are creating maps to a solid-state drive (SSD), the data that is stored on the VDisk is not protected against SSD failures or node failures. To avoid data loss, add a VDisk copy that maps to an SSD on another node.

This task assumes that the cluster has been set up and that you have created managed disk (MDisk) groups. You can establish an empty MDisk group to hold the MDisks that are used for image mode VDisks.

Note: If you want to keep the data on an MDisk, create image mode VDisks. This task describes how to create a VDisk with striped virtualization.

Perform the following steps to create VDisks:

1. Issue the **svcinfo lsmdiskgrp** CLI command to list the available MDisk groups and the amount of free storage in each group.

   The following is an example of the CLI command you can issue to list MDisk groups:

   ```
   svcinfo lsmdiskgrp -delim :
   ```

   The following is an example of the output that is displayed:

   ```
   id:name:status:mdisk_count:vdisk_count:capacity:extent_size:free_capacity:
   virtual_capacity:used_capacity:real_capacity:overallocation:warning
   0:mdiskgrp0:degraded:4:0:34.2GB:16:34.2GB:0:0:0:0:0
   1:mdiskgrp1:online:4:6:200GB:16:100GB:400GB:75GB:100GB:200:80
   ```

2. Decide which MDisk group you want to provide the storage for the VDisk.
3. Issue the **svcinfo lsiogrp** CLI command to show the I/O groups and the number of VDisks 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. You can use FlashCopy to make copies of VDisks regardless of whether the source and target VDisk are in the same I/O group. If you plan to use intracluster Metro Mirror or Global Mirror, both the master and auxiliary VDisk must be in the same I/O group.

   The following is an example of the CLI command you can issue to list I/O groups:

   ```
   svcinfo lsiogrp -delim :
   ```

   The following is an example of the output that is displayed:

```
id:name:node_count:vdisk_count:host_count
0:io_grp0:2:0:2
1:io_grp1:2:0:1
2:io_grp2:0:0:0
3:io_grp3:0:0:0
4:recovery_io_grp:0:0:0
```

4. Decide which I/O group you want to assign the VDisk 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, 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.
5. Issue the **svctask mkvdisk** CLI command to create a VDisk. See the *IBM System Storage SAN Volume Controller Command-Line Interface User's Guide* for more information on this command.

   The following is an example of the CLI command you can issue to create a VDisk using the I/O group ID and MDisk group ID:

   ```
   svctask mkvdisk -name mainvdisk1 -iogrp 0
    -mdiskgrp 0 -vtype striped -size 256 -unit gb
   ```

   where *mainvdisk1* is the name that you want to call the VDisk, *0* is the ID of the I/O group that want the VDisk to use, *0* is the ID of the MDisk group that you want the VDisk to use, and *256* is the capacity of the VDisk.

   The following is an example of the CLI command that you can issue to create a VDisk using the I/O group and MDisk group name:

   ```
   svctask mkvdisk -name bkpvdisk1 -iogrp io_grp1
   -mdiskgrp bkpmdiskgroup -vtype striped -size 256 -unit gb
   ```

   where *bkpvdisk1* is the name that you want to call the VDisk, *io_grp1* is the name of the I/O group that want the VDisk to use, *bkpmdiskgroup* is the name of the MDisk group that you want the VDisk to use, and *256* is the capacity of the VDisk.

   The following is an example of the the CLI command that you can issue to create a space-efficient VDisk using the I/O group and MDisk group name:

   ```
   svctask mkvdisk  -iogrp io_grp1 -mdiskgrp bkpmdiskgroup -vtype striped
   -size 10 unit gb -rsize 20% -autoexpand -grainsize 32
   ```

   where *io_grp1* is the name of the I/O group that you want the VDisk to use and *20%* is how much real storage to allocate to the VDisk, as a proportion of its virtual size. In this example, the size is 10 GB so that 2 GB will be allocated.

   The following is an example of the CLI command that you can issue to create a VDisk with two copies using the I/O group and MDisk group name:

   ```
   svctask mkvdisk -iogrp io_grp1 -mdiskgrp grpa:grpb
   -size 500 -vtype striped -copies 2
   ```

   where *io_grp1* is the name of the I/O group that you want the VDisk to use, *grpa* is the name of the MDisk group for the primary copy of the VDisk and *grpb* is the name of the MDisk group for the second copy of the VDisk, and 2 is the number of VDisk copies.

   Note: If you want to create two VDisk copies of different types, create the first copy using the mkvdisk command and then add the second copy using the addvdiskcopy command.
6. Issue the **svcinfo lsvdisk** CLI command to list all the VDisks that have been created.

Adding a copy to a VDisk using the CLI

You can use the command-line interface (CLI) to add a mirrored copy to a virtual disk (VDisk). Each VDisk can have a maximum of two copies.

The addvdiskcopy command adds a copy to an existing VDisk, which changes a nonmirrored VDisk into a mirrored VDisk.

Use the **-copies** parameter to specify the number of copies to add to the VDisk; this is currently limited to the default value of **1** copy. Use the **-mdiskgrp** parameter to specify the managed disk group that will provide storage for the copy; the svcinfo lsmdiskgrp CLI command lists the available managed disk groups and the amount of available storage in each group.

For image copies, you must specify the virtualization type using the **-vtype** parameter and an MDisk that is in unmanaged mode using the **-mdisk** parameter. This MDisk must be in the unmanaged mode. The **-vtype** parameter is optional for sequential (seq) and striped VDisks. The default virtualization type is **striped**.

Issue the addvdiskcopy CLI command to add a mirrored copy to a VDisk:

```
svctask addvdiskcopy -mdiskgrp 0 vdisk8
```

where *0* is the name of the managed disk group and *vdisk8* is the VDisk to which the copy will be added.

The command returns the IDs of the newly created VDisk copies.

Deleting a copy from a VDisk using the CLI

You can use the command-line interface (CLI) to delete a mirrored copy from a virtual disk (VDisk).

If you are using solid-state drives (SSDs) that are inside a SAN Volume Controller node, always use VDisk mirroring with these SSDs. Data stored on SSDs inside SAN Volume Controller nodes is not protected against SSD failures or node failures. Therefore, if you are deleting a VDisk copy that uses a SSD, ensure that the data that is stored on the copy is protected with another VDisk copy.

The rmvdiskcopy CLI command deletes the specified copy from the specified VDisk. The command fails if all other copies of the VDisk are not synchronized; in this case, you must specify the **-force** parameter, delete the VDisk, or wait until the copies are synchronized. You must specify the **vdisk_name | vdisk_id** parameter last on the command line.

Issue the rmvdiskcopy CLI command to delete a mirrored copy from a VDisk:

```
svctask rmvdiskcopy -copy 1 vdisk8
```

where *1* is the ID of the copy to delete and *vdisk8* is the virtual disk to delete the copy from.

The command does not return any output.

Configuring host objects using the CLI

You can use command-line interface (CLI) to create host objects.

If you are configuring a host object on a fibre-channel attached host, ensure that you have completed all zone and switch configuration. Also test the configuration to ensure that zoning was created correctly.

If you are configuring a host object on the cluster that uses iSCSI connections, ensure that you have completed the necessary host-system configurations and have configured the cluster for iSCSI connections.

At least one WWPN or iSCSI name must be specified.

Perform the following steps to create host objects:

1. Issue the svctask mkhost CLI command to create a logical host object for a fibre-channel attached host. Assign your worldwide port name (WWPN) for the host bus adapters (HBAs) in the hosts.

   The following is an example of the CLI command that you can issue to create a fibre-channel attached host:

   ```
   svctask mkhost -name new_name -hbawwpn wwpn_list
   ```

   where *new_name* is the name of the host and *wwpn_list* is the WWPN of the HBA.

2. To create an iSCSI-attached host, issue the following CLI command:

   ```
   svctask mkhost -iscsiname iscsi_name_list
   ```

   where *iscsi_name_list* specifies one or more iSCSI qualified names (IQNs) of this host. Up to 16 names can be specified, provided that the command-line limit is not reached. Each name should comply with the iSCSI standard, RFD 3720.

3. To add ports to a fibre-channel attached host, issue the svctask addhostport CLI command.

   For example, issue the following CLI command:

   ```
   svctask addhostport -hbawwpn wwpn_list new_name
   ```

   This command adds another HBA WWPN *wwpn_list* to the host that was created in step 1.

4. To add ports to an iSCSI-attached host, issue the svctask addhostport CLI command.

   For example, issue the following CLI command:

   ```
   svctask addhostport -iscsiname iscsi_name_list new_name
   ```

   where *iscsi_name_list* Specifies the comma-separated list of IQNs to add to the host. This command adds an IQN to the host that was created in step 2.

5. To set the Challenge Handshake Authentication Protocol (CHAP) secret that is used to authenticate the host for iSCSI I/O, issue the svctask chhost CLI command. This secret is shared between the host and the cluster. For example, issue the following CLI command:

   ```
   svctask chhost -chapsecret chap_secret
   ```

   where *chap_secret* is the CHAP secret that is used to authenticate the host for iSCSI I/O. To list the CHAP secret for each host, use the svcinfo lsiscsiauth command. To clear any previously set CHAP secret for a host, use the svctask chhost -nochapsecret command.

Creating VDisk-to-host mappings using the CLI

You can use the command-line interface (CLI) to create virtual disk (VDisk)-to-host mappings.

Perform the following steps to create VDisk-to-host mappings:

Issue the **svctask mkvdiskhostmap** CLI command to create VDisk-to-host mappings.
The following is an example of the CLI command you can issue to create VDisk-to-host mappings:

```
svctask mkvdiskhostmap -host demohost1 mainvdisk1
```

Where *demohost1* is the name of the host and *mainvdisk1* is the name of the VDisk.

Creating FlashCopy mappings using the CLI

You can use the command-line interface (CLI) to create FlashCopy mappings.

A FlashCopy mapping specifies the source and target virtual disk (VDisk). Source VDisks and target VDisks must meet the following requirements:

- They must be the same size.
- They must be managed by the same cluster.

A VDisk can be the source in up to 256 mappings. A mapping is started at the point in time when the copy is required.

Perform the following steps to create FlashCopy mappings:

1. The source and target VDisk must be the exact same size. Issue the **svcinfo lsvdisk -bytes** CLI command to find the size (capacity) of the VDisk in bytes.
2. Issue the **svctask mkfcmap** CLI command to create a FlashCopy mapping.

   The following CLI command example creates a FlashCopy mapping and sets the copy rate:

   ```
   svctask mkfcmap -source mainvdisk1 -target bkpvdisk1
    -name main1copy -copyrate 75
   ```

   Where *mainvdisk1* is the name of the source VDisk, *bkpvdisk1* is the name of the VDisk that you want to make the target VDisk, *main1copy* is the name that you want to call the FlashCopy mapping, and *75* is the copy rate.

   The following is an example of the CLI command you can issue to create FlashCopy mappings without the copy rate parameter:

   ```
   svctask mkfcmap -source mainvdisk2 -target bkpvdisk2
    -name main2copy
   ```

   Where *mainvdisk2* is the name of the source VDisk, *bkpvdisk2* is the name of the VDisk that you want to make the target VDisk, *main2copy* is the name that you want to call the FlashCopy mapping.

   Note: The default copy rate of 50 is used if you do not specify a copy rate.

   If the specified source and target VDisks are also the target and source VDisks of an existing mapping, the mapping that is being created and the existing mapping become partners. If one mapping is created as incremental, its partner is automatically incremental. A mapping can have only one partner.
3. Issue the **svcinfo lsfcmap** CLI command to check the attributes of the FlashCopy mappings that have been created:

The following is an example of a CLI command that you can issue to view the attributes of the FlashCopy mappings:

```
svcinfo lsfcmap -delim :
```

Where **-delim** species the delimiter. The following is an example of the output that is displayed:

```
id:name:source_vdisk_id:source_vdisk_name:target_vdisk_id:target_vdisk_name:
group_id:group_name:status:progress:copy_rate:clean_progress:incremental
0:main1copy:77:vdisk77:78:vdisk78:::idle_or_copied:0:75:100:off
1:main2copy:79:vdisk79:80:vdisk80:::idle_or_copied:0:50:100:off
```

Preparing and starting a FlashCopy mapping using the CLI

Before you start the FlashCopy process using the command-line interface (CLI), you must prepare a FlashCopy mapping.

Starting a FlashCopy mapping creates a point-in-time copy of the data on the source virtual disk (VDisk) and writes it to the target VDisk for the mapping.

Perform the following steps to prepare and start a FlashCopy mapping:

1. Issue the **svctask prestartfcmap** CLI command to prepare the FlashCopy mapping.

   To run the following command, the FlashCopy mapping cannot belong to a consistency group.

   ```
   svctask prestartfcmap -restore main1copy
   ```

   Where *main1copy* is the name of the FlashCopy mapping.

   This command specifies the optional **restore** parameter, which forces the mapping to be prepared even if the target VDisk is being used as a source in another active FlashCopy mapping.

   The mapping enters the preparing state and moves to the prepared state when it is ready.
2. Issue the **svcinfo lsfcmap** CLI command to check the state of the mapping.

   The following is an example of the output that is displayed:

   ```
   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** CLI command to start the FlashCopy mapping.

   The following is an example of the CLI command you can issue to start the FlashCopy mapping:

   ```
   svctask startfcmap -restore main1copy
   ```

   Where *main1copy* is the name of the FlashCopy mapping.

   This command specifies the optional **restore** parameter, which forces the mapping to be started even if the target VDisk is being used as a source in another active FlashCopy mapping.
4. Issue the **svcinfo lsfcmapprogress** CLI command with the FlashCopy mapping name or ID to check the progress of the mapping.

   The following is an example of the output that is displayed; the FlashCopy mapping ID 0 is 47% completed.

```
svcinfo lsfcmapprogress -delim :
id:progress
0:47
```

You have created a point-in-time copy of the data on a source VDisk and written that data to a target VDisk. The data on the target VDisk is only recognized by the hosts that are mapped to it.

Stopping FlashCopy mappings using the CLI

You can use the command-line interface (CLI) to stop a FlashCopy mapping.

Follow these steps to stop a single stand-alone FlashCopy mapping.

1. To stop a FlashCopy mapping, issue the following svctask stopfcmap command:

   `svctask stopfcmap` *`fc_map_id or fc_map_name`*

   where *fc_map_id or fc_map_name* is the ID or name of the mapping to stop.
2. To stop immediately all processing associated with the mapping and break the dependency on the source VDisk of any mappings that are also dependent on the target disk, issue the following command:

   `svctask stopfcmap -force -split` *`fc_map_id or fc_map_name`*

   When you use the **force** parameter, all FlashCopy mappings that depend on this mapping (as listed by the lsfcmapdependentmaps command) are also stopped. The **split** parameter can be specified only when stopping a map that has a progress of 100 as shown by the svcinfo lsfcmap command. The **split** parameter removes the dependency of any other mappings on the source VDisk. It might be used prior to starting another FlashCopy mapping whose target disk is the source disk of the mapping being stopped. After the mapping is stopped with the **split** option, you can start the other mapping without the **restore** option.

Deleting a FlashCopy mapping using the CLI

You can use the command-line interface (CLI) to delete a FlashCopy mapping.

The svctask rmfcmap CLI command deletes an existing mapping if the mapping is in the idle_or_copied or stopped state. If it is in the stopped state, the **-force** parameter is required to specify that the target VDisk is brought online. If the mapping is in any other state, you must stop the mapping before you can delete it.

If deleting the mapping splits the tree that contains the mapping, none of the mappings in either resulting tree can depend on any mapping in the other tree. To display a list of dependent FlashCopy mappings, use the svcinfo lsfcmapdependentmaps command.

1. To delete an existing mapping, issue the rmfcmap CLI command:

   `svctask rmfcmap` *`fc_map_id or fc_map_name`*

   where *fc_map_id or fc_map_name* is the ID or name of the mapping to delete.
2. To delete an existing mapping and bring the target VDisk online, issue the following command:

   `svctask rmfcmap -force` *`fc_map_id or fc_map_name`*

   where *fc_map_id or fc_map_name* is the ID or name of the mapping to delete.

The command does not return any output.

Creating a FlashCopy consistency group and adding mappings using the CLI

You can use the command-line interface (CLI) to create and add mappings to a FlashCopy consistency group.

If you have created several FlashCopy mappings for a group of virtual disks (VDisks) that contain elements of data for the same application, it can be convenient to assign these mappings to a single FlashCopy consistency group. You can then issue a single prepare or start command for the whole group. For example, you can copy all of the files for a database at the same time.

Perform the following steps to add FlashCopy mappings to a new FlashCopy consistency group:

1. Issue the **svctask mkfcconsistgrp** CLI command to create a FlashCopy consistency group.

   The following is an example of the CLI command you can issue to create a FlashCopy consistency group:

   ```
   svctask mkfcconsistgrp -name FCcgrp0 -autodelete
   ```

   Where **FCcgrp0** is the name of the FlashCopy consistency group. The **-autodelete** parameter specifies to delete the consistency group when the last FlashCopy mapping is deleted or removed from the consistency group.
2. Issue the **svcinfo lsfcconsistgrp** CLI command to display the attributes of the group that you have created.

   The following is an example of the CLI command you can issue to display the attributes of a FlashCopy consistency group:

   ```
   svcinfo lsfcconsistgrp -delim : FCcgrp0
   ```

   The following is an example of the output that is displayed:

   ```
   id:1
   name:FCcgrp0
   status:idle_or_copied
   autodelete:on
   FC_mapping_id:0
   FC_mapping_name:fcmap0
   FC_mapping_id:1
   FC_mapping_name:fcmap1
   ```

   Note: For any group that has just been created, the status reported is `empty`
3. Issue the **svctask chfcmap** CLI command to add FlashCopy mappings to the FlashCopy consistency group:

   The following are examples of the CLI commands you can issue to add Flash Copy mappings to the FlashCopy consistency group:

   ```
   svctask chfcmap -consistgrp FCcgrp0 main1copy
   svctask chfcmap -consistgrp FCcgrp0 main2copy
   ```

   Where **FCcgrp0** is the name of the FlashCopy consistence group and *main1copy, main2copy* are the names of the FlashCopy mappings.
4. Issue the **svcinfo lsfcmap** CLI command to display the new attributes of the FlashCopy mappings.

   The following is an example of the output that is displayed:

```
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:28:maindisk1:29:bkpdisk1:1:FCcgrp0:idle_copied::75
1:main2copy:30:maindisk2:31:bkpdisk2:1:FCcgrp0:idle_copied::50
```

5. Issue the **svcinfo lsfcconsistgrp** CLI command to display the detailed attributes of the group.

   The following is an example of a CLI command that you can issue to display detailed attributes:

   ```
   svcinfo lsfcconsistgrp -delim : FCcgrp0
   ```

   Where **FCcgrp0** is the name of the FlashCopy consistency group, and **-delim** specifies the delimiter.

   The following is an example of the output that is displayed:

   ```
   id:1
   name:FCcgrp0
   status:idle_or_copied
   autodelete:off
   FC_mapping_id:0
   FC_mapping_name:main1copy
   FC_mapping_id:1
   FC_mapping_name:main2copy
   ```

Preparing and starting a FlashCopy consistency group using the CLI

You can use the command-line interface (CLI) to prepare and start a FlashCopy consistency group to start the FlashCopy process.

Successful completion of the FlashCopy process creates a point-in-time copy of the data on the source virtual disk (VDisk) and writes it to the target VDisk for each mapping in the group. When you have assigned several mappings to a FlashCopy consistency group, you only have to issue a single prepare command to prepare every FlashCopy mapping in the group, and you only have to issue a single start command to start every FlashCopy mapping in the group.

Perform the following steps to prepare and start a FlashCopy consistency group:

1. Issue the **svctask prestartfcconsistgrp** CLI command to prepare the FlashCopy consistency group before the copy process can be started.

   Remember: You only have to issue a single prepare command for the whole group to prepare all of the mappings simultaneously.

   The following is an example of the CLI command you can issue to prepare the FlashCopy consistency group:

   ```
   svctask prestartfcconsistgrp -restore maintobkpfcopy
   ```

   Where *maintobkpfcopy* is the name of the FlashCopy consistency group.

   The optional **restore** parameter forces the consistency group to be prepared even if the target VDisk of one of the mappings in the consistency group is being used as a source VDisk of another active mapping. An active mapping is in the copying, suspended, or stopping state. The group enters the preparing state, and then moves to the prepared state when it is ready.

2. Issue the **svcinfo lsfcconsistgrp** command to check the status of the FlashCopy consistency group.

The following is an example of the CLI command that you can issue to check the status of the FlashCopy consistency group.

```
svcinfo lsfcconsistgrp -delim :
```

The following is an example of the output that is displayed:

```
id:name:status
1:maintobkpfcopy:prepared
```

3. Issue the **svctask startfcconsistgrp** CLI command to start the FlashCopy consistency group to make the copy.

   Remember: You only have to issue a single start command for the whole group to start all the mappings simultaneously.

   The following is an example of the CLI command that you can issue to start the FlashCopy consistency group mappings:

   ```
   svctask startfcconsistgrp -prep -restore maintobkpfcopy
   ```

   Where *maintobkpfcopy* is the name of the FlashCopy consistency group.

   When you use the **prep** parameter, the system automatically issues the prestartfcconsistgrp command for the group that you specify. When the **restore** parameter is combined with the prep option, you force the consistency group to be started even if the target VDisk of one of the mappings in the consistency group is being used as a source VDisk in another active mapping. An active mapping is in the copying, suspended, or stopping state. The FlashCopy consistency group enters the copying state and returns to the idle_copied state when complete.

4. Issue the **svcinfo lsfcconsistgrp** command to check the status of the FlashCopy consistency group.

   The following is an example of the CLI command that you can issue to check the status of the FlashCopy consistency group:

   ```
   svcinfo lsfcconsistgrp -delim : maintobkpfcopy
   ```

   Where *maintobkpfcopy* is the name of the FlashCopy consistency group.

   The following is an example of the output that is displayed when the process is still copying:

   ```
   id:name:status
   1:maintobkpfcopy:copying
   ```

   The following is an example of the output that is displayed when the process has finished copying:

   ```
   id:1
   name:maintobkpfcopy
   status:idle_copied
   autodelete:off
   FC_mapping_id:0
   FC_mapping_name:main1copy
   FC_mapping_id:1
   FC_mapping_name:main2copy
   ```

Stopping a FlashCopy consistency group using the CLI

You can use the command-line interface (CLI) to stop a FlashCopy consistency group.

The svctask stopfcconsistgrp CLI command stops all processing that is associated with a FlashCopy consistency group that is in one of the following processing states: prepared, copying, stopping, or suspended.

1. To stop a FlashCopy consistency group, issue the stopfcconsistgrp CLI command:

   ```
   svctask stopfcconsistgrp fc_map_id or fc_map_name
   ```

   where *fc_map_id or fc_map_name* is the ID or name of the mapping to delete.
2. To stop a consistency group and break the dependency on the source VDisks of any mappings that are also dependent on the target VDisk, issue the following command:

   ```
   svctask stopfcconsistgrp -split fc_map_id or fc_map_name
   ```

   You can specify the **-split** parameter when all the maps in the group have a progress of 100. It removes the dependency of any other maps on the source VDisks. You can use this option before you start another FlashCopy consistency group whose target disks are the source disks of the mappings that are being stopped. After the consistency group is stopped with the split option, you can start the other consistency group without the restore option

The command does not return any output.

Deleting a FlashCopy consistency group using the CLI

You can use the command-line interface (CLI) to delete a FlashCopy consistency group.

The svctask rmfcconsistgrp CLI command deletes an existing FlashCopy consistency group. The **-force** parameter is required only when the consistency group that you want to delete contains mappings.

1. To delete an existing consistency group that does not contain mappings, issue the svctask rmfcconsistgrp CLI command:

   ```
   svctask rmfcconsistgrp fc_map_id or fc_map_name
   ```

   where *fc_map_id or fc_map_name* is the ID or name of the consistency group to delete.
2. To delete an existing consistency group that contains mappings that are members of the consistency group, issue the following command:

   ```
   svctask rmfcconsistgrp -force fc_map_id or fc_map_name
   ```

   where *fc_map_id or fc_map_name* is the ID or name of the mapping to delete.

   All the mappings that are associated with the consistency group are removed from the group and changed to stand-alone mappings. To delete a single mapping in the consistency group, you must use the svctask rmfcmap command.

The command does not return any output.

Creating Metro Mirror or Global Mirror relationships using the CLI

You can use the command-line interface (CLI) to create Metro Mirror or Global Mirror relationships.

To create Metro Mirror or Global Mirror relationships, perform these steps:

1. To create a Metro Mirror relationship, run the **svctask mkrcrelationship** command. For example, enter:

```
svctask mkrcrelationship -master master_vdisk_id
-aux aux_vdisk_id -cluster cluster_id
```

Where *master_vdisk_id* is the ID of the master VDisk, *aux_vdisk_id* is the ID of the auxiliary VDisk, and *cluster_id* is the ID of the remote cluster.

2. To create a new Global Mirror relationship, run the **svctask mkrcrelationship** command with the **-global** parameter. For example, enter:

```
svctask mkrcrelationship -master master_vdisk_id
-aux aux_vdisk_id -cluster cluster_id -global
```

Where *master_vdisk_id* is the ID of the master VDisk, *aux_vdisk_id* is the ID of the auxiliary VDisk, and *cluster_id* is the ID of the remote cluster.

Modifying Metro Mirror or Global Mirror relationships using the CLI

You can use the command-line interface (CLI) to modify certain attributes of Metro Mirror or Global Mirror relationships. You can change only one attribute at a time for each command submission.

To modify Metro Mirror or Global Mirror relationships, run the **svctask chrcrelationship** command.

1. Run the **svctask chrcrelationship** command to change the name of a Metro Mirror or Global Mirror relationship. For example, to change the relationship name, enter:

```
svctask chrcrelationship -name new_rc_rel_name previous_rc_rel_name
```

Where *new_rc_rel_name* is the new name of the relationship and *previous_rc_rel_name* is the previous name of the relationship.

2. Run the **svctask chrcrelationship** command to remove a relationship from whichever consistency group it is a member of. For example, enter:

```
svctask chrcrelationship -force rc_rel_name/id
```

Where *rc_rel_name/id* is the name or ID of the relationship.

Starting and stopping Metro Mirror or Global Mirror relationships using the CLI

You can use the command-line interface (CLI) to start and stop stand-alone Metro Mirror and Global Mirror relationships. Relationships that are members of consistency groups must be started and stopped using the consistency group CLI commands.

To start and stop Metro Mirror or Global Mirror relationships, perform these steps:

1. To start a Metro Mirror or Global Mirror relationship, run the **svctask startrcrelationship** command. For example, enter:

```
svctask startrcrelationship rc_rel_id
```

Where *rc_rel_id* is the ID of the relationship that you want to start in a stand-alone relationship.

2. To stop a Metro Mirror or Global Mirror relationship, run the **svctask stoprcrelationship** command. This command applies to a stand-alone relationship.

For example, enter:

```
svctask stoprcrelationship rc_rel_id
```

Where *rc_rel_id* is the ID of the stand-alone relationship that you want to stop mirroring I/O.

Displaying the progress of Metro Mirror or Global Mirror relationships using the CLI

You can use the command-line interface (CLI) to display the background copy of Metro Mirror or Global Mirror relationships as a percentage. When the initial background copy process for a relationship has completed, `null` is displayed for the progress of that relationship.

To display the progress of the background copy of Metro Mirror or Global Mirror relationships, run the **lsrcrelationshipprogress** command.

1. To display data progress without headings for columns of data or for each item of data in a Metro Mirror or Global Mirror relationship, run the **lsrcrelationshipprogress -nohdr** command. For example, to display data of the relationship with headings suppressed, enter:

   ```
   lsrcrelationshipprogress -nohdr rc_rel_name
   ```

   Where *rc_rel_name* is the name of the specified object type.
2. To display the progress of a background copy of a Metro Mirror or Global Mirror relationship as a percentage, run the **lsrcrelationshipprogress -delim** command. The colon character (:) separates all items of data in a concise view, and the spacing of columns does not occur. In a detailed view, the data is separated from its header by the specified delimiter. For example, enter:

   ```
   svcinfo lsrcrelationshipprogress -delim : 0
   ```

   The resulting output is displayed, such as in this example:

   ```
   id:progress
   0:58
   ```

Switching Metro Mirror or Global Mirror relationships using the CLI

You can use the command-line interface (CLI) to reverse the roles of primary and secondary virtual disks in a stand-alone Metro Mirror or Global Mirror relationship when that relationship is in a consistent state. Relationships that are members of consistency groups must be switched by using the consistency group CLI commands.

To switch the roles of primary and secondary VDisks in Metro Mirror or Global Mirror relationships, follow these steps:

1. To make the master disk in a Metro Mirror or Global Mirror relationship to be the primary, run the **svctask switchrcrelationship -primary master** command. For example, enter:

   ```
   svctask switchrcrelationship -primary master rc_rel_id
   ```

   Where *rc_rel_id* is the ID of the relationship to switch.
2. To make the auxiliary disk in a Metro Mirror or Global Mirror relationship to be the primary, run the **svctask switchrcrelationship -primary aux** command. For example, enter:

   ```
   svctask switchrcrelationship -primary aux rc_rel_id
   ```

   Where *rc_rel_id* is the ID of the relationship to switch.

Deleting Metro Mirror and Global Mirror relationships using the CLI

You can use the command-line interface (CLI) to delete Metro Mirror and Global Mirror relationships.

To delete Metro Mirror and Global Mirror relationships, run the **svctask rmrcrelationship** command. For example, enter:

```
svctask rmrcrelationship rc_rel_name/id
```

Where *rc_rel_name/id* is the name or ID of the relationship.

Creating Metro Mirror or Global Mirror consistency groups using the CLI

You can use the command-line interface (CLI) to create Metro Mirror or Global Mirror consistency groups.

To create Metro Mirror or Global Mirror consistency groups, perform these steps:

1. To create a Metro Mirror or Global Mirror consistency group, run the **svctask mkrcconsistgrp** command. For example, enter:

   ```
   svctask mkrcconsistgrp -name new_name -cluster cluster_id
   ```

   Where *new_name* is the name of the new consistency group and *cluster_id* is the ID of the remote cluster for the new consistency group. If **-cluster** is not specified, a consistency group is created only on the local cluster. The new consistency group does not contain any relationships and will be in the empty state.
2. To add Metro Mirror or Global Mirror relationships to the group, run the **svctask chrcrelationship** command. For example, enter:

   ```
   svctask chrcrelationship -consistgrp consist_group_name rc_rel_id
   ```

   Where *consist_group_name* is the name of the new consistency group to assign the relationship to and *rc_rel_id* is the ID of the relationship.

Modifying Metro Mirror or Global Mirror consistency groups using the CLI

You can use the command-line interface (CLI) to assign a new name or modify the name of an existing Metro Mirror or Global Mirror consistency group.

To assign or modify the name of a Metro Mirror or Global Mirror consistency group, run the **svctask chrcconsistgrp** command.

1. Run the **svctask chrcconsistgrp** command to assign a new name of a Metro Mirror or Global Mirror consistency group. For example, enter:

   ```
   svctask chrcconsistgrp -name new_name_arg
   ```

   Where *new_name_arg* is the assigned new name of the consistency group.
2. Run the **svctask chrcconsistgrp** command to change the name of the consistency group. For example, enter:

   ```
   svctask chrcconsistgrp -name new_consist_group_name previous_consist_group_name
   ```

   Where *new_consist_group_name* is the assigned new name of the consistency group and *previous_consist_group_name* is the previous name of the consistency group.

Starting and stopping Metro Mirror or Global Mirror consistency-group copy processes using the CLI

You can use the command-line interface (CLI) to start and stop Metro Mirror or Global Mirror consistency-group copy processes.

To start and stop Metro Mirror or Global Mirror consistency-group copy processes, perform these steps:

1. To start a Metro Mirror or Global Mirror consistency-group copy process, set the direction of copy if it is undefined and optionally mark the secondary VDisks of the consistency group as clean. Run the **svctask startrcconsistgrp** command. For example, enter:

   ```
   svctask startrcconsistgrp rc_consist_group_id
   ```

   Where *rc_consist_group_id* is the ID of the consistency group to start processing.
2. To stop the copy process for a Metro Mirror or Global Mirror consistency group, run the **svctask stoprcconsistgrp** command.

   For example, enter:

   ```
   svctask stoprcconsistgrp rc_consist_group_id
   ```

   Where *rc_consist_group_id* is the ID of the consistency group that you want to stop processing.

   If the group is in a consistent state, you can also use this command to enable write access to the secondary virtual disks (VDisks) in the group.

Deleting Metro Mirror or Global Mirror consistency groups using the CLI

You can use the command-line interface (CLI) to delete Metro Mirror or Global Mirror consistency groups.

To delete existing Metro Mirror or Global Mirror consistency groups, follow these steps:

1. To delete a Metro Mirror or Global Mirror consistency group, run the **svctask rmrcconsistgrp** command. For example, enter:

   ```
   svctask rmrcconsistgrp rc_consist_group_id
   ```

   Where *rc_consist_group_id* is the ID of the consistency group to delete.
2. If a Metro Mirror or Global Mirror consistency group is not empty, you must use the **-force** parameter to delete the consistency group. For example, enter:

   ```
   svctask rmrcconsistgrp -force rc_consist_group_id
   ```

   Where *rc_consist_group_id* is the ID of the consistency group to delete. This command causes all relationships that are members of the deleted group to become stand-alone relationships.

Creating Metro Mirror and Global Mirror partnerships using the CLI

You can use the command-line interface (CLI) to create Metro Mirror and Global Mirror partnerships between two clusters.

Perform the following steps to create Metro Mirror and Global Mirror partnerships:

1. To create Metro Mirror and Global Mirror partnerships, run the svctask mkpartnership command. For example, enter:

   ```
   svctask mkpartnership -bandwidth bandwidth_in_mbps remote_cluster_id
   ```

where *bandwidth_in_mbps* specifies the bandwidth (in megabytes per second) that is used by the background copy process between the clusters and *remote_cluster_id* is the ID of the remote cluster.

2. Run the svctask mkpartnership command from the remote cluster. For example, enter:

   ```
   svctask mkpartnership -bandwidth bandwidth_in_mbps local_cluster_id
   ```

   where *bandwidth_in_mbps* specifies the bandwidth (in megabytes per second) that is used by the background copy process between the clusters and *local_cluster_id* is the ID of the local cluster.

Modifying Metro Mirror and Global Mirror partnerships using the CLI

You can use the command-line interface (CLI) to modify Metro Mirror and Global Mirror partnerships.

Perform the following steps to modify Metro Mirror and Global Mirror partnerships:

1. To modify Metro Mirror and Global Mirror partnerships, run the svctask chpartnership command. For example, enter:

   ```
   svctask chpartnership -bandwidth bandwidth_in_mbps remote_cluster_id
   ```

   where *bandwidth_in_mbps* is the new bandwidth (in megabytes per second) from the local cluster to the remote cluster, and *remote_cluster_id* is the ID of the remote cluster.

2. Run the svctask chpartnership command from the remote cluster. For example, enter:

   ```
   svctask chpartnership -bandwidth bandwidth_in_mbps local_cluster_id
   ```

   where *bandwidth_in_mbps* is the new bandwidth (in megabytes per second) from the remote cluster to the local cluster, and *local_cluster_id* is the ID of the local cluster.

Starting and stopping Metro Mirror and Global Mirror partnerships using the CLI

You can use the command-line interface (CLI) to start and stop Metro Mirror and Global Mirror partnerships.

Perform the following steps to start and stop Metro Mirror and Global Mirror partnerships:

1. To start a Metro Mirror or Global Mirror partnership, run the **svctask chpartnership** command from either cluster. For example, enter:

   ```
   svctask chpartnership -start remote_cluster_id
   ```

   Where *remote_cluster_id* is the ID of the remote cluster. The **svctask mkpartnership** command starts the partnership by default.

2. To stop a Metro Mirror or Global Mirror partnership, run the **svctask chpartnership** command from either cluster.

   For example, enter:

   ```
   svctask chpartnership -stop remote_cluster_id
   ```

   Where *remote_cluster_id* is the ID of the remote cluster.

Deleting Metro Mirror and Global Mirror partnerships using the CLI

You can use the command-line interface (CLI) to delete Metro Mirror and Global Mirror partnerships.

Perform the following steps to delete Metro Mirror and Global Mirror partnerships:

1. If a Metro Mirror or Global Mirror partnership has configured relationships or groups, you must stop the partnership before you can delete it. For example, enter:

   ```
   svctask chpartnership -stop remote_cluster_id
   ```

   Where *remote_cluster_id* is the ID of the remote cluster.
2. To delete a Metro Mirror and Global Mirror partnership, run the **svctask rmpartnership** command from either cluster. For example, enter:

   ```
   svctask rmpartnership remote_cluster_id
   ```

   Where *remote_cluster_id* is the ID of the remote cluster.

Determining the WWPNs of a node using the CLI

You can determine the worldwide port names (WWPNs) of a node using the command-line interface (CLI).

Perform the following steps to determine the WWPNs of a node:

1. Issue the **svcinfo lsnode** CLI command to list the nodes in the cluster.
2. Record the name or ID of the node for which you want to determine the WWPNs.
3. Issue the **svcinfo lsnode** CLI command and specify the node name or ID that was recorded in step 2.

   The following is an example of the CLI command you can issue:

   ```
   svcinfo lsnode node1
   ```

   Where *node1* is the name of the node for which you want to determine the WWPNs.
4. Record the four port IDs (WWPNs).

Listing node-dependent VDisks using the CLI

You can use the command-line interface (CLI) to list the virtual disks (VDisks) that are dependent on the status of a node.

If a node goes offline or is removed from a cluster, all VDisks that are dependent on the node go offline. Before taking a node offline or removing a node from a cluster, run the lsnodedependentvdisks command to identify any node-dependent VDisks.

By default, the lsnodedependentvdisks command also checks all available quorum disks. If the quorum disks are accessible only through the specified node, the command returns an error.

Various scenarios can produce node-dependent VDisks. The following examples are common scenarios in which the lsnodedependentvdisks command will return node-dependent VDisks:

1. The node contains solid-state drives (SSDs) and also contains the only synchronized copy of a mirrored VDisk.

2. The node is the only node that can access an MDisk on the SAN fabric.
3. The other node in the I/O group is offline (all VDisks in the I/O group are returned).
4. Pinned data in the cache is stopping the partner node from joining the I/O group.

To resolve (1), allow VDisk mirror synchronizations between SSD MDisks to complete. To resolve (2-4), bring any offline MDisks online and repair any degraded paths.

Note: The command lists the node-dependent VDisks at the time the command is run; subsequent changes to a cluster require running the command again.

1. Issue the **svcinfo lsnodedependentvdisks** CLI command.

   The following example shows the CLI format for listing the VDisks that are dependent on node01:

   ```
   svcinfo lsnodedependentvdisks -node01 :
   ```

   The following example shows the output that is displayed:

   ```
   vdisk_id       vdisk_name
   0              vdisk0
   1              vdisk1
   ```

2. If the **svcinfo lsnodedependentvdisks** command returns an error, you must move your quorum disks to MDisks that are accessible through all nodes. Rerun the command until no errors are returned.
3. Reissue the **svcinfo lsnodedependentvdisks** command. When the command returns no VDisks, the cluster is free from any node-dependent VDisks.

   The following example shows the command syntax for listing the VDisks that are dependent on node01:

   ```
   svcinfo lsnodedependentvdisks -node01 :
   ```

   The following example shows the command output if there are no node-dependent VDisks in the cluster:

   ```
   vdisk_id       vdisk_name
   ```

Determining the VDisk name from the device identifier on the host

You can use the command-line interface (CLI) to determine the virtual disk (VDisk) name from the device identifier on the host.

Each VDisk that is exported by the SAN Volume Controller is assigned a unique device identifier. The device identifier uniquely identifies the VDisk and can be used to determine which VDisk corresponds to the volume that the host sees.

Perform the following steps to determine the VDisk name from the device identifier:

1. Find the device identifier. For example, if you are using the subsystem device driver (SDD), the disk identifier is referred to as the virtual path (vpath) number. You can issue the following SDD command to find the vpath serial number:

   ```
   datapath query device
   ```

   For other multipathing drivers, refer to the documentation that is provided with your multipathing driver to determine the device identifier.
2. Find the host object that is defined to the SAN Volume Controller and corresponds with the host that you are working with.

a. Find the worldwide port numbers (WWPNs) by looking at the device definitions that are stored by your operating system. For example, on AIX the WWPNs are in the ODM and if you use Windows you have to go into the HBA Bios.
b. Verify which host object is defined to the SAN Volume Controller for which these ports belong. The ports are stored as part of the detailed view, so you must list each host by issuing the following CLI command:

   ```
   svcinfo lshost id | name
   ```

   Where *name/id* is the name or ID of the host.
c. Check for matching WWPNs.

3. Issue the following command to list the VDisk-to-host mappings:

   ```
   svcinfo lshostvdiskmap hostname
   ```

   Where *hostname* is the name of the host.
4. Find the VDisk UID that matches the device identifier and record the VDisk name or ID.

Determining the host that a VDisk is mapped to

You can determine the host that a virtual disk (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 want to check.
2. Issue the following CLI command to list the hosts that this VDisk is mapped:

   ```
   svcinfo lsvdiskhostmap vdiskname/id
   ```

   Where *vdiskname/id* is the name or ID of the VDisk.
3. Find 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 virtual disks (VDisks) and managed disks (MDisks) using the command-line interface (CLI).

Select one or more of the following options to determine the relationship between VDisks and MDisks:

- To display a list of the IDs that correspond to the MDisks that comprise the VDisk, issue the following CLI command:

  ```
  svcinfo lsvdiskmember vdiskname/id
  ```

  where *vdiskname/id* is the name or ID of the VDisk.
- To display a list of IDs that correspond to the VDisks that are using this MDisk, issue the following CLI command:

  ```
  svcinfo lsmdiskmember mdiskname/id
  ```

  where *mdiskname/id* is the name or ID of the MDisk.
- To display a table of VDisk IDs and the corresponding number of extents that are being used by each VDisk, issue the following CLI command:

  ```
  svcinfo lsmdiskextent mdiskname/id
  ```

  where *mdiskname/id* is the name or ID of the MDisk.

- To display a table of MDisk IDs and the corresponding number of extents that each MDisk provides as storage for the given VDisk, issue the following CLI command:

  ```
  svcinfo lsvdiskextent vdiskname/id
  ```

  where *vdiskname/id* is the name or ID of the VDisk.

Determining the relationship between MDisks and RAID arrays or LUNs using the CLI

You can determine the relationship between managed disks (MDisks) and RAID arrays or LUNs using the command-line interface (CLI).

Each MDisk corresponds with a single RAID array, or with a single partition on a given RAID array. Each RAID controller defines a LUN number for this disk. The LUN number and controller name or ID are needed 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. Issue the following command to display a detailed view of the MDisk:

   ```
   svcinfo lsmdisk mdiskname
   ```

   Where *mdiskname* is the name of the MDisk for which you want to display a detailed view.
2. Record the controller name or controller ID and the controller LUN number.
3. Issue the following command to display a detailed view of the controller:

   ```
   svcinfo lscontroller controllername
   ```

   Where *controllername* is the name of the controller that you recorded in step 2.
4. Record the vendor ID, product ID, and WWNN. You can use this information to determine what is being presented to the MDisk.
5. From the native user interface for the given controller, list the LUNs it is presenting and match the LUN number with that noted in step 1. This tells you the exact RAID array or partition that corresponds with the MDisk.

Increasing the size of your cluster using the CLI

You can increase throughput by adding more nodes to the cluster. The nodes must be added in pairs and assigned to a new I/O group.

Perform the following steps to increase the size of your cluster:

1. Add a node to your cluster and repeat this step for the second node.
2. If you want to balance the load between the existing I/O groups and the new I/O groups, you can migrate your virtual disks (VDisks) to new I/O groups. Repeat this step for all VDisks that you want to assign to the new I/O group.

Adding a node to increase the size of a cluster using the CLI

You can use the command-line interface (CLI) to increase the size of a cluster by adding two nodes to create a new I/O group.

Attention: If you are adding a node that was previously removed from a cluster, ensure that either the following two conditions have been met:

- The WWPN for the removed node is swapped with the node that replaces it.
- All hosts that accessed the removed node through its WWPNs have been reconfigured to use the WWPN for the new node.

Failure to do either of these action can result in data corruption.

Complete the following steps to add a node and increase the size of a cluster:

1. Install the new nodes and connect to the fibre channel.
2. Issue the following command to verify that the node is detected on the fabric:
   ```
   svcinfo lsnodecandidate
   ```
3. Using the front panel of the node, record the WWNN.
4. Issue the following command to determine the I/O group where the node should be added:
   ```
   svcinfo lsiogrp
   ```
5. Record the 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.

   Note: You only need to do this step for the first node that is added. The second node of the pair uses the same I/O group number.
6. 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.
7. Issue the following command to add the node to the cluster:
   ```
   svctask addnode  -wwnodename WWNN -iogrp newiogrpname/id  [-name newnodename]
   ```
   Where *WWNN* is the WWNN of the node, *newiogrpname/id* is the name or ID of the I/O group that you want to add the node to and *newnodename* is the name that you want to assign to the node. If you do not specify a new node name, a default name is assigned; however, it is recommended you specify a meaningful name.
8. 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.
9. Issue the following command to verify that the node is online:
   ```
   svcinfo lsnode
   ```

You may need to reconfigure your storage systems to allow the new I/O group nodes to access them. If the disk controller uses mapping to present RAID arrays or partitions to the cluster and the WWNNs or the worldwide port names have changed, you must modify the port groups that belong to the cluster.

Migrating a VDisk to a new I/O group using the CLI

You can use the command-line interface (CLI) to migrate a virtual disk (VDisk) to a new I/O group to increase the size of your cluster.

You can migrate a VDisk to a new I/O group to manually balance the workload across the nodes in the cluster. However, you might end up with a pair of nodes that are overworked and another pair that are not 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 is lost while you follow this procedure. Migration commands fail if the target or source VDisk is offline, or if there is insufficient quorum disk space to store the metadata. Correct the offline or quorum disk condition and reissue the command.

Perform the following steps to migrate a single VDisk:

1. Quiesce all I/O operations for the VDisk. You might have to determine the hosts that are using this VDisk in advance.
2. Before migrating the VDisk, it is essential that for each device identifier that is presented by the VDisk you intend to move, the subsystem device driver (SDD) or other multipathing driver configuration is updated to remove the device identifiers. Failure to do this can result in data corruption. See the *IBM System Storage Multipath Subsystem Device Driver User's Guide* or the documentation that is provided with your multipathing driver for details about how to dynamically reconfigure device identifiers for the given host operating system.
3. 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.

   a. Find the **FC_id** and **RC_id** fields. If these are not blank, the VDisk is part of a mapping or relationship.
   b. Stop or delete any FlashCopy mappings, Global Mirror, or Metro Mirror relationships that use this VDisk.
4. Issue the following command to migrate the VDisk:

   ```
   svctask chvdisk -iogrp newiogrpname/id -node preferred_node vdiskname/id
   ```

   where *preferred_node* is the name of the node that you want to move the VDisk, *newiogrpname/id* is the name or ID of the I/O group where you want to migrate the VDisk and *vdiskname/id* is the name or ID of the VDisk that you want to migrate.
5. Discover the new device identifiers and check that each device identifier presents the correct number of paths. See the *IBM System Storage Multipath Subsystem Device Driver User's Guide* or the documentation that is provided with your multipathing driver for details about how to discover device identifiers for the given host operating system.

Validating and repairing mirrored VDisk copies using the CLI

You can use the **repairvdiskcopy** command from the command-line interface (CLI) to validate and repair mirrored VDisk copies.

Attention: Run the **repairvdiskcopy** command only if all VDisk copies are synchronized.

When you issue the **repairvdiskcopy** command, you must use only one of the **-validate**, **-medium**, or **-resync** parameters. You must also specify the name or ID of the VDisk to be validated and repaired as the last entry on the command line. After you issue the command, no output is displayed.

-validate
Use this parameter if you only want to verify that the mirrored VDisk copies are identical. If any difference is found, the command stops and logs an error that includes the logical block address (LBA) and the length of the first difference. You can use this parameter, starting at a different LBA each time to count the number of differences on a VDisk.

-medium
Use this parameter to convert sectors on all VDisk copies that contain different contents into virtual medium errors. Upon completion, the command logs an event, which indicates the number of differences that were found, the number that were converted into medium errors, and the number that were not converted. Use this option if you are unsure what the correct data is, and you do not want an incorrect version of the data to be used.

-resync
Use this parameter to overwrite contents from the specified primary VDisk copy to the other VDisk copy. The command corrects any differing sectors by copying the sectors from the primary copy to the copies being compared. Upon completion, the command process logs an event, which indicates the number of differences that were corrected. Use this action if you are sure that either the primary VDisk copy data is correct or that your host applications can handle incorrect data.

-startlba *lba*
Optionally, use this parameter to specify the starting Logical Block Address (LBA) from which to start the validation and repair. If you previously used the **validate** parameter, an error was logged with the LBA where the first difference, if any, was found. Reissue repairvdiskcopy with that LBA to avoid reprocessing the initial sectors that compared identically. Continue to reissue repairvdiskcopy using this parameter to list all the differences.

Issue the following command to validate and, if necessary, automatically repair mirrored copies of the specified VDisk:

```
svctask repairvdiskcopy -resync -startlba 20 vdisk8
```

Notes:

1. Only one **repairvdiskcopy** command can run on a VDisk at a time.
2. Once you start the **repairvdiskcopy** command, you cannot use the command to stop processing.
3. The primary copy of a mirrored VDisk cannot be changed while the **repairvdiskcopy -resync** command is running.
4. If there is only one mirrored copy, the command returns immediately with an error.
5. If a copy being compared goes offline, the command is halted with an error. The command is not automatically resumed when the copy is brought back online.
6. In the case where one copy is readable but the other copy has a medium error, the command process automatically attempts to fix the medium error by writing the read data from the other copy.
7. If no differing sectors are found during **repairvdiskcopy** processing, an informational error is logged at the end of the process.

Checking the progress of validation and repair of VDisk copies using the CLI

Use the lsrepairvdiskcopyprogress command to display the progress of mirrored VDisk validation and repairs. You can specify a VDisk copy using the **-copy** *id* parameter. To display the VDisks that have two or more copies with an active task, specify the command with no parameters; it is not possible to have only one VDisk copy with an active task.

To check the progress of validation and repair of mirrored VDisks, issue the following command:

```
svcinfo lsrepairvdiskcopyprogress –delim :
```

The following example shows how the command output is displayed:

```
vdisk_id:vdisk_name:copy id:task:progress:estimated_completion_time
0:vdisk0:0:medium:50:070301120000
0:vdisk0:1:medium:50:070301120000
```

Repairing a space-efficient VDisk using the CLI

You can use the **repairsevdiskcopy** command from the command-line interface to repair the metadata on a space-efficient virtual disk (VDisk).

The **repairsevdiskcopy** command automatically detects and repairs corrupted metadata. The command holds the VDisk offline during the repair, but does not prevent the disk from being moved between I/O groups.

If a repair operation completes successfully and the volume was previously offline because of corrupted metadata, the command brings the volume back online. The only limit on the number of concurrent repair operations is the number of virtual disk copies in the configuration.

When you issue the **repairsevdiskcopy** command, you must specify the name or ID of the VDisk to be repaired as the last entry on the command line. Once started, a repair operation cannot be paused or cancelled; the repair can only be terminated by deleting the copy.

Attention: Use this command only to repair a space-efficient VDisk that has reported corrupt metadata.

Issue the following command to repair the metadata on a space-efficient VDisk:

```
svctask repairsevdiskcopy vdisk8
```

After you issue the command, no output is displayed.

Notes:

1. Because the volume is offline to the host, any I/O that is submitted to the volume while it is being repaired fails.
2. When the repair operation completes successfully, the corrupted metadata error is marked as fixed.
3. If the repair operation fails, the volume is held offline and an error is logged.

Checking the progress of the repair of a space-efficient VDisk using the CLI

Issue the lsrepairsevdiskcopyprogress command to list the repair progress for space-efficient VDisk copies of the specified VDisk. If you do not specify a VDisk, the command lists the repair progress for all space-efficient copies in the cluster.

Note: Only run this command after you run the svctask repairsevdiskcopy command, which you must only run as required by the Directed Maintenance Procedures or by IBM support.

Recovering from offline VDisks using the CLI

If a node or an I/O group fails, you can use the command-line interface (CLI) to recover offline virtual disks (VDisks).

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, you must perform one of the following procedures to regain access to your VDisks. Depending on the failure type, you might have lost data that was cached for these VDisks and the VDisks are now offline.

Data loss scenario 1

One node in an I/O group has failed and failover has started on the second node. During the failover process, the second node in the I/O group fails before the data in the write cache is written to hard disk. The first node is successfully repaired but its hardened data is not the most recent version that is committed to the data store; therefore, it cannot be used. The second node is repaired or replaced and has lost its hardened 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 when one node has down-level hardened data and the other node has lost hardened data:

1. Recover the node and add it back into the cluster.
2. Delete all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the offline VDisks.
3. Run the recovervdisk, recovervdiskbyiogrp or recovervdiskbycluster command.
4. Recreate all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the VDisks.

Data loss scenario 2

Both nodes in the I/O group have failed and have been repaired. The nodes have lost their hardened data, therefore, the nodes have no way of recognizing that they are part of the cluster.

Perform the following steps to recover from an offline VDisk when both nodes have lost their hardened data and cannot be recognized by the cluster:

1. Delete all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the offline VDisks.
2. Run the recovervdisk, recovervdiskbyiogrp or recovervdiskbycluster command.
3. Recreate all FlashCopy mappings and Metro Mirror or Global Mirror relationships that use the VDisks.

Recovering a node and returning it to the cluster using the CLI

After a node or an I/O group fails, you can use the command-line interface (CLI) to recover a node and return it to the cluster.

Perform the following steps to recover a node and return it to the cluster:

1. Issue the following command to verify that the node is offline:

   ```
   svcinfo lsnode
   ```

2. Issue the following command to remove the old instance of the offline node from the cluster:

   ```
   svctask rmnode nodename/id
   ```

   Where *nodename/id* is the name or ID of the node.

3. Issue the following command to verify that the node can be seen on the fabric:

   ```
   svcinfo lsnodecandidate
   ```

   Note: Remember the worldwide node names (WWNNs) for each node because you will need them in the following step.

4. If the nodes are repaired by replacing the service controller, or the node is replaced, be sure to follow the replacement instructions for the specific node or controller. You will be instructed to reset the WWNN of the node to that of the original node. If you do not do that, you may need to reconfigure your SAN fabric, your hosts, and your storage systems.

   Attention: If more than one I/O group is affected, ensure that you are adding the node to the same I/O group from which it was removed. 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 the IBM Support Center 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, you must record the following information:

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

5. Issue the following command to add the node back into the cluster:

   ```
   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 and *NODENAME* is the name of the node.

   In a service situation, a node should normally be added back into a cluster using the original node name. As long as the partner node in the I/O group has not been deleted too, this is the default name used if **-name** is not specified.

6. Issue the following command to verify that the node is online:

   ```
   svcinfo lsnode
   ```

Recovering offline VDisks using the CLI

You can recover offline virtual disks (VDisks) using the command-line interface (CLI).

Perform the following steps to recover offline VDisks:

1. Issue the following CLI command to list all VDisks that are offline and belong to an I/O group, enter:

   ```
   svcinfo lsvdisk -filtervalue IO_group_name=
    IOGRPNAME/ID:status=offline
   ```

   where *IOGRPNAME/ID* is the name of the I/O group that failed.
2. To acknowledge data loss for a VDisk with a *fast_write_state* of **corrupt** and bring the VDisk back online, enter:

   ```
   svctask recovervdisk vdisk_id | vdisk_name
   ```

   where *vdisk_id | vdisk_name* is the name or ID of the VDisk.

   Notes:

   - If the specified VDisk is space-efficient or has space-efficient copies, the **recovervdisk** command starts the space-efficient repair process.
   - If the specified VDisk is mirrored, the **recovervdisk** command starts the resynchronization process.
3. To acknowledge data loss for all virtual disks in an I/O group with a *fast_write_state* of **corrupt** and bring them back online, enter:

   ```
   svctask recovervdiskbyiogrp io_group_id | io_group_name
   ```

   where *io_group_id | io_group_name* is the name or ID of the I/O group.

   Notes:

   - If any VDisk is space-efficient or has space-efficient copies, the **recovervdiskbyiogrp** command starts the space-efficient repair process.
   - If any VDisk is mirrored, the **recovervdiskbyiogrp** command starts the resynchronization process.
4. To acknowledge data loss for all VDisks in the cluster with a *fast_write_state* of **corrupt** and bring them back online, enter:

   ```
   svctask recovervdiskbycluster
   ```

   Notes:

   - If any VDisk is space-efficient or has space-efficient copies, the **recovervdiskbycluster** command starts the space-efficient repair process.
   - If any VDisk is mirrored, the **recovervdiskbycluster** command starts the resynchronization process.

Moving offline VDisks to their original I/O group using the CLI

You can move offline virtual disks (VDisks) to their original I/O group using the command-line interface (CLI).

Beginning with SAN Volume Controller 4.3.1, the recovery I/O group is no longer used for VDisk recovery, but it is possible that VDisks were moved to the I/O group before the upgrade.

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: Do not move VDisks to an offline I/O group. Ensure that the I/O group is online before you move the VDisks back to avoid any further data loss.

Perform the following steps to move offline VDisks to their original I/O group:

1. Issue the following command to move the VDisk back into the original I/O group:

   ```
   svctask chvdisk -iogrp IOGRPNAME/ID -force
    vdiskname/ID
   ```

   where *IOGRPNAME/ID* is the name or ID of the original I/O group and *vdiskname/ID* is the name or ID of the offline VDisk.
2. Issue the following command to verify that the VDisks are now online:

   ```
   svcinfo lsvdisk -filtervalue IO_group_name=
    IOGRPNAME/ID
   ```

   where *IOGRPNAME/ID* is the name or ID of the original I/O group.

Informing the SAN Volume Controller of changes to host HBAs using the CLI

You can use the command-line interface (CLI) to inform the SAN Volume Controller of a change to a defined host object.

Because it is sometimes necessary to replace the HBA that connects the host to the SAN, you must inform the SAN Volume Controller of the new worldwide port names (WWPNs) that this HBA contains.

Ensure that your switch is zoned correctly.

Perform the following steps to inform the SAN Volume Controller of a change to a defined host object:

1. Issue the following CLI command to list the candidate HBA ports:

   ```
   svcinfo lshbaportcandidate
   ```

   You should see a list of the HBA ports that are available for addition to host objects. One or more of these HBA ports should correspond with the one or more WWPNs that belong to the new HBA port.
2. Locate the host object that corresponds with the host in which you have replaced the HBA. The following CLI command lists all the defined host objects:

   ```
   svcinfo lshost
   ```
3. Issue the following CLI command to list the WWPNs that are currently assigned to the host object:

   ```
   svcinfo lshost hostobjectname
   ```

   where *hostobjectname* is the name of the host object.
4. Issue the following CLI command to add the new ports to the existing host object:

   ```
   svctask addhostport -hbawwpn one or more existing WWPNs
    separated by : hostobjectname/ID
   ```

where *one or more existing WWPNs separated by :* is the WWPNs that are currently assigned to the host object and *hostobjectname/ID* is the name or ID of the host object.

5. Issue the following CLI command to remove the old ports from the host object:

   ```
   svctask rmhostport -hbawwpn one or more existing WWPNs
    separated by : hostobjectname/ID
   ```

   where *one or more existing WWPNs separated by :* is the WWPNs that are currently assigned to the host object and *hostobjectname/ID* is the name or ID of the host object.

Any mappings that exist between the host object and the virtual disks (VDisks) are automatically applied to the new WWPNs. Therefore, the host sees the VDisks as the same SCSI LUNs as before.

See the *IBM System Storage Multipath Subsystem Device Driver User's Guide* or the documentation that is provided with your multipathing driver for additional information about dynamic reconfiguration.

Expanding VDisks

You can use the command-line interface (CLI) or the SAN Volume Controller Console to expand a virtual disk (VDisk).

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 before you attempt to expand a VDisk that is mapped to a Windows host.

Determine the exact size of the source or master VDisk by issuing the following CLI command:

```
svcinfo lsvdisk -bytes vdiskname
```

where *vdiskname* is the name of the VDisk for which you want to determine the exact size.

VDisks can be expanded under Windows concurrently with I/O operations.

You can expand VDisks for the following reasons:

- To increase the available capacity 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 so that it can be used in a FlashCopy mapping or Metro Mirror relationship.

A VDisk that is not mapped to any hosts and does not contain any data can be expanded at any time. If the VDisk contains data that is in use, you can expand the VDisks if your host has a supported AIX or Microsoft Windows operating system.

The following table provides the supported operating systems and requirements for expanding VDisks that contain data:

Operating system	Supported	Requirement
AIX	Yes	AIX version 5.2 or later
HP-UX	No	-
Linux	No	-
SUN Solaris	No	-
Microsoft Windows NT®	No	-
Microsoft Windows	Yes	Windows version 2000 or later

Expanding a VDisk that is mapped to an AIX host

The SAN Volume Controller supports the ability to dynamically expand the size of a virtual disk (VDisk) if the AIX host is using AIX version 5.2 or later.

The **chvg** command options provide the ability to expand the size of a physical volume that the Logical Volume Manager (LVM) uses, without interruptions to the use or availability of the system. See the *AIX System Management Guide Operating System and Devices* for more information.

Expanding a VDisk that is mapped to a Microsoft Windows host using the CLI

You can use the command-line interface (CLI) to dynamically expand the size of a virtual disk (VDisk) that is mapped to a Microsoft Windows host.

Perform the following steps to expand a VDisk that is mapped to a Windows host:

1. Issue the following CLI command to expand the VDisk:

   ```
   svctask expandvdisksize -size disk_size -unit
    b | kb | mb | gb | tb | pb vdisk_name/vdisk_id
   ```

   where *disk_size* is the capacity by which you want to expand the VDisk, *b | kb | mb | gb | tb | pb* is the data unit to use in conjunction with the capacity and *vdisk_name/vdisk_id* is the name of the VDisk or the ID of the VDisk to expand.
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.

You can expand dynamic disks without stopping I/O operations in most cases. However, in some applications the operating system might report I/O errors. When this problem occurs, either of the following entries might 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 in 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.

If the Computer Management application was open before you expanded the VDisk, use the Computer Management application to issue a rescan command.

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.

Shrinking a virtual disk using the CLI

You can reduce the size of a virtual disk (VDisk) using the command-line interface (CLI).

VDisks can be reduced in size, if it is necessary. You can make a target or auxiliary VDisk the same size as the source or master VDisk when you create FlashCopy® mappings, Metro Mirror relationships, or Global Mirror relationships. However, if the VDisk contains data, do not shrink the size of the disk.

Attention:

1. The SAN Volume Controller arbitrarily reduces the capacity of the VDisk by removing one or more extents from those that are allocated to the VDisk. You cannot control which extents are removed so you cannot guarantee that it is unused space that is removed.
2. If the VDisk contains data that is being used, *do not attempt under any circumstances to shrink a VDisk without first backing up your data*.
3. For performance reasons, some operating systems or file systems use the outer edge of the disk.

You can use the **shrinkvdisksize** command to shrink the physical capacity that is allocated to the particular VDisk by the specified amount. You can also shrink the virtual capacity of a space-efficient VDisk without altering the physical capacity assigned to the VDisk.

For more information about the command parameters, see the *IBM System Storage SAN Volume Controller Command-Line Interface User's Guide*.

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 concurrently 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 highly utilized MDisks.
2. Migrating VDisks from one MDisk group to another. This can be used to remove highly utilized MDisk groups. For example, you can reduce the utilization of a group of MDisks.
3. Migration commands fail if the target or source VDisk is offline, or if there is insufficient quorum disk space to store the metadata. Correct the offline or quorum disk condition and reissue the command.

Notes:

1. The source MDisk must not currently be the source MDisk for any other migrate extents operation.
2. 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 input/output (I/O) statistics about nodes, MDisks, and VDisks. After you have gathered this data, you can analyze it to determine which MDisks are highly utilized. 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 CLI command:

   ```
   svctask startstats -interval 15
   ```

   This command generates a new I/O statistics dump file approximately every 15 minutes.
2. Wait for at least 15 minutes after issuing the **svctask startstats** command and then issue the following command:

   ```
   svcinfo lsiostatsdumps
   ```

   This command lists the I/O statistics files generated on a per-node basis. These are prefixed with `Nm` for MDisk statistics, `Nv` for VDisk statistics, and `Nn` for node statistics.
3. Use secure copy (scp) to retrieve the dumps files for analysis. For example, issue the following AIX CLI command:

   ```
   scp clusterip:/dumps/iostats/m_*
   ```

This command copies all the MDisk statistics files to the AIX host in the current directory.

4. Analyze the dumps to determine which MDisks are highly utilized. You can also determine which VDisks are being highly utilized so that you can spread data more evenly across all the MDisks in the group.
5. Stop the statistics collection by issuing the following command:

   ```
   svctask stopstats
   ```

After you have determined the MDisks that are highly utilized, you can migrate some of the data onto other MDisks within the same MDisk group.

1. Determine the number of extents that are in use by each VDisk for the given MDisk by issuing the following CLI command:

   ```
   svcinfo lsmdiskextent mdiskname
   ```

   This command returns 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 CLI command:

      ```
      svcinfo lsmdisk mdiskname | ID
      ```

   b. List the MDisks in the group by issuing the following CLI 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 CLI 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 does not already own a large set of extents on the target MDisk.
4. For each set of extents, issue the following CLI command to move them to another MDisk:

   ```
   svctask migrateexts -source mdiskname | ID -exts num_extents
   -target newmdiskname | ID -threads 4 vdiskid
   ```

   Where *num_extents* is the number of extents on the *vdiskid*. The *newmdiskname | ID* value is the name or ID of the MDisk to migrate this set of extents to.

   Note: The number of threads indicates the priority of the migration processing, where **1** is the lowest priority and **4** is the highest priority.
5. Repeat the previous steps for each set of extents that you are moving.
6. You can check the progress of the migration by issuing the following CLI command:

   ```
   svcinfo lsmigrate
   ```

Migrating VDisks between MDisk groups using the CLI

You can migrate virtual disks (VDisks) between managed disk (MDisk) groups using the command-line interface (CLI).

You can determine the usage of particular MDisks by gathering input/output (I/O) statistics about nodes, MDisks, and VDisks. After you have gathered this data, you can analyze it to determine which VDisks or MDisks are hot. You can then migrate VDisks from one MDisk group to another.

Perform the following steps to gather statistics about MDisks and VDisks:

1. Isolate any VDisks that are overused. You can determine this by requesting an I/O statistics dump and analyzing the output.

   To start I/O statistics gathering, issue the following CLI command:

   ```
   svctask startstats -interval 15
   ```

   This command generates a new I/O statistics dump file approximately every 15 minutes.
2. Wait for at least 15 minutes after issuing the svctask startstats command and then issue the following command:

   ```
   svcinfo lsiostatsdumps
   ```

   This command lists the I/O statistics files are generated on a per-node basis. These are prefixed with `Nm` for MDisk statistics, `Nv` for VDisk statistics, and `Nn` for node statistics.
3. Use secure copy (**scp** command) to retrieve the dump files for analyzing. For example, issue the following:

   ```
   scp clusterip:/dumps/iostats/v_*
   ```

   This copies all the VDisk statistics files to the AIX host in the current directory.
4. Analyze the dumps to determine which VDisks are hot. It might be helpful to also determine which MDisks are being used heavily as you can spread the data that 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
   ```

After you analyze the I/O statistics data, you can determine which VDisks are hot. You also need to determine the MDisk group that you want to move this VDisk to. Either create a new MDisk group or determine an existing group that is not yet overly used. To do this, check the I/O statistics files that you generated and then ensure that the MDisks or VDisks in the target MDisk group are used less than those in the source group.

You can use data migration or VDisk mirroring to migrate data between MDisk groups. Data migration uses the command **svctask migratevdisk**. VDisk mirroring uses the commands **svctask addvdiskcopy** and **svctask rmvdiskcopy**.

When you issue the **svctask migratevdisk** command, a check is made to ensure that the destination of the migration has enough free extents to satisfy the command. If it does, the command proceeds. The command takes some time to complete.

Notes:

- You cannot use the SAN Volume Controller data migration function to move a VDisk between MDisk groups that have different extent sizes.
- Migration commands fail if the target or source VDisk is offline, or if there is insufficient quorum disk space to store the metadata. Correct the offline or quorum disk condition and reissue the command.

When you use data migration, it is possible for the free destination extents to be consumed by another process; for example, if a new VDisk is created in the destination MDisk group or if more migration commands are started. In this scenario, after all the destination extents are allocated, the migration commands suspend and an error is logged (error ID 020005). To recover from this situation, use either of the following methods:

- Add additional MDisks to the target MDisk group. This provides additional extents in the group and allows the migrations to be restarted. You must mark the error as fixed before you reattempt the migration.
- Migrate one or more VDisks that are already created from the MDisk group to another group. This frees up extents in the group and allows the original migrations to be restarted.

Perform the following steps to use the **svctask migratevdisk** command to migrate VDisks between MDisk groups:

1. After you determine the VDisk that you want to migrate and the new MDisk group you want to migrate it to, issue the following CLI command:

   ```
   svctask migratevdisk -vdisk vdiskname/ID -mdiskgrp
    newmdiskgrname/ID -threads 4
   ```

2. You can check the progress of the migration by issuing the following CLI command:

   ```
   svcinfo lsmigrate
   ```

When you use data migration, the VDisk goes offline if either MDisk group fails. VDisk mirroring can be used to minimize the impact to the VDisk because the VDisk goes offline only if the source MDisk group fails.

Perform the following steps to use VDisk mirroring to migrate VDisks between MDisk groups:

1. After you determine the VDisk that you want to migrate and the new MDisk group that you want to migrate it to, issue the following command:

   ```
   svctask addvdiskcopy -mdiskgrp newmdiskgrname/ID vdiskname/ID
   ```

2. The copy ID of the new copy is returned. The copies now synchronize such that the data is stored in both MDisk groups. You can check the progress of the synchronization by issuing the following command:

   ```
   svcinfo lsvdisksyncprogress
   ```

3. After the synchronization is complete, remove the copy from the original I/O group to free up extents and decrease the utilization of the MDisk group. To remove the original copy, issue the following command:

   ```
   svctask rmvdiskcopy -copy original copy id vdiskname/ID
   ```

Migrating a VDisk between I/O groups using the CLI

Ensure that you are familiar with migrating a virtual disk (VDisk) between I/O groups.

Attention:

- These migration tasks are disruptive. The cached data that is held within the cluster must first be written to disk before the allocation of the VDisk can be changed.
- Migration commands fail if the target or source VDisk is offline, or if there is insufficient quorum disk space to store the metadata. Correct the offline or quorum disk condition and reissue the command.

Modifying the I/O group that services the VDisk cannot be done concurrently with I/O operations. It also requires a rescan at the host level to ensure that the multipathing driver is notified that the allocation of the preferred node has changed and the ports by which the VDisk 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. Synchronize all file systems that are mounted on the given VDisk.
2. Stop all I/O operations to the VDisk.
3. Issue the following CLI command to migrate the VDisk into a new I/O group:

   ```
   svctask chvdisk -iogrp iogrp_name_or_id -node preferred_node vdisk
   ```

   where *iogrp_name_or_id* is the name or ID of the I/O group that you want to migrate the VDisk to, *preferred_node* is the name of the node that you want to move the VDisk to, and *vdisk* is the name of the VDisk that you want to migrate.
4. Resynchronize the VDisk to host mapping. See the *IBM System Storage Multipath Subsystem Device Driver User's Guide* or the documentation that is provided with your multipathing driver for more information.
5. Restart the I/O operations to the VDisk.

Creating an image mode VDisk using the CLI

You can use the command-line interface (CLI) to import storage that contains existing data and continue to use this storage. You can also 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 MDisks to the cluster by adding these disks one at a time. For example, map a single LUN from your RAID controller to the cluster and refresh the view of MDisks. The newly detected MDisk is displayed.
2. *Do not* manually add an unmanaged-mode MDisk that contains existing data to an MDisk group. If you do, the data is lost. When you use the command to convert an image mode VDisk from an unmanaged-mode disk, you will select the MDisk group where it should be added.

See the following Web site for more information:

www.ibm.com/storage/support/2145

Perform the following steps to create an image mode VDisk:

1. Stop all I/O operations from the hosts. Unmap the logical disks that contain the data from the hosts.
2. Create one or more MDisk groups.
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 appears as an unmanaged-mode MDisk to the SAN Volume Controller.
4. Issue 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. Issue the **svctask detectmdisk** command to scan the fibre-channel network for the unmanaged-mode MDisks.

   Note: The **svctask detectmdisk** command also rebalances MDisk access across the available controller device ports.
5. Convert the unmanaged-mode MDisk to an image mode virtual disk.

   Note: If the VDisk that you are converting maps to a solid-state drive (SSD), the data that is stored on the VDisk is not protected against SSD failures or node failures. To avoid data loss, add a VDisk copy that maps to an SSD on another node.

   Issue the **svctask mkvdisk** command to create an image mode virtual disk object.
6. Map the new VDisk to the hosts that were previously using the data that the MDisk now contains. You can use the **svctask mkvdiskhostmap** command to create a new mapping between a VDisk and a host. This makes the image mode VDisk accessible for I/O operations to the host.

After the VDisk is mapped to a host object, the VDisk 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 VDisk. Migrate the data on the image mode VDisk to managed-mode disks in another MDisk group. Issue the **svctask migratevdisk** command to migrate an entire image mode VDisk from one MDisk group to another MDisk group.

Migrating to an image mode virtual disk using the CLI

You can use the command-line interface (CLI) to migrate data to an image mode virtual disk (VDisk).

The **svctask migratetoimage** CLI command allows you to migrate the data from an existing VDisk onto a different managed disk (MDisk).

When the **svctask migratetoimage** CLI command is issued, it migrates the data of the user specified source VDisk onto the specified target MDisk. When the command completes, the VDisk is classified as an image mode VDisk.

Note: Migration commands fail if the target or source VDisk is offline, or if there is insufficient quorum disk space to store the metadata. Correct the offline or quorum disk condition and reissue the command.

The MDisk specified as the target must be in an unmanaged state at the time the command is run. Issuing this command results in the inclusion of the MDisk into the user specified MDisk group.

Issue the following CLI command to migrate data to an image mode VDisk:

```
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.

Deleting a node from a cluster using the CLI

You can use the command-line interface (CLI) to remove a node from a cluster.

After the node is deleted, the other node in the I/O group enters write-through mode until another node is added back into the I/O group.

By default, the **rmnode** command flushes the cache on the specified node before taking the node offline. When operating in a degraded state, the SAN Volume Controller ensures that data loss does not occur as a result of deleting the only node with the cache data.

Attention:

- If you are removing a single node and the remaining node in the I/O group is online, the data can be exposed to a single point of failure if the remaining node fails.
- If both nodes in the I/O group are online and the VDisks are already degraded prior to deleting the node, redundancy to the VDisks is already degraded and loss of access to data and loss of data might occur if the **-force** option is used.
- Removing the last node in the cluster destroys the cluster. Before you delete the last node in the cluster, ensure that you want to destroy the cluster.
- To take the specified node offline immediately without flushing the cache or ensuring data loss does not occur, run the **rmnode** command with the **-force** parameter. The **force** parameter forces continuation of the command even though any node-dependent VDisks will be taken offline. Use the **force** parameter with caution; access to data on node-dependent VDisks will be lost.

Perform the following steps to delete a node:

1. If you are deleting the last node in an I/O group, determine the VDisks that are still assigned to this I/O group:
   a. Issue the following CLI command to request a filtered view of the VDisks:
      ```
      svcinfo lsvdisk -filtervalue IO_group_name=name
      ```
      Where *name* is the name of the I/O group.
   b. Issue the following CLI command to list the hosts that this VDisk is mapped to:
      ```
      svcinfo lsvdiskhostmap vdiskname/id
      ```
      Where *vdiskname/id* is the name or ID of the VDisk.

- If VDisks are assigned to this I/O group that contain data that you want to continue to access, back up the data or migrate the VDisks to a different (online) I/O group.

2. If this is *not* the last node in the cluster, turn the power off to the node that you intend to remove. This step ensures that the multipathing device driver, such as the subsystem device driver (SDD), does not rediscover the paths that are manually removed before you issue the delete node request.

   Attention:

   a. If you are removing the configuration node, the **rmnode** command causes the configuration node to move to a different node within the cluster. This process may take a short time, typically less than a minute. The cluster IP address remains unchanged, but any SSH client attached to the configuration node might need to reestablish a connection. The SAN Volume Controller Console reattaches to the new configuration node transparently.
   b. If you turn on the power to the node that has been removed and it is still connected to the same fabric or zone, it attempts to rejoin the cluster. At this point, the cluster causes the node to remove itself from the cluster and the node becomes a candidate for addition to this cluster or another cluster.
   c. If you are adding this node into the cluster, ensure that you add it to the same I/O group that it was previously a member of. Failure to do so can result in data corruption.
   d. In a service situation, a node should normally be added back into a cluster using the original node name. As long as the partner node in the I/O group has not been deleted too, this is the default name used if **-name** is not specified.

3. Before you delete the node, update the multipathing device driver configuration on the host to remove all device identifiers that are presented by the VDisk that you intend to remove. If you are using the subsystem device driver (SDD), the device identifiers are referred to as virtual paths (vpaths).

   Attention: Failure to perform this step can result in data corruption.

   See the *IBM System Storage Multipath Subsystem Device Driver User's Guide* for details about how to dynamically reconfigure SDD for the given host operating system.

4. Issue the following CLI command to delete a node from the cluster:

   Attention: Before you delete the node: The rmnode command checks for node-dependent VDisks, which are not mirrored at the time that the command is run. If any node-dependent VDisks are found, the command stops and returns a message. To continue removing the node despite the potential loss of data, run the rmnode command with the -force parameter. Alternatively, follow these steps before you remove the node to ensure that all VDisks are mirrored:

   a. Run the lsnodedependentvdisks command.
   b. For each node-dependent VDisk that is returned, run the lsvdisk command.
   c. Ensure that each VDisk returns in-sync status.

   ```
   svctask rmnode node_name_or_id
   ```

   Where *node_name_or_id* is the name or ID of the node.

   Note: Before removing a node, the command checks for any node-dependent VDisks that would go offline. If the node that you selected to delete contains a solid-state drive (SSD) that has dependent VDisks, VDisks that use the SSDs go offline and become unavailable if the node is deleted. To maintain access to VDisk data, mirror these VDisks before

removing the node. To continue removing the node without mirroring the VDisks, specify the **force** parameter.

Performing the cluster maintenance procedure using the CLI

You can use the command-line interface (CLI) to perform the cluster maintenance procedure.

Perform the following steps for cluster maintenance:

1. 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.
2. Issue the `svctask dumperrlog` command to dump the contents of the error log to a text file.
3. Locate and fix the error.
4. Issue the `svctask clearerrlog` command to clear all entries from the error log, including status events and any unfixed errors. Only issue 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.

   Note: Clearing the error log does not fix the errors.
5. Issue the `svctask cherrstate` command to toggle the state of an error between `unfixed` and `fixed`.

Modifying the cluster IP addresses using the CLI

You can use the command-line interface (CLI) to change the IP addresses that are associated with a cluster.

Attention: When you specify a new IP address for a cluster, the existing communication with the cluster is broken. You must reconnect to the cluster with the new IP address.

Perform the following steps to change the cluster IP address:

1. Issue the svcinfo lsclusterip command to list the current IP addresses that are used by the cluster.
2. Record the current IP addresses for future reference.
3. To change an IPv4 cluster IP address, issue the following command:

   `svctask chclusterip -clusterip` *`cluster_ip_address`* `-port` *`cluster_port`*

   where *cluster_ip_address* is the new IP address for the cluster and *cluster_port* specifies which port (1 or 2) to apply changes to.
4. To change an IPv4 cluster IP address to an IPv6 cluster IP address, issue the following command:

   `svctask chclusterip -clusterip_6` *`cluster_ip_address`* `-port` *`cluster_port`*

   where *cluster_ip_address* is the new IPv6 address for the cluster and *cluster_port* specifies which port (1 or 2) to apply changes to.
5. To change an IPv4 default gateway IP address, issue the following command:

   `svctask chclusterip -gw` *`cluster_gateway_address`* `-port` *`cluster_port`*

   where *cluster_gateway_address* is the new gateway address for the cluster and *cluster_port* specifies which port (1 or 2) to apply changes to.
6. To change an IPv6 default gateway address, issue the following command:

```
svctask chclusterip -gw_6 cluster_gateway_address -port cluster_port
```

where *cluster_gateway_address* is the new gateway address for the cluster and *cluster_port* specifies which port (1 or 2) to apply changes to.

7. Issue the following command to change an IPv4 cluster subnet mask:

```
svctask chclusterip -mask cluster_subnet_mask -port cluster_port
```

where *cluster_subnet_mask* is the new subnet mask for the cluster and *cluster_port* specifies which port (1 or 2) to apply changes to.

8. For IPv6 addresses, you can issue the following command to set the prefix for the cluster:

```
svctask chclusterip -prefix_6 -port cluster_port
```

where *cluster_port* specifies which port (1 or 2) to apply changes to.

9. Optionally, if you want to delete all of the IPv4 addresses in the cluster after you have changed all addresses to IPv6, issue the following command:

```
svctask chcluster -noip
```

10. Optionally, if you want to delete all of the IPv6 addresses in the cluster after you have changed all addresses to IPv4, issue the following command:

```
svctask chcluster -noip_6
```

11. The IP routing table provides details of the gateway that is used for IP traffic to a range of IP addresses for each Ethernet port. This information can be used to diagnose configuration node accessibility problems. To display the IP routing table, enter the following CLI command:

```
svcinfo lsroute
```

12. The ping command can be used to diagnose IP configuration problems by checking whether a given IP address is accessible from the configuration node. The command can be useful for diagnosing problems where the configuration node cannot be reached from a specific management server. For example, enter the following CLI command:

```
svctask ping ipv4_address | ipv6_address
```

where *ipv4_address | ipv6_address* is either the IPv4 address or the IPv6 address.

Changing the cluster gateway address using the CLI

You can use the command-line interface (CLI) to change the gateway address for a cluster.

Perform the following steps to change the cluster gateway address:

1. Issue the svcinfo lsclusterip command to list the current gateway address of the cluster.
2. Record the current gateway address for future reference.
3. Issue the following command to change an IPv4 cluster gateway address:

```
svctask chclusterip -gw cluster_gateway_address -port cluster_port
```

where *cluster_gateway_address* is the new gateway address for the cluster. The **port** parameter specifies which port (1 or 2) to apply changes to.

4. Issue the following command to change an IPv6 cluster gateway address:

```
svctask chclusterip -gw_6 cluster_gateway_address -port cluster_port
```

where *cluster_gateway_address* is the new gateway address for the cluster. The **port** parameter specifies which port (1 or 2) to apply changes to.

Changing the relationship bandwidth for a cluster using the CLI

You can use the command-line interface (CLI) to change the relationship bandwidth for a cluster.

The relationship bandwidth limit controls the maximum rate at which any one remote-copy relationship can synchronize. The overall limit is controlled by the **bandwidth** parameter of each cluster partnership. The default value for the relationship bandwidth limit is 25 megabytes per second (MBps), but you can change this by following these steps:

1. Issue the svcinfo lscluster command to list the current relationship bandwidth limit of the cluster. For example:

   ```
   svcinfo lscluster cluster_id_or_cluster_name
   ```

   Where *cluster_id_or_cluster_name* is the ID or name of the cluster.
2. For future reference, record the current relationship bandwidth limit that is displayed. For example: `relationship_bandwidth_limit 25`
3. To change the relationship bandwidth limit of the cluster, issue the following command:

   ```
   svctask chcluster -relationshipbandwidthlimit
    cluster_relationship_bandwidth_limit
   ```

   Where *cluster_relationship_bandwidth_limit* is the new limit for the cluster. Issue the command on both clusters in a relationship.

Configuring the cluster for iSCSI using the CLI

You need to complete several tasks to configure the cluster to work with iSCSI-attached hosts. The tasks include general tasks on the host system before you configure cluster on the SAN Volume Controller.

Before completing any iSCSI-configuration tasks on the cluster, it is important that you complete all the necessary iSCSI-related configuration on the host. Because the SAN Volume Controller supports a variety of host systems, consult the documentation for specific instructions and requirements for a particular host. For a list of supported hosts, see the Support for SAN Volume Controller (2145) Web site:

www.ibm.com/storage/support/2145

To configure a cluster for iSCSI, follow these general tasks on the host system:

1. Select a software-based iSCSI initiator, such as Microsoft Windows iSCSI Software Initiator and verify the iSCSI driver installation.
2. If required, install and configure a multipathing driver for the host system.

In addition, determine a naming convention for iSCSI names such as iSCSI qualified names (IQNs) for the iSCSI initiators that you are planning for the SAN Volume Controller cluster. Hosts use iSCSI names to connect to the SAN Volume Controller node. Each node, for example, has a unique IQN, and the cluster name and node name are used as part of that IQN.

Port IP addresses are the IP addresses that are used by iSCSI-attached hosts to perform I/O.

1. To configure a new port IP address to a specified Ethernet port of a node with an IPv4 address, enter the following command-line interface (CLI) command:

```
svctask cfgportip -node node_name | node_id -ip ipv4addr
-gw ipv4gw -mask subnet_mask -failover port_id
```

where *node_name* | *node_id* specifies the name or ID of the node that is being configured, *ipv4addr* is the IPv4 address for the Ethernet port, *ipv4gw* is the IPv4 gateway IP address, *subnet_mask* is the IPv4 subnet mask, and *port_id* specifies the Ethernet port ID (1 or 2). To view a list of ports, use the svcinfo lsportip command.

The optional **-failover** parameter specifies that the data is failover data, which is data that is related to the partner node. If the node that is specified is the only online node in the I/O group, the address is configured and presented by this node. When another node in the I/O group comes online, the failover address is presented by that node. If two nodes in the I/O group are online when the command is issued, the address is presented by the other node to that specified.

2. To configure a new port IP address that belongs to a partner node with an IPv6 address in the I/O group, enter the following CLI command:

```
svctask cfgportip -node node_name | node_id -ip_6 ipv6addr
-gw_6 ipv6gw -prefix_6 prefix -failover port_id
```

where *node_name* | *node_id* specifies the name or ID of the node that is being configured, *ipv6addr* is the IPv6 address for the iSCSI port, *ipv6gw* if the gateway address for the given IP address, *prefix* is the IPv6 prefix for the gateway, and *port_id* specifies the Ethernet port ID (1 or 2). To view a list of ports, use the svcinfo lsportip command. If the partner node is offline, the address is configured and presented by this node. When another node comes online in the I/O group, the failover address is presented by that node.

The optional **-failover** parameter specifies that the data is failover data, which is data that is related to the partner node. If the node that is specified is the only online node in the I/O group, the address is configured and presented by this node. When another node in the I/O group comes online, the failover address is presented by that node. If two nodes in the I/O group are online when the command is issued, the address is presented by the other node to that specified.

3. To remove an iSCSI IP address from a node Ethernet port, enter either of these CLI commands. The following command deletes an IPv4 configuration for the specified iSCSI Ethernet port:

```
svctask rmportip -failover
-node node_name | node_id port_id
```

where *node_name* | *node_id* specifies the name or ID of the node with the Ethernet port that the IP address is being removed from and *port_id* specifies the Ethernet port ID. To list the valid values for the Ethernet port, enter the svcinfo lsportip command. The optional **-failover** parameter indicates that the specified data is failover data.

The following command deletes an IPv6 configuration for the specified iSCSI Ethernet port:

```
svctask rmportip -ip_6 -failover
-node node_name | node_id port_id
```

where **-ip_6** indicates that this command will remove an IPv6 configuration, *node_name* | *node_id* specifies the name or ID of the node with the Ethernet port that the IP address is being removed from, and *port_id* specifies the Ethernet port ID. To list the valid values for the Ethernet port, enter the svcinfo lsportip command. The optional **-failover** parameter indicates that the specified data is failover data.

After you configure your IP addresses, you can optionally create iSCSI aliases.

Configuring or modifying an iSCSI alias using the CLI

You can use the command-line interface (CLI) to optionally create or change the iSCSI alias for the selected node. An iSCSI alias is a user-assigned name that identifies the SAN Volume Controller node to the iSCSI-attached host. It is possible to change the iSCSI alias of a node even if the node is offline.

To configure or modify an iSCSI alias, follow these steps:

1. To configure a new port IP address to a specified Ethernet port of a node, enter the following CLI command:

   ```
   svctask chnode -iscsialias  alias node_name | node_id
   ```

   where *alias node_name | node_id* specifies the name or ID of the node.
2. To specify that the name or iSCSI alias that is being set is the name or alias of the partner node in the I/O group, enter the following CLI command. When there is no partner node, the values set are applied to the partner node when it is added to the cluster. If this parameter is used when there is a partner node, the name or alias of that node changes

   ```
   svctask chnode -iscsialias alias -failover node_name | node_id
   ```

   where *alias* specifies the iSCSI name of the node and *node_name | node_id* specifies the node to be modified.

After you create iSCSI aliases, you can optionally configure the address for the Internet Storage Name Service (iSNS) server for the cluster.

Configuring the iSNS server address using the CLI

If you are using iSCSI-attached hosts with the SAN Volume Controller cluster, you can the command-line interface (CLI) to optionally configure the address for the Internet Storage Name Service (iSNS) server for the cluster. Host systems use the iSNS server to manage iSCSI targets and for iSCSI discovery.

1. To specify an IPv4 address for the iSCSI storage name service (SNS), enter the following CLI command:

   ```
   svctask chcluster -isnsip sns_server_address
   ```

   where *sns_server_address* is the IP address of the iSCSI storage name service in IPv4 format.
2. To specify an IPv6 address for the iSCSI storage name service (SNS), enter the following CLI command:

   ```
   svctask chcluster -isnsip_6 ipv6_sns_server_address
   ```

   where *ipv6_sns_server_address* is the IP address of the iSCSI storage name service in IPv6 format.

After you configure the iSNS server address for the cluster, you can configure cluster iSCSI authentication.

Note: To help in problem determination, this step can be delayed until after the first one or two hosts have been configured and their connectivity has been tested without authentication configured.

Configuring cluster iSCSI authentication using the CLI

You can use the command-line interface (CLI) to configure the Challenge-Handshake Authentication Protocol (CHAP) to authenticate the SAN Volume Controller cluster to the iSCSI-attached hosts. After the CHAP is set for the

cluster, all attached hosts must be configured to authenticate this way. To help in problem determination, this step can be delayed until after the first one or two hosts have been configured and their connectivity has been tested without authentication configured.

To configure authentication between the SAN Volume Controller cluster and the iSCSI-attached hosts, follow these steps:

1. To set the authentication method for the iSCSI communications of the cluster, enter the following CLI command:

   ```
   svctask chcluster -iscsiauthmethod chap -chapsecret chap_secret
   ```

   where *chap* sets the authentication method for the iSCSI communications of the cluster and *chap_secret* sets the CHAP secret to be used to authenticate the cluster via iSCSI. This parameter is required if the **iscsiauthmethod chap** parameter is specified. The specified CHAP secret cannot begin or end with a space.
2. To clear any previously set CHAP secret for iSCSI authentication, enter the following CLI command:

   ```
   svctask chcluster -nochapsecret
   ```

   The **nochapsecret** parameter is not allowed if the **chapsecret** parameter is specified.
3. The lsiscsiauth command lists the Challenge Handshake Authentication Protocol (CHAP) secret that is configured for authenticating an entity to the SAN Volume Controller cluster. The command also displays the configured iSCSI authentication method. For example, enter the following CLI command:

   ```
   svcinfo lsiscsiauth
   ```

After you configure the CHAP secret for the SAN Volume Controller cluster, ensure that the cluster CHAP secret is added to each iSCSI-attached host. On all iSCSI-attached hosts, specify a CHAP secret that the hosts use to authenticate to the SAN Volume Controller cluster.

Configuring remote authentication service using CLI

You can use the command-line interface (CLI) to configure the SAN Volume Controller to use remote authentication.

To use the SAN Volume Controller with a remote authentication service, follow these steps:

1. Configure the cluster with the location of the remote authentication server.

   To change settings, issue the svctask chauthservice command. To view settings, issue the svcinfo lscluster command.

   You can use either an `http` or `https` connection to the server. If you use an `http` option, the user and password information is transmitted in clear text over the IP network.
2. Configure user groups on the cluster by matching those that are used by the authentication service.

   For each group of interest known to the authentication service, a SAN Volume Controller user group must be created with the same name and with the remote setting enabled. If, for example, members of a group called *sysadmins* require the SAN Volume Controller Administrator (admin) role, issue the following command:

```
svctask mkusergrp -name sysadmins -remote -role Administrator
```

If none of the groups for a user match any of the SAN Volume Controller user groups, the user is not permitted to access the cluster.

3. Configure users who do not require Secure Shell (SSH) access.

   SAN Volume Controller users who are to use the remote authentication service and do not require SSH access should be deleted from the system. The superuser cannot be deleted and cannot use the remote authentication service.

4. Configure users who require SSH access.

   All SAN Volume Controller users who are to use the remote authentication service and require SSH access must have their remote settings enabled and the same password set both on the cluster and on the authentication service.

   The remote setting instructs SAN Volume Controller to check the authentication service for group information for determining the role of the user.

5. Configure the system time.

   The current time of both the SAN Volume Controller cluster and the system that is running the remote authentication service must match. The easiest way to do this is to use the same Network Time Protocol (NTP) server for both.

   Attention: Failure to follow this step could result in either poor interactive performance of the SAN Volume Controller user interface or in incorrect user-role assignments.

Creating and working with user groups using the CLI

You can use the command-line interface (CLI) to create and work with users and user groups.

To create and work with user groups, follow these steps:

1. Issue the svctask mkusergrp CLI command to create a new user group. For example:

   ```
   svctask mkusergrp -name group_name -role role_name -remote
   ```

   where *group_name* specifies the name of the user group and *role_name* specifies the role that is associated with any users that belong to this group. The **remote** parameter specifies that the group is visible to the remote authentication service.

   The command returns the ID of the user group that was created.

2. Issue the svctask chusergrp CLI command to change attributes of an existing user group. For example:

   ```
   svctask chusergrp -role role_name -remote yes | no group_id_or_name
   ```

   where *role_name* specifies the role that is associated with any users that belong to this group and *group_id_or_name* specifies the group to be changed. The **remote** parameter specifies whether the group is visible to the authentication server.

3. Issue the svctask rmusergrp CLI command to delete a user group: For example:

   ```
   svctask rmusergrp -force group_id_or_name
   ```

   where *group_id_or_name* specifies the group to delete. The **force** parameter specifies to delete the group even if there are users in the user group. All users that were assigned to this group are assigned to the Monitor group.

4. Issue the svcinfo lsusergrp CLI command to display the user groups that have been created on the cluster. For example:

```
svcinfo lsusergrp usergrp_id_or_name
```

where *group_id_or_name* specifies the user group to view. If you do not specify a user group ID or name, all user groups on the cluster are displayed.

Creating and working with users using the CLI

You can use the command-line interface (CLI) to create and work with users.

To create and work with users, follow these steps:

1. Issue the svctask mkuser CLI command to create either a local user or a remote user to access the SAN Volume Controller. For example:

   ```
   svctask mkuser -name user_name -remote
   ```

   where *user_name* specifies the name of the user. The **remote** parameter specifies that the user authenticates to the remote authentication service.

   ```
   svctask mkuser -name user_name -usergrp group_name_or_id
   ```

   where *user_name* specifies the name of the user and *group_name_or_id* specifies the name or ID of the user group with which the local user is to be associated. The **usergrp** parameter specifies that the user authenticates to the cluster using cluster authentication methods.
2. Issue the svctask chuser CLI command to change the attributes of an existing user. For example:

   ```
   svctask chuser -usergrp group_id_or_name user_id_or_name
   ```

   where the *group_id_or_name* specifies the new group for the user and *user_id_or_name* specifies the user to be changed.
3. Issue the svctask chcurrentuser CLI command to change the attributes of the current user. For example:

   ```
   svctask chcurrentuser -nokey
   ```

   where the nokey parameter specifies that the SSH key of the user is to be deleted.
4. Issue the svctask rmuser CLI command to delete a user: For example:

   ```
   svctask rmuser user_id_or_name
   ```

   where *user_id_or_name* specifies the user to be removed.
5. Issue the svcinfo lsuser CLI command to display a list of users that have been created on the cluster. For example:

   ```
   svcinfo lsuser user_id_or_name
   ```

   where *user_id_or_name* specifies the ID or name of the user view. If you do not specify an ID or name, the concise view is displayed. If you do not specify a user ID or name, all users on the cluster are displayed.
6. Issue the svcinfo lscurrentuser CLI command to display the name and role of the logged-in user. For example:

   ```
   svcinfo lscurrentuser
   ```

   The name and the role of the user are displayed.

Setting up SNMP notifications using the CLI

You can set up event notifications using the command-line interface (CLI).

The notification settings apply to the entire cluster. You can specify the types of events that cause the cluster to send a notification. The cluster sends a Simple Network Management Protocol (SNMP) notification. The SNMP setting represents the type of notification.

The possible types of event notifications are error, warning, and information. Event notifications are reported to the SNMP destinations of your choice. To specify an SNMP destination, you *must* provide a valid IP address and SNMP community string.

Note: A valid community string can contain up to 60 letters or digits (most characters). A maximum of six SNMP destinations can be specified.

In configurations that use SNMP, the SAN Volume Controller uses the notifications settings to call home if errors occur. You must specify `Error` and send the trap to the IBM System Storage Productivity Center or the master console if you want the SAN Volume Controller to call home when errors occur.

To configure the SNMP notification settings, use the following commands:

1. To create a new SNMP server to receive notifications, use the svctask mksnmpserver CLI command. For example, enter one of the following commands:

   ```
   svctask mksnmpserver -ip 9.11.255.634
   ```

   where *9.11.255.634* is the IP addresses for this server.

   ```
   svctask mksnmpserver -ip 9.11.255.634 -port  remoteportnumber
   ```

   where *9.11.255.634* is the IP addresses for this server and *remoteportnumber* is the port number for the remote SNMP server.
2. To change the settings of an existing SNMP server, enter the svctask chsnmpserver command. For example:

   ```
   svctask chsnmpserver -name newserver snmp_server_name_or_id
   ```

   where *newserver* is the new name or ID of the server and *snmp_server_name_or_id* is the name or ID of the server to be modified.
3. To remove an existing SNMP server from the system, enter the svctask rmsnmpserver command. For example:

   ```
   svctask rmsnmpserver snmp_server_name_or_id
   ```

   where *snmp_server_name_or_id* is either the name or the ID of the SNMP server to be deleted.
4. To display either a concise list or a detailed view of the SNMP servers that are detected by the cluster, enter the svcinfo lssnmpserver command. For example, to display a concise view, enter the following command:

   ```
   svcinfo lssnmpserver -delim :
   ```

   To display a detailed view of an SNMP server, enter the following command:

   ```
   svcinfo lssnmpserver snmp_server_name_or_id
   ```

Setting up syslog notifications using the CLI

You can set up syslog event notifications using the command-line interface (CLI).

The syslog protocol is a standard protocol for forwarding log messages from a sender to a receiver on an IP network. The IP network can be either IPv4 or IPv6. SAN Volume Controller can send syslog messages that notify personnel about an event. SAN Volume Controller can transmit syslog messages in either expanded or concise format. You can use a syslog manager to view the syslog messages that SAN Volume Controller sends. SAN Volume Controller uses the User Datagram Protocol (UDP) to transmit the syslog message. You can use the SAN Volume Controller Console or the SAN Volume Controller command-line interface to configure and modify your syslog settings.

The syslog event notification settings apply to the entire cluster. You can specify the types of events that cause the cluster to send a notification. The possible types of notifications are error, warning, or information.

To specify a syslog destination, you *must* provide a valid IP address.

Note: Servers that are configured with facility values of 0 - 3 receive syslog messages in concise format. Servers that are configured with facility values of 4 - 7 receive syslog messages in fully expanded format.

The SAN Volume Controller uses the notifications settings to call home if errors occur.

To configure and work with notification settings, use the following commands:

1. Issue the svctask mksyslogserver CLI command to specify the action that you want to take when a syslog error or event is logged to the error log. For example, you can issue the following CLI command to set up a syslog notification:

   ```
   svctask mksyslogserver -ip 9.11.255.634
   ```

   where *9.11.255.634* is the IP address of the syslog server.
2. To modify a syslog notification, issue the svctask chsyslogserver command. For example:

   ```
   svctask chsyslogserver -name -facility facility_number syslog_server_name_or_id
   ```

   where *facility number* is a facility number to identify the origin of the message to the receiving server and *syslog_server_name_or_id* is the name or ID of the server to be modified.
3. To delete a syslog notification, issue the svctask rmsyslogserver command. For example:

   ```
   svctask rmsyslogserver syslog_server_name_or_id
   ```
4. To display either a concise list or a detailed view of syslog servers that are configured on the cluster, issue the svcinfo lssyslogserver command. For example, to display a concise view, enter the following command:

   ```
   svcinfo lssyslogserver -delim :
   ```

   To display a detailed view of a syslog server, enter the following command:

   ```
   svcinfo lssyslogserver snmp_server_name_or_id
   ```

Setting up e-mail event notifications and inventory reports using the CLI

You can use the command-line interface (CLI) to set up your system to send event notification and inventory reports to specified recipients and the IBM Support Center.

To set up, manage, and activate e-mail event and inventory notifications, complete the following steps:

1. Enable your system to use the e-mail notification function. To do this, issue the **svctask mkemailserver** CLI command. Up to six SMTP e-mail servers can be configured to provide redundant access to the external e-mail network.

   The following example creates an e-mail server object. It specifies the name, IP address, and port number of the SMTP e-mail server. After you issue the command, you see a message that indicates that the e-mail server was successfully created.

   ```
   svctask mkemailserver -ip ip_address -port port_number
   ```

   where *ip_address* specifies the IP address of a remote e-mail server and *port_number* specifies the port number for the e-mail server.

2. Add recipients of e-mail event and inventory notifications to the e-mail event notification facility. To do this, issue the svctask mkemailuser CLI command. You can add up to twelve recipients, one recipient at a time.

   The following example adds e-mail recipient **manager2008** and designates that **manager2008** receive e-mail error-type event notifications.

   ```
   svctask mkemailuser -address manager2008@ibm.com
    -error on -usertype local
   ```

3. Set the contact information that is used by the e-mail event notification facility. To do this, issue the svctask chemail CLI command. If you are starting the e-mail event notification facility, the **reply**, **contact**, **primary**, and **location** parameters are required. If you are modifying contact information used by the e-mail event notification facility, at least one of the parameters must be specified.

   The following example sets the contact information for the e-mail recipient **manager2008**.

   ```
   svctask chemail -reply manager2008@ibm.com -contact manager2008
    -primary 0441234567 -location 'room 256 floor 1 IBM'
   ```

4. Optionally, generate a report that lists e-mail event notification settings for all e-mail recipients, or change or delete e-mail recipients.
   - To generate a report that lists the e-mail event notification settings for all e-mail recipients, an individual e-mail recipient, or a specified type of e-mail recipient (local or support), issue the svctask lsemailuser CLI command.
   - To change the settings that are defined for a recipient, issue the svctask chemailuser CLI command. You must specify the user ID or name of the e-mail recipient for whom you are modifying settings.
   - To remove a previously defined e-mail recipient, issue the svctask rmemailuser CLI command. You must specify the user ID or name of the e-mail recipient that you want to remove.

5. Activate the e-mail and inventory notification function. To do this, issue the svctask startemail CLI command. There are no parameters for this command.

   Note: Inventory information is automatically reported to IBM when you activate error reporting.

6. Optionally, test the e-mail notification function to ensure that it is operating correctly and send an inventory e-mail notification.
   - To send a test e-mail notification to one or more recipients, issue the svctask testemail CLI command. You must either specify **all** or the user ID or user name of an e-mail recipient that you want to send a test e-mail to.

- To send an inventory e-mail notification to all recipients that are enabled to receive inventory e-mail notifications, issue the svctask sendinventoryemail CLI command. There are no parameters for this command.

Setting up e-mail servers using the CLI

You can set up e-mail server objects using the command-line interface (CLI).

You can specify a server object that describes a remote Simple Mail Transfer Protocol (SMTP) e-mail server to receive event notifications from the cluster. You can specify up to six servers to receive notifications. To configure and work with e-mail servers, use the following commands:

1. Issue the svctask mkemailserver CLI command to create an e-mail server object that describes a remote Simple Mail Transfer Protocol (SMTP) e-mail server. For example, issue the following CLI command to set up an e-mail server:

   ```
   svctask mkemailserver -ip ip_address
   ```

   where *ip_address* is the IP address of a remote e-mail server. This must be a valid IPv4 or IPv6 address.
2. To change the parameters of an existing e-mail server object, issue the svctask chemailserver command. For example:

   ```
   svctask chemailserver -ip ip_address email_server_name_or_id
   ```

   where *ip_address* is the IP address of the e-mail server object and *email_server_name_or_id* is the name or ID of the server object to be changed.
3. To delete a specified e-mail server object, issue the svctask rmemailserver command. For example:

   ```
   svctask rmemailserver email_server_name_or_id
   ```
4. To display either a concise list or a detailed view of e-mail servers that are configured on the cluster, issue the svcinfo lsemailserver command. For example, to display a concise view, enter the following command:

   ```
   svcinfo lsemailserver -delim :
   ```

   To display a detailed view of an e-mail server, enter the following command:

   ```
   svcinfo lsemailserver email_server_name_or_id
   ```

Changing cluster passwords using the CLI

You can use the command-line interface (CLI) to change the superuser and service passwords.

Passwords only affect the SAN Volume Controller Console that accesses the cluster. To restrict access to the CLI, you must control the list of SSH client keys that are installed on the cluster.

Perform the following steps to change the superuser and service passwords:

1. Issue the following command to change the superuser password:

   ```
   svctask chuser -password superuser_password superuser
   ```

   Where *superuser_password* is the new superuser password that you want to use.
2. Issue the following command to change the service password:

   ```
   svtask chcluster -servicepwd service_password
   ```

   Where *service_password* is the new service password that you want to use.

Changing the locale setting using the CLI

You can use the command-line interface (CLI) to specify the locale for a SAN Volume Controller cluster. The language that you select as your locale setting is used to display command results and error messages in the CLI.

The following locales are available:

- 0 US English (default)
- 3 Japanese

Issue the svctask setlocale CLI command with the ID for the locale.

For example, issue the following CLI command to change the locale setting from US English to Japanese:

```
svctask setlocale 3
```

where *3* is the ID for the Japanese locale setting.

Viewing the feature log using the CLI

You can use the command-line interface (CLI) to view the feature log.

Perform the following steps to view the feature log:

1. 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.
2. Issue the **svcservicemodeinfo lsfeaturedumps** command to return a list of the files that exist of the type specified on the given node.

Analyzing the error log using the CLI

You can use the command-line interface (CLI) to analyze the error log.

Perform the following steps to analyze the error log:

Issue any of the following CLI commands to list error log entries by file type:

- **svcinfo lserrlogbymdisk**
- **svcinfo lserrlogbymdiskgroup**
- **svcinfo lserrlogbyvdisk**
- **svcinfo lserrlogbyhost**
- **svcinfo lserrlogbynode**
- **svcinfo lserrlogbyiogrp**
- **svcinfo lserrlogbyfcconsistgrp**
- **svcinfo lserrlogbyfcmap**
- **svcinfo lserrlogbyrcconsistgrp**
- **svcinfo lserrlogbyrcrelationship**

These commands list the error log entries by type. For example, the **svcinfo lserrlogbymdisk** command displays the error log by managed disks (MDisks).

You can display the whole log or filter the log so that only errors, events, or unfixed errors are displayed. You can also request that the output is sorted either by error priority or by time. For error priority, the most serious errors are the lowest-numbered errors. Therefore, the most serious errors are displayed first in the table. For time, either the older or the latest entry can be displayed first in the output.

Shutting down a cluster using the CLI

You can use the command-line interface (CLI) to shut down a cluster.

If you want to remove all input power to a cluster (for example, the machine room power must 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 Controller nodes detect the loss of power and continue to run on battery power until all data that is held in memory is saved to the internal disk drive. This increases the time that is required to make the cluster operational when input power is restored and severely increases the time that is 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 supply units, they start to recharge. However, the SAN Volume Controller nodes do not permit any I/O activity to be performed to the virtual disks (VDisks) 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 two hours. Shutting down the cluster prior to removing input power to the uninterruptible power supply units prevents the battery power from being drained and makes it possible for I/O activity to resume as soon as input power is restored.

Before shutting down a cluster, quiesce all I/O operations that are destined for this cluster. Failure to do so can result in failed I/O operations being reported to your host operating systems.

Attention: If you are shutting down the entire cluster, you lose access to all VDisks that are provided by this cluster. Shutting down the cluster also shuts down all SAN Volume Controller nodes. This shutdown causes the hardened data to be dumped to the internal hard drive.

Begin the following process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDisks that are provided by the cluster.

1. Determine which hosts are using the VDisks that are provided by the cluster.
2. Repeat the previous step for all VDisks.

If input power is lost and subsequently restored, you must press the power button on the uninterruptible power supply units before you press the power buttons on the SAN Volume Controller nodes.

Perform the following steps to shut down a cluster:

1. Issue the following command to shut down a cluster:

   ```
   svctask stopcluster
   ```

The following output is displayed:

```
Are you sure that you want to continue with the shut down?
```

2. Type y to shut down the entire cluster.

Chapter 7. Backing up and restoring the cluster configuration

You can back up and restore the cluster configuration data after preliminary tasks are completed.

Cluster configuration data provides information about your cluster and the objects that are defined in it. The backup and restore functions of the svcconfig command can only back up and restore your cluster configuration data. You must regularly back up your application data using the appropriate backup methods.

Information about the following objects is included in the cluster configuration data:

- Storage system
- Hosts
- I/O groups
- Managed disks (MDisks)
- MDisk groups
- Nodes
- Virtual disks (VDisks)
- VDisk-to-host mappings
- Users and user groups
- FlashCopy mappings
- FlashCopy consistency groups
- Metro Mirror relationships
- Metro Mirror consistency groups
- Global Mirror relationships
- Global Mirror consistency groups

You can maintain your cluster configuration data by complete the following tasks:

- Backing up the cluster configuration data
- Restoring the cluster configuration data
- Deleting unwanted backup configuration data files

Before you backup your cluster configuration data, the following prerequisites must be met:

- No independent operations that change the cluster configuration can be running while the backup command is running.
- No object name can begin with an underscore.

Note:

- The default object names for controllers, I/O groups and managed disks (MDisks) do not restore correctly if the ID of the object is different than what is recorded in the current cluster configuration data file.
- All other objects with default names are renamed during the restore process. The new names appear in the format *name*_r. Where *name* is the name of the object in your cluster.

Before you restore your cluster configuration data, the following prerequisites must be met:

- You have the Security Administrator role associated with your username and password.
- You have a copy of your backup cluster configuration files on a server that is accessible to the cluster.
- You have a backup copy of your application data.
- You know the current license settings for your cluster.
- You have not removed any hardware since the last backup of your cluster configuration. If you had to replace a faulty node, the new node must use the same worldwide node name (WWNN) as the faulty node that it replaced.

  Note: You can add new hardware, but you must not remove any hardware because the removal can cause the restore process to fail.
- No zoning changes have been made on the fibre-channel fabric which would prevent communication between the SAN Volume Controller and with other nodes or storage controllers which are present in the configuration.

The restore must be performed to a single node cluster. You can restore the configuration using any node as the configuration node. However, if you do not use the node that was the configuration node when the cluster was first created, the unique identifier (UID) of the VDisks that are within the I/O groups can change. This can affect IBM TotalStorage Productivity Center for Fabric, VERITAS Volume Manager, and any other programs that record this information.

The SAN Volume Controller analyzes the backup configuration data file and the cluster to verify that the required disk controller system nodes are available.

Before you begin, hardware recovery must be complete. The following hardware must be operational: hosts, SAN Volume Controller, disk controller systems, disks, the Ethernet network, and the SAN fabric.

Backing up the cluster configuration

You can backup your cluster configuration data from the Backing up a Cluster Configuration panel.

This task assumes that you have already launched the SAN Volume Controller Console.

The backup function is designed to back up information about your cluster configuration, such as virtual disks (VDisks), local Metro Mirror information, local Global Mirror information, managed disk (MDisk) groups, and nodes. All other data that you have written to the VDisks is *not* backed up. Any application that uses the VDisks on the cluster as storage, must back up its application data using appropriate backup methods.

You must regularly back up your cluster configuration data and your application data to avoid data loss. If a cluster is lost after a severe failure occurs, both cluster configuration and application data is lost. You must reinstate the cluster to the exact state it was in prior to the failure and then recover the application data.

The backup function creates three files that provide information about the backup process and cluster configuration. When you use the SAN Volume Controller

Console to perform the backup, these files are created in the \console\ ClusterConfiguration.backup\cluster directory of the IBM System Storage Productivity Center or the master console. Where *console* is the directory where the SAN Volume Controller Console is installed and *cluster* is the name of the cluster for which you want to back up the cluster configuration data.

The following table describes the file that is created by the backup process:

File name	Description
svc.config.backup.xml	This file contains your cluster configuration data.

If the svc.config.backup.xml file already exists in the directory, it is renamed to svc.config.backup.bak. After the file is renamed the new svc.config.backup.xml is written.

Complete the following steps to backup your cluster configuration data:

1. Click **Service and Maintenance** → **Backup Configuration** in the portfolio. The Backing up a Cluster Configuration panel is displayed.
2. Click **Backup**.

Backing up the cluster configuration using the CLI

You can backup your cluster configuration data using the command-line interface (CLI).

Before you backup your cluster configuration data, the following prerequisites must be met:

- No independent operations that change the cluster configuration can be running while the backup command is running.
- No object name can begin with an underscore.
- All objects should have non-default names, that is, names that are not assigned by the SAN Volume Controller.

  Note:

  – The default object names for controllers, I/O groups and managed disks (MDisks) do not restore correctly if the ID of the object is different than what is recorded in the current cluster configuration data file.

  – All other objects with default names are renamed during the restore process. The new names appear in the format *name*_r. Where *name* is the name of the object in your cluster.

The backup feature of the **svcconfig** CLI command is designed to back up information about your cluster configuration, such as virtual disks (VDisks), local Metro Mirror information, local Global Mirror information, managed disk (MDisk) groups, and nodes. All other data that you have written to the VDisks is *not* backed up. Any application that uses the VDisks on the cluster as storage, must back up its application data using the appropriate backup methods.

You must regularly back up your cluster configuration data and your application data to avoid data loss. If a cluster is lost after a severe failure occurs, both cluster

configuration and application data is lost. You must reinstate the cluster to the exact state it was in prior to the failure and then recover the application data.

Perform the following steps to backup your cluster configuration data:

1. Back up all of the application data that you have stored on your VDisks using your preferred backup method.
2. Open a command prompt.
3. Issue the following command to log onto the cluster:

   ```
   ssh -l admin your_cluster_name -p 22
   ```

   Where *your_cluster_name* is the name of the cluster for which you want to backup cluster configuration data.
4. Issue the following CLI command to remove all of the existing cluster configuration backup and restore files that are located on your configuration node in the **/tmp** directory.

   ```
   svcconfig clear -all
   ```
5. Issue the following CLI command to backup your cluster configuration:

   ```
   svcconfig backup
   ```

   The following output is an example of the messages that are displayed during the backup process:

   ```
   CMMVC6112W io_grp io_grp1 has a default name
   CMMVC6112W io_grp io_grp2 has a default name
   CMMVC6112W mdisk mdisk14 ...
   CMMVC6112W node node1 ...
   CMMVC6112W node node2 ...
   .................................................
   ```

   The **svcconfig backup** CLI command creates three files that provide information about the backup process and cluster configuration. These files are created in the /tmp directory of the configuration node.

   The following table describes the three files that are created by the backup process:

File name	Description
svc.config.backup.xml	This file contains your cluster configuration data.
svc.config.backup.sh	This file contains the names of the commands that were issued to create the backup of the cluster.
svc.config.backup.log	This file contains details about the backup, including any error information that might have been reported.

6. Issue the following command to exit the cluster:

   ```
   exit
   ```
7. Issue the following command to copy the backup files to a location that is not in your cluster:

   ```
   scp -P 22 admin@your_cluster:/tmp/svc.config.backup.*
    /offclusterstorage/
   ```

   Where *your_cluster* is the name of your cluster and *offclusterstorage* is the location where you want to store the backup files.

You must copy these files to a location outside of your cluster because the /tmp directory on this node becomes inaccessible if the configuration node changes. The configuration node might change in response to an error recovery action or to a user maintenance activity.

Tip: To maintain controlled access to your cluster configuration data, copy the backup files to a location that is password protected.

8. Ensure that the copies of the backup files are stored in the location that you specified in step 7 on page 298.

You can rename the backup files to include the configuration node name either at the start or end of the file names so you can easily identify these files when you are ready to restore your configuration.

Issue the following command to rename the backup files that are stored on a Linux or AIX host:

```
mv /offclusterstorage/svc.config.backup.xml
 /offclusterstorage/svc.config.backup.xml_myconfignode
```

Where *offclusterstorage* is the name of the directory where the backup files are stored and *myconfignode* is the name of your configuration node.

To rename the backup files that are stored on a Windows host, right-click on the name of the file and select **Rename**.

Downloading backup configuration data files

You can use the SAN Volume Controller Console to download backup configuration data files to your IBM System Storage Productivity Center or master console.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to download the backup configuration data files to your IBM System Storage Productivity Center or master console:

1. Click **Service and Maintenance** → **List Dumps** in the portfolio. The List Dumps panel is displayed.
2. Click **Software Dumps**. The Software Dumps panel is displayed.
3. Find the name of your backup configuration data file.
4. Right-click on your backup configuration data file and click **Save Target As**.
5. Select the location where you want to save the file and click **Save**.

Restoring the cluster configuration using the CLI

You can restore your cluster configuration data using the command-line interface (CLI).

Important: There are two phases during the restore process: prepare and execute. You must not make any changes to the fabric or cluster between these two phases.

Complete the following steps to restore your cluster configuration data:

1. Select delete cluster from the front panel on each node in the cluster that does *not* display `Cluster :` on the front panel. If the front panel of the node displays `Cluster :`, the node is already a candidate node.
2. Create a new cluster from the front panel of any node in the cluster. If possible, use the node that was originally the configuration node for the cluster.
3. Generate an SSH key pair for all of the hosts to use to access the CLI.
4. Start the SAN Volume Controller Console.
5. On the Viewing Clusters panel, select the cluster that you are recovering from the list, select **Remove the Cluster** from the task list and click **Go**. The Remove Cluster panel displays. Click **Yes** to confirm the removal of the cluster. The Viewing Cluster panel displays.
6. On the Viewing Cluster panel, select **Add a Cluster** from the task list and click **Go**. The Adding a Cluster panel displays. Enter the IP address for the cluster that you are recovering. Do **not** select **Create (Initialize) Cluster**. Click **OK**.
7. To work with the command-line interface to finish restoring the cluster, you also need to assign an SSH key to the user that has Security Administrator role on the cluster by completing these steps:
   a. On the Viewing Cluster panel, select the new cluster and select Launch SAN Volume Controller Console from the task list and click **Go**.
   b. Click **Manage Authentication** → **Users** in the portfolio. The Viewing Users panel is displayed.
   c. Select the user that you want to change and select **Modify a User** from the task list. Click **Go**. The Modifying a User panel is displayed.
   d. To assign the SSH key that you generated in Step 3 to the user, enter the name of SSH key file in the SSH Key Public File field or click **Browse** to select the file.
   e. Click **OK**.
8. Using the command-line interface, issue the following command to log onto the cluster:

   ```
   ssh -l admin your_cluster_name -p 22
   ```

   Where *your_cluster_name* is the name of the cluster for which you want to restore the cluster configuration.

   Note: Because the RSA host key has changed, a warning message displays when connecting to the cluster using SSH.
9. Issue the following CLI command to ensure that only the configuration node is online.

   ```
   svcinfo lsnode
   ```

   The following is an example of the output that is displayed:

   ```
   id  name  status  IO_group_id   IO_group_name   config_node
    1  node1 online  0  io_grp0  yes
   ```
10. Verify that the most recent version of your **/tmp/svc.config.backup.xml** configuration file has been copied to your SSPC. The most recent file is located on your configuration node in the **/tmp** or **/dumps** directory. In addition, a **/dumps/svc.config.cron.xml**_*node_name* configuration file is created daily on the configuration node. In certain cases, you might prefer to copy an earlier configuration file. If necessary, back up your configuration file as described in "Backing up the cluster configuration using the CLI" on page 297.

11. Issue the following CLI command to remove all of the existing backup and restore cluster configuration files that are located on your configuration node in the **/tmp** directory:

    ```
    svcconfig clear -all
    ```

12. Copy the `svc.config.backup.xml` file from the IBM System Storage Productivity Center or master console to the /tmp directory of the cluster using the PuTTY pscp program. Perform the following steps to use the PuTTY pscp program to copy the file:
    a. Open a command prompt from the IBM System Storage Productivity Center or 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 <private key location> source [source...] [user@]host:target
       ```

13. Issuing the following CLI command to compare the current cluster configuration with the backup configuration data file:

    ```
    svcconfig restore -prepare
    ```

    This CLI command creates a log file in the /tmp directory of the configuration node. The name of the log file is `svc.config.restore.prepare.log`.

    Note: It can take up to a minute for each 256-MDisk batch to be discovered. If you receive error message CMMVC6119E for an MDisk after you enter this command, all the managed disks (MDisks) might not have been discovered yet. Allow a suitable time to elapse and try the `svcconfig restore -prepare` command again.

14. Issue the following command to copy the log file to another server that is accessible to the cluster:

    ```
    pscp -i <private key location> [user@]host:source target
    ```

15. Open the log file from the server where the copy is now stored.
16. Check the log file for errors.
    - If there are errors, correct the condition which caused the errors and reissue the command. You must correct all errors before you can proceed to step 17.
    - If you need assistance, contact the IBM Support Center.
17. Issue the following CLI command to restore the cluster configuration:

    ```
    svcconfig restore -execute
    ```

    Note: Issuing this CLI command on a single node cluster adds the other nodes and hosts to the cluster.

    This CLI command creates a log file in the /tmp directory of the configuration node. The name of the log file is `svc.config.restore.execute.log`.

18. Issue the following command to copy the log file to another server that is accessible to the cluster:

    ```
    scp -P 22 admin@your_cluster:/tmp/svc.config.restore.execute.log
     /offclusterstorage/
    ```

    Where *your_cluster* is the name of your cluster and *offclusterstorage* is the location where you want to store the log file.

19. Open the log file from the server where the copy is now stored.
20. Check the log file to ensure that no errors or warnings have occurred.

Note: You might receive a warning that states that a licensed feature is not enabled. This means that after the recovery process, the current license settings do not match the previous license settings. The recovery process continues normally and you can enter the correct license settings in the SAN Volume Controller Console at a later time.

The following output is displayed after a successful cluster configuration restore operation:

```
IBM_2145:your_cluster_name:admin>
```

21. After the cluster configuration is restored, verify that the quorum disks are restored to the MDisks that you want by using `svcinfo lsmdisk` *mdisk_id or mdisk_name*. To restore the quorum disks to the correct MDisks, issue the appropriate svctask setquorum CLI commands.

You can remove any unwanted configuration backup and restore files from the **/tmp** directory on your configuration by issuing the `svcconfig clear -all` CLI command.

Note: The recovery process does not recreate the superuser password and SSH keys. Ensure those are recreated before managing the recovered cluster.

Deleting backup configuration files

You can delete a backup cluster configuration from the Deleting a Cluster Configuration panel.

This task assumes that you have already launched the SAN Volume Controller Console.

Perform the following steps to delete backup configuration files:

1. Click **Service and Maintenance** → **Delete Backup** in the portfolio. The Deleting a Cluster Configuration panel is displayed.
2. Click **OK**.

Deleting backup configuration files using the CLI

You can use the command-line interface (CLI) to delete backup configuration files.

Perform the following steps to delete backup configuration files:

1. Issue the following command to log onto the cluster:

   ```
   ssh -l admin your_cluster_name -p 22
   ```

   Where *your_cluster_name* is the name of your cluster.
2. Issue the following CLI command to erase all of the files that are stored in the /tmp directory:

   ```
   svconfig clear -all
   ```

Chapter 8. Upgrading the SAN Volume Controller software

The SAN Volume Controller software can be upgraded while you run day-to-day operations.

However, performance is degraded during the software upgrade process. Only the following commands can be issued during the software upgrade:

- All svcinfo commands
- svctask rmnode

Attention:

- You must upgrade the SAN Volume Controller Console before you upgrade the SAN Volume Controller software.
- Applying a software upgrade takes a varying length of time. Plan for at least one hour because there is a 30-minute delay that allows the multipathing 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 increases each time a new release is made. The package includes Linux, Apache and the SAN Volume Controller software.

Some software levels support upgrades only from specific previous levels, or the software can be installed only on certain hardware types. If you upgrade to more than one level above your current level, you might be required to install an intermediate level. For example, if you are upgrading from level 1 to level 3, you might have to install level 2 before you can install level 3. For information about the prerequisites for each software level, see the Support for SAN Volume Controller (2145) Web site:

www.ibm.com/storage/support/2145

Attention:

- If you apply the software upgrade while a node is in service mode, the node is deleted from the cluster. All status information that is stored on the node is deleted and data loss can occur if the cluster depends solely on this node.
- Ensure that you have no unfixed errors in the log and that the cluster date and time are correctly set. Start the Directed Maintenance Procedures (DMPs) and ensure that you fix any outstanding errors before you attempt to concurrently upgrade the software.

Metro Mirror and Global Mirror relationships

When you upgrade software where the cluster participates in one or more intercluster relationships, update the clusters one at a time. Do not upgrade the clusters concurrently because you can lose synchronization and availability.

You can create new Metro Mirror or Global Mirror partnerships between clusters with different software levels. If the partnerships are between a SAN Volume Controller version 5.1.0 cluster and a cluster that is earlier than version 5.1.0, each

cluster can participate in a single partnership with another cluster. If the clusters are all SAN Volume Controller version 5.1.0, each cluster can participate in up to four cluster partnerships.

Attention: If you want to upgrade a cluster to SAN Volume Controller version 5.1.0 and the partner is running version 4.2.0 or earlier, you must first upgrade the partner cluster to SAN Volume Controller 4.2.1 or later before you upgrade the first cluster to version 5.1.0.

Installing or upgrading the SAN Volume Controller software

The SAN Volume Controller software can be installed or upgraded after you download the software package from the SAN Volume Controller Web site.

Software package

The software installation or upgrade procedure copies the new software level to the cluster and starts an automatic installation process. During the installation process, each node is restarted. While each node restarts, there might be some degradation in the maximum I/O rate that can be sustained by the cluster. The amount of time that is needed to install or upgrade the software is dependent on the size of the cluster and the size of the software update package. The size of the software update package is determined by the number of components that are being replaced. After all the nodes in the cluster are successfully restarted with the new software level, the new software level is automatically committed.

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, these restrictions are documented on the SAN Volume Controller Web site that you use to download the software packages. During the software upgrade procedure, only the following SAN Volume Controller commands are operational from the time the install process starts to the time that the new software level is committed, or until the process has been backed-out.

- All `svcinfo` commands
- `svctask rmnode`

All other commands fail with a message that indicates a software upgrade is in progress.

To determine when your software upgrade process has completed, you will be notified through the SAN Volume Controller Console or, if you are using the command-line interface, issue the svcinfo lssoftwareupgradestatus command to display the status of the upgrade.

Because of the operational limitations that occur during the software upgrade process, the software installation is a user task.

Copying the SAN Volume Controller software upgrade files using PuTTY scp

PuTTY scp (pscp) provides a file transfer application for secure shell (SSH) to copy files either between two directories on the configuration node or between the configuration node and another host.

To use the pscp application, you must have the appropriate permissions on the source and destination directories on your respective hosts.

The pscp application is available when you install an SSH client on your host system. You can access the pscp application through a Microsoft Windows command prompt.

Perform the following steps to use the pscp application:

1. Start a PuTTY session.
2. Configure your PuTTY session to access your SAN Volume Controller cluster.
3. Save your PuTTY configuration session. For example, you can name your saved session SVCPUTTY.
4. Open a command prompt.
5. Issue the following command to set the path environment variable to include the PuTTY directory:

   ```
   set path=C:\Program Files\putty;%path%
   ```

   where *Program Files* is the directory where PuTTY is installed.
6. Issue the following command to copy the package onto the node where the CLI runs:

   ```
   pscp -load saved_putty_configuration
   directory_software_upgrade_files/software_upgrade_file_name
   admin@cluster_ip_address:/home/admin/upgrade
   ```

   where *saved_putty_configuration* is the name of the PuTTY configuration session, *directory_software_upgrade_files is the location of the software upgrade files*, *software_upgrade_file_name* is the name of the software upgrade file, and *cluster_ip_address* is an IP address of your cluster.

   If there is insufficient space to store the software upgrade file on the cluster, the copy process fails. Perform the following steps:

   a. Use pscp to copy data that you want to preserve from the /dumps directory. For example, issue the following command to copy all error logs from the cluster to the IBM System Storage Productivity Center or the master console:

      ```
      pscp -unsafe -load saved_putty_configuration
      admin@cluster_ip_address:/dumps/elogs/*
      your_preferred_directory
      ```

      where *saved_putty_configuration* is the name of the PuTTY configuration session, *cluster_ip_address* is the IP address of your cluster, and *your_preferred_directory* is the directory where you want to transfer the error logs.

   b. Issue the svctask cleardumps command to free space on the cluster:

      ```
      svctask cleardumps -prefix /dumps
      ```

   c. Then repeat step 6.

Upgrading the SAN Volume Controller software automatically

When new nodes are added to the cluster, the software upgrade file is automatically downloaded to the new nodes from the SAN Volume Controller cluster.

If you add a new node that has or requires a software level that is higher than the software level available on the cluster, the new node is *not* configured into the cluster. The new node must be downgraded to the software level of the cluster before it can join the cluster. If a node is added to the cluster that does not have software installed or has an old software level that cannot be recognized by the cluster, a node rescue must be performed to force a reinstallation of the software.

If the new node requires a level of software that is higher than the software level that is available on the cluster, the entire cluster must be upgraded before the new node can be added to the cluster.

Error counts

During the software upgrade if you are using IBM Subsystem device driver (SDD) as the multipathing software on the host, increased I/O error counts are displayed by the **datapath query device** or **datapath query adapter** commands if active I/O operations exist between the hosts and the SANs during a software upgrade. See the *IBM System Storage Multipath Subsystem Device Driver User's Guide* for more information about the **datapath query** commands.

During the software upgrade, each node of a working pair is upgraded sequentially. The node that is being upgraded is temporarily unavailable and all I/O operations to that node fails. As a result, the I/O error counts increase and the failed I/O operations are directed to the partner node of the working pair. Applications should not see any I/O failures.

Upgrading the SAN Volume Controller cluster software using the SAN Volume Controller Console

You can upgrade the cluster software using the SAN Volume Controller Console.

Attention: Before you start a software upgrade, you must check for offline or degraded VDisks. An offline VDisk can cause write data that has been modified to be pinned in the SAN Volume Controller cache. This prevents VDisk failover and causes a loss of I/O access during the software upgrade. If the fast_write_state is empty, a VDisk can be offline and not cause errors during the software upgrade.

The software upgrade files can be quite large. To shorten upload times, you should disable proxies on the Web browser from where you will upload the file. If you disable proxies, you might not be able to connect to external Web sites. Therefore, you must make a record of your existing settings before you disable proxies in case you have to restore access to other Web sites. To disable proxy settings if you are using Internet Explorer, complete the following steps:

1. Click **Tools** in the menu bar.
2. Select the **Internet Options** → **Connections** tab.
3. Click on **LAN Settings...** and ensure that the box marked **Use a proxy server for your LAN (These settings will not apply to dial-up or VPN connections)** is unchecked.

4. Click **OK** twice to accept the settings.

To disable proxy settings if you are using Netscape, complete the following steps:

1. Click **Edit** in the menu bar.
2. Click on **Preferences...**. Expand the Advanced section and select **Proxies**.
3. Select the **Direct connection to the Internet** button and click **OK** to accept the settings.

To disable proxy setting if you are using Mozilla Firefox, complete the following steps:

1. Open your Firefox browser and click **Tools** → **Options** → **Advanced**. The Advanced window displays.
2. Select the Network tab and click **Settings** under **Connections** heading. The Connection Setting panel displays
3. Under Configure Proxies to Access the Internet, ensure that **No Proxies** is selected.
4. Click **OK**.

When using Firefox 3 to navigate the SAN Volume Controller Console, and you select **Service and Maintenance** → **Upgrade Software** for the first time, you might see the following pop-up error:

```
hostname.ibm.com:443 uses an invalid security certificate

The certificate is not trusted because it is self signed.
The certificate is only valid for <a id="cert_domain_link" title="2145">2145</a>

(Error code: sec_error_ca_cert_invalid)
```

Click **Or you can add an exception** if the browser presents that option and complete the following steps:

1. Click **Add Exception...**
2. Click **Get Certificate**.
3. Click **Confirm Security Exception**

If the browser does not present this option, you can add the exception manually by completing the following steps:

1. Click **Tools** → **Options** → **Advanced**. The Advanced window displays.
2. Click the **Encryption** tab and click **View Certificates**.
3. Click **Servers** tab.
4. Click **Add Exception...** and enter `https://hostname.ibm.com:443` for the location and click **Get Certificate**.
5. Click **Confirm Security Exception**.

If you are using Internet Explorer 8, you see the message: `Content was blocked because it was not signed by a valid security certificate.` To display blocked content, complete the following steps:

1. Click the Internet Explorer Information Bar, and select **Display Blocked Content**.
2. Close the Information Bar. The Software Upgrade panel now displays.

This task assumes that you have already launched the SAN Volume Controller Console.

Complete the following steps to upgrade the software:

1. Download the SAN Volume Controller code from the following Web site:

   www.ibm.com/storage/support/2145

   - If you want to write the SAN Volume Controller code to a CD, you must download the CD image.
   - If you do not want to write the SAN Volume Controller code to a CD, you must download the install image.
2. Click **Service and Maintenance** in the portfolio.
3. Click **Upgrade Software** to install a new level of software on the cluster. The Software Upgrade panel is displayed. The Software Upgrade panel displays different information depending upon the current state of the cluster. If the cluster is currently in the middle of a cluster software upgrade, the panel displays a **Check Upgrade Status** button that allows you to view information about the current state of the upgrade. See step 8 on page 309 for details of upgrade status information. Otherwise, you are presented with options to upgrade the SAN Volume Controller cluster software or upgrade the firmware on solid-state drives (SSDs) in the SAN Volume Controller nodes.
4. From the task list, select **Upload File** and click **Go**. The File Upload panel displays.
5. Click **Browse** and select the SAN Volume Controller cluster software file that you downloaded in step 1.
6. Click **OK** to copy the SAN Volume Controller software file to the cluster. After the file uploads, the Software Upgrade panel is displayed. The new uploaded file should be displayed in the list.

   Before you begin the software upgrade, you must be aware of the following:

   - The installation process fails under the following conditions:
     - If the software that is installed on the remote cluster is not compatible with the new software or if there is an intercluster communication error that does not allow the software to check that the software is compatible.
     - If any node in the cluster has a hardware type that is not supported by the new software.
     - If the SAN Volume Controller software determines that one or more virtual disks (VDisks) in the cluster would be taken offline by rebooting the nodes as part of the upgrade process. Details about which VDisks would be affected can be found by using the svcinfo lsnodedependentvdisks command or the **View Dependent VDisks** action from the Viewing Nodes panel. You can use the force flag to override this restriction if you are prepared to lose access to data during the upgrade.
   - The software upgrade 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 run the new software, concurrently with normal cluster activity.
   - While the node is updated, it does not participate in I/O activity in the I/O group. As a result, all I/O activity for the VDisks in the I/O group is directed to the other node in the I/O group by the host multipathing software.
   - There is a 30 minute delay between node updates. The delay allows time for the host multipathing software to rediscover paths to the nodes which have been upgraded, so that there is no loss of access when another node in the I/O group is updated.

- The software update is not committed until all nodes in the cluster have been successfully updated to the new software level. If all nodes successfully restart with the new software level, the new level is committed. When the new level is committed, the cluster vital product data (VPD) is updated to reflect the new software level.
- You cannot invoke the new functions of the upgraded software until all member nodes are upgraded and the update has been committed.
- Because the software upgrade process takes some time, the install command completes as soon as the software level is verified by the cluster. To determine when the upgrade has completed, you can use the **Check Upgrade Status** button available from the Upgrade Software panel. See step 8 for more details.
- During a software upgrade the version number of each node is updated when the software has been installed and the node has been restarted. The cluster software version number is updated when the new software level is committed.
- When the software upgrade starts an entry is made in the error or event log and another entry is made when the upgrade completes or fails.

7. Select the cluster software file that you uploaded from the list of files and select Click **Apply Cluster Software** from the task list and click **Go**. The Applying Software Upgrade panel is displayed. The Applying Software Upgrade panel enables you to select the upgrade and apply it to the cluster. A list of the software levels that you can apply to the cluster is displayed.
8. To monitor the upgrade, select Check Upgrade Status on the Upgrade Software panel. The panel displays the current software running on every node in the cluster, as well as the software level of the cluster. The upgrade is complete when every node in the cluster is running the new level of software, and the cluster software version has been updated to match that of the node software versions.

   Attention: If the cluster upgrade process stalls and has the status of stalled or stalled_non_redundant you should contact IBM Support for guidance about the best procedure to restore your cluster to a fully operational state with the least disruption to hosts. Starting the abort procedure without guidance from IBM Support may cause a preventable loss of I/O access.

When a new software level is applied, it is automatically installed on all the nodes that are in the cluster.

Note: The software upgrade can take up to 30 minutes per node.

Upgrading solid-state drive (SSD) software

The Upgrade Software panel in the SAN Volume Controller Console provides a method to upgrade cluster software and SSD software. This panel displays currently uploaded packages, the status of upgrades, and the current version levels.

This task assumes that you have already launched the SAN Volume Controller Console.

Software upgrades include cluster software versions and upgrades to SSDs that are located on SAN Volume Controller 2145-CF8 nodes. For software updates to SSDs, each SAN Volume Controller 2145-CF8 node can contain up to four SSDs, which

are presented to SAN Volume Controller Console as managed disks (MDisks). Each of these MDisks contains two programmable components:

- The drive firmware
- The Field Programmable Gate Array (FPGA)

These programmable components can be upgraded individually.

Note: Do not upgrade the FPGA unless directed to do so by an IBM service representative.

If you want to upgrade the SSD software, you can download the upgrade files from the Support for SAN Volume Controller (2145) Web site:

www.ibm.com/storage/support/2145

1. Download the appropriate SSD software from the following Web site:
   www.ibm.com/storage/support/2145
   - If you want to write the SAN Volume Controller code to a CD, you must download the CD image.
   - If you do not want to write the SAN Volume Controller code to a CD, you must download the install image.
2. In the portfolio, click **Service and Maintenance**.
3. Click **Upgrade Software** to upgrade SSD software. The Software Upgrade panel is displayed. This panel displays all currently uploaded packages that contain cluster software upgrades and SSD software. Cluster software upgrade files are automatically deleted when they have been successfully uploaded and applied to the all nodes within the cluster. SSD upgrade files remain listed on the panel because they can be applied to multiple drives within the cluster. When you have successfully upgraded the SSD drives in the cluster, you can delete these files.
4. If you have not uploaded the SSD software file, select **Upload File** from the task list, and click **Upload File**. The File Upload panel displays.
5. Click **Browse** and select the SSD software file that you downloaded in step 1.
6. Click **OK** to copy the SSD software file to the cluster. After the file uploads, the Software Upgrade panel is displayed. The new uploaded file should be displayed in the list.
7. On the Software Upgrade panel, select the SSD software file that you uploaded and select **Apply MDisk Software** from the task list and click **Go**. The Select Solid State Drives panel displays.
8. On the Select Solid State Drives panel, you can select to either upgrade the SSD firmware or upgrade the FPGA. SSD drives are represented as MDisks. To upgrade SSD firmware, select the SSD drive to be upgraded and select **Upgrade Firmware** from the task list. Click **Go**.
   a. On the Upgrading Firmware panel, verify that the correct SSD drive is being upgraded with the correct firmware software package. Click **Confirm**.
9. To upgrade the FPGA firmware, select the SSD drive to be upgraded and select **Upgrade FPGA** from the task list. Click **Go**.

   Note: Do not upgrade the FPGA of an SSD unless you have been directed to do so by an IBM service representative.

   a. On the Upgrading FPGA panel, verify that the correct SSD drive is being upgraded with the correct FPGA firmware package. Click **Confirm**.

Upgrading solid-state drive (SSD) firmware using the CLI

You can upgrade a solid-state drive (SSD) by downloading and applying firmware updates using the SAN Volume Controller command-line interface (CLI).

This procedure upgrades firmware on a solid-state drive (SSD) that is internal to a supported SAN Volume Controller node. If the upgrade could cause any VDisks to go offline, the **-force** option is required. For example, a firmware update to a managed MDisk requires the **-force** option.

Perform the following steps to upgrade SSD firmware:

1. Run the lsnodedependentvdisks command for the node that you are upgrading. If any VDisks are returned, continuing with this procedure will take the VDisks offline. To avoid losing access to data, you must mirror the VDisks before continuing with this upgrade procedure.
2. Locate the firmware upgrade file at the IBM Support for SAN Volume Controller (2145) Web site: www.ibm.com/storage/support/2145
3. Download the firmware upgrade file to the /home/admin/upgrade directory.
4. Run the applymdisksoftware command as follows. You must specify the firmware file name and the MDisk name or ID of the SSD:

   ```
   svctask applymdisksoftware -file filename mdisk name | id
   ```

   To apply the upgrade even if it will cause one or more VDisks to go offline, specify the `-force` option.

   Attention: Only use the `-type fpga` option, which upgrades Field Programmable Gate Array (FPGA) firmware, under the direction of an IBM service representative.

Upgrading the SAN Volume Controller software using the CLI

You can use the command-line interface (CLI) to install software upgrades.

Attention: Before you start a software upgrade, you must check for offline or degraded VDisks. An offline VDisk can cause write data that has been modified to be pinned in the SAN Volume Controller cache. This prevents VDisk failover and causes a loss of I/O access during the software upgrade. If the fast_write_state is empty, a VDisk can be offline and not cause errors during the software upgrade.

Perform the following steps to upgrade the software:

1. Download the SAN Volume Controller code from the Support for SAN Volume Controller (2145) Web site:

   www.ibm.com/storage/support/2145

   - If you want to write the SAN Volume Controller code to a CD, you must download the CD image.
   - If you do not want to write the SAN Volume Controller code to a CD, you must download the install image.
2. Use PuTTY scp (pscp) to copy the software upgrade files to the node.
3. Ensure that the software upgrade file has been successfully copied.

   Before you begin the software upgrade, you must be aware of the following:

   - The install process fails under the following conditions:

- If the software that is installed on the remote cluster is not compatible with the new software or if there is an intercluster communication error that does not allow the software to check that the software is compatible.
- If any node in the cluster has a hardware type that is not supported by the new software.
- .If the SAN Volume Controller software determines that one or more virtual disks (VDisks) in the cluster would be taken offline by rebooting the nodes as part of the upgrade process. Details about which VDisks would be affected can be found by using the svcinfo lsnodedependentvdisks command or the **View Dependent VDisks** action from the Viewing Nodes panel. You can use the force flag to override this restriction if you are prepared to lose access to data during the upgrade.

- The software upgrade 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 run the new software, concurrently with normal cluster activity.
- While the node is updated, it does not participate in I/O activity in the I/O group. As a result, all I/O activity for the VDisks in the I/O group is directed to the other node in the I/O group by the host multipathing software.
- There is a 30-minute delay between node updates. The delay allows time for the host multipathing software to rediscover paths to the nodes which have been upgraded, so that there is no loss of access when another node in the I/O group is updated.
- The software update is not committed until all nodes in the cluster have been successfully updated to the new software level. If all nodes successfully restart with the new software level, the new level is committed. When the new level is committed, the cluster vital product data (VPD) is updated to reflect the new software level.
- You cannot invoke the new functions of the upgraded software until all member nodes are upgraded and the update has been committed.
- Because the software upgrade process takes some time, the install command completes as soon as the software level is verified by the cluster. To determine when the upgrade has completed, you must either display the software level 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 software level or fails at any other time during the process, the software level is backed-off.
- During a software upgrade the version number of each node is updated when the software has been installed and the node has been restarted. The cluster software version number is updated when the new software level is committed.
- When the software upgrade starts an entry is made in the error or event log and another entry is made when the upgrade completes or fails.

4. Issue the following CLI command to start the software upgrade process:
   - If you are upgrading from SAN Volume Controller version 4.3.1 to version 5.1, enter this command:

     ```
     svcservicetask applysoftware -file software_upgrade_file
     ```

   - If you are upgrading from SAN Volume Controller version 5.1, enter this command:

     ```
     svctask applysoftware -file software_upgrade_file
     ```

where *software_upgrade_file* is the name of the software upgrade file.The software upgrade does not start if the cluster identifies any VDisks that would go offline as a result of rebooting the nodes as part of the cluster upgrade. An optional force parameter can be used to indicate that the upgrade should continue in spite of the problem identified. Use the svcinfo lsnodedependentvdisks command to identify the cause for the failed upgrade. If you use this parameter, you are prompted to confirm that you want to continue. The behavior of the force parameter has changed, and is no longer required when applying an upgrade to a cluster with errors in the error log.

5. Issue the following CLI command to check the status of the software upgrade process:

   ```
   svcinfo lssoftwareupgradestatus
   ```

   Note: If a status of `stalled_non_redundant` is displayed, proceeding with the remaining set of node upgrades might result in offline VDisks. Contact an IBM service representative to complete the upgrade.

6. Perform the following steps to verify that the software upgrade successfully completed:
   a. Issue the **svctask dumperrlog** CLI command to send the contents of the error log to a text file.

      The following output is displayed in the text file if the software is successfully upgraded:

      ```
      Upgrade completed successfully
      ```

   b. Issue the **svcinfo lsnodevpd** CLI command for each node that is in the cluster. The software version field displays the new software level.

When a new software level is applied, it is automatically installed on all the nodes that are in the cluster.

Note: The software upgrade can take up to 30 minutes per node.

Performing a disruptive software upgrade using the CLI

You can use the command-line interface (CLI) to perform a disruptive software upgrade.

The SAN Volume Controller only supports concurrent software upgrades. To ensure that a software upgrade is coordinated across all nodes in the cluster, the nodes must be able to communicate with each other across the fibre-channel SAN. However, if this is not possible, you can perform a disruptive software upgrade.

Perform the following steps to complete the disruptive software upgrade process:

1. Stop any host applications and unmount the file systems that use storage that is managed by the SAN Volume Controller. If you are shutting down your hosts, this occurs as the host is shutdown. If you are not shutting down your hosts, you must manually stop host applications and unmount the file systems on each host. This step ensures that the hosts stop issuing I/O operations and that any data in the file system caches is flushed.
2. Issue the **svctask stopcluster** CLI command to shutdown the cluster. This CLI command stops the SAN Volume Controller from issuing I/O to backend controllers and flushes data from the SAN Volume Controller nodes 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 during 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 SAN Volume Controller has been isolated from the backend storage, errors that indicate the backend storage is unavailable are logged.
5. Perform the software upgrade in the same manner as for a concurrent software upgrade.
6. Restore the original switch configuration.
7. Clear any error logs that were produced in step 4 indicating that backend storage is unavailable. Check that all backend storage is now online and accessible to the SAN Volume Controller nodes.
8. Remount file systems and start host applications.

Performing the node rescue

If it is necessary to replace the hard disk drive or if the software on the hard disk drive is corrupted, you can use the node rescue procedure to reinstall the SAN Volume Controller software.

Similarly, if you have replaced the service controller, you should use the node rescue procedure to ensure that the service controller has the correct software.

Attention: If you recently replaced both the service controller and the disk drive as part of the same repair operation, node rescue fails.

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 node cannot boot and the hardware boot indicator remains on the front panel display or the boot operation does not progress. If this occurs, use the node rescue procedure to reinstall the SAN Volume Controller software.

Node rescue works by booting the operating system from the service controller and running a program that copies all the SAN Volume Controller software from any other node that can be found on the fibre-channel fabric.

Attention: When running node rescue operations, only run one node rescue operation on the same SAN, at any one time. Wait for one node rescue operation to complete before starting another.

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 node is connected to the fibre-channel fabric.
3. Ensure that the SAN zoning allows a connection between at least one port of this node and one port of another node. It is better if multiple ports can connect. This is particularly important if the zoning is by worldwide port name (WWPN) and you are using a new service controller. In this case, you might

need to use SAN monitoring tools to determine the WWPNs of the node. If you need to change the zoning, remember to set it back when the service procedure is complete.

4. Turn off the node.
5. Press and hold the left and right buttons on the front panel.
6. Press the power button.
7. Continue to hold the left and right buttons until the node-rescue-request symbol is displayed on the front panel (Figure 38).

Figure 38. Node rescue display

The node rescue request symbol displays on the front panel display until the node starts to boot from the service controller. If the node rescue request symbol displays for more than two minutes, go to the hardware boot MAP to resolve the problem. When the node rescue starts, 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 is now 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, it is not possible to add the node back into the cluster until the upgrade or downgrade process has completed. This can take up to four hours for an eight-node cluster.

Recovering from software upgrade problems automatically

The cluster automatically stops the software upgrade process if 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 are downgraded to the original software level. If a node fails to restart during this downgrade process, the process is suspended. The following scenarios can cause the downgrade process to suspend:

- A node (other than the node that is currently upgrading) is offline, restarted, or has detected an error condition.
- A node fails to update to the new software level.
- A node is deleted while it is in the process of updating.

You must check the error log to determine the reason for the failure before you attempt to upgrade the cluster again.

Recovering from software upgrade problems manually

When a new software level is committed, you might not be able to return to a previous software level because some data structures might have been changed such that they cannot be used with the previous software level. Therefore, if you have any problems, you must install the newest level of the software.

Attention: This procedure causes a loss of *all* data that is currently configured in the cluster. This procedure must only be used as 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 level, you can use the following procedure.

Attention: This procedure causes the total loss of the SAN Volume Controller cluster. This procedure must only be used 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 service access mode.
3. Use the service access mode functions to force the download of the older software level.
4. Repeat the action for each of the failed nodes.
5. Use a node with a new software level to create a new cluster.

Chapter 9. Upgrading the SAN Volume Controller Console

You can download the SAN Volume Controller Console software and upgrade or reinstall an existing SAN Volume Controller Console installation.

Overview of the upgrade or reinstallation process

The following list provides an overview of the upgrade or reinstallation tasks, as well as any configuration tasks that you must perform after you upgrade or reinstall the SAN Volume Controller Console:

1. Upgrade or reinstall the SAN Volume Controller Console in graphical mode with the help of an installation wizard. If errors were generated during the upgrade or reinstallation process, you must remove and reinstall the SAN Volume Controller Console.
2. Verify that the IBM WebSphere Application Server V6 - SVC is installed and started on your system.
3. Use a Web browser to access the SAN Volume Controller Console.
4. Identify the clusters that are to be managed by the SAN Volume Controller Console.

Using the IBM System Storage SAN Volume Controller Installer to upgrade SAN Volume Controller Console

You can upgrade the SAN Volume Controller Console using the IBM System Storage SAN Volume Controller Installer. You can also use this process for reinstalling an existing installation.

Before you can upgrade or reinstall the SAN Volume Controller Console and PuTTY in graphical mode, you must ensure that you have performed the following tasks:

- Ensure that your system meets the IBM System Storage Productivity Center hardware and software requirements provided in the *IBM System Storage Productivity Center Introduction and Planning Guide*
- Download the SAN Volume Controller Console zip file from the following Web site:

  www.ibm.com/storage/support/2145

  After you have downloaded the zip file, you can extract the contents and write it to a CD or you can extract the contents to a directory on your system and perform the installation tasks from that directory.

If you are upgrading, you need to know the following information before initiating the upgrade of the SAN Volume Controller Console:

- The SAN Volume Controller Console superuser password
- The administrator password for each SAN Volume Controller cluster that is being managed by the SAN Volume Controller Console. You need the administrator passwords to ensure that you do not lose access to the SAN Volume Controller Console during the upgrade to 5.1.0. When you upgrade from SAN Volume Controller Console version 4.3.1 to 5.1.0, the cluster superuser password is initially set to the cluster's administrator password from 4.3.1.

If you are upgrading the SAN Volume Controller Console from version 4.3.1 to version 5.1.0 and do not know the administrator password, you can reset it using the front panel of a node, or you can reset the password from the command-line interface by issuing the svctask chcluster -admpwd command.

Note: The svctask chcluster -admpwd command is only available on 4.3.1 and earlier releases.

The SAN Volume Controller Console version 5.1.0 and later does not share the superuser password between all clusters that it manages. Instead you access the same level of function for a particular cluster by using a password that is unique to that cluster. This password is called the *cluster superuser* password. When upgrading from 5.1.0 to a later release, cluster superuser passwords are kept during the upgrade process.

During the upgrade or reinstallation process, you use the IBM System Storage SAN Volume Controller Console Launchpad application. The Launchpad allows you to select from the following options:

Console overview
Provides information about the SAN Volume Controller Console and its components.

Readme file
Provides any last minute product information that is not provided in the topics for installing the SAN Volume Controller Console.

Configuration guide
Provides instructions for installing and configuring the SAN Volume Controller Console.

License agreement
Provides information about the license for the SAN Volume Controller Console.

SAN Volume Controller Web site
Opens the SAN Volume Controller product Web site.

Installation wizard
Starts the 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 SAN Volume Controller Console LaunchPad program.

The SAN Volume Controller Console installation program determines if this is a reinstallation or an upgrade 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 that is 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.

Note: You cannot upgrade to SAN Volume Controller Console version 5.1.0 from releases before version 4.3.1. If you want to upgrade from a 4.3.0 version of the SAN Volume Controller Console, you must first upgrade to version 4.3.1, then upgrade to version 5.1.0.

- If the new code level is lower than the level on the system, the installation is not valid.

Complete the following steps to upgrade the SAN Volume Controller Console:

1. Log on to the system as a local system administrator.
2. Perform one of the following steps:
   - If you wrote the contents of the zip file to a CD and you have autorun mode set on your system, insert the CD into the drive. The IBM System Storage SAN Volume Controller Console Launchpad application starts.
   - If you wrote the contents of the zip file to a CD and you do not have autorun mode set on your system, insert the CD into the drive. Open a command prompt window and change to the W2K directory on the CD.

     Issue the following command:

     ```
     Launchpad
     ```

     The IBM System Storage SAN Volume Controller Console Launchpad application panel is displayed.
   - If you did not write the contents of the zip file to a CD, open a command prompt window and change to the following directory:

     extract_directory\W2K

     Where *extract_directory* is the directory where you extracted the zip file.

     Issue the following command:

     ```
     Launchpad
     ```

     The IBM System Storage SAN Volume Controller Console Launchpad application panel is displayed.
3. Click **Readme file** in the LaunchPad window to read installation information that is specific to this SAN Volume Controller Console software level.
4. Click **Installation wizard** in the LaunchPad window 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. You can click **Exit** if you want to close LaunchPad.

   There might be a slight delay while the software loads on your system. After the software loads, a command prompt window opens to display the following message:

   ```
   Initializing InstallShield Wizard...
   Preparing Java <tm> Virtual Machine ...................
   ...................................................
   ...................................................
   ```

   The Welcome panel for the installation wizard is displayed. The Welcome panel provides the names of the documentation that you should read before you continue with the installation.
5. Click **Next** to continue or **Cancel** to exit the installation. If you click Next, the license agreement panel is displayed.
6. Read the license agreement information and perform one of the following steps:

- Select **I accept the terms of the license agreement** and click **Next** to accept the license agreement.
- Select **I do not accept the terms of the license agreement** and click **Cancel** to exit the installation.

7. Wait while the installation wizard verifies that your system meets all of the requirements. You might have to perform additional steps before the installation process can start if any of the following apply:
   - The installation wizard checks for previous versions of the Service Location Protocol (SLP) service that was installed with SAN Volume Controller Console. If other versions of SLP are installed on the system that are different from the version installed with the SAN Volume Controller Console, they should be removed before completing the installation wizard.
   - During an upgrade from SAN Volume Controller Console version 4.3.1 to version 5.1.0, if the SLP, the IBM System Storage SAN Volume Controller Pegasus Server, or IBM WebSphere Application Server V6 - SVC services are started, the installation program stops these services during the installation. If you are installing a 5.1.0 version of the SAN Volume Controller Console, then the installation program only installs IBM WebSphere Application Server V6 - SVC.

   Note: If you or any of your users plan to use the SAN Volume Controller command-line interface (CLI) in addition to the SAN Volume Controller Console, you can use the **putty-<version>-installer.exe** file that is located in the SSHClient/PuTTY folder that is included as part of the SAN Volume Controller Console zip file to install PuTTY on your system.

   If you are upgrading from SAN Volume Controller Console version 4.3.1 to version 5.1.0, the Product Installation panel displays with the Preserve Configuration option. If you chose to preserve the current configuration under in the backup directory, the installation program skips the next steps and goes directly to the Installation Confirmation panel. If you do not preserve the current configuration or are installing the SAN Volume Controller Console for the first time, the Destination Directory panel is displayed.
8. Select one of the following options from the Destination Directory panel:
   - Click **Next** to accept the default directory.
   - Click **Browse** to select a different directory for installation and then click **Next** to continue the installation process.
   - Click **Cancel** to exit the installation process.

   Note:
   - The directory name, including the drive letter, can be a maximum of 40 characters. When specifying a directory, the following characters cannot be used: \ / : * ?< > !.

   After you click **Next**, the program checks to ensure there is adequate space available on the system for the SAN Volume Controller Console. If the program detects insufficient space for the 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 click **Back**, and select a different destination. After the system verifies the space requirements for the destination directory, the Updating Embedded WAS Ports panel is displayed.

9. Update the default ports assignments by typing unique port numbers (SOAP port, RMI port, HTTP port, and HTTPS port) for the IBM WebSphere Application Server V6 - SVC. The installation program checks for available ports; however, you can check the availability of ports manually. To check the ports manually that are in use, issue the `netstat -a` command and view the `C:\WINNT\system32\drivers\etc\services` file. Click **Next** to continue. The Install Confirmation panel displays.
10. Click **Install** to confirm the installation location and file size and to start the installation. Click **Cancel** to exit the installation wizard or click **Back** to go to the previous panel. If you click **Cancel** a pop-up panel opens and asks you to confirm the cancellation of the installation wizard. Click **Yes** to confirm the cancellation or click **No** to continue the installation. If you confirm the cancellation, the information that you entered or selected in the previous panel is not saved and you must restart the installation process. If you choose to continue the installation, the Installation Progress panel indicates how much of the installation has been completed. Installation usually takes 3 - 10 minutes depending on the configuration of your workstation. If you are upgrading from SAN Volume Controller Console version 4.3.1 to version 5.1.0, the Migration Input panel is displayed. Use this panel to provide user account information to migrate users on the existing clusters to the upgraded version of the clusters. For each cluster IP address, provide the administrator password that is used. The migration status for the cluster displays whether the migration failed or succeeded for each user account that was specified. If the there are errors in the connection to the cluster or with password authorization, they are displayed on the Migration Results panel. Because the Migration Results panel does not display detailed error messages if the migration fails, you can run the migration tool manually after the installation program completes to get detailed error results for the migration.
11. After the installation completes successfully, the Finish panel displays. Click **Finish** to complete the wizard. The wizard displays the post installation task file that provides an overview of the task that need to be complete to complete the configuration of a cluster. You need to reboot your system before launching the SAN Volume Controller Console. When your system restarts, the program automatically starts IBM WebSphere Application Server V6 - SVC.
12. Review the log file for error messages. The log file is located in *install_directory*\logs\install.log, where *install_directory* is the directory where the SAN Volume Controller Console was installed. The install.log file contains a trace of the installation process. After the wizard finishes the Post Installation Readme file displays. This file contains important information on the next steps for completing your configuration. Follow the instructions in this file to complete the post installation tasks for the SAN Volume Controller Console. You can now launch the
13. Use the Services component of the Windows Computer Management utility to verify that the following services Status is set to Started and Startup Type is set to Automatic:
    - IBM WebSphere Application Server V6 - SVC
14. You can now launch the by selecting **Start** → **All Programs** → **IBM System Storage SAN Volume Controller** → **Launch SVC Console**.

    Note: The SAN Volume Controller Console version 5.1.0 does not prompt you for a password when it first starts. Instead it immediately displays the Welcome panel from which you can access the list of available clusters. After selecting a cluster, you are prompted for a password. You can complete either of the following actions to access the cluster:

- Enter `superuser` for the user name and the cluster superuser password, which was the cluster administrator password on the 4.3.1 version of the SAN Volume Controller Console.
- Enter a user name and password of a SAN Volume Controller Console user account that was successfully migrated during the installation. The superuser password that was used to access SAN Volume Controller Console version 4.3.1 is no longer used by the system.

Migrating user accounts manually

During the upgrade from SAN Volume Controller Console version 4.3.1 to version 5.1.0, the installation program attempts to migrate user accounts currently defined on the cluster. If the migration of those accounts fail with the installation program, you can manually migrate the user accounts by using the pmigrate.bat file in the support folder in the directory that you specified to install the SAN Volume Controller Console. The manual method provides detailed messages to indicate the possible cause of the failure.

You must run the manual migration command while the cluster is still running version 4.3.1. After you upgrade the SAN Volume Controller cluster to version 5.1.0, you cannot migrate users.

Note: If the cluster that you are migrating user accounts from is behind a firewall, ensure that you authenticate to the firewall before you attempt to migrate user accounts manually.

The pmigrate.bat file is launched from a Windows command prompt and has the following format:

```
pmigrate PasswordFile=<Location of proxy CIMOM password file>
RoleFile=<Location of proxy CIMOM role file>
CimomIPAddress=<eCIMOM address> CimomPassword=<eCIMOM password>
[CimomUser=<eCIMOM username>]
```

In this format, the CIMOM password is the administrator password for the cluster and the CIMOM IP address is the cluster IP address. The installation program preserves the CIMOM password in the cimserver.passwd file and preserves the CIMOM role definitions in the roles.txt file. Both of these files are stored in the latest backup directory. When you specify these files, include the absolute path to the directory that they are contained in. The `eCIMOM username` must be `superuser`. To migrate user accounts manually, follow these steps:

At a Windows command prompt, enter the following command:

```
"C:\Program Files\IBM\svcconsole\support\pmigrate.bat"
     PasswordFile="C:\Program Files\IBM\svcconsole\backup009
     \cimserver.passwd"
     RoleFile="C:\Program Files\IBM\svcconsole\backup009\roles.txt"
     CimomUser=superuser
     CimomIPAddress=1.1.1.1
     CimomPassword=mypassword
```

In this example, the backup directory is backup009, the CIMOM password file is named cimserver.passwrd, and the CIMOM roles file is named roles.txt.

This example migrates the user accounts successfully with the following results:

```
SVC Command Line Utility Version IBM-SVC-SMIS-Agent-5.1.0
 Returning default value AgentDirectory=<C:\Program Files\IBM\svcconsole\cimom>
 Customization property file=<C:\Program
    Files\IBM\svcconsole\cimom\svcutil.properties> is not present.
 Found CimomIPAddress=<1.1.1.1>  on command line
 Returning default value CimomIPPort=<5989>
 Returning default value CimomProtocol=<https://>
 Found CimomUser=<superuser> on command line
 Connected to: https://1.1.1.1:5989
 Found PasswordFile=<C:\Program Files\IBM\svcconsole\backup009\cimserver.passwd>
    on command line.
 Found  RoleFile=<C:\Program Files\IBM\svcconsole\backup009\roles.txt> on
    command line.
 Enumerating instances of:IBMTSSVC_ObjectManager
 Enumerating instances of:IBMTSSVC_Cluster
 CIM_AGENT_INFORMATION: Connected to CIM Agent running on cluster <MyCluster> at
    software version <4.3. 1.7>
 Enumerating instances of:IBMTSSVC_User
 >>> Invoking method: makeEncrypted
 with input parameters:
 [0] : CIMArgument(name=name, value="newadmin")
 [1] : CIMArgument(name=encryptedPassword, value="admin1")
 [2] : CIMArgument(name=role, value=0)
 <<< Method <makeEncrypted> returned value: 0  ADD_USER_SUCCESS: Successfully
    added user <newadmin> with role <Administrator(0)>
>>> Invoking method: makeEncrypted
 with input parameters:
 [0] : CIMArgument(name=name, value="newmonitor")
 [1] : CIMArgument(name=encryptedPassword, value="monitor1")
 [2] : CIMArgument(name=role, value=3)
 <<< Method <makeEncrypted> returned value: 0
 ADD_USER_SUCCESS: Successfully added user <newmonitor> with role <Monitor(3)>
 Restarting CIM Agent....
 Sending shutdown message to CIM Agent
 >>> Invoking method: BeginShutdown
 CIMOM appears to have exited during shutdown call
```

If the manual migration fails, error messages are displayed in the migration results. Table 27 describes the all the possible error messages that are generated when user account migration fails.

Table 27. Error messages for user account migration and resolutions

Migration Error	Resolution
Customization property file <****> is not present.	Indicates a minor migration problem that does not need to be resolved.
Invalid hostname: ″https://1.1.1.1.1″	Enter the current IP address and protocol for the cluster.
ERROR: Could not find the file <***>	Indicates that the password or role file does not exist in the specified backup directory.
Exception in thread ″main″ org.sblim.wbem.cim.CIMTransportException: EXT_ERR_UNBLE_TO_CONNECT; nested exception is:	Indicates the connection problem to the cluster. This error can occur if there is a network connection problem or if the username or password is not valid. Verify your connection and username and password information.

Verifying the IBM WebSphere Application Server V6 - SVC

You must verify that the IBM WebSphere Application Server V6 - SVC service that is associated with your SAN Volume Controller Console is correctly installed and started.

Perform the following steps to verify that the services are correctly installed:

Important: Do not close the Services window until you are instructed to close it.

Verify the installation of the IBM WebSphere Application Server V6 - SVC service.

1. Verify that the IBM WebSphere Application Server V6 - SVC is started. Select **Start** → **Settings** → **Control Panel**.
2. Double-click the **Administrative Tools** icon.
3. Double-click the **Services** icon.
4. Find IBM WebSphere Application Server V6 - SVC in the **Services** list. For this component, the **Status** column is marked `Started`.
5. If the **IBM WebSphere Application Server V6 - SVC** service is not started, right click **IBM WebSphere Application Server V6 - SVC** and select **Start** from the pop-up menu. Wait for the **Status** column to change to `Started`.
6. Close the Services window.
7. Close the Administrative Tools window.

Uninstalling the SAN Volume Controller Console

You can uninstall the SAN Volume Controller Console from your system.

1. Log on to the system as a local system administrator.
2. Stop the IBM WebSphere Application Server V6 - SVC service:
   a. Click **Start** → **Settings** → **Control Panel**.
   b. In the Control Panel window, double-click on the **Administrative Tools** icon.
   c. Double-click the **Services** icon. The Services window opens.
   d. In the Services window, find `IBM WebSphere Application Server V6 - SVC`. Click on the service to select it.
   e. If the **Status** column shows `Started`, right-click the service and click **Stop** on the menu.
   f. Wait for the service to stop.
3. Use the Windows Add/Remove Programs facility to remove the SAN Volume Controller Console.
   a. From the Windows menu bar, click **Start** → **Settings** → **Control Panel**. Double-click **Add/Remove Programs**.
   b. Click **IBM System Storage SAN Volume Controller Console** from the list of currently installed programs and click **Remove** to remove the product. The Welcome panel for the Uninstaller opens.
4. Click **Next** to continue or click **Cancel** to stop the removal of the SAN Volume Controller Console. The program detects whether the IBM WebSphere Application Server V6 - SVC is running. If this service is running, the uninstaller stops it before proceeding with the uninstallation.
5. Click **Next** to have the program stop the service for you or click **Cancel** to exit the removal process if you wish to manually stop the service. Instructions for

stopping the service begin in step 2 on page 324. You must then restart the removal process from the Windows Add/Remove facility. The Confirmation panel opens.

6. Click **Remove** to continue or click **Cancel** to stop the removal of the SAN Volume Controller Console. Click **Back** to return to the previous panel. The Uninstallation Progress panel opens.
7. Wait for the program to remove the SAN Volume Controller Console product. The Finish panel for the Uninstaller opens.
8. 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 cannot remove information from the system, a **Next** button is displayed 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. Click **Finish** to complete the removal process and exit the wizard.
9. Close the Add/Remove Programs window.
10. If the system has not been restarted since the SAN Volume Controller Console was removed, restart the system now.
11. Log on to the system as a local system administrator.

    Note: The removal process saves files uniquely related to the configuration in a backup directory under the destination path where you installed the SAN Volume Controller Console. Save these files if you plan to reinstall the application. Otherwise you can remove the backup folder and files. An example of the default destination path is: C:\Program Files\IBM\svcconsole.
12. Empty your Windows Recycle Bin to reclaim the disk space that was made available during the removal process.

Chapter 10. Replacing or adding nodes to an existing cluster

You can replace cluster nodes to upgrade to newer hardware models; you can also add cluster nodes to increase your cluster's workload capability.

Replacing nodes nondisruptively

You can replace SAN Volume Controller 2145-8F2, SAN Volume Controller 2145-8F4, or SAN Volume Controller 2145-8G4 nodes with SAN Volume Controller 2145-8A4 or SAN Volume Controller 2145-CF8 nodes. You can also use these procedures if you are replacing a SAN Volume Controller 2145-4F2 node with a SAN Volume Controller 2145-8A4 or earlier node.

These procedures do not include replacing a SAN Volume Controller 2145-4F2 node with a SAN Volume Controller 2145-CF8 node.

For information about replacing a SAN Volume Controller 2145-4F2 node with a SAN Volume Controller 2145-CF8 node, see *Procedures for Replacing SAN Volume Controller 2145-4F2 Nodes with SAN Volume Controller 2145-CF8 Nodes* on the Support for SAN Volume Controller (2145) Web site:

www.ibm.com/storage/support/2145

The link to the procedures is located under **Related Reading** on the page that contains the V5.1.x *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*.

The following procedures are nondisruptive, because changes to your SAN environment are not required. This is because the replacement (new) node uses the same worldwide node name (WWNN) as the node you are replacing.

This task assumes that the following conditions have been met:

- The existing cluster software must be at a version that supports the new node. If a node is being replaced by a SAN Volume Controller 2145-CF8 node, the cluster software version must be 5.1.0. If a node is being replaced by a SAN Volume Controller 2145-8A4 node, the cluster software version must be 4.3.1 or later.
- All nodes that are configured in the cluster are present and online.
- All errors in the cluster error log are addressed and marked as fixed.
- There are no virtual disks (VDisks), managed disks (MDisks), or controllers with a status of degraded or offline.
- The replacement node is not powered on.
- The replacement node is not connected to the SAN.
- You have a 2145 UPS-1U unit (feature code 8115) for each new SAN Volume Controller 2145-CF8 or 2145-8A4 node.
- You have backed up the cluster configuration and saved the svc.config.backup.xml file.
- The replacement node must be able to operate at the fibre-channel or Ethernet connection speed of the node it is replacing.
- If the node being replaced contains solid-state drives (SSDs), all SSDs and SAS adapters should be transferred to the new node if it supports the drives. If the

new node does not support the existing SSDs, you must transfer the data off of the SSDs before replacing the node to avoid losing access to the data.

Important:

1. Do not continue this task if any of the conditions listed above are not met unless you are instructed to do so by the IBM Support Center.
2. Review all of the steps listed below before you perform this task.
3. Do not perform this task if you are not familiar with SAN Volume Controller environments or the procedures described in this task.
4. If you plan to reuse the node that you are replacing, ensure that the WWNN of the node is set to a unique number on your SAN. If you do not ensure that the WWNN is unique, the WWNN and WWPN are duplicated in the SAN environment and can cause issues.

   Tip: You can change the WWNN of the node you are replacing to the factory default WWNN of the replacement node to ensure that the number is unique.
5. The node ID and possibly the node name change during this task. After the cluster assigns the node ID, the ID cannot be changed. However, you can change the node name after this task is complete.

Perform the following steps to replace active nodes in a cluster:

1. (If the cluster software version is at 5.1, complete this step.)

   Confirm that no hosts have dependencies on the node.

   When shutting down a node that is part of a cluster, or when deleting the node from a cluster, use the Show Dependent VDisks menu option on the Viewing Nodes panel in the SAN Volume Controller Console to display all the VDisks that are dependent on a node, or use the svcinfo lsnodedependentvdisk command to view dependent VDisks.

   If dependent VDisks exist, determine if the VDisks are being used. If the VDisks are being used, either restore the redundant configuration or suspend the host application. If a dependent quorum disk is reported, repair the access to the quorum disk or modify the quorum disk configuration.
2. Perform the following steps to determine the cluster configuration node, and the ID, name, I/O group ID, and I/O group name for the node that you want to replace. If you already know the physical location of the node that you want to replace, you can skip this step and proceed to step 3 on page 329.

   Tip: If any of the nodes you want to replace are the cluster configuration node, replace it last.

   a. Issue the following command from the command-line interface (CLI):

      ```
      svcinfo lsnode -delim :
      ```

      The following is an example of the output that is displayed for this command:

      ```
      id:name:UPS_serial_number:WWNN:status:IO_group_id:IO_group_name:
      config_node:UPS_unique_id:hardware:iscsi_name:iscsi_alias
      3:dvt113294:100089J137:5005076801005A07:online:0:io_grp0:yes:
      20400002096810C7:8A4:iqn.1986-03.com.ibm:2145.ldcluster-80.dvt113294:
      14:des113004:10006BR010:5005076801004F0F:online:0:io_grp0:no:
      2040000192880040:8G4:iqn.1986-03.com.ibm:2145.ldcluster-80.des113004:
      ```

b. In the config_node column, find the value yes and record the values in the id and name columns.
c. Record the values in the id and the name columns for each node in the cluster.
d. Record the values in the IO_group_id and the IO_group_name columns for each node in the cluster.
e. Issue the following command from the CLI for each node in the cluster to determine the front panel ID:

   ```
   svcinfo lsnodevpd node_name or node_id
   ```

   where *node_name or node_id* is the name or ID of the node for which you want to determine the front panel ID.
f. Record the value in the front_panel_id column. The front panel ID is displayed on the front of each node. You can use this ID to determine the physical location of the node that matches the node ID or node name that you want replace.

3. Perform the following steps to record the WWNN or iSCSI name of the node that you want to replace:
   a. Issue the following command from the CLI:

      ```
      svcinfo lsnode -delim : node_name or node_id
      ```

      where *node_name or node_id* is the name or ID of the node for which you want to determine the WWNN or iSCSI name.
   b. Record the WWNN or iSCSI name of the node that you want to replace. Also record the order of the fibre-channel and Ethernet ports.
4. Issue the following command from the CLI to power off the node:

   ```
   svctask stopcluster -node node_name
   ```

   Important:
   a. Record and mark the order of the fibre-channel or Ethernet cables with the node port number (port 1 to 4 for fibre-channel, or port 1 to 2 for ethernet) before you remove the cables from the back of the node. The fibre-channel ports on the back of the node are numbered 1 to 4 from left to right. You must reconnect the cables in the exact order on the replacement node to avoid issues when the replacement node is added to the cluster. If the cables are not connected in the same order, the port IDs can change, which impacts the ability of the host to access VDisks. See the hardware documentation specific to your model to determine how the ports are numbered.
   b. Do not connect the replacement node to different ports on the switch or director. The SAN Volume Controller can have 4 or 8 Gbps HBAs; however, do not move them to faster switch or director ports at this time to avoid issues when the replacement node is added to the cluster.
   c. Do not move the fibre-channel cables of the node to faster or different ports on the switch or director at this time. This is a separate task that must be planned independently of replacing nodes in a cluster.
5. Issue the following CLI command to delete this node from the cluster and I/O group:

   ```
   svctask rmnode node_name or node_id
   ```

Where *node_name or node_id* is the name or ID of the node that you want to delete. You can use the CLI to verify that the deletion process has completed.

6. Issue the following CLI command to ensure that the node is no longer a member of the cluster:

   ```
   svcinfo lsnode
   ```

   A list of nodes is displayed. Wait until the removed node is not listed in the command output.

7. Perform the following steps to change the WWNN or iSCSI name of the node that you just deleted from the cluster to FFFFF:

   For SAN Volume Controller V4.3 or later:

   a. Power on the node.
   b. With the Node WWNN: panel displayed, press and hold the down button, press and release the select button, and then release the down button. The display switches into edit mode. `Edit WWNN` is displayed on line 1. Line 2 of the display contains the last five numbers of the WWNN.
   c. Change the displayed number to `FFFFF`. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.
   d. Press the select button to save your changes and apply `FFFFF` as the new WWNN for the node.

   For SAN Volume Controller versions prior to V4.3:

   a. Power on the node.
   b. Press and release the right button until the Status: panel is displayed.
   c. 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 numbers of the WWNN.
   d. With the WWNN displayed on the front panel; press and hold the down button; press and release the select button; release the down button. The display switches into edit mode.
   e. Change the displayed number to `FFFFF`. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.
   f. Press the select button to retain the numbers that you have updated and return to the WWNN panel.
   g. Press the select button to apply the numbers as the new WWNN for the node.
   h. Remove the node to be replaced and optionally the uninterruptible power supply from the rack.

8. Install the replacement node and the uninterruptible power supply in the rack and connect the uninterruptible power supply cables. See the *IBM System Storage SAN Volume Controller Model 2145-XXX Hardware Installation Guide* to determine how to connect the node and the uninterruptible power supply.

   Important: Do not connect the fibre-channel or Ethernet cables during this step.

9. Power on the replacement node.

10. Record the WWNN or iSCSI nameof the replacement node. You can use this name if you plan to reuse the node that you are replacing.
11. Perform the following steps to change the WWNN or iSCSI nameof the replacement node to match the name that you recorded in step 3 on page 329:

    For SAN Volume Controller V4.3 or later:

    a. With the Node WWNN: panel displayed, press and hold the down button, press and release the select button, and then release the down button. The display switches into edit mode. `Edit WWNN` is displayed on line 1. Line 2 of the display contains the last five numbers of the WWNN.
    b. Change the WWNN that is displayed to match the last five numbers of the WWNN that you recorded in step 3 on page 329. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.
    c. Press the select button to apply the numbers as the new WWNN for the node.

    For SAN Volume Controller versions prior to V4.3:

    a. Press and release the right button until the Status: panel is displayed.
    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 numbers of the WWNN.
    c. With the WWNN displayed on the front panel; press and hold the down button; press and release the select button; release the down button. The display switches into edit mode. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.
    d. When the five numbers match the last five numbers of the WWNN that you recorded in step 3 on page 329, press the select button to retain the numbers that you have updated and return to the WWNN panel.
    e. Press the select button to apply the numbers as the new WWNN for the node.

    Wait one minute. If `Cluster:` is displayed on the front panel, this indicates that the node is ready to be added to the cluster. If `Cluster:` does not display, see the *IBM System Storage SAN Volume Controller Troubleshooting Guide* to determine how to address this problem or contact the IBM Support Center before you continue with the next step.
12. Connect the fibre-channel or Ethernet cables to the same port numbers that you recorded for the original node in step 4 on page 329.
13. Issue the following CLI command to verify that the last five characters of the WWNN are correct:

    ```
    svcinfo lsnodecandidate
    ```

    Important: If the WWNN is not what you recorded in step 3 on page 329, you must repeat step 11.
14. Issue the following CLI command to add the node to the cluster and ensure that the node has the same name as the original node and is in the same I/O group as the original node. See the svctask addnode CLI command documentation for more information.

    ```
    svctask addnode -wwnodename WWNN -iogrp iogroupname/id
    ```

where *WWNN* and *iogroupname/id* are the values that you recorded for the original node.

The SAN Volume Controller V5.1 automatically reassigns the node with the name that was used originally. For versions prior to V5.1, use the **name** parameter with the svctask addnode command to assign a name. If the original node's name was automatically assigned by SAN Volume Controller, it is not possible to reuse the same name. It was automatically assigned if its name starts with `node`. In this case, either specify a different name that does not start with `node` or do not use the **name** parameter so that SAN Volume Controller automatically assigns a new name to the node.

If necessary, the new node is updated to the same SAN Volume Controller software version as the cluster. This update can take up to 20 minutes.

Important:

a. Both nodes in the I/O group cache data; however, the cache sizes are asymmetric. The replacement node is limited by the cache size of the partner node in the I/O group. Therefore, it is possible that the replacement node does not utilize the full cache size until you replace the other node in the I/O group.

b. You do not have to reconfigure the host multipathing device drivers because the replacement node uses the same WWNN and WWPN as the previous node. The multipathing device drivers should detect the recovery of paths that are available to the replacement node.

c. The host multipathing device drivers take approximately 30 minutes to recover the paths. Do not upgrade the other node in the I/O group until for at least 30 minutes after you have successfully upgraded the first node in the I/O group. If you have other nodes in different I/O groups to upgrade, you can perform those upgrades while you wait.

15. See the documentation that is provided with your multipathing device driver for information on how to query paths to ensure that all paths have been recovered before proceeding to the next step. If you are using the IBM System Storage Multipath Subsystem Device Driver (SDD), the command to query paths is **datapath query device**.
16. Repeat steps 3 on page 329 to 15 for each node that you want to replace.

Replacing nodes disruptively (rezoning the SAN)

You can replace SAN Volume Controller 2145-8F2, SAN Volume Controller 2145-8F4, SAN Volume Controller 2145-8A4, or SAN Volume Controller 2145-8G4 nodes with SAN Volume Controller 2145-8A4 or SAN Volume Controller 2145-CF8 nodes. The following procedures are disruptive, because you do not use the same WWNN and WWPNs for the new node. You must rezone your SAN and the host multipathing device drivers must discover new paths. Access to virtual disks (VDisks) is lost during this task.

This task assumes that the following conditions exist:

- The cluster software is at 5.1.0 or higher
- All nodes that are configured in the cluster are present
- All errors in the cluster error log are fixed
- All managed disks (MDisks) are online

- You have a 2145 UPS-1U unit for each new SAN Volume Controller 2145-8G4 node.

Perform the following steps to replace nodes:

1. (If the cluster software version is at 5.1, complete this step.)

   Confirm that no hosts have dependencies on the node.

   When shutting down a node that is part of a cluster, or when deleting the node from a cluster, use the Show Dependent VDisks menu option on the Viewing Nodes panel in the SAN Volume Controller Console to display all the VDisks that are dependent on a node, or use the svcinfo lsnodedependentvdisk command to view dependent VDisks.

   If dependent VDisks exist, determine if the VDisks are being used. If the VDisks are being used, either restore the redundant configuration or suspend the host application. If a dependent quorum disk is reported, repair the access to the quorum disk or modify the quorum disk configuration.
2. Quiesce all I/O from the hosts that access the I/O group of the node that you are replacing.
3. Delete the node that you want to replace from the cluster and I/O group.

   Notes:

   a. The node is not deleted until the SAN Volume Controller cache is destaged to disk. During this time, the partner node in the I/O group transitions to write through mode.
   b. You can use the command-line interface (CLI) or the SAN Volume Controller Console to verify that the deletion process has completed.
4. Ensure that the node is no longer a member of the cluster.
5. Power-off the node and remove it from the rack.
6. Install the replacement (new) node in the rack and connect the uninterruptible power supply cables and the fibre-channel cables.
7. Power-on the node.
8. Rezone your switch zones to remove the ports of the node that you are replacing from the host and storage zones. Replace these ports with the ports of the replacement node.
9. Add the replacement node to the cluster and I/O group.

   Important: Both nodes in the I/O group cache data; however, the cache sizes are asymmetric if the remaining partner node in the I/O group is a SAN Volume Controller 2145-4F2 node. The replacement node is limited by the cache size of the partner node in the I/O group. Therefore, the replacement node does not use the full 8 GB cache size until you replace the other SAN Volume Controller 2145-4F2 node in the I/O group.
10. From each host, issue a rescan of the multipathing software to discover the new paths to VDisks.

    Notes:

    a. If your system is inactive, you can perform this step after you have replaced all nodes in the cluster.
    b. The host multipathing device drivers take approximately 30 minutes to recover the paths.

11. See the documentation that is provided with your multipathing device driver for information on how to query paths to ensure that all paths have been recovered before proceeding to the next step.
12. Repeat steps 2 on page 333 to 11 for the partner node in the I/O group.
13. Repeat steps 2 on page 333 to 12 for each node in the cluster that you want to replace.
14. Resume host I/O.

Replacing nodes disruptively (moving VDisks to new I/O group)

You can replace SAN Volume Controller 2145-8F2, SAN Volume Controller 2145-8F4, or SAN Volume Controller 2145-8A4 nodes with SAN Volume Controller 2145-8A4 or SAN Volume Controller 2145-CF8 nodes. The following procedures are disruptive, because you move VDisks to a new I/O group.

This task assumes the following:

- The cluster software is at 5.1.0 or higher
- Your cluster contains six or less nodes
- All nodes that are configured in the cluster are present
- All errors in the cluster error log are fixed
- All managed disks (MDisks) are online
- You have a 2145 UPS-1U unit for each new SAN Volume Controller 2145-8G4 node.

Perform the following steps to replace nodes:

1. (If the cluster software version is at 5.1, complete this step.)

   Confirm that no hosts have dependencies on the node.

   When shutting down a node that is part of a cluster, or when deleting the node from a cluster, use the Show Dependent VDisks menu option on the Viewing Nodes panel in the SAN Volume Controller Console to display all the VDisks that are dependent on a node, or use the svcinfo lsnodedependentvdisk command to view dependent VDisks.

   If dependent VDisks exist, determine if the VDisks are being used. If the VDisks are being used, either restore the redundant configuration or suspend the host application. If a dependent quorum disk is reported, repair the access to the quorum disk or modify the quorum disk configuration.
2. Quiesce all I/O from the hosts that access the I/O groups of the nodes that you are replacing.
3. Add two replacement nodes to the cluster to create a new I/O group.
4. Rezone your switch zones to add the ports of the new nodes to the host and storage zones.
5. Move all of the VDisks from the I/O group of the nodes that you are replacing to the new I/O group.
6. From each host, issue a rescan of the multipathing software to discover the new paths to VDisks. The host multipathing device drivers take approximately 30 minutes to recover the paths.
7. See the documentation that is provided with your multipathing device driver for information on how to query paths to ensure that all paths have been recovered before proceeding to the next step.
8. Delete the nodes that you are replacing from the cluster and remove the ports from the switch zones.

9. Repeat steps 2 on page 334 to 8 on page 334 for each node in the cluster that you want to replace.

Adding SAN Volume Controller 2145-CF8 nodes to an existing cluster

You can add SAN Volume Controller 2145-CF8 nodes to your cluster.

This task requires that the following conditions are met:

- The cluster SAN Volume Controller software version is 5.1.0 or later.

  Note: SAN Volume Controller 2145-4F2 nodes are not supported on this version or later versions.
- All nodes that are configured in the cluster are present.
- All errors in the cluster error log are fixed.
- All managed disks (MDisks) are online.

1. Install the SAN Volume Controller 2145-CF8 nodes and the 2145 UPS-1U units in the rack.
2. Connect the SAN Volume Controller 2145-CF8 nodes to the LAN.
3. Connect the SAN Volume Controller 2145-CF8 nodes to the SAN fabric.
4. Power on the SAN Volume Controller 2145-CF8 nodes and the 2145 UPS-1U units.
5. Zone the SAN Volume Controller 2145-CF8 node ports in the existing SAN Volume Controller zone. The SAN Volume Controller zone exists in each fabric with only node ports.
6. Zone the SAN Volume Controller 2145-CF8 node ports in the existing SAN Volume Controller and storage zone. A storage zone contains all of the SAN Volume Controller 2145-CF8 node ports and controller ports that are in the fabric and used to access the physical disks.
7. For each system that is used with the SAN Volume Controller cluster, use the system management application to map the LUNs that are currently used by the cluster to all of the WWPNs of the SAN Volume Controller 2145-CF8 nodes that you want to add to the cluster. The SAN Volume Controller 2145-CF8 nodes must recognize the same LUNs that the existing nodes in the cluster can recognize before they can be added to the cluster. If the SAN Volume Controller 2145-8A4 nodes cannot recognize the same LUNs, the system is marked degraded.
8. Add the SAN Volume Controller 2145-CF8 nodes to the cluster.
9. Check the status of the systems and MDisks to ensure that status has not been marked degraded. If the status is degraded, there is a configuration problem that must be resolved before any further cluster configuration tasks can be performed. If the problem cannot be resolved, remove the newly added SAN Volume Controller 2145-CF8 nodes from the cluster and contact the IBM Support Center for assistance.

Adding SAN Volume Controller 2145-8A4 nodes to an existing cluster

You can add SAN Volume Controller 2145-8A4 nodes to your cluster.

This task requires that the following conditions are met:

- The cluster software version is 4.3.1 or higher
- All nodes that are configured in the cluster are present
- All errors in the cluster error log are fixed

- All managed disks (MDisks) are online

1. Install the SAN Volume Controller 2145-8A4 nodes and the 2145 UPS-1U units in the rack.
2. Connect the SAN Volume Controller 2145-8A4 nodes to the LAN.
3. Connect the SAN Volume Controller 2145-8A4 nodes to the SAN fabric.
4. Power on the SAN Volume Controller 2145-8A4 nodes and the 2145 UPS-1U units.
5. Zone the SAN Volume Controller 2145-8A4 node ports in the existing SAN Volume Controller zone. The SAN Volume Controller zone exists in each fabric with only node ports.
6. Zone the SAN Volume Controller 2145-8A4 node ports in the existing SAN Volume Controller and storage zone. A storage zone contains all of the SAN Volume Controller 2145-8A4 node ports and controller ports that are in the fabric and used to access the physical disks.
7. For each system that is used with the SAN Volume Controller cluster, use the system management application to map the LUNs that are currently used by the cluster to all of the WWPNs of the SAN Volume Controller 2145-8A4 nodes that you want to add to the cluster. The SAN Volume Controller 2145-8A4 nodes must recognize the same LUNs that the existing nodes in the cluster can recognize before they can be added to the cluster. If the SAN Volume Controller 2145-8A4 nodes cannot recognize the same LUNs, the system is marked degraded.
8. Add the SAN Volume Controller 2145-8A4 nodes to the cluster.
9. Check the status of the systems and MDisks to ensure that status has not been marked degraded. If the status is degraded, there is a configuration problem that must be resolved before any further cluster configuration tasks can be performed. If the problem cannot be resolved, remove the newly added SAN Volume Controller 2145-8A4 nodes from the cluster and contact the IBM Support Center for assistance.

Adding SAN Volume Controller 2145-8G4 nodes to an existing cluster

You can add SAN Volume Controller 2145-8G4 to increase the size of your cluster.

This task assumes that the following conditions exist:

- The cluster software version is 4.2.0 or higher
- All nodes that are configured in the cluster are present
- All errors in the cluster error log are fixed
- All managed disks (MDisks) are online

1. Install the SAN Volume Controller 2145-8G4 nodes and the 2145 UPS-1U units in the rack.
2. Connect the SAN Volume Controller 2145-8G4 nodes to the LAN.
3. Connect the SAN Volume Controller 2145-8G4 nodes to the SAN fabric.
4. Power on the SAN Volume Controller 2145-8G4 nodes and the 2145 UPS-1U units.
5. Zone the SAN Volume Controller 2145-8G4 node ports in the existing SAN Volume Controller zone. A SAN Volume Controller zone should exist in each fabric with only node ports.
6. Zone the SAN Volume Controller 2145-8G4 node ports in the existing SAN Volume Controller and storage zone. A storage zone should contain all of the

SAN Volume Controller 2145-8G4 node ports and controller ports that are in the fabric and used to access the physical disks.

7. For each system that is used with the SAN Volume Controller cluster, use the system management application to map the LUNs that are currently used by the cluster to all of the WWPNs of the SAN Volume Controller 2145-8G4 nodes that you want to add to the cluster. The SAN Volume Controller 2145-8G4 nodes must see the same LUNs that the existing nodes in the cluster can see before they can be added to the cluster. If the SAN Volume Controller 2145-8G4 nodes cannot see the same LUNs, the system is marked degraded.
8. Add the SAN Volume Controller 2145-8G4 nodes to the cluster.
9. Check the status of the systems and MDisks to ensure that status has not been marked degraded. If the status is degraded, there is a configuration problem that must be resolved before any further cluster configuration tasks can be performed. If the problem cannot be resolved, remove the newly added SAN Volume Controller 2145-8G4 nodes from the cluster and contact the IBM Support Center for assistance.

Adding SAN Volume Controller 2145-8F4 nodes to an existing cluster

You can add SAN Volume Controller 2145-8F4 to increase the size of your cluster.

This task assumes that the following conditions exist:

- The cluster software version is 4.2.0 or higher
- All nodes that are configured in the cluster are present
- All errors in the cluster error log are fixed
- All managed disks (MDisks) are online

1. Install the SAN Volume Controller 2145-8F4 nodes and the 2145 UPS-1U units in the rack.
2. Connect the SAN Volume Controller 2145-8F4 nodes to the LAN.
3. Connect the SAN Volume Controller 2145-8F4 nodes to the SAN fabric.
4. Power on the SAN Volume Controller 2145-8F4 nodes and the 2145 UPS-1U units.
5. Zone the SAN Volume Controller 2145-8F4 node ports in the existing SAN Volume Controller zone. A SAN Volume Controller zone should exist in each fabric with only node ports.
6. Zone the SAN Volume Controller 2145-8F4 node ports in the existing SAN Volume Controller and storage zone. A storage zone should contain all of the SAN Volume Controller 2145-8F4 node ports and controller ports that are in the fabric and used to access the physical disks.
7. For each system that is used with the SAN Volume Controller cluster, use the system management application to map the LUNs that are currently used by the cluster to all of the WWPNs of the SAN Volume Controller 2145-8F4 nodes that you want to add to the cluster. The SAN Volume Controller 2145-8F4 nodes must see the same LUNs that the existing nodes in the cluster can see before they can be added to the cluster. If the SAN Volume Controller 2145-8F4 nodes cannot see the same LUNs, the system is marked degraded.
8. Add the SAN Volume Controller 2145-8F4 nodes to the cluster.
9. Check the status of the systems and MDisks to ensure that status has not been marked degraded. If the status is degraded, there is a configuration problem that must be resolved before any further cluster configuration tasks can be

performed. If the problem cannot be resolved, remove the newly added SAN Volume Controller 2145-8F4 nodes from the cluster and contact theIBM Support Center for assistance.

Adding SAN Volume Controller 2145-8F2 nodes to an existing cluster

You can add SAN Volume Controller 2145-8F2 nodes to increase the size of your cluster.

This task assumes that the following conditions exist:

- The cluster software version is 3.1.0 or higher
- All nodes that are configured in the cluster are present
- All errors in the cluster error log are fixed
- All managed disks (MDisks) are online

1. Install the SAN Volume Controller 2145-8F2 nodes and the uninterruptible power supply units in the rack.
2. Connect the SAN Volume Controller 2145-8F2 nodes to the LAN.
3. Connect the SAN Volume Controller 2145-8F2 nodes to the SAN fabric.
4. Power on the SAN Volume Controller 2145-8F2 nodes and the uninterruptible power supply units.
5. Zone the SAN Volume Controller 2145-8F2 node ports in the existing SAN Volume Controller zone. A SAN Volume Controller zone should exist in each fabric with only node ports.
6. Zone the SAN Volume Controller 2145-8F2 node ports in the existing SAN Volume Controller and storage zone. A storage zone should contain all of the SAN Volume Controller 2145-8F2 node ports and controller ports that are in the fabric and used to access the physical disks.
7. For each system that is used with the SAN Volume Controller cluster, use the system management application to map the LUNs that are currently used by the cluster to all of the WWPNs of the SAN Volume Controller 2145-8F2 nodes that you want to add to the cluster. The SAN Volume Controller 2145-8F2 nodes must see the same LUNs that the existing nodes in the cluster can see before they can be added to the cluster. If the SAN Volume Controller 2145-8F2 nodes cannot see the same LUNs, the system is marked degraded.
8. Add the SAN Volume Controller 2145-8F2 nodes to the cluster.
9. Check the status of the systems and MDisks to ensure that status has not been marked degraded. If the status is degraded, there is a configuration problem that must be resolved before any further cluster configuration tasks can be performed. If the problem cannot be resolved, remove the newly added SAN Volume Controller 2145-8F2 nodes from the cluster and contact theIBM Support Center for assistance.

Replacing a faulty node in the cluster using the CLI

You can use the command-line interface (CLI) and the SAN Volume Controller front panel to replace a faulty node in the cluster.

Before you attempt to replace a faulty node with a spare node you must ensure that you meet the following requirements:

- 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 must make a record of the last five characters of the original worldwide node name (WWNN) of the spare node. If you repair a faulty node, and you want to make it a spare node, you can use the WWNN of the node. You do not want to duplicate the WWNN because it is unique. It is easier to swap in a node when you use the WWNN.

Attention: Never connect a node with a WWNN of `00000` to the cluster. If this node is no longer required as a spare and is to be used for normal attachment to a cluster, you must change the WWNN to the number you recorded when a spare was created. Using any other number might cause data corruption.

If a node fails, the cluster continues to operate with degraded performance until the faulty node is repaired. If the repair operation takes an unacceptable amount of time, it is useful to replace the faulty node with a spare node. However, the appropriate procedures must be followed and precautions must be taken so you do *not* interrupt I/O operations and compromise the integrity of your data.

The following table describes the changes that are made to your configuration when you replace a faulty node in the cluster:

Node attributes	Description
Front panel ID	This is the number that is printed on the front of the node and is used to select the node that is added to a cluster.
Node ID	This is the ID that is assigned to the node. 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. This ID will change during this procedure.
Node name	This is the name that is assigned to the node. If you are using SAN Volume Controller version 5.1.0 nodes, the SAN Volume Controller automatically re-adds nodes that have failed back to the cluster. If the cluster reports an error for a node missing (error code 1195) and that node has been repaired and restarted, the cluster automatically re-adds the node back into the cluster. For releases prior to 5.1.0, if you do not specify a name, the SAN Volume Controller assigns a default name. The SAN Volume Controller creates a new default name each time a node is added to a cluster. If you choose to assign your own names, you must type the node name on the Adding a node to a cluster panel. You cannot manually assign a name that matches the naming convention used for names assigned automatically by SAN Volume Controller. If you are using scripts to perform management tasks on the cluster and those scripts use the node name, you can avoid the need to make changes to the scripts by assigning the original name of the node to a spare node. This name might change during this procedure.
Worldwide node name	This is the WWNN that is assigned to the node. The WWNN is used to uniquely identify the node and the fibre-channel ports. During this procedure, the WWNN of the spare node changes to that of the faulty node. The node replacement procedures must be followed exactly to avoid any duplication of WWNNs. This name does not change during this procedure.

Node attributes	Description

Worldwide port names	These are the WWPNs that are assigned to the node. WWPNs are derived from the WWNN that is written to the spare node as part of this procedure. For example, if the WWNN for a node is `50050768010000F6`, the four WWPNs for this node are 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`
	These names do not change during this procedure.

Complete the following steps to replace a faulty node in the cluster:

1. Verify the name and ID of the node that you want to replace.

   Complete the following step to verify the name and ID:

   a. Issue the **svcinfo lsnode** CLI command to ensure that the partner node in the I/O group is online.

   - If the other node in the I/O group is offline, start Directed Maintenance Procedures (DMPs) to determine the fault.
   - 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 after a node or an I/O group failed.
   - If you are replacing the node for other reasons, determine the node you want to replace and ensure that the partner node in the I/O group is online.
   - If the partner node is offline, you will lose access to the VDisks that belong to this I/O group. Start the DMPs and fix the other node before proceeding to the next step.

2. Find and record the following information about the faulty node using Steps 2a through 2h:

   - Node serial number
   - Worldwide node name
   - All of the worldwide port names
   - Name or ID of the I/O group that contains the node
   - Front panel ID
   - Uninterruptible power supply serial number

   a. Issue the **svcinfo lsnode** CLI command to find and record the node name and I/O group name. The faulty node will be offline.

   b. Issue the following CLI command:

      `svcinfo lsnodevpd` *nodename*

      Where *nodename* is the name that you recorded in step 2a.

   c. Find the `WWNN` field in the output.

   d. Record the last five characters of the WWNN.

   e. Find the `front_panel_id` field in the output.

   f. Record the front panel ID.

   g. Find the `UPS_serial_number` field in the output.

   h. Record the uninterruptible power supply serial number.

3. Ensure that the faulty node has been powered off.
4. Issue the following CLI command to remove the faulty node from the cluster:

   `svctask rmnode` *nodename/id*

   Where *nodename/id* is the name or ID of the faulty node.
5. Disconnect all four fibre-channel cables from the node.

   Important: Do not plug the fibre-channel cables into the spare node until the spare node is configured with the WWNN of the faulty node.
6. Connect the power and signal cables from the spare node to the uninterruptible power supply that has the serial number you recorded in step 2h on page 340.

   Note: For 2145 UPS-1U units, you must disconnect the cables from the faulty node.
7. Disconnect the faulty node's power and serial cable from the 2145 UPS-1U and connect the new node's power and signal cable in their place.
8. Power on the spare node.
9. Display the node status on the front-panel display.
10. You must change the WWNN of the spare node to that of the faulty node. The procedure for doing this depends on the SAN Volume Controller version that is installed on the spare node. Press and release the down button until the Node: panel displays. Then press and release the right button until the WWNN: panel displays. If repeated pressing of the right button returns you to the Node: panel, without displaying a Node WWNN: panel, go to step 12; otherwise, continue with step 11.
11. Change the WWNN of the spare node (with SAN Volume Controller V4.3 and above installed) to match the WWNN of the faulty node by completing the following steps:
    a. With the Node WWNN: panel displayed, press and hold the down button, press and release the select button, and then release the down button.The display switches into edit mode. `Edit WWNN` is displayed on line 1. Line 2 of the display contains the last five numbers of the WWNN.
    b. Change the WWNN that is displayed to match the last five numbers of the WWNN that you recorded in step 13 on page 342. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.
    c. When the five numbers match the last five numbers of the WWNN that you recorded in step 2d on page 340, press the select button to accept the numbers.
12. Change the WWNN of the spare node (with SAN Volume Controller versions prior to V4.3 installed) to match the WWNN of the faulty node by performing the following steps:
    a. Press and release the right button until the Status: panel is displayed.
    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 numbers of the WWNN.
    c. With the WWNN displayed on the front panel; press and hold the down button; press and release the select button; release the down button. The display switches into edit mode.

d. Change the WWNN that is displayed to match the last five numbers of the WWNN that you recorded in step 2d on page 340. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.

e. When the five numbers match the last five numbers of the WWNN that you recorded in step 2d on page 340, press the select button to retain the numbers that you have updated and return to the WWNN panel.

f. Press the select button to apply the numbers as the new WWNN for the node.

13. Connect the four fibre-channel cables that you disconnected from the faulty node to the spare node.

    If the spare node has less Ethernet cables connected than the faulty node, move the Ethernet cables from the faulty node to the spare node. Ensure you connect the cable into the same port on the spare node as it was in on the faulty node.

14. Issue the following command to add the spare node to the cluster:

    ```
    svctask addnode -wwnodename WWNN -iogrp iogroupname/id
    ```

    where *WWNN* and *iogroupname/id* are the values that you recorded for the original node.

    The SAN Volume Controller V5.1 automatically reassigns the node with the name that was used originally. For versions prior to V5.1, use the **name** parameter with the svctask addnode command to assign a name. If the original node's name was automatically assigned by SAN Volume Controller, it is not possible to reuse the same name. It was automatically assigned if its name starts with `node`. In this case, either specify a different name that does not start with `node` or do not use the **name** parameter so that SAN Volume Controller automatically assigns a new name to the node.

    If necessary, the new node is updated to the same SAN Volume Controller software version as the cluster. This update can take up to 20 minutes.

15. Use the tools that are provided with your multipathing device driver on the host systems to verify that all paths are now online. See the documentation that is provided with your multipathing device driver for more information. For example, if you are using the subsystem device driver (SDD), see the *IBM System Storage Multipath Subsystem Device Driver User's Guide* for instructions on how to use the SDD management tool on host systems. It might take up to 30 minutes for the paths to come online.

16. Repair the faulty node.

    Attention: When the faulty node is repaired, do not connect the fibre-channel cables to it. Connecting the cables might cause data corruption because the spare node is using the same WWNN as the faulty node.

    If you want to use the repaired node as a spare node, perform the following steps.

    For SAN Volume Controller V4.3 and later versions:

    a. With the Node WWNN: panel displayed, press and hold the down button, press and release the select button, and then release the down button.

    b. The display switches into edit mode. `Edit WWNN` is displayed on line 1. Line 2 of the display contains the last five numbers of the WWNN.

    c. Change the displayed number to `00000`. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.

d. Press the select button to accept the numbers.

   This node can now be used as a spare node.

For SAN Volume Controller versions prior to V4.3:

a. Press and release the right button until the Status: panel is displayed.

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 numbers of the WWNN.

c. With the WWNN displayed on the front panel; press and hold the down button; press and release the select button; release the down button. The display switches into edit mode.

d. Change the displayed number to `00000`. To edit the highlighted number, use the up and down buttons to increase or decrease the numbers. The numbers wrap F to 0 or 0 to F. Use the left and right buttons to move between the numbers.

e. Press the select button to accept the numbers.

f. Press the select button to retain the numbers that you have updated and return to the WWNN panel.

   This node can now be used as a spare node.

Chapter 11. Configuring and servicing storage systems

To avoid performance issues, you must ensure that your storage systems and switches are correctly configured to work with the SAN Volume Controller.

Virtualization provides many benefits over direct-attached or direct SAN-attached storage systems. However, virtualization is more susceptible to performance hot spots than direct-attached storage. Hot spots can cause I/O errors on your hosts and can potentially cause a loss of access to data.

Identifying your storage system

The serial number that is presented by the command-line interface (CLI) and the SAN Volume Controller Console for the SAN Volume Controller is the serial number of the device.

The serial numbers can be viewed on your storage system. If the serial numbers are not displayed, the worldwide node name (WWNN) or worldwide port name (WWPN) is displayed. The WWNN or WWPN can be used to identify the different storage systems.

Configuration guidelines for storage systems

You must follow the guidelines and procedures for your storage system to maximize performance and to avoid potential I/O problems.

General guidelines

You must follow these general guidelines when configuring your storage systems.

- Avoid splitting arrays into multiple logical disks at the storage system level. Where possible, create a single logical disk from the entire capacity of the array.
- Depending on the redundancy that is required, create RAID-5 arrays using between 5 and 8 plus parity components. That is 5 + P, 6 + P, 7 + P or 8 + P.
- Do not mix managed disks (MDisks) that greatly vary in performance in the same MDisk group. The overall MDisk group performance is limited by the slowest MDisk. Because some disk controllers can sustain much higher I/O bandwidths than others, do not mix MDisks that are provided by low-end storage systems with those that are provided by high-end storage systems. You must consider the following factors:
  - The underlying RAID type that the storage system is using to implement the MDisk.
  - The number of physical disks in the RAID array and the physical disk type. For example: 10K/15K rpm, FC/SATA.
- When possible, include similarly sized MDisks in an MDisk group. This makes it easier to balance the MDisks in the group. If the MDisks in an MDisk group are significantly different sizes, you can balance the proportion of space that is allocated on each MDisk by including the larger MDisk multiple times in the MDisk list. This is specified when you create a new VDisk. For example, if you have two 400 MB disks and one 800 MB disk that are identified as MDisk 0, 1, and 2, you can create the striped VDisk with the MDisk IDs of 0:1:2:2. This

doubles the number of extents on the 800 MB drive, which accommodates it being double the size of the other MDisks.

- Avoid leaving VDisks in image mode. Only use image mode to import existing data into the cluster. To optimize the benefits of virtualization, migrate this data across the other MDisks in the group as soon as possible.
- Follow the FlashCopy feature requirements before you set up the storage. Balance the spread of the FlashCopy VDisks across the MDisk groups and then across the storage systems. The I/O characteristics of the application that is writing to the source VDisk also affects the impact that FlashCopy operations have on your overall I/O throughput.
- Perform the appropriate calculations to ensure that your storage systems are configured correctly.
- If any controller that is associated with an MDisk has the **allowquorum** parameter set to **no**, the **setquorum** command will fail for that MDisk. Before setting the **allowquorum** parameter to **yes** on any controller, check the following Web site for controller configuration requirements.

  http://www.ibm.com/storage/support/2145

Logical disk configuration guidelines for storage systems

Most storage systems provide some mechanism to create multiple logical disks from a single array. This is useful when the storage system presents storage directly to the hosts.

However, in a virtualized SAN, use a one-to-one mapping between arrays and logical disks so that the subsequent load calculations and the managed disk (MDisk) and MDisk group configuration tasks are simplified.

Scenario: the logical disks are uneven

In this scenario, you have two RAID-5 arrays and both contain 5 + P components. Array A has a single logical disk that is presented to the SAN Volume Controller cluster. This logical disk is seen by the cluster as mdisk0. Array B has three logical disks that are presented to the cluster. These logical disks are seen by the cluster as mdisk1, mdisk2 and mdisk3. All four MDisks are assigned to the same MDisk group that is named mdisk_grp0. When a virtual disk (VDisk) is created by striping across this group, array A presents the first extent and array B presents the next three extents. As a result, when the system reads and writes to the VDisk, the loading is split 25% on the disks in array A and 75% on the disks in array B. The performance of the VDisk is about one third of what array B can sustain.

The uneven logical disks cause performance degradation and complexity in a simple configuration. You can avoid uneven logical disks by creating a single logical disk from each array.

RAID array configuration guidelines for storage systems

With virtualization, ensure that the storage devices are configured to provide some type of redundancy against hard disk failures.

A failure of a storage device can affect a larger amount of storage that is presented to the hosts. To provide redundancy, storage devices can be configured as RAID arrays that 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. If you use a

large amount of disks, you can reduce the number of disks that are required to provide availability for the same total capacity (1 per array). However, more disks mean that it takes a longer time to rebuild a replacement disk after a disk failure, and during this period a second disk failure causes a loss of all array data. More data is affected by a disk failure for a larger number of member disks because performance is reduced while you rebuild onto a hot spare (a redundant disk) and more data is exposed if a second disk fails before the rebuild operation is complete. The smaller the number of disks, the more likely it is that write operations span an entire stripe (stripe size, multiplied by the number of members, minus one). In this case, write performance is improved. The number of disk drives required to provide availability can be unacceptable if arrays are too small.

Notes:

1. For optimal performance, use arrays with between 6 and 8 member disks.
2. When creating RAID arrays with mirroring, the number of component disks in each array does not affect redundancy or performance.

Optimal MDisk group configuration guidelines for storage systems

A managed disk (MDisk) group provides the pool of storage from which virtual disks (VDisks) are created. You must ensure that the entire pool of storage provides the same performance and reliability characteristics.

Notes:

1. The performance of an MDisk group is generally governed by the slowest MDisk in the group.
2. The reliability of an MDisk group is generally governed by the weakest MDisk in the group.
3. If a single MDisk in a group fails, access to the entire group is lost.

Use the following guidelines when you group similar disks:

- Group equally performing MDisks in a single group.
- Group similar arrays in a single group. For example, configure all 6 + P RAID-5 arrays in one group.
- Group MDisks from the same type of storage system in a single group.
- Group MDisks that use the same type of underlying physical disk in a single group. For example, group MDisks by fibre-channel or SATA.
- Do not use single disks. Single disks do not provide redundancy. Failure of a single disk results in total data loss of the MDisk group to which it is assigned.

Scenario: Similar disks are not grouped together

Under one scenario, you could have two storage systems that are attached behind your SAN Volume Controller. One device is an IBM TotalStorage Enterprise Storage Server (ESS), which contains ten 6 + P RAID-5 arrays and MDisks 0 through 9. The other device is an IBM System Storage DS5000, 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 MDisk group, and the JBOD MDisk11 fails, you lose access to all of the IBM ESS arrays, even though

they are online. The performance is limited to the performance of the JBOD in the IBM DS5000 storage system, therefore slowing down the IBM ESS arrays.

To fix this problem, you can create three groups. The first group must contain the IBM ESS arrays, MDisks 0 through 9, the second group must contain the RAID-1 array, and the third group must contain the large RAID-5 array.

Attention: Do not combine internal SAN Volume Controller solid-state drives (SSDs) and external HDD-based storage systems in the same MDisk group.

FlashCopy mapping guidelines for storage systems

Ensure that you have considered the type of I/O and frequency of update before you create the virtual disks (VDisks) that you want to use in FlashCopy mappings.

FlashCopy operations perform in direct proportion to the performance of the source and target disks. If you have a fast source disk and slow target disk, the performance of the source disk is reduced because it has to wait for the write operation to occur at the target before it can write to the source.

The FlashCopy implementation that is provided by the SAN Volume Controller copies at least 256 K every time a write is made to the source. This means that *any* write involves at minimum a read of 256 K from the source, write of the same 256 K at the target, and a write of the original change at the target. Therefore, when an application performs small 4 K writes, this is translated into 256 K.

Because of this overhead, consider the type of I/O that your application performs during a FlashCopy operation. Ensure that you do not overload the storage. The calculations contain a heavy weighting when the FlashCopy feature is active. The weighting depends on the type of I/O that is performed. Random writes have a much higher overhead than sequential writes. For example, the sequential write would have copied the entire 256 K.

You can spread the FlashCopy source VDisks and the FlashCopy target VDisks between as many managed disk (MDisk) groups as possible. This limits the potential bottle-necking of a single storage system, (assuming that the MDisk groups contain MDisks from different storage systems). However, this can still result in potential bottle-necking if you want to maintain all your target VDisks on a single storage system. You must ensure that you add the appropriate weighting to your calculations.

Image mode VDisks and data migration guidelines for storage systems

Image mode virtual disks (VDisks) enable you to import and then migrate existing data that is managed by the SAN Volume Controller cluster.

Ensure that you follow the guidelines for using image mode VDisks. This might be difficult because a configuration of logical disks and arrays that performs well in a direct SAN-attached environment can contain hot spots or hot component disks when they are connected through the cluster.

If the existing storage systems do not follow the configuration guidelines, consider stopping I/O operations on the host systems while you migrate the data into the cluster. If I/O operations are continued and the storage system does not follow the guidelines, I/O operations can fail at the hosts and ultimately loss of access to the data can occur.

Attention: Migration commands fail if the target or source VDisk is offline, or if there is insufficient quorum disk space to store the metadata. Correct the offline or quorum disk condition and reissue the command.

The procedure for importing managed disks (MDisks) that contain existing data depends on the amount of free capacity you have in the cluster. You must have the same amount of free space in the cluster as the data that you want to migrate into the cluster. If you do not have this amount of available capacity, the migration causes the MDisk group to have an uneven distribution of data because some MDisks are more heavily loaded than others. Further migration operations are required to ensure an even distribution of data and subsequent I/O loading.

Migrating data with an equivalent amount of free capacity

To prevent managed disks (MDisks) from having an uneven distribution of data, ensure that the cluster has the same amount of free space as the data that you want to migrate.

Perform the following steps to migrate data:

1. Stop all I/O operations from the hosts. Unmap the logical disks that contain the data from the hosts.
2. Create one or more MDisk groups with free capacity. Ensure that the MDisk groups have enough free capacity to contain all of the migrating data and that they have balanced data distribution.
3. Create an empty MDisk group. This temporarily contains the data that is imported.
4. Perform the following steps to create an image mode virtual disk (VDisk) from the first unmanaged-mode MDisk that contains the data to be imported:
   a. Map one logical disk from the storage system to the SAN Volume Controller ports.
   b. Issue the **svctask detectmdisk** command-line interface (CLI) command on the cluster or use the SAN Volume Controller Console to discover MDisks. The new unmanaged-mode MDisk that is found corresponds with the logical disk mapped in the previous step.
   c. Create an image mode VDisk from this unmanaged-mode MDisk and assign it to the empty MDisk group that you just created.
   d. Repeat steps 4a through 4c for all logical disks as required.
5. You can restart host applications by mapping the image mode VDisks to the hosts and restarting the hosts. Alternatively, you can wait until the migration is complete to map the VDisks to the hosts.
6. Perform the following steps to migrate the data to the MDisk groups that you created in step 2:
   a. Select the first image mode VDisk to be migrated.
   b. Using the **svctask migratevdisk** or **svctask migratetoimage** command, migrate this VDisk from its current MDisk group to one of the MDisk groups that you created in step 2. This migrates all of the data from the logical disk into the new free space.
   c. Run the **svcinfo lsmigrate** command to verify that the migration has completed.
   d. Select the next image mode VDisk and repeat the previous step after the migration completes.

7. When all the VDisks have been migrated, the MDisk groups that you created in step 2 on page 349 contain the data that was on the image mode VDisks. The data is striped across the new groups and is virtualized.
8. Destroy the temporary MDisk group that contained the original image mode VDisks.
9. Go back to the storage system and reconfigure 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 VDisks.

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 that is imported, you can still migrate data.

Scenario

Under one scenario, you could have one managed disk (MDisk) in the destination MDisk group. You add image mode logical units from an array on the storage system and migrate these logical units to the destination MDisk group. The logical units are then striped across the one managed-mode disk. Next, you add another logical unit to the destination MDisk group. The MDisk now contains two managed-mode disks, but all of the data is on the first managed-mode disk. As a result, some of the data must be migrated from the overloaded managed-mode disks to the underused managed-mode disks.

Attention: The migration causes an uneven distribution of data across the MDisks in the MDisk group. The impact of this depends on the number of MDisks that are initially in the MDisks group and how many of these have free capacity.

This task might require subsequent migration of data within the MDisk group in order to balance the distribution of data across the MDisks in the group.

Perform the following steps to migrate data:

1. Select an MDisk group that contains enough free capacity to migrate *all* of the logical disks on the first array that you want to migrate to the cluster.
2. Create an empty MDisk group that can temporarily contain the data that is imported.
3. Stop all I/O operations to the logical disks that you want to migrate first, and unmap these disks from their hosts.
4. Perform the following steps to create an image mode virtual disk (VDisk) from the first unmanaged-mode MDisk that contains the data that you want to import:
   a. Map one logical disk from the storage system to the SAN Volume Controller ports.
   b. Issue the **svctask detectmdisk** command-line interface (CLI) command on the cluster or use the SAN Volume Controller Console to discover MDisks. The new unmanaged-mode MDisk that is found corresponds with the logical disk that was mapped in the previous step.
   c. Create an image mode VDisk from this unmanaged-mode MDisk and assign it to use the empty MDisk group just created.
   d. Repeat steps 4a through 4c for all logical disks.

5. You can restart host applications by mapping the image mode VDisks to the hosts and restarting the hosts, using the SAN Volume Controller. Alternatively, you can wait until the migration is complete to map the VDisks to the hosts.
6. Perform the following steps to migrate the data into the MDisk groups that you created in step 1 on page 350:
   a. Select the first image mode VDisk that you want to migrate.
   b. Using the **svctask migratevdisk** or **svctask migratetoimage** command, migrate this VDisk from its current MDisk group to one of the MDisk groups that you created in step 2 on page 350. This migrates all of the data from the logical disk into the new free space.
   c. Run the **svcinfo lsmigrate** command to verify that the migration has completed.
   d. Select the next image mode VDisk and repeat the previous step when the migration completes.
7. Perform the following steps to reconfigure the RAID array that contains the logical disks and add it to the MDisk group that you selected in step 1 on page 350:
   a. Remove the MDisks from the temporary MDisk group.
   b. At the storage system, unmap the logical disks that have been migrated from the SAN Volume Controller cluster and delete them from the array (if more than one existed).
   c. Create a single logical disk that uses the entire array capacity.
   d. Map this new logical disk to the SAN Volume Controller ports.
   e. Issue the **svctask detectmdisk** CLI command on the cluster or use the SAN Volume Controller Console to discover MDisks. The new managed-mode MDisk that is found corresponds with the new logical disk that you created.
   f. Add this managed-mode MDisk to the MDisk group that you selected in step 1 on page 350.
8. Repeat steps 3 on page 350 through 7 for the next array.

Configuring a balanced storage system

The attachment of a storage system to a SAN Volume Controller requires that specific settings are applied to the device.

There are two major steps to attaching a storage system to a SAN Volume Controller:

1. Setting the characteristics of the SAN Volume Controller to storage connections
2. Mapping logical units to these storage connections that allow the SAN Volume Controller to access the logical units

The virtualization features of the SAN Volume Controller enable you to choose how your storage is divided and presented to hosts. While virtualization provides you with a great deal of flexibility, it also offers the potential to set up an overloaded storage system. A storage system is overloaded if the quantity of I/O transactions that are issued by the host systems exceeds the capability of the storage to process those transactions. If a storage system is overloaded, it causes delays in the host systems and might cause I/O transactions to time out in the host. If I/O transactions time out, the host logs errors and I/Os fail to the applications.

Scenario: You have an overloaded storage system

Under one scenario, you have used the SAN Volume Controller to virtualize a single RAID array and to divide the storage across 64 host systems. If all host systems attempt to access the storage at the same time, the single RAID array is overloaded.

Perform the following steps to configure a balanced storage system:

1. Use Table 28 to calculate the I/O rate for each RAID array in the storage system.

   Note: The actual number of I/O operations per second that can be processed depends 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 28. Calculate the I/O rate

Type of RAID array	Number of component disks in the RAID array	Approximate I/O rate per second
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 (MDisk).
   - If there is a one-to-one relationship between backend arrays and MDisks, the I/O rate for an MDisk is the same as the I/O rate of the corresponding array.
   - If an array is divided into multiple MDisks, the I/O rate per MDisk is the I/O rate of the array divided by the number of MDisks that are using the array.
3. Calculate the I/O rate for an MDisk group. The I/O rate for an MDisk group is the sum of the I/O rates of the MDisk that is in the MDisk group. For example, an MDisk group contains eight MDisks and each MDisk corresponds to a RAID-1 array. Using Table 28, the I/O rate for each MDisk is calculated as 300. The I/O rate for the MDisk group is 300×8 = 2400.
4. Use Table 29 to calculate the impact of FlashCopy mappings. If you are using the FlashCopy feature that is provided by the SAN Volume Controller, you must consider the additional amount of I/O that FlashCopy operations generate because it reduces the rate at which I/O from host systems can be processed. When a FlashCopy mapping copies write I/Os from the host systems to areas of the source or target virtual disk (VDisk) that are not yet copied, the SAN Volume Controller generates extra I/Os to copy the data before the write I/O is performed. The effect of using the FlashCopy feature depends on the type of I/O workload that is generated by an application.

Table 29. Calculate the impact of FlashCopy mappings

Type of application	Impact to I/O rate	Additional weighting for FlashCopy
Application is not performing I/O	Insignificant impact	0
Application is only reading data	Insignificant impact	0

Table 29. Calculate the impact of FlashCopy mappings (continued)

Type of application	Impact to I/O rate	Additional weighting for FlashCopy
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 VDisk that is the source or target of an active FlashCopy mapping, consider the type of application that you want to use the VDisk and record the additional weighting for the VDisk.

Example

For example, a FlashCopy mapping is used to provide point-in-time backups. During the FlashCopy process, a host application generates an I/O workload of random read and write operations to the source VDisk. A second host application reads the target VDisk and writes the data to tape to create a backup. The additional weighting for the source VDisk is 14. The additional weighting for the target VDisk is 0.

5. Calculate the I/O rate for VDisks in an MDisk group by performing the following steps:
   a. Calculate the number of VDisks in the MDisk group.
   b. Add the additional weighting for each VDisk that is the source or target of an active FlashCopy mapping.
   c. Divide the I/O rate of the MDisk group by this number to calculate the I/O rate per VDisk.

   Example 1

   An MDisk group has an I/O rate of 2400 and contains 20 VDisks. There are no FlashCopy mappings. The I/O rate per VDisk is 2400 / 20 = 120.

   Example 2

   An MDisk group has an I/O rate of 5000 and contains 20 VDisks. There are two active FlashCopy mappings that have source VDisks in the MDisk group. Both source VDisks are accessed by applications that issue random read and write operations. As a result, the additional weighting for each VDisk is 14. The I/O rate per VDisk is 5000 / (20 + 14 + 14) = 104.
6. Determine if the storage system is overloaded. The figure that was determined in step 4 on page 352 provides some indication of how many I/O operations per second can be processed by each VDisk in the MDisk group.
   - If you know how many I/O operations per second that 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 that your host applications generate, you can use the I/O statistics facilities that are provided by the SAN Volume Controller to measure the I/O rate of your virtual disks, or you can use Table 30 as a guideline.

Table 30. Determine if the storage system is overloaded

Type of Application	I/O rate per VDisk
Applications that generate a high I/O workload	200

Table 30. Determine if the storage system is overloaded (continued)

Type of Application	I/O rate per VDisk
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 that is generated by the application exceeds the I/O rate per VDisk that you calculated, you might be overloading your storage system. You must carefully monitor the storage system to determine if the backend storage limits the overall performance of the storage system. It is also possible that the previous calculation is too simplistic to model your storage use after. For example, the calculation assumes that your applications generate the same I/O workload to all VDisks, which might not be the case.

   You can use the I/O statistics facilities that are provided by the SAN Volume Controller to measure the I/O rate of your MDisks. You can also use the performance and I/O statistics facilities that are provided by your storage systems.

If your storage system is overloaded there are several actions that you can take to resolve the problem:

- Add more backend storage to the system to increase the quantity of I/O that can be processed by the storage system. The SAN Volume Controller provides virtualization and data migration facilities to redistribute the I/O workload of VDisks across a greater number of MDisks without having to take the storage offline.
- Stop unnecessary FlashCopy mappings to reduce the amount of I/O operations that are submitted to the backend storage. If you perform FlashCopy operations in parallel, consider reducing the amount of FlashCopy mappings that start in parallel.
- Adjust the queue depth to limit the I/O workload that is generated by a host. Depending on the type of host and type of host bus adapters (HBAs), it might be possible to limit the queue depth per VDisk or limit the queue depth per HBA, or both. The SAN Volume Controller also provides I/O governing features that can limit the I/O workload that is generated by hosts.

Note: Although these actions can be used to avoid I/O time-outs, performance of your storage system is still limited by the amount of storage that you have.

Storage system requirements

The performance of applications at the local cluster can be limited by the performance of the storage systems at the remote cluster.

Your set up must meet the following requirements to maximize the amount of I/O operations that applications can run on Global Mirror VDisks:

- The Global Mirror VDisks at the remote cluster must be in dedicated MDisk groups that only contain other Global Mirror VDisks.
- Configure storage systems to support the Global Mirror workload that is required of them. The following guidelines can be used to fulfill this requirement:
  - Dedicate storage systems to only Global Mirror VDisks

  - Configure the storage system to guarantee sufficient quality of service for the disks that are being used by Global Mirror operations
  - Ensure that physical disks are not shared between Global Mirror VDisks and other I/O operations. For example, do not split an individual RAID array.
- For Global Mirror MDisk groups, use MDisks with the same characteristics. For example, use MDisks that have the same RAID level, physical disk count, and disk speed. This requirement is important to maintain performance when you use the Global Mirror feature.

You must provision the storage systems that are attached to the remote cluster to accommodate the following:

- The peak application workload to the Global Mirror VDisks
- The specified background copy level
- All I/O operations that run on the remote cluster

Storage system requirements for FlashCopy, VDisk mirroring, and space-efficient VDisks

Application performance on a local cluster can be affected by the use of FlashCopy, VDisk mirroring, and space-efficient VDisks for storage systems.

The FlashCopy, VDisk mirroring, and space-efficient VDisk features can all have a negative impact on cluster performance. The impact depends on the type of I/O taking place, and is estimated using a weighting factor from Table 31.

A FlashCopy mapping effectively adds a number of loaded VDisks to the MDisk group. The effect of mirrored and space-efficient VDisks is also estimated in Table 31. The estimates assume that space-efficient VDisks are running at approximately 80% capacity of a fully allocated VDisk, and that mirrored VDisks read from one copy and write to all copies.

Table 31. Performance impact estimates for FlashCopy, VDisk mirroring, and space-efficient VDisks

Type of I/O (to VDisk)	Impact on I/O weighting	FlashCopy weighting	VDisk mirroring weighting	Space-efficient weighting
None or minimal	Insignificant	0	0	0
Read Only	Insignificant	0	0	0.25 * Sv
Sequential read and write	Up to 2 x I/O	2 * F	C–V	0.25 * Sc
Random read and write	Up to 15 x I/O	14 * F	C–V	0.25 * Sc
Random write	Up to 50 x I/O	49 * F	C–V	0.25 * Sc

Table 31. Performance impact estimates for FlashCopy, VDisk mirroring, and space-efficient VDisks (continued)

Type of I/O (to VDisk)	Impact on I/O weighting	FlashCopy weighting	VDisk mirroring weighting	Space-efficient weighting
Notes: • In an MDisk group with two FlashCopy mappings and random read/write to those VDisks, the weighting factor is 14 * 2 = 28. • In an MDisk group with ten copies, five of which are primary copies of a VDisk, a weighting factor of 10–5 = 5 applies. If the copies are space-efficient, an additional weighting factor of 0.25 * 10 = 2.5 applies. **Key:** **C** Number of VDisk Copies in this MDisk Group **V** Number of VDisks with their primary copy in this MDisk Group **F** Number of FlashCopy mappings affecting VDisks that have copies in this MDisk Group **Sv** Number of space-efficient VDisk Copies in this MDisk Group that are the primary copy of a VDisk **Sc** Number of space-efficient VDisk Copies in this MDisk Group				

To calculate the average I/O rate per VDisk, use the following equation:

```
I/O rate = (I/O capacity) / ( V + weighting factor for FlashCopy +
weighting factor for VDisk mirroring + weighting factor for space-efficient)
```

For example, consider 20 VDisks with an I/O capacity of 5250, a FlashCopy weighting of 28, a mirroring weighting of 5, and a space-efficient weighting of 0.25. The I/O rate per VDisk is 5250 / (20 + 28 + 5 + 2.5) = 94.6. This estimate is an average I/O rate per VDisk; for example, half of the VDisks could be running at 200 IOPs, and the other half could be running at 20 IOPs. This would not overload the system however, because the average load is 94.6.

If the average I/O rate to the VDisks in the example exceeds 94.6, the system would be overloaded. As approximate guidelines, a heavy I/O rate is 200, a medium I/O rate is 80, and a low I/O rate is 10.

With VDisk mirroring, a single VDisk can have multiple copies in different MDisk groups. The I/O rate for such a VDisk is the minimum I/O rate calculated from each of its MDisk Groups.

If system storage is overloaded, you can migrate some of the VDisks to MDisk groups with available capacity.

Note: Solid-state drives (SSDs) are exempt from these calculations, with the exception of overall node throughput, which increases substantially for each additional SSD in the node.

Discovering logical units

The SAN Volume Controller initialization includes a process called discovery.

The discovery process systematically recognizes all visible ports on the SAN for devices that identify themselves as storage systems and the number of logical units (LUs) that they export. The LUs can contain new storage or a new path for previously discovered storage. The set of LUs forms the SAN Volume Controller managed disk (MDisk) view.

The discovery process runs when ports are added to or deleted from the SAN and when certain error conditions occur. You can also manually run the discovery process using the **svctask detectmdisk** command-line interface (CLI) command or the **Discover MDisks** function from the SAN Volume Controller Console. The **svctask detectmdisk** CLI command and the **Discover MDisks** function have the cluster rescan the fibre-channel network. The rescan discovers any new MDisks that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

Note: Some storage systems do not automatically export LUs to the SAN Volume Controller.

Guidelines for exporting LUs

Ensure that you are familiar with the following guidelines for exporting LUs to the SAN Volume Controller:

- When you define the SAN Volume Controller as a host object to the storage systems, you must include *all* ports on *all* nodes and candidate nodes.
- When you first create an LU, you *must* wait until it is initialized before you export it to the SAN Volume Controller.

  Attention: Failure to wait for the LUs to initialize can result in excessive discovery times and an unstable view of the SAN.
- Do not present new LUs to the SAN Volume Controller until the array initialization and format is complete. If you add a LUN to an MDisk group before the array initialization format is complete, the MDisk group goes offline. While the MDisk group is offline, you cannot access the VDisks that are in the MDisk group.
- When you export an LU to the SAN Volume Controller, the LU *must* be accessible through all ports on the storage system that are visible to the SAN Volume Controller.

  Important: The LU *must* be identified by the same logical unit number (LUN) on all ports.

Expanding a logical unit using the CLI

You can use the command-line interface (CLI) to expand a logical unit.

Some storage systems enable you to expand the size of a logical unit (LU) using vendor-specific disk-configuration software that is provided. The steps in this procedure are required for the SAN Volume Controller to use extra capacity that is provided in this way.

To ensure that this additional capacity is available to the SAN Volume Controller, follow these steps:

1. Issue the svctask rmmdisk CLI command to remove the managed disk (MDisk) from the MDisk group. Use the **force** parameter to migrate data on the specified MDisk to other MDisks in the group. The command completes

asynchronously if `-force` is specified. You can check the progress of active migrations by running the svcinfo lsmigrate command.

2. Use the vendor-specific, disk-configuration software to expand the size of the logical unit on the storage system.
3. Issue the svctask detectmdisk CLI command to rescan the fibre-channel network. The rescan discovers any changes to existing MDisks and any new MDisks that have been added to the cluster. This command completes asynchronously and might take a few minutes. To determine whether a discovery operation is still in progress, use the svcinfo lsdiscoverystatus command.
4. Issue the svcinfo lsmdisk CLI command to display the additional capacity that has been expanded.
5. Issue the svctask addmdisk CLI command to add the MDisk back to the group.

The extra capacity is available for use by the SAN Volume Controller.

Modifying a logical unit mapping using the CLI

You can modify a logical unit (LU) mapping using the command-line interface (CLI).

Perform the following steps to modify an LU mapping:

1. Migrate all of the data from the managed disk (MDisk) by performing the following steps:
   a. If the MDisk is in managed mode or image mode and the virtual disk (VDisk) must be kept online, issue the following CLI command and then proceed to step 2:

      ```
      svctask rmmdisk -mdisk MDisk number -force MDisk group number
      ```

      Where *MDisk number* is the number of the MDisk that you want to modify and *MDisk group number* is the number of the MDisk group for which you want to remove the MDisk.

      Note:

      - The VDisk becomes a striped MDisk *not* an image-mode VDisk.
      - All data that is stored on this MDisk is migrated to the other MDisks in the MDisk group.
      - This CLI command can fail if there are not enough free extents in the MDisk group.

   b. If the MDisk is in image mode and you do not want to convert the VDisk to a striped VDisk, stop all I/O to the image mode VDisk.
   c. Issue the following CLI command to remove the host mapping and any SCSI reservation that the host has on the VDisk:

      ```
      svctask rmvdiskhostmap -host host name VDisk name
      ```

      Where *host name* is the name of the host for which you want to remove the VDisk mapping and *VDisk name* is the name of the VDisk for which you want to remove mapping.
   d. Issue the following command to delete the VDisk:

      ```
      svctask rmvdisk VDisk name
      ```

      Where *VDisk name* is the name of the VDisk that you want to delete.
2. Remove the LU mapping on the storage system so that the LUN is not visible to the SAN Volume Controller.

3. Issue the following CLI command to clear all error counters on the MDisk:
   ```
   svctask includemdisk MDisk number
   ```
   Where *MDisk number* is the number of the MDisk that you want to modify.
4. Issue the following CLI command to rescan the fibre-channel network and detect that the LU is no longer there.
   ```
   svctask detectmdisk MDisk number
   ```
   Where *MDisk number* is the number of the MDisk that you want to modify. The MDisk is removed from the configuration.
5. Issue the following CLI command to verify that the MDisk is removed:
   ```
   svcinfo lsmdisk MDisk number
   ```
   Where *MDisk number* is the number of the MDisk that you want to modify.
   - If the MDisk is still displayed, repeat steps 3 and 4.
6. Configure the mapping of the LU to the new LUN on the storage system.
7. Issue the following CLI command:
   ```
   svctask detectmdisk
   ```
8. Issue the following CLI command to check that the MDisk now has the correct LUN:
   ```
   svcinfo lsmdisk
   ```

The MDisk has the correct LUN.

Accessing controller devices with multiple remote ports

If a managed disk (MDisk) logical unit (LU) is accessible through multiple controller device ports, the SAN Volume Controller ensures that all nodes that access this LU coordinate their activity and access the LU through the same controller device port.

Monitoring LU access through multiple controller device ports

When the SAN Volume Controller can access an LU through multiple controller device ports, the SAN Volume Controller uses the following criteria to determine the accessibility of these controller device ports:

- The SAN Volume Controller node is a member of a cluster.
- The SAN Volume Controller node has fibre-channel connections to the controller device port.
- The SAN Volume Controller node has successfully discovered the LU.
- Slandering has not caused the SAN Volume Controller node to exclude access to the MDisk through the controller device port.

An MDisk path is presented to the cluster for all SAN Volume Controller nodes that meet these criteria.

Controller device port selection

When an MDisk is created, the SAN Volume Controller selects one of the controller device ports to access the MDisk.

Table 32 on page 360 describes the algorithm that the SAN Volume Controller uses to select the controller device port.

Table 32. Controller device port selection algorithm

Criteria	Description
Accessibility	Creates an initial set of candidate controller device ports. The set of candidate controller device ports include the ports that are accessible by the highest number of nodes.
Slandering	Reduces the set of candidate controller device ports to those with the lowest number of nodes.
Preference	Reduces the set of candidate controller device ports to those that the controller device uses as preferred ports.
Load balance	Selects the port from the set of candidate controller device ports that has the lowest MDisk access count.

After the initial device port selection is made for an MDisk, the following events can cause the selection algorithm to rerun:

- A new node joins the cluster and has a different view of the controller device than the other nodes in the cluster.
- The **svctask detectmdisk** command-line interface (CLI) command is run or the **Discover MDisks** SAN Volume Controller Console function is used. The **svctask detectmdisk** CLI command and the **Discover MDisks** function have the cluster rescan the fibre-channel network. The rescan discovers any new MDisks that might have been added to the cluster and rebalances MDisk access across the available controller device ports.
- Error recovery procedures (ERPs) are started because a controller device has changed its preferred port.
- New controller device ports are discovered for the controller device that is associated with the MDisk.
- The controller device port that is currently selected becomes inaccessible.
- Slandering has caused the SAN Volume Controller to exclude access to the MDisk through the controller device port.

Determining a storage system name from its SAN Volume Controller name

You can determine a storage system name from its SAN Volume Controller name.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to determine the name of the storage system:

1. Click **Work with Managed Disks → Disk Controller Systems**. The Viewing Disk Controller Systems panel is displayed.
2. Select the link for the name of the storage system for which you want to determine the name.
3. Record the worldwide node name (WWNN). You can launch the native user interface for the storage system, or use the command-line tools to verify the name of the storage system that uses this WWNN.

Determining a storage system name from its SAN Volume Controller name using the CLI

You can determine a storage system name from its SAN Volume Controller name using the command-line interface (CLI).

1. Issue the following CLI command to list the storage system:

   ```
   svcinfo lscontroller
   ```

2. Record the name or ID for the storage system that you want to determine.
3. Issue the following CLI command:

   ```
   svcinfo lscontroller controllername/id
   ```

   where *controllername/id* is the name or ID that you recorded in step 2.
4. Record the worldwide node name (WWNN) for the device. The WWNN can be used to determine the actual storage system by launching the native user interface or using the command-line tools it provides to verify the actual storage system that has this WWNN.

Renaming a storage system

You can rename a storage system from the Renaming a Disk Controller System panel.

This task assumes that you have already launched the SAN Volume Controller application.

Perform the following steps to rename a storage system:

1. Click **Work with Managed Disks → Disk Controller Systems** in the portfolio. The Viewing Disk Controller Systems panel is displayed.
2. Select the storage system to rename and select **Rename a Disk Controller System** from the list. Click **Go**. The Renaming Disk Controller System panel is displayed.

Renaming a storage system using the CLI

You can use the command-line interface (CLI) to rename a storage system.

Perform the following step to rename a storage system:

Issue the **svctask chcontroller -name** ***new_name controller_id*** command.

Changing the configuration of an existing storage system using the CLI

You can use the command-line interface (CLI) to change the configuration of an existing storage system. You must change the configuration for a storage system when you want to delete and replace logical units (LUs).

Perform the following steps to delete existing LUs and replace them with new LUs:

1. Issue the following CLI command to delete the managed disks (MDisks) that are associated with the LUs from their MDisk groups:

```
svctask rmmdisk -mdisk MDisk name1:MDisk name2 -force MDisk group name
```

Where *MDisk name1:MDisk name2* are the names of the MDisks to delete.

2. Delete the existing LUs using the configuration software of the storage system.
3. Issue the following command to delete the associated MDisks from the cluster:

```
svctask detectmdisk
```

4. Configure the new LUs using the configuration software of the storage system.
5. Issue the following command to add the new LUs to the cluster:

```
svctask detectmdisk
```

Adding a new storage controller to a running configuration

You can add a new storage controller to your SAN at any time.

You must follow the zoning guidelines for your switch and also ensure that the controller is setup correctly for use with the SAN Volume Controller.

You must create one or more arrays on the new controller.

If your controller provides array partitioning, create a single partition from the entire capacity available in the array. You must record the LUN number that you assign to each partition. You must 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.

Perform the following steps to add a new storage controller:

1. Ensure that the cluster has detected the new storage (MDisks).
   a. Click **Work with Managed Disks** → **Managed Disks**. The Viewing Managed Disks panel is displayed.
   b. Select **Discover MDisks** from the task list and click **Go**.
2. Determine the storage controller name to validate that this is the correct controller. The controller will have automatically been assigned a default name.
   - If you are unsure which controller is presenting the MDisks perform the following steps:
   a. Click **Work with Managed Disks** → **Disk Controller Systems**. The Viewing Disk Controller Systems panel is displayed.
   b. Find the new controller in the list. The new controller has the highest numbered default name.
3. Record the field controller LUN number. The controller LUN number corresponds with the LUN number that you assigned to each of the arrays or partitions.
   a. Click **Work with Managed Disks** → **Managed Disks**. The Viewing Managed Disks panel is displayed.
   b. Find the MDisks that are not in managed mode. These MDisks should correspond with the RAID arrays or partitions that you created.
4. Create a new MDisk group and add only the RAID arrays that belong to the new controller to this MDisk group. To avoid mixing RAID types, create a new MDisk group for each set of RAID array types (for example, RAID-5, RAID-1).
   a. Click **Work with Managed Disks** → **Managed Disk Groups**.
   b. Select **Create an MDisk Group** from the task list and click **Go**. The Create Managed Disk Group wizard begins.

c. Complete the wizard to create a new MDisk group.

Tip: Give each MDisk group that you create a descriptive name. For example, if your controller is named FAST650-fred, and the MDisk group contains RAID-5 arrays, name the MDisk Group F600-fred-R5.

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 command-line interface (CLI).

You must follow the zoning guidelines for your switch and also ensure that the controller is setup correctly for use with the SAN Volume Controller.

You must create one or more arrays on the new controller.

If your controller provides array partitioning, create a single partition from the entire capacity available in the array. You must record the LUN number that you assign to each partition. You must 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.

Perform the following steps to add a new storage controller:

1. Issue the following CLI command to ensure that the cluster has detected the new storage (MDisks):

   ```
   svctask detectmdisk
   ```

2. Determine the storage controller name to validate that this is the correct controller. The controller is automatically assigned a default name.
   - If you are unsure which controller is presenting the MDisks, issue the following command to list the controllers:

   ```
   svcinfo lscontroller
   ```

3. Find the new controller in the list. The new controller has the highest numbered default name.
4. Record the name of the controller and follow the instructions in the section about determining a disk controller system name.
5. Issue the following command to change the controller name to something that you can easily use to identify it:

   ```
   svctask chcontroller -name newname oldname
   ```

   Where *newname* is the name that you want to change the controller to and *oldname* is the name that you are changing.
6. Issue the following command to list the unmanaged MDisks:

   ```
   svcinfo lsmdisk -filtervalue mode=unmanaged:controller_name=new_name
   ```

   These MDisks should correspond with the RAID arrays or partitions that you have created.
7. Record the field controller LUN number. This number corresponds with the LUN number that you assigned to each of the arrays or partitions.
8. Create a new MDisk group and add only the RAID arrays that belong to the new controller to this MDisk group. To avoid mixing RAID types, create a new

MDisk group for each set of RAID array types (for example, RAID-5, RAID-1). Give each MDisk group that you create a descriptive name. For example, if your controller is named FAST650-fred, and the MDisk group contains RAID-5 arrays, name the MDisk Group F600-fred-R5.

```
svctask mkmdiskgrp -ext 16 -name mdisk_grp_name
-mdisk colon separated list of RAID-x mdisks returned
in step 4
```

This creates a new MDisk group with an extent size of 16MB.

Removing a storage system

You can replace or decommission a storage system.

This task assumes that you have already launched the SAN Volume Controller Console.

During this procedure, you will add a new device, migrate data off of the storage system and remove the old MDisks.

An alternative to following this procedure is to migrate all of the virtual disks (VDisks) that are using storage in this MDisk group to another MDisk group. This allows you to consolidate the VDisks in a single or new group. However, you can only migrate one VDisk at a time. The procedure outlined below migrates all the data through a single command.

You can also use this procedure to remove or replace a single MDisk in a group. If an MDisk experiences 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 replace just that MDisk. Steps 1 and 3 detail how you can add or remove a single MDisk rather than a list of MDisks.

Perform the following steps to remove a storage system:

1. Add the new MDisks to the MDisk group by performing the following steps:
   a. Click **Work with Managed Disks** → **Managed Disks**. The Viewing Managed Disk Groups panel is displayed.
   b. Select the MDisk group that you want to add the new MDisks to and select **Add MDisks** from the task list. Click **Go**. The Adding Managed Disks to Managed Disk Group panel is displayed.
   c. Select the new MDisks and click **OK**. The MDisk group should now contain both the old and new MDisks.
2. Ensure that the capacity of the new MDisks is the same or exceeds that of the old MDisks before proceeding to step 3.
3. Force the deletion of the old MDisks from the MDisk group to migrate all the data from the old MDisks to the new MDisks.
   a. Click **Work with Managed Disks** → **Managed Disks**. The Viewing Managed Disk Groups panel is displayed.
   b. Select the MDisk group that you want to add the new MDisks to and select **Remove MDisks** from the task list. Click **Go**. The Deleting Managed Disks from Managed Disk Group panel is displayed.
   c. Select the old MDisks and click **OK**. The migration process begins.

Note: The amount of time this process runs depends on the number and size of MDisks and the number and size of the VDisks that are using the MDisks.

4. Check the progress of the migration process by issuing the following command from the command-line interface (CLI): `svcinfo lsmigrate`
5. When all the migration tasks are complete, for example, the command in step 4 returns no output, verify that the MDisks are unmanaged.
6. Access the storage system and unmap the LUNs from the SAN Volume Controller ports.

   Note: You can delete the LUNs if you no longer want to preserve the data that is on the LUNs.
7. Perform the following steps to have the cluster rescan the fibre-channel network:
   a. Click **Work with Managed Disks → Managed Disks**.
   b. Select **Discover MDisks** from the task list and click **Go**. The Discovering Managed Disks panel is displayed. The rescan discovers that the MDisks have been removed from the cluster and also rebalances MDisk access across the available controller device ports.
8. Verify that there are no MDisks for the storage system that you want to decommission.
9. Remove the storage system from the SAN so that the SAN Volume Controller ports can no longer access the storage system.

Removing a storage system using the CLI

You can replace or decommission a storage system using the command-line interface (CLI).

During this procedure, you will add a new device, migrate data off of the storage system and remove the old MDisks.

An alternative to following this procedure is to migrate all of the virtual disks (VDisks) that are using storage in this MDisk group to another MDisk group. This allows you to consolidate the VDisks in a single or new group. However, you can only migrate one VDisk at a time. The procedure outlined below migrates all the data through a single command.

You can also use this procedure to remove or replace a single MDisk in a group. If an MDisk experiences 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 replace just that MDisk.

Perform the following steps to remove a storage system:

1. Add the new storage system 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 decommissioned MDisks and *mdisk_grp_name* is the name of the MDisk group that contains the MDisks that you want to decommission.

You should now have an MDisk group that you want to decommission and the new MDisks.

3. Ensure that the capacity of the new MDisks is the same or exceeds that of the old MDisks before you proceed to step 4.
4. Issue the following command to force delete the old MDisks from the group:

   ```
   svctask rmmdisk -force -mdisk mdiskx:mdisky:mdiskz... mdisk_grp_name
   ```

   Where *mdiskx:mdisky:mdiskz...* are the old MDisks that you want to delete and *mdisk_grp_name* is the name of the MDisk group that contains the MDisks that you want to 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 takes some time to complete, even though the command returns immediately.
5. Check the progress of the migration process by issuing the following command:

   ```
   svcinfo lsmigrate
   ```
6. When all the migration tasks are complete, for example, the command in step 5 returns no output, verify that the MDisks are unmanaged.
7. Access the storage system and unmap the LUNs from the SAN Volume Controller ports.

   Note: You can delete the LUNs if you no longer want to preserve the data that is on the LUNs.
8. Issue the following CLI command:

   ```
   svctask detectmdisk
   ```
9. Verify that there are no MDisks for the storage system that you want decommission.
10. Remove the storage system from the SAN so that the SAN Volume Controller ports can no longer access the storage system.

Removing MDisks that represent unconfigured LUs using the CLI

You can use the command-line interface (CLI) to remove MDisks from the cluster.

When you remove LUs from your storage system, the managed disks (MDisks) that represent those LUs might still exist in the cluster. However, the cluster cannot access these MDisks because the LUs that these MDisks represent have been unconfigured or removed from the storage system. You must remove these MDisks.

Perform the following steps to remove MDisks:

1. Run the **svctask includemdisk** CLI command on all the affected MDisks.
2. Run the **svctask rmmdisk** CLI command on all affected MDisks. This puts the MDisks into the unmanaged mode.
3. Run the **svctask detectmdisk** CLI command. The cluster detects that the MDisks no longer exist in the storage system.

All of the MDisks that represent unconfigured LUs are removed from the cluster.

Creating a quorum disk

A quorum disk is used to resolve tie-break situations when the voting set of nodes disagree on the current state of the cluster.

Quorum disk creation and extent allocation

The use of a quorum disk allows the cluster to manage a SAN fault that splits the cluster exactly in half. One half of the cluster continues to operate and the other half stops until the SAN connectivity is restored.

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 quorum candidate disks.

An LU must meet the following criteria to be considered a candidate for a quorum disk:

- It must be in managed space mode.
- It must be visible to all nodes in the cluster.
- It must be presented by a storage system 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 are presented by different devices. After the quorum candidate disks are selected, the cluster selects one of the candidate quorum disks to become the active quorum disk, which means it is used first to break a tie in the event of a cluster partition. After the active quorum disk is selected, the cluster does not attempt to ensure that the candidate quorum disks are presented by different devices. However, you can also manually select the active quorum disk if you want to ensure the active quorum disk is presented by a different device. Selecting the active quorum disk is useful in split-site cluster configurations and ensures that the most highly available quorum disk is used. You can set the **-active** parameter on the setquorum command to set a disk as an active quorum disk. You can also use the Set Active Quorum Disk panel in the SAN Volume Controller Console by selecting **Work with Managed Disks** → **Quorum Disks** to specify an active quorum disk. The quorum disk candidates can be updated by configuration activity if other eligible LUs are available.

To view a list of current quorum disk candidates, use the svcinfo lsquorum command.

If no quorum disk candidates are found after the discovery, one of the following situations has occurred:

- No LUs exist in managed space mode. An error is logged when this situation occurs.
- LUs exist in managed space mode, but they do not meet the eligibility criteria. An error is logged when this situation occurs.

Manual discovery

When you create or remove LUNs on a storage system, the managed disk (MDisk) view is not automatically updated.

You must issue the **svctask detectmdisk** command-line interface (CLI) command or use the **Discover MDisks** function from the SAN Volume Controller Console to have the cluster rescan the fibre-channel network. The rescan discovers any new MDisks that might have been added to the cluster and rebalances MDisk access across the available controller device ports.

Servicing storage systems

Storage systems that are supported for attachment to the SAN Volume Controller are designed with redundant components and access paths to enable concurrent maintenance. Hosts have continuous access to their data during component failure and replacement.

The following guidelines apply to all storage systems that are attached to the SAN Volume Controller:

- Always follow the service instructions that are provided in the documentation for your storage system.
- Ensure that there are no unfixed errors in the SAN Volume Controller error log before you perform any maintenance procedures.
- After you perform a maintenance procedure, check the SAN Volume Controller error log and fix any errors. Expect to see the following types of errors:
  - MDisk error recovery procedures (ERPs)
  - Reduced paths

The following categories represent the types of service actions for storage systems:

- Controller code upgrade
- Field replaceable unit (FRU) replacement

Controller code upgrade

Ensure that you are familiar with the following guidelines for upgrading controller code:

- Check to see if the SAN Volume Controller supports concurrent maintenance for your storage system.
- Allow the storage system to coordinate the entire upgrade process.
- If it is not possible to allow the storage system to coordinate the entire upgrade process, perform the following steps:
  1. Reduce the storage system workload by 50%.
  2. Use the configuration tools for the storage system to manually failover all logical units (LUs) from the controller that you want to upgrade.
  3. Upgrade the controller code.
  4. Restart the controller.
  5. Manually failback the LUs to their original controller.
  6. Repeat for all controllers.

FRU replacement

Ensure that you are familiar with the following guidelines for replacing FRUs:

- If the component that you want to replace is directly in the host-side data path (for example, cable, fibre-channel port, or controller), disable the external data paths to prepare for upgrade. To disable external data paths, disconnect or disable the appropriate ports on the fabric switch. The SAN Volume Controller ERPs reroute access over the alternate path.
- If the component that you want to replace is in the internal data path (for example, cache, or disk drive) and did not completely fail, ensure that the data is backed up before you attempt to replace the component.

- If the component that you want to replace is not in the data path, (for example, uninterruptible power supply units, fans, or batteries) the component is generally dual-redundant and can be replaced without additional steps.

Configuring Bull FDA systems

This section provides information about configuring Bull StoreWay FDA systems for attachment to a SAN Volume Controller.

Supported firmware levels for the Bull FDA

The Bull FDA system must use a firmware level that is supported by the SAN Volume Controller.

See the following Web site for specific firmware levels and the latest supported hardware:www.ibm.com/storage/support/2145

Logical unit creation and deletion for Bull FDA

You can create or delete logical units for the Bull FDA. See the storage configuration guidelines that are specified in the Bull FDA documentation that is provided for this system.

Platform type for Bull FDA

You must set all logical units that the SAN Volume Controller accesses to platform type AX (AIX).

Access control methods for Bull FDA

You can use access control to restrict access from hosts and SAN Volume Controller clusters. You do not need to use access control to allow a SAN Volume Controller cluster to use all of the defined logical units on the system.

The following table lists the access control methods that are available:

Method	Description
Port Mode	Allows access to logical units that you want to define on a per storage controller port basis. SAN Volume Controller visibility (through switch zoning, physical cabling, etc.) must allow the SAN Volume Controller cluster to have the same access from all nodes and the accessible controller ports have been assigned the same set of logical units with the same logical unit number. This method of access control is not recommended for SAN Volume Controller connection.
WWN Mode	Allows access to logical units using the WWPN of each of the ports of an accessing host device. All WWPNs of all the SAN Volume Controller nodes in the same cluster must be added to the list of linked paths in the controller configuration. This becomes the list of host (SAN Volume Controller) ports for an LD Set or group of logical units. This method of access control allows sharing because different logical units can be accessed by other hosts.

Setting cache allocations for Bull FDA

Cache allocations can be set manually; however, changes to the default settings can adversely effect performance and cause you to lose access to the system.

Snapshot Volume and Link Volume for Bull FDA

You cannot use Copy Services logical volumes with logical units that are assigned to the SAN Volume Controller.

Configuring EMC CLARiiON systems

This section provides information about configuring the EMC CLARiiON storage system for attachment to a SAN Volume Controller.

Access Logix

Access Logix is an optional feature of the firmware code that provides the functionality that is known as LUN Mapping or LUN Virtualization.

You can use the software tab in the storage systems properties page of the EMC Navisphere GUI to determine if Access Logix is installed.

After Access Logix is installed it can be disabled but not removed. The following are the two modes of operation for Access Logix:

- **Access Logix not installed:** In this mode of operation, 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 that are assigned to the storage group are allowed to access these LUNs.

Configuring the EMC CLARiiON controller with Access Logix installed

The SAN Volume Controller does not have access to the storage controller logical units (LUs) if Access Logix is installed on the EMC CLARiiON controller. You must use the EMC CLARiiON configuration tools to associate the SAN Volume Controller and LU.

The following prerequisites must be met before you can configure an EMC CLARiiON controller with Access Logix installed:

- The EMC CLARiiON controller is not connected to the SAN Volume Controller
- You have a RAID controller with LUs and you have identified which LUs you want to present to the SAN Volume Controller

You must complete the following tasks to configure an EMC CLARiiON controller with Access Logix installed:

- Register the SAN Volume Controller ports with the EMC CLARiiON
- Configure storage groups

The association between the SAN Volume Controller and the LU is formed when you create a storage group that contains both the LU and the SAN Volume Controller.

Registering the SAN Volume Controller ports with the EMC CLARiiON

You must register the SAN Volume Controller ports with an EMC CLARiiON controller if Access Logix is installed.

The following prerequisites must be met before you can register the SAN Volume Controller ports with an EMC CLARiiON controller that has Access Logix installed:

- The EMC CLARiiON controller is not connected to the SAN Volume Controller
- You have a RAID controller with LUs and you have identified which LUs you want to present to the SAN Volume Controller

Each initiator port [worldwide port name (WWPN)] must be registered against a host name and against a target port to which access is granted. If a host has multiple initiator ports, multiple table entries with the same host name are listed. If a host is allowed access using multiple target ports, multiple table entries are listed. For SAN Volume Controller hosts, all WWPN entries should carry the same host name.

The following table lists the associations:

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	Array

1. Connect the fibre channel and zone the fabric as required.
2. Issue the **svctask detectmdisk** command-line interface (CLI) command.
3. Right-click on the storage system from the Enterprise Storage window.
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. Use the WWPN to Identify them. 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.
9. Click the right arrow to move them to the selected box.
10. Fill in the **HBA WWN** field. You must know the following information:
    - WWNN of each SAN Volume Controller in the cluster
    - 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 is an example of the output:

    ```
    50:05:07:68:01:00:8B:D8:50:05:07:68:01:20:8B:D8
    ```
11. Select A in the field marked SP and 0 in the SP Port field.
12. Select **CLARiiON Open** in the drop down list of the **Initiator Type** field.
13. Deselect the ArrayCommPath checkbox if it has been selected.
14. Select **2** in the drop down list of the **Failover Mode** field.

Attention: Failure to select failover mode 2 prevents the SAN Volume Controller from being able to failover I/O. Your data might become unavailable in the event of a single failure.

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.

15. Select **Array** in the drop down list of the **Unit Serial Number** field.
16. Assign a host name in the Host Name field.
17. Click **OK**.
18. Specify the IP address of your switch. The EMC CLARiiON does not use this IP address. However it must be unique (within the EMC CLARiiON) to prevent errant behavior by Navisphere.
19. Repeat step 11 on page 371 for all possible combinations. The following example shows the different combinations of a system with four ports:
    - SP: A SP Port: 0
    - SP: A SP Port: 1
    - SP: B SP Port: 0
    - SP: B SP Port: 1
20. Repeat steps 1 on page 371 to 19 to register the rest of your SAN Volume ControllerWWPNs.

All your WWPNs are registered against the host name that you specified.

Configuring your storage groups

Storage groups can only be configured if Access Logix is installed and enabled.

Access Logix provides the following LUN mapping:

Notes:

1. A subset of logical units (LUs) can form a storage group.
2. An LU can be in multiple storage groups.
3. A host can be added to a storage group. This host has access to all LUs in the storage group.
4. A host *cannot* be added to a second storage group.

1. Right-click on the storage system from the Enterprise Storage window.
2. Select **Create Storage Group**. The Create Storage Group window is displayed.
3. Enter a name for your storage group in the **Storage Group Name** field.
4. If available, select **Dedicated** in the **Sharing State** field.
5. Click **OK**. The storage group is created.
6. Right-click the storage group in the Enterprise Storage window.
7. Select **Properties**. The Storage Group Properties window is displayed.
8. Perform the following steps from the Storage Group Properties window:
   a. Select the **LUNs** tab.
   b. Select the LUNs that you want the SAN Volume Controller to manage in the Available LUNs table.

      Attention: Ensure that the LUs 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 that you created when you registered the SAN Volume Controller ports with the EMC CLARiiON.

   Attention: Ensure that 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 without Access Logix installed

If Access Logix is not installed on an EMC CLARiiON controller, all logical units (LUs) that were created on the controller can 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.

Supported models of the EMC CLARiiON

The SAN Volume Controller supports models of the EMC CLARiiON.

See the following Web site for the latest supported models:

www.ibm.com/storage/support/2145

Supported firmware levels for the EMC CLARiiON

The EMC CLARiiON must use a firmware level that is supported by the SAN Volume Controller.

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

www.ibm.com/storage/support/2145

Concurrent maintenance on EMC CLARiiON systems

Concurrent maintenance is the ability to perform I/O operations to a controller while simultaneously performing maintenance on it.

Important: An EMC Field Engineer must perform all maintenance procedures.

The EMC CLARiiON FC series and the SAN Volume Controller cluster allow concurrent replacement of the following components:

- Disk drives
- Controller fans (fans must be replaced within 2 minutes or controllers are shut down.)
- Disk enclosure fans (fans must be replaced within 2 minutes or controllers are shut down.)

- Controller (service processor: you must first disable cache)
- Fibre Channel Bypass cards (LCC)
- Power supplies (you must first remove fans.)
- Uninterruptible power supply battery (SPS)

EMC CLARiiON FC devices require that the I/O is quiesced during code upgrade. Consequently, the SAN Volume Controller cluster does not support concurrent upgrade of the FC controller code.

The EMC CLARiiON CX series and the SAN Volume Controller cluster allow 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 are shut down.)
- Uninterruptible power supply battery (SPS)

The SAN Volume Controller cluster and EMC CLARiiON CX devices support concurrent code upgrade of the CX controllers.

Note:

- EMC CLARiiON procedures for concurrent upgrade must be followed in all cases.
- 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* a concurrent operation.

EMC CLARiiON user interfaces

Ensure that you are familiar with the user interface applications that EMC CLARiiON systems use.

Navisphere or Navicli

The following user interface applications are available with EMC CLARiiON systems:

- Navisphere is the Web-based application that can be accessed from any Web browser.
- Navicli is the command-line interface (CLI) that 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 EMC CLARiiON in both cases is out-of-band. Therefore, the host does not need to be connected to the storage over fibre-channel and cannot be connected without Access Logix.

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.

Switch zoning limitations for the EMC CLARiiON systems

There are limitations in switch zoning for SAN Volume Controller clusters and EMC CLARiiON systems.

FC4500 and CX200 models

The EMC CLARiiON FC4500 and CX200 systems limit the number of initiator HBAs to only allow 15 connections for each controller port. This limit is less than the 16 initiator ports that are required to connect to an 8–node cluster in a dual fabric configuration. To use EMC CLARiiON FC4500 and CX200 systems with an 8–node cluster, you must zone the system to use one SAN Volume Controller port for each node in each fabric. This reduces the initiator HBA count to eight.

FC4700 and CX400 models

EMC CLARiiON FC4700 and CX400 systems provide 4 target ports and allow 64 connections. Using a single SAN fabric, a 4-node cluster requires 64 connections ($4 \times 4 \times 4$), which is equal to the number of connections that are allowed. If split support with other hosts is required, this can cause issues. You can reduce either the number of initiator ports or target ports so that only 32 of the available 64 connections are used.

CX600 models

EMC CLARiiON CX600 systems provide 8 target ports and allow 128 connections. A 4-node cluster consumes all 128 connections ($4 \times 4 \times 8$). An 8-node cluster exceeds the connection limit and no reduction methods can be used.

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.

Advanced functions for the EMC CLARiiON

Some advanced functions of the EMC CLARiiON are not supported by the SAN Volume Controller.

Advanced copy functions

Advanced copy functions for EMC CLARiiON, for example, SnapView, MirrorView and SANcopy, are not supported for disks that are managed by the SAN Volume Controller because the copy function does *not* extend to the SAN Volume Controller cache.

MetaLUN

MetaLUN allows a logical unit (LU) to be expanded using LUs in other RAID groups. The SAN Volume Controller only supports MetaLUN for the migration of image mode virtual disks.

Logical unit creation and deletion on the EMC CLARiiON

Binding a logical unit (LU) to a RAID group can take a significant amount of time on EMC CLARiiON systems.

Do not add the LU to a storage group until binding is complete. If the LU is mapped to a SAN Volume Controller cluster during the binding process, the LU might be identified with the wrong capacity. If this occurs, run the following procedure to rediscover the LU with the correct capacity:

1. Unmap the LU from the SAN Volume Controller cluster.
2. Run detectmdisk and wait for the managed disk to be deconfigured.
3. Wait for any binding activity to complete.
4. Remap the LU to the SAN Volume Controller cluster.
5. Run detectmdisk.

Configuring settings for the EMC CLARiiON

A number of settings and options are available through the EMC CLARiiON configuration interface.

The following settings and options are supported by the SAN Volume Controller:

- System
- Port
- Logical unit

Global settings for the EMC CLARiiON

Global settings apply across an EMC CLARiiON system. Not all options are available on all EMC CLARiiON models.

Table 33 lists the global settings that are supported by the SAN Volume Controller.

Table 33. 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

Table 33. EMC CLARiiON global settings supported by the SAN Volume Controller (continued)

Option	EMC CLARiiON default setting	SAN Volume Controller required setting
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

Controller settings for the EMC CLARiiON

The controller settings for the EMC CLARiiON are the settings that apply across one EMC CLARiiON system.

Table 34 lists the options that can be set by the EMC CLARiiON.

Table 34. 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	Default recommended
Statistics Logging	Disable	Either Enable or Disable

Note: The SAN Volume Controller cannot obtain or change the configuration options that are listed above. You must configure the options that are listed above.

Port settings for the EMC CLARiiON

Port settings are configurable at the port level.

Table 35 lists port settings, the EMC CLARiiON defaults, and the required settings for SAN Volume Controller clusters.

Table 35. EMC CLARiiON port settings

Option	EMC CLARiiON default setting	SAN Volume Controller required setting
Port speed	Depends on the model	Any

Note: The SAN Volume Controller cluster cannot obtain or change the configuration options that are listed above. You must configure the options that are listed above.

Logical unit settings for the EMC CLARiiON

Logical unit (LU) settings are configurable at the LU level.

Table 36 on page 378 lists the options that must be set for each LU that is accessed by the SAN Volume Controller. LUs that are accessed by hosts can be configured differently.

Table 36. 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
Write Aside	16 to 65534	16 to 65534

Note: The SAN Volume Controller cannot obtain or change the configuration options that are listed above. You must configure the options that are listed above.

Configuring EMC Symmetrix and Symmetrix DMX systems

This section provides information about configuring the EMC Symmetrix and Symmetrix DMX for attachment to a SAN Volume Controller.

Supported models of the EMC Symmetrix and Symmetrix DMX controllers

The SAN Volume Controller supports models of the EMC Symmetrix and Symmetrix DMX controllers.

See the following Web site for the latest supported models:

www.ibm.com/storage/support/2145

Supported firmware levels for the EMC Symmetrix and Symmetrix DMX

The EMC Symmetrix and Symmetrix DMX must use a firmware level that is supported by the SAN Volume Controller.

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

www.ibm.com/storage/support/2145

Concurrent maintenance on the EMC Symmetrix and Symmetrix DMX

Concurrent maintenance is the capability to perform I/O operations to the EMC Symmetrix or Symmetrix DMX while simultaneously performing maintenance operations on it.

Important: Service actions and upgrade procedures can only be performed by an EMC Field Engineer.

The EMC Symmetrix and Symmetrix DMX are Enterprise class devices that support nondisruptive 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

The SAN Volume Controller and EMC Symmetrix/Symmetrix DMX support concurrent upgrade of the EMC Symmetrix/Symmetrix DMX firmware.

User interfaces on EMC Symmetrix and Symmetrix DMX

Ensure that you are familiar with the user interface applications that support the EMC Symmetrix and Symmetrix DMX systems.

EMC Control Center

A basic EMC Symmetrix or Symmetrix DMX configuration is performed by an EMC Field Engineer (FE) using the EMC Symmetrix service processor. After the initial configuration, you can configure and control the exported storage. The FE defines the storage device types and sets the configurable options.

You can configure and control the exported storage as described below.

You can use the EMC Control Center to manage and monitor the EMC Symmetrix and Symmetrix DMX systems.

You can use Volume Logix for volume configuration management. Volume Logix allows you to control access rights to the storage when multiple hosts share target ports.

SYMCLI

The EMC Symmetrix Command Line Interface (SYMCLI) allows the server to monitor and control the EMC Symmetrix and Symmetrix DMX.

Sharing the EMC Symmetrix or Symmetrix DMX system between a host and a SAN Volume Controller cluster

There are restrictions for sharing EMC Symmetrix and Symmetrix DMX systems between a host and a SAN Volume Controller cluster.

An EMC Symmetrix or Symmetrix DMX system can be shared between a host and a SAN Volume Controller under the following conditions:

- When possible, avoid sharing target ports between the SAN Volume Controller cluster and other hosts. If this cannot be avoided, you must regularly check the combined I/O workload that is generated by the SAN Volume Controller cluster and the other hosts. The performance of either the SAN Volume Controller cluster or the hosts is impacted if the workload exceeds the target port capabilities.
- A single host must not be connected to a SAN Volume Controller and an EMC Symmetrix or Symmetrix DMX because the multipathing drivers (for example, subsystem device driver (SDD) and PowerPath) cannot coexist.
- If the EMC Symmetrix or Symmetrix DMX is configured in such a way that other hosts cannot access the LUs that are managed by the SAN Volume Controller cluster, other hosts can be directly connected to an EMC Symmetrix or Symmetrix DMX system at the same time as a SAN Volume Controller cluster.

Switch zoning limitations for the EMC Symmetrix and Symmetrix DMX

There are limitations in switch zoning for the SAN Volume Controller and the EMC Symmetrix and Symmetrix DMX systems.

Switch zoning

The SAN Volume Controller switch zone must include at least one target port on two or more fibre-channel adapters to avoid a single point of failure.

The EMC Symmetrix and Symmetrix DMX must be configured to present logical units (LUs) to all SAN Volume Controller initiator ports that are in the fabric zone.

Only SAN Volume Controller initiator ports that are LUN masked on the EMC Symmetrix or Symmetrix DMX controller should be present in the fabric zone.

Note: The EMC Symmetrix and Symmetrix DMX systems present themselves to a SAN Volume Controller cluster as separate controllers for each port zoned to the SAN Volume Controller. For example, if one of these storage systems has 4 ports zoned to the SAN Volume Controller, each port appears as a

separate controller rather than one controller with 4 WWPNs. In addition, a given logical unit (LU) must be mapped to the SAN Volume Controller through all controller ports zoned to the SAN Volume Controller using the same logical unit number (LUN).

Connecting to the SAN

You can connect a maximum of 16 EMC Symmetrix or Symmetrix DMX ports to the SAN Volume Controller cluster. There are no further special zoning requirements. Configurations that are setup to adhere to the requirements that are described in previous SAN Volume Controller releases are also supported, but should not be followed for new installations.

Quorum disks on EMC Symmetrix and Symmetrix DMX

The SAN Volume Controller chooses managed disks (MDisks) that are presented by the EMC Symmetrix or Symmetrix DMX as quorum disks.

The SAN Volume Controller uses a logical unit (LU) that is presented by an EMC Symmetrix or Symmetrix DMX as a quorum disk. The SAN Volume Controller provides a quorum disk even if the connection is through a single port.

Advanced functions for EMC Symmetrix and Symmetrix DMX

SAN Volume Controller cache-disabled virtual disks (VDisks) can be used as the source or target in Symmetrix advanced copy functions (for example, SRDF and TimeFinder).

LU creation and deletion on EMC Symmetrix and Symmetrix DMX

A logical unit (LU) that is exported by an EMC Symmetrix or Symmetrix DMX, meaning it is visible to a host, is either a *Symmetrix device* or a *Meta device*.

Symmetrix device

Restriction: An LU with a capacity of 32 MB or less is ignored by the SAN Volume Controller.

Symmetrix device is an EMC term for an LU that is hosted by an EMC Symmetrix. These are all emulated devices and have exactly the same characteristics. The following are the characteristics of a Symmetrix device:

- 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 EMC Symmetrix Command Line Interface (SYMCLI). The configuration of an LU can be changed using the **convert dev** command from the SYMCLI. Each physical storage device in an EMC Symmetrix is partitioned into 1 to 128 hyper-volumes (hypers). Each hyper can be up to 16 GB. A Symmetrix device maps to one or more hypers, depending on how it is configured. The following are examples of hyper configurations:

- Hypers can be mirrored (2-way, 3-way, 4-way)
- Hypers can be formed into RAID-S groups

Meta device

Meta device is an EMC term for a concatenated chain of EMC Symmetrix devices. This enables the EMC Symmetrix to provide LUs that are larger than a hyper. Up to 255 hypers can be concatenated to form a single meta device. Meta devices can be created using the **form meta** and **add dev** commands from the SYMCLI. This allows an extremely large LU to be created, however, if exported to the SAN Volume Controller, only the first 2 TB is used.

Do not extend or reduce meta devices that are used for managed disks (MDisks). Reconfiguration of a meta device that is used for an MDisk causes unrecoverable data-corruption.

Configuring settings for the EMC Symmetrix and Symmetrix DMX

A number of settings and options are available through the EMC Symmetrix configuration interface.

The settings and options can have a scope of the following:

- System
- Port
- Logical unit (LU)

Global settings for the EMC Symmetrix and Symmetrix DMX

Global settings apply across the EMC Symmetrix and Symmetrix DMX systems.

You can specify EMC Symmetrix and Symmetrix DMX settings with the **set Symmetrix** command from the Symmetrix Command Line Interface (SYMCLI). The settings can be viewed using the **symconfigure** command from the SYMCLI.

Table 37 lists the EMC Symmetrix global settings that can be used with SAN Volume Controller clusters.

Table 37. EMC Symmetrix and Symmetrix DMX global settings

Option	EMC Symmetrix and Symmetrix DMX default setting	SAN Volume Controller required setting
max_hypers_per_disk	-	Any
dynamic_rdf	Disable	Any
fba_multi_access_cache	Disable	N/A
Raid_s_support	Disable	Enable or Disable

Port settings for the EMC Symmetrix and Symmetrix DMX

Target port characteristics can be set using the **set port** command from the Symmetrix Command Line Interface (SYMCLI).

The target port characteristics can be viewed using the **symcfg** command from the SYMCLI.

Table 38 on page 383 lists the EMC Symmetrix and Symmetrix DMX port settings that can be used with the SAN Volume Controller cluster.

Table 38. EMC Symmetrix and Symmetrix DMX port settings that can be used with the SAN Volume Controller

Option	EMC Symmetrix and Symmetrix DMX default setting	SAN Volume Controller required setting
Disk_Array	Enabled	Disabled
Volume_Set_Addressing	Enabled	Disabled
Hard_Addressing	Enabled	Enabled
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 or Enabled
Environ_Set	Disabled	Disabled
Unique_WWN	Enabled	Enabled
Point_to_Point	Disabled	Enabled
VCM_State	Disabled	Disabled or Enabled
OpenVMS	Disabled	Disabled

Logical unit settings for the EMC Symmetrix and Symmetrix DMX

Logical unit (LU) settings are configurable at the LU level.

LU characteristics can be set using the **set device** command from the Symmetrix Command Line Interface (SYMCLI).

Table 39 lists the options that must be set for each LU that is accessed by the SAN Volume Controller.

Table 39. EMC Symmetrix and Symmetrix DMX LU settings supported by the SAN Volume Controller

Option	EMC Symmetrix and Symmetrix DMX default setting	SAN Volume Controller required setting
emulation	-	FBA
attribute	-	Set all attributes to disabled.

Mapping and virtualization settings for the EMC Symmetrix and Symmetrix DMX

Mapping a logical unit (LU) to a host is a function of the EMC Control Center.

LUs can be mapped to a particular director or target port using the **map dev** command from the Symmetrix Command Line Interface (SYMCLI). LUs can be unmapped using the **unmap dev** command from the SYMCLI.

Volume Logix and masking

Volume Logix allows you to restrict access to particular WWPNs on the fabric for Symmetrix Volumes.

This function can be switched on and off by changing the VMC_State port setting. The SAN Volume Controller requires that you do not share target ports between a host and a SAN Volume Controller. However, you can still use Volume Logix to protect the system from errors that can occur if the SAN is not correctly configured.

To mask a volume to the SAN Volume Controller, you must first identify the SAN Volume Controller ports that are connected to each system. This can be done using the EMC Symmetrix symmask command.

The SAN Volume Controller automatically logs into any EMC Symmetrix system it sees on the fabric. You can use the SAN Volume Controller svcinfo lsnode CLI command to find the correct port identifiers.

After you have identified the ports, you can map each volume on each port to each WWPN. The EMC Symmetrix stores the LUN masking in a database, so you must apply the changes you have made to refresh the contents of the database to view the changes.

Configuring Fujitsu ETERNUS systems

This section provides information about configuring the Fujitsu ETERNUS systems for attachment to a SAN Volume Controller.

Supported models of the Fujitsu ETERNUS

The SAN Volume Controller supports models of the Fujitsu ETERNUS series of systems.

See the following Web site for the latest supported models:www.ibm.com/storage/support/2145

Supported firmware levels for the Fujitsu ETERNUS

The Fujitsu ETERNUS must use a firmware level that is supported by the SAN Volume Controller.

See the following Web site for specific firmware levels and the latest supported hardware:www.ibm.com/storage/support/2145

User interfaces on the Fujitsu ETERNUS

Ensure that you are familiar with the user interface application that is used by the Fujitsu ETERNUS.

You can use the ETERNUSmgr web-based configuration utility. See the documentation that is provided with the Fujitsu ETERNUS system for more information.

Configuring the Fujitsu ETERNUS to use with the SAN Volume Controller

Ensure that you use the settings that are required to use the Fujitsu ETERNUS with the SAN Volume Controller. It is important that you use the correct settings to avoid data access problems.

Use the following sequence of steps to configure the Fujitsu ETERNUS system:

1. Configure the SAN Volume Controller host response pattern.
2. Register the host world wide names (WWNs) and associate them with the host response pattern.
3. Setup the affinity group for SAN Volume Controller volumes or setup LUN mapping.
4. Create or reassign storage to the SAN Volume Controller.

For all other settings and procedures, consider the SAN Volume Controller a host. See the documentation that is provided with the Fujitsu ETERNUS system.

CA parameters

The following table lists the port settings that are required. See the documentation that is provided with your Fujitsu ETERNUS system for more information because some options are only available on certain models.

Option	Fujitsu ETERNUS default setting	SAN Volume Controller required setting
Connection Topology/FC Connection Settings	FC-AL Connection	Fabric Connection
Service Class	Class 3	Class 3
FC Transfer Rate	Auto Setting	Any
Reset Scope/Scope of LUR Actions	T_L	T_L **Note:** If this option is not set correctly, data corruption can occur.
Release Reservation upon Chip Reset	Enable/valid	Enable/valid
HP-UX Connection Setting	Disable	Disable
Frame Size Setting	2048	Any
Affinity/Addressing Mode	Off	Any

Host response pattern

The SAN Volume Controller requires that a new host response pattern is created. If the Host Affinity/Host Table Settings Mode is used, this host response pattern must be associated with each WWN. If the Host Affinity/Host Table Settings Mode is not used, this host response pattern must be associated with the target port.

The following table lists the settings that are required. See the documentation that is provided with your Fujitsu ETERNUS system for more information because some options are only available on certain models.

Option	Fujitsu ETERNUS default setting	SAN Volume Controller required setting
Command timeout interval	Depends on the Fujitsu ETERNUS model	Default
Response status in overload	Unit Attention	Unit Attention
Byte 0 of Inquiry response/Response to inquiry commands	Default	Default
Inquiry Standard Data NACA Function	Disable	Disable
Inquiry Standard Data Version	Depends on the Fujitsu ETERNUS model	Default
Inquiry Command Page 83/Inquiry VPD ID Type	Depends on the Fujitsu ETERNUS model	Type 01
Reservation Conflict Response to Test Unit Ready Commands	Disable/Normal Response	Enable/Conflict Response
Target Port Group Access Support	Disable	Enable
Host Specific Mode	Normal Mode	Normal Mode
Response Sense at Firmware Hot Switching	Enable	Enable
Change LUN mapping	No Report	Report
LUN Capacity Expansion	No Report	Report
Aymmetric / Symmetric Logical Unit Access	Active/Active	Active/Active
Pattern of Sense Code Conversion	No Conversion	No Conversion

Notes:

1. If you set Inquiry VPD ID Type option to Type 3 on E4000 or E8000 range, the MDisks go offline.
2. If you set the Target Port Group Access Support option to Disabled on E3000 range, a 1370 error is shown in the error log.

Host WWNs

After the SAN Volume Controller is zoned on the fabric to see the Fujitsu ETERNUS, the system might not initially appear in the list of controllers when you issue the **lscontroller** CLI command. This is normal and expected behavior.

See the documentation that is provided with the Fujitsu ETERNUS system to add all SAN Volume Controller WWPNs as host WWNs. The following restrictions apply:

- The SAN Volume Controller WWNs must be associated with a host response pattern. The host response pattern must be defined prior to registration. If you use an incorrect or default host response pattern, you can lose access to data.
- All SAN Volume Controller WWNs must be registered on all Fujitsu ETERNUS ports on the same fabric. If the WWNs are not registered, you can lose access to data.

Affinity groups/zones

Use the affinity groups/zones mode to protect the SAN Volume Controller LUs if the SAN is incorrectly configured. The affinity group mode is setup in the CA configuration. See the documentation that is provided with your Fujitsu ETERNUS system for more information about using the affinity groups/zones mode. The following restrictions apply:

- Each SAN Volume Controller must have exactly one affinity group/zone.
- The SAN Volume Controller affinity group/zone must be associated with all SAN Volume Controller WWNs.

LUN mapping

You can use the LUN mapping mode (also called the zone settings mode for some models) with the following restrictions:

- The SAN zoning must only allow a single SAN Volume Controller access to this target port.
- The host response pattern must be set in CA configuration using the required SAN Volume Controller settings.

Note: If you use the LUN mapping mode, you cannot use the host affinity mode. The host affinity mode is set to OFF.

Assigning storage to the SAN Volume Controller

Ensure that you understand all SAN Volume Controller and Fujitsu ETERNUS restrictions before you assign storage to the SAN Volume Controller. See the documentation that is provided with the Fujitsu ETERNUS system for more information.

Zoning configuration for the Fujitsu ETERNUS

If LUN mapping mode is used for a Fujitsu ETERNUS port, you must exclusively zone the SAN Volume Controller with this target port.

Migrating logical units from the Fujitsu ETERNUS to the SAN Volume Controller

You can use the standard migration procedure with the following restrictions:

- The SAN Volume Controller must have software level 4.2.0 or higher installed before you start migration. Upgrades from previous SAN Volume Controller software levels to software level 4.2.0 or higher causes all Fujitsu ETERNUS systems that are attached to be excluded.
- You must configure the Fujitsu ETERNUS system to work with the SAN Volume Controller before you start migration.
- The subsystem device driver (SDD) and Fujitsu Multipath driver cannot coexist.
- The SAN Volume Controller must support all host code levels.

Concurrent maintenance on the Fujitsu ETERNUS

Concurrent maintenance is the capability to perform I/O operations to a Fujitsu ETERNUS while simultaneously performing maintenance operations on it.

You can perform nondisruptive maintenance procedures concurrently on the following components:

- Fujitsu ETERNUS controller module
- Fujitsu ETERNUS controller cache
- Fujitsu ETERNUS cache battery pack
- Fan
- Power supply
- Disk drive
- SFP transceiver

See the documentation that is provided with the Fujitsu ETERNUS system for more information.

Advanced functions for the Fujitsu ETERNUS

The Fujitsu ETERNUS system provides several Advanced Copy functions. Do not use these Advanced Copy functions for storage that is managed by the SAN Volume Controller, even if the VDisk cache is disabled.

Configuring IBM TotalStorage ESS systems

This section provides information about configuring the IBM TotalStorage Enterprise Storage Server (ESS) for attachment to a SAN Volume Controller.

Configuring the IBM ESS

The IBM Enterprise Storage Server (ESS) provides functionality that is compatible with the SAN Volume Controller.

Perform the following steps to configure the IBM ESS:

1. Enter the IP address of the IBM ESS in a Web browser to access the ESS Specialist.
2. Login with your 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 each initiator port on each SAN Volume Controller node in your cluster. Complete the following fields:
   a. Enter a unique name for each port in the **Nickname** field. For example, enter `knode` or `lnode`.
   b. Select **IBM SAN Volume Controller** in the **Host Type** field. If this option is not available, select **RS/6000**.
   c. Select **Fibre Channel attached** in the **Host Attachment** field.
   d. Leave the **Hostname/IP address** field blank.
   e. Select the WWPN from the list or enter it manually into the **WWPN** field. A configuration command fails if you use WWPN 0 in the command string.
8. Click **Perform Configuration Update** after you are finished adding all of the ports.
9. Click **Add Volumes** to add the volumes that you want the SAN Volume Controller to use. The Add Volumes panel is displayed.
10. Perform the following steps in the Add Volumes panel:

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. Click **Perform Configuration Update** after you have created all the volumes.

11. Perform the following steps to map the volumes to all of your SAN Volume Controller ports:
    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. Click **Perform Configuration Update**.

    Important: If you are adding SAN Volume Controller ports to a volume that is already assigned to other SAN Volume Controller ports, you must select the **Use same ID/LUN in source and target** check box.

Supported models of the IBM ESS

The SAN Volume Controller supports models of the IBM Enterprise Storage Server (ESS).

See the following Web site for the latest supported models:

www.ibm.com/storage/support/2145

Supported firmware levels for the IBM ESS

The SAN Volume Controller supports the IBM Enterprise Storage Server (ESS).

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

www.ibm.com/storage/support/2145

Concurrent maintenance on the IBM ESS

Concurrent maintenance is the capability to perform I/O operations to an IBM Enterprise Storage Server (ESS) while simultaneously performing maintenance operations on it.

All IBM ESS concurrent maintenance procedures are supported.

User interface on the IBM ESS

Ensure that you are familiar with the user interface application that supports the IBM Enterprise Storage Server (ESS) system.

Web Server

A Web server runs on each of the controllers on the system. During normal operation, the user interface application allows only basic monitoring of the system

and displays an error log. If you press the reset button on the controller to put the controller into diagnostic mode, the user interface application allows firmware upgrades and system configuration resets.

Sharing the IBM ESS between a host and the SAN Volume Controller

The IBM Enterprise Storage Server (ESS) can be shared between a host and a SAN Volume Controller.

The following restrictions apply when you share the IBM ESS between a host and a SAN Volume Controller:

- If an IBM ESS port is in the same zone as a SAN Volume Controller port, that same IBM ESS port should not be in the same zone as another host.
- A single host can have both IBM ESS direct-attached and SAN Volume Controller virtualized disks configured to it.
- If a LUN is managed by the SAN Volume Controller, it *cannot* be mapped to another host.

See the following Web site for the latest supported configurations:

www.ibm.com/storage/support/2145

Switch zoning limitations for the IBM ESS

Consider the following limitations when you zone the IBM Enterprise Storage Server (ESS) to the SAN Volume Controller.

To avoid a single point of failure on the IBM ESS, you must have a minimum of two SAN connections from two separate adapter bays. The maximum number of IBM ESS SAN connections in the SAN Volume Controller switch zone is 16.

Note: The IBM ESS provides ESCON®, FICON® and Ultra SCSI connectivity; however, only a 1 or 2 Gb fibre-channel SAN attachment is supported by the SAN Volume Controller.

Quorum disks on the IBM ESS

The SAN Volume Controller can choose managed disks (MDisks) that are presented by the IBM Enterprise Storage Server (ESS) controller as quorum disks.

Advanced functions for the IBM ESS

SAN Volume Controller cache-disabled virtual disk (VDisks) can be used as the source or target for IBM Enterprise Storage Server (ESS) advanced copy functions (for example, FlashCopy, MetroMirror, GlobalCopy).

Logical unit creation and deletion on the IBM ESS

Certain IBM Enterprise Storage Server (ESS) types are supported for use with the SAN Volume Controller.

Before you delete or unmap a logical unit (LU) from the SAN Volume Controller, remove the LU from the managed disk (MDisk) group. The following is supported:

- LU size of 1 GB to 2 TB.
- RAID 5 and RAID 10 LUs.
- LUs can be added dynamically.

Attention: When adding additional SAN Volume Controller ports to an existing LU, you must select the **Use same ID/LUN in source and target** check box. Failure to select the **Use same ID/LUN in source and target** checkbox can cause loss in redundancy or a loss of data. If this checkbox is not available, the option is not required. The detect MDisks task in the SAN Volume Controller Console or the **svctask detectmdisk** command-line interface (CLI) command must be run for the SAN Volume Controller to detect the new disks.

Configuring IBM System Storage DS5000, IBM DS4000 and IBM DS3000 systems

This section provides information about configuring IBM System Storage DS5000, IBM DS4000 and IBM DS3000 systems for attachment to a SAN Volume Controller cluster. Some IBM System Storage DS4000 controllers are equivalent to StorageTek models; SAN Volume Controller also supports certain StorageTek FlexLine series and StorageTek D series. The information in this section also applies to the supported models of the StorageTek FlexLine series and StorageTek D series.

IBM System Storage DS5000, IBM DS4000, and IBM DS3000 are similar systems. The concepts in this section apply generally to all three systems; however, some options may not be available. See the documentation that is provided with your system for specific information.

Configuring IBM System Storage DS5000 and IBM DS4000 systems for the storage server

IBM System Storage DS5000 and IBM DS4000 disk controllers are supported with the SAN Volume Controller cluster.

The following steps provide the supported options and impact on the SAN Volume Controller cluster:

1. Perform the following steps for the host type option:
   a. Depending on your IBM System Storage DS5000 or IBM DS4000 model, you must set either the default host type of the system, or the host type of the chosen partition to one of the following:

      ```
      IBM TS SAN VCE
      SAN Volume Contr
      ```

      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. Perform the following steps for the worldwide node name (WWNN) option:
   a. Set the system so that both controllers have the same WWNN.
   b. See the following Web site for the scripts that are available to change the setup of the IBM System Storage DS5000 or IBM DS4000 system:

      www.ibm.com/storage/support/
3. Perform the following steps for the auto volume transfer (AVT) option:
   a. Ensure that the AVT option is enabled. The host type selection should have already enabled the AVT option.
   b. View the storage system profile data to confirm that you have the AVT option enabled. This storage profile is presented as a text view in a separate window.
   c. See the following Web site for the scripts that are available to enable the AVT option:

www.ibm.com/storage/support/

The following limitations apply to partitions:

- Only one IBM System Storage DS5000 or IBM DS4000 system storage partition that contains any of the ports of any of the nodes in a single SAN Volume Controller cluster can be created.
- Only map one partition to any of the ports on any of the nodes that are in the SAN Volume Controller cluster to avoid unexpected behavior. For example, you can lose access to your storage or you might not receive warning messages, even if there are errors logged in the SAN Volume Controller error log.

The following limitation applies to IBM System Storage DS5000 and IBM DS4000 Copy Services:

- Do not use IBM System Storage DS5000 or IBM DS4000 Copy Services when the SAN Volume Controller cluster is attached to an IBM System Storage DS5000 or IBM DS4000 system.
- You can use partitioning to allow IBM System Storage DS5000 or IBM DS4000 Copy Services usage for other hosts.

The following information applies to the access LUN (also known as the Universal Transport Mechanism (UTM) LUN):

- The access/UTM LUN is a special LUN that allows a IBM System Storage DS5000 or IBM DS4000 system to be configured through software over the fibre-channel connection. The access/UTM LUN does not have to be in the partition that contains the SAN Volume Controller ports because the access/UTM LUN is not required by the SAN Volume Controller cluster. No errors are generated if the access/UTM LUN is not in the partition.

The following information applies to the logical unit (LU):

- The SAN Volume Controller cluster attempts to follow the preferred ownership that is specified by the IBM System Storage DS5000 or IBM DS4000 system. You can specify which controller (A or B) is used for I/O operations to an LU.
- If the SAN Volume Controller cluster can see the ports of the preferred controller and error conditions do not exist, the SAN Volume Controller cluster accesses the LU through one of the ports on the preferred controller.
- If error conditions exist, the SAN Volume Controller cluster ignores the preferred ownership of the IBM System Storage DS5000 or IBM DS4000 system.

Supported options for IBM System Storage DS5000 and IBM DS4000 controllers

IBM System Storage DS5000 and IBM DS4000 series disk controllers provide functions that can be used with the SAN Volume Controller.

The storage manager for IBM System Storage DS5000 and IBM DS4000 systems has several options and actions that you can perform.

Controller run diagnostics

Diagnostics are automatically recovered by the SAN Volume Controller software. After the controller run diagnostics option is used, check your managed disks (MDisks) to ensure that they have not been set to degraded mode.

Controller disable data transfer

The controller disable data transfer option is not supported when a SAN Volume Controller is attached to IBM System Storage DS5000 or IBM DS4000 systems.

Setting an array Offline

Do not set an array offline because you can lose access to the MDisk group.

Array increase capacity

The array increase capacity option is supported but the new capacity is not usable until the MDisk is removed from the MDisk group and re-added to the MDisk group. You might have to migrate data to increase the capacity.

Redistribute logical drives or change ownership of the preferred path

You can redistribute logical drives or change ownership of the preferred path; however, these options might not take effect until a discovery is started on the SAN Volume Controller cluster. You can use the **svctask detectmdisk** command-line interface (CLI) command to restart a cluster discovery process. The discovery process rescans the fibre-channel network to discover any new MDisks that might have been added to the cluster and to rebalance MDisk access across the available controller device ports.

Controller reset

You must only use the controller reset option if you are directed to do so by IBM Service and the alternate controller is functional and available to the SAN. The SAN Volume Controller reset is automatically recovered by the SAN Volume Controller software.

Check your MDisks to ensure that they have not been set to the degraded state during the controller reset process. You can issue the **svctask includemdisk** CLI command to repair degraded MDisks.

Supported models of IBM System Storage DS5000, IBM DS4000 and IBM DS3000 systems

The SAN Volume Controller supports models of the IBM System Storage DS5000, IBM DS4000 and IBM DS3000 systems. Some IBM System Storage DS4000 series controllers are equivalent to Sun StorageTek and StorageTek models; SAN Volume Controller also supports some Sun StorageTek, StorageTek FlexLine and D series models.

See the following Web site for the latest supported models:

www.ibm.com/storage/support/2145

Note: Some levels of IBM System Storage DS5000 microcode support a maximum of 32 LUNs per host partition, newer versions allow up to 256 LUNs per host partition.

Supported firmware levels for IBM System Storage DS5000, IBM DS4000 and IBM DS3000 systems

You must ensure that the firmware level of the system can be used with the SAN Volume Controller cluster.

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

www.ibm.com/storage/support/2145

The Web site includes the maximum number of LUNs per partition that are supported by the firmware level.

Concurrent maintenance on IBM System Storage DS5000 and IBM DS4000 systems

Concurrent maintenance is the capability to perform I/O operations to an IBM System Storage DS5000 or IBM DS4000 series controller while simultaneously performing maintenance operations on the system.

See your IBM System Storage DS5000 or IBM DS4000 series documentation for information about concurrent maintenance.

IBM System Storage DS5000, IBM DS4000 and IBM DS3000 systems user interface

Ensure that you are familiar with the user interface that supports the IBM System Storage DS5000, IBM DS4000 and IBM DS3000 systems.

Web Server

A Web server is running on each of the controllers in the system. During normal operation, the user interface only allows basic monitoring of the system and displays an error log. If you press the reset button to put a controller in diagnostic mode, the user interface allows firmware upgrades and system configuration resets.

Sharing an IBM System Storage DS5000, IBM DS4000 or IBM DS3000 system between a host and the SAN Volume Controller

You can share an IBM System Storage DS5000, IBM DS4000 or IBM DS3000 system between a host and a SAN Volume Controller cluster.

The IBM System Storage DS5000 and IBM DS4000 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 logical units that are accessed by the SAN Volume Controller cluster.

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 controller ports. At least one port from each controller must be visible by the SAN Volume Controller cluster.

Quorum disks on IBM System Storage DS5000, IBM DS4000 and IBM DS3000 systems

The SAN Volume Controller can choose managed disks (MDisks) that are presented by a IBM System Storage DS5000, IBM DS4000 or IBM DS3000 system as quorum disks.

Note: The FASsT series 200 does not support quorum disks.

Advanced functions for IBM System Storage DS5000, IBM DS4000 and IBM DS3000 systems

SAN Volume Controller cache-disabled virtual disks (VDisks) can be used as the source or target for IBM System Storage DS5000, IBM DS4000 and IBM DS3000 advanced copy functions: for example, FlashCopy and Metro Mirror.

Data migration on partitioned IBM System Storage DS5000 and IBM DS4000 systems

You can migrate data on partitioned IBM System Storage DS5000 and IBM DS4000 systems.

You can enable the SAN Volume Controller to be introduced to an existing SAN environment, so that you have the option of using image mode LUNs to import the existing data into the virtualization environment without requiring a backup and restore cycle. Each partition can only access a unique set of HBA ports, as defined by the worldwide port names (WWPNs). For a single host to access multiple partitions, unique host fibre ports (WWPNs) must be assigned to each partition. All LUNs within a partition are identified to the assigned host fibre ports (no subpartition 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, you must 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 backup and restore cycles requires a set of unique SAN Volume Controller HBA port WWPNs for each partition. This allows the IBM System Storage DS5000 or IBM DS4000 system to make the LUNs known to the SAN Volume Controller, which then configures these LUNs as image-mode LUNs and identifies them to the required hosts. This violates a requirement that all SAN Volume Controller nodes must be able to see all back-end storage. For example, to fix this problem for an IBM DS4000 system, change the configuration to allow more than 32 LUNs in one storage partition, so that you can move all the LUNs from all the other partitions into one partition and map to the SAN Volume Controller cluster.

Scenario: the SAN Volume Controller nodes cannot see all back-end storage

The IBM DS4000 series has eight partitions with 30 LUNs in each.

Perform the following steps to allow the SAN Volume Controller nodes to see all back-end storage:

1. Change the mappings for the first four partitions on the IBM DS4000 system such that each partition is mapped to one port on each node. This maintains redundancy across the cluster.
2. Create a new partition on the system that is mapped to all four ports on all the nodes.
3. Gradually migrate the data into the managed disks (MDisks) in the target partition. As storage is freed from the source partitions, it can be reused as new storage in the target partition. As partitions are deleted, new partitions that must be migrated can be mapped and migrated in the same way. The host side data access and integrity is maintained throughout this process.

Logical unit creation and deletion on IBM System Storage DS5000 and IBM DS4000 systems

You can create or delete logical units on IBM System Storage DS5000 and IBM DS4000 systems.

Some IBM System Storage DS5000 and IBM DS4000 controllers are supported for use with SAN Volume Controller clusters.

To create a logical disk, you must set either the default host type of the IBM System Storage DS5000 or IBM DS4000 system, or the host type of the chosen partition to one of the following settings, depending on the model:

```
IBM TS SAN VCE
SAN Volume Contr
```

Perform one of the following tasks to set the host type:

- Click **Storage Subsystem → Change → Default Host Type**
- For each host port, specify the host type of that port or modify existing ports.

Configuration interface for IBM System Storage DS5000 and IBM DS4000 systems

IBM System Storage DS5000 and IBM DS4000 systems include a configuration application.

The access LUN, also known as the Universal Transport Mechanism (UTM) LUN, is the configuration interface for IBM System Storage DS5000 and IBM DS4000 systems.

The access LUN might not be in a partition that contains the SAN Volume Controller ports because it is not required by the SAN Volume Controller cluster. The UTM LUN is a special LUN that allows IBM System Storage DS5000 and IBM DS4000 systems to be configured through suitable software over the fibre-channel connection. Because the SAN Volume Controller does not require the UTM LUN, it does not generate errors either way. IBM System Storage DS5000 and IBM DS4000 systems *must not* have the Access UTM LUN that is presented as LUN 0 (zero).

It is possible to use in-band (over fibre channel) and out-of-band (over Ethernet) to allow the configuration software to communicate with more than one IBM System Storage DS5000 or IBM DS4000 system. If using in-band configuration, the Access UTM LUN must be configured in a partition that does not include any logical units that are accessed by the SAN Volume Controller cluster.

Note: In-band is not supported for access to the LUN while in the SAN Volume Controller partition.

Controller settings for IBM System Storage DS5000 and IBM DS4000 systems

Controller settings are the settings that apply across one IBM System Storage DS5000 or IBM DS4000 system.

You must configure the following settings for IBM System Storage DS5000 and IBM DS4000 systems:

- Depending on your model, you must set either the default host type of your IBM System Storage DS5000 or IBM DS4000 system, or the host type of the chosen partition, to one of the following:

  ```
  IBM TS SAN VCE
  SAN Volume Contr
  ```

  Perform one of the following tasks to set the host type:
  - Click **Storage Subsystem** → **Change** → **Default Host Type**
  - For each host port, specify the host type of that port or modify existing ports.
- Set the system so that both controllers have the same worldwide node name (WWNN). See the following Web site for the scripts that are available to change the setup for IBM System Storage DS5000 and IBM DS4000 systems:

  www.ibm.com/storage/support/
- Ensure that the AVT option is enabled. The host type selection should have already enabled the AVT option. View the storage system profile data to confirm that you have the AVT option enabled. This storage profile is presented as a text view in a separate window. See the following Web site for the scripts that are available to enable the AVT option:

  www.ibm.com/storage/support/
- You must have the following options enabled on any logical units that are mapped to IBM System Storage DS5000 and IBM DS4000 systems:
  - read caching
  - write caching
  - write cache mirroring
- You must *not* have caching without batteries enabled.

Configuration settings for IBM System Storage DS5000, IBM DS4000 and IBM DS3000 systems

The system configuration interface provides configuration settings and options that can be used with the SAN Volume Controller cluster.

These settings and options can have the following scope:

- System
- Logical unit (LU)
  - The SAN Volume Controller cluster attempts to follow preferred ownership that is specified by the system. You can specify which controller (A or B) is used to perform I/O operations to a given LU. If the SAN Volume Controller cluster can see the ports of the preferred controller and no error conditions exist, the SAN Volume Controller cluster accesses that LU through one of the ports on that controller. Under error conditions, the ownership is ignored.

– You must have the following options enabled on any LUs that are mapped to the SAN Volume Controller cluster:
  - read caching
  - write caching
  - write cache mirroring
– You must *not* have caching without batteries enabled.

Global settings for IBM System Storage DS5000 or IBM DS4000 systems

Global settings apply across IBM System Storage DS5000 or IBM DS4000 systems.

Table 40 lists the global settings that can be used with SAN Volume Controller clusters.

Table 40. IBM System Storage DS5000 and DS4000 system global options and required settings

Option	IBM DS5000 or IBM DS4000 system default setting
Start flushing	80%
Stop flushing	80%
Cache block size	4 Kb

Attention: Do not modify these settings unless you are directed by the IBM Support Center

Depending on the IBM DS5000 or IBM DS4000 model, use a host type of IBM TS SAN VCE or SAN Volume Contr to establish the correct global settings for the SAN Volume Controller cluster. Either set this as the system default host type or, if partitioning is enabled, associate each SAN Volume Controller port with this host type.

Logical unit settings for IBM System Storage DS5000, IBM DS4000 and IBM DS3000 systems

Logical unit (LU) settings are configurable at the LU level.

LUs that are accessed by hosts can be configured differently.

The read ahead cache multiplier is typically set to 0 or 1. Do not modify this setting unless you are directed to do so by the IBM Support Center.

The following options must be enabled on any LUs that are mapped to the SAN Volume Controller cluster:

- read caching
- write caching
- write cache mirroring

You must not have caching without batteries enabled.

Depending on your system model, set the host type to the one of following when you create a new LU:

```
IBM TS SAN VCE
SAN Volume Contr
```

Miscellaneous settings for IBM System Storage DS5000, IBM DS4000 or IBM DS3000 systems

The SAN Volume Controller cluster supports all media scan settings that are provided by the system. Set the background media scan to enabled and set the frequency to 30 days. These settings are enabled at both the system level and the individual logical drive level.

See the documentation that is provided with your system for information about other settings.

Configuring IBM System Storage DS6000 systems

This section provides information about configuring the IBM System Storage DS6000™ system for attachment to a SAN Volume Controller.

Configuring the IBM DS6000

The IBM DS6000 provides functions that are compatible with the SAN Volume Controller.

After you have defined at least one storage complex, storage unit, and I/O port, you can define the SAN Volume Controller as a host and create host connections. If you have not defined all of these required storage elements, use the IBM System Storage DS6000 Storage Manager or the IBM DS6000 command-line interface (CLI) to define these elements and return to this topic after they are configured.

This task assumes that you have already launched the IBM System Storage DS6000 Storage Manager.

Perform the following steps to configure the IBM DS6000:

1. Click **Real-time manager** → **Manage hardware** → **Host systems**.
2. Select **Create** from the **Select Action** list. The Create Host System wizard is displayed.
3. Perform the following steps to select a host type:
   a. Select **IBM SAN Volume Controller (SVC)** from the **Host Type** list.
   b. Enter a unique name of up to 16 characters for each port in the **Nickname** field. The value that you enter in this field appears in other panels when you select defined hosts. This is a required field.
   c. Optionally, enter a detailed description of up to 256 characters in the **Description** field.
   d. Click **Next**. The Define host wizard panel is displayed.
4. Perform the following steps in the Define host panel:
   a. Enter the number of ports that you are defining for the SAN Volume Controller node in the **Quantity** field.

      Note: You must add all of the SAN Volume Controller node ports.
   b. Select **FC Switch fabric (P-P)** from the **Attachment Port Type** list.
   c. Click **Add**.
   d. Select **Group ports to share a common set of volumes**.
   e. Click **Next**. The Define host WWPN panel is displayed.
5. Specify a WWPN for each SAN Volume Controller node port that you are configuring. After you have defined all SAN Volume Controller node port WWPNs, click **Next**.

6. Perform the following steps in the Specify storage units panel:
   a. Select all the available storage units that use the ports that you defined in step 5 on page 399.
   b. Click **Add** to move the selected storage units to the **Selected storage units** field.
   c. Click **Next**. The Specify storage units parameters panel is displayed
7. Perform the following steps in the Specify storage units parameters panel:
   a. Select a host attachment identifier from the table.
   b. Click **the following specific storage unit I/O ports** in the **This host attachment can login to** field. The available ports are displayed in the Available storage unit I/O ports table.
   c. Select each port in the Available storage unit I/O ports table.

      Note: The **Type** for each port should be **FcSf**. If the listed type is not FcSf, click **Configure I/O Ports**. The Configure I/O Ports panel is displayed. Click the port that you want to configure and select **Change to FcSf** from the **Select Action** list.
   d. Click **Apply assignment**.
   e. Click **OK**. The Verification panel is displayed.
8. Verify that the attributes and values that are displayed in the table are correct.
9. Click **Finish** if the values that are displayed in the table are correct. Otherwise, click **Back** to return to the previous panels and change the values that are not correct.

Supported firmware levels for the IBM DS6000

The IBM DS6000 must use a firmware level that is supported by the SAN Volume Controller.

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

www.ibm.com/storage/support/2145

Supported models of the IBM DS6000 series

The SAN Volume Controller supports models of the IBM DS6000 series of controllers.

See the following Web site for the latest supported models:

www.ibm.com/storage/support/2145

User interfaces on the IBM DS6000

Ensure that you are familiar with the user interfaces that support the IBM DS6000.

Web server

You can manage, configure, and monitor the IBM DS6000 through the IBM System Storage DS6000 Storage Manager.

CLI

You can also manage, configure, and monitor the IBM DS6000 through the IBM System Storage DS command-line interface.

Concurrent maintenance on the IBM DS6000

Concurrent maintenance is the capability to perform I/O operations to an IBM DS6000 while simultaneously performing maintenance operations on it.

All IBM DS6000 concurrent maintenance procedures are supported.

Target port groups on the IBM DS6000

The IBM DS6000 uses the SCSI Target Port Groups feature to indicate a preferred path for each logical unit (LU).

Sharing an IBM System Storage DS6000 system between a host and the SAN Volume Controller

You can share an IBM System Storage DS6000 system between a host and a SAN Volume Controller cluster.

Quorum disks on IBM System Storage DS6000 systems

The SAN Volume Controller can choose managed disks (MDisks) that are presented by a IBM System Storage DS6000 system as quorum disks.

Configuring IBM System Storage DS8000 systems

This section provides information about configuring the IBM System Storage DS8000 system for attachment to a SAN Volume Controller.

Configuring the IBM DS8000

The IBM DS8000 provides functions that are compatible with the SAN Volume Controller.

After you have defined at least one storage complex, storage unit, and I/O port, you can define the SAN Volume Controller as a host and create host connections. If you have not defined all of these required storage elements, use the IBM System Storage DS8000 Storage Manager or the IBM System Storage DS® command-line interface to define these elements and return to this topic after they are configured.

This task assumes that you have already launched the IBM System Storage DS8000 Storage Manager.

Perform the following steps to configure the IBM DS8000:

1. Click **Real-time manager → Manage hardware → Host systems**.
2. Select **Create** from the **Select Action** list. The Create Host System wizard begins.
3. Perform the following steps to select a host type:
   a. Select **IBM SAN Volume Controller (SVC)** from the **Host Type** list.
   b. Enter a unique name of up to 16 characters for each port in the **Nickname** field. The value that you enter in this field is displayed in other panels when you select defined hosts. This is a required field.

c. Optionally, enter a detailed description of up to 256 characters in the **Description** field.
d. Click **Next**. The Define host wizard panel is displayed.

4. Perform the following steps in the Define host panel:
   a. Enter the number of ports that you are defining for the SAN Volume Controller node in the **Quantity** field.

      Note: You must add all of the SAN Volume Controller node ports.
   b. Select **FC Switch fabric (P-P)** from the **Attachment Port Type** list.
   c. Click **Add**.
   d. Select **Group ports to share a common set of volumes**.
   e. Click **Next**. The Define host WWPN panel is displayed.
5. Specify a WWPN for each SAN Volume Controller node port that you are configuring. When you have defined all SAN Volume Controller node port WWPNs, click **Next**.
6. Perform the following steps in the Select storage images panel:
   a. Select all the available storage units that use the ports that you defined in the previous step.
   b. Click **Add** to move the selected storage units to the **Select storage images** field.
   c. Click **Next**. The Specify storage image parameters panel is displayed
7. Perform the following steps in the Specify storage image parameters panel:
   a. Select a host attachment identifier from the table.
   b. Click **the following specific storage image I/O ports** in the **This host attachment can login to** field. The available ports are displayed in the Available storage unit I/O ports table.
   c. Select each port in the Available storage unit I/O ports table.

      Note: The **Type** for each port should be **FcSf**. If the listed type is not FcSf, click **Configure I/O Ports**. The Configure I/O Ports panel is displayed. Click the port that you want to configure and select **Change to FcSf** from the **Select Action** list.
   d. Click **Apply assignment**.
   e. Click **OK**. The Verification panel is displayed.
8. Verify that the attributes and values that are displayed in the table are correct.
9. Click **Finish** if the values that are displayed in the table are correct. Otherwise, click **Back** to return to the previous panels and change the incorrect values.

Supported firmware levels for the IBM DS8000

The SAN Volume Controller supports the IBM DS8000 series.

See the following Web site for specific firmware levels and the latest supported hardware:www.ibm.com/storage/support/2145

Supported models of the IBM DS8000

The SAN Volume Controller supports models of the IBM DS8000 series of controllers.

See the following Web site for the latest supported models:

www.ibm.com/storage/support/2145

User interfaces on the IBM DS8000

Ensure that you are familiar with the user interfaces that support the IBM DS8000.

Web server

You can manage, configure, and monitor the IBM DS8000 through the IBM System Storage DS8000 Storage Manager.

CLI

You can also manage, configure, and monitor the IBM DS8000 through the IBM System Storage DS command-line interface.

Concurrent maintenance for the IBM DS8000

Concurrent maintenance is the capability to perform I/O operations to an IBM DS8000 while simultaneously performing maintenance operations on it.

All IBM DS8000 concurrent maintenance procedures are supported.

Sharing an IBM System Storage DS8000 system between a host and the SAN Volume Controller

You can share an IBM System Storage DS8000 system between a host and a SAN Volume Controller cluster.

Quorum disks on IBM System Storage DS8000 systems

The SAN Volume Controller can choose managed disks (MDisks) that are presented by a IBM System Storage DS8000 system as quorum disks.

Configuring HDS Lightning series systems

This section provides information about configuring the Hitachi Data Systems (HDS) Lightning series system for attachment to a SAN Volume Controller.

The information in this section also applies to the supported models of the Sun StorEdge series and the HP XP series.

Supported models of the HDS Lightning

The SAN Volume Controller supports models of the HDS Lightning. Certain models of the HDS Lightning are equivalent to Sun StorEdge and HP XP models; therefore, the SAN Volume Controller also supports models of the Sun StorEdge and the HP XP.

See the following Web site for the latest supported models:

www.ibm.com/storage/support/2145

Supported firmware levels for HDS Lightning

The SAN Volume Controller supports the HDS Lightning.

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

www.ibm.com/storage/support/2145

Note: Concurrent upgrade of the controller firmware is supported with the SAN Volume Controller.

Concurrent maintenance on the HDS Lightning

Concurrent maintenance is the capability to perform I/O operations to an HDS Lightning while simultaneously performing maintenance operations on it.

Important: An HDS Field Engineer must perform all maintenance procedures.

User interface on HDS Lightning

Ensure that you are familiar with the user interface application that supports the HDS Lightning system.

Service Processor (SVP)

HDS Lightning has a laptop in the controller frame. The laptop runs the Service Processor (SVP) as the primary configuration user interface. You can use SVP to perform most configuration tasks and to monitor the controller.

HiCommand

The HiCommand is a graphical user interface that allows basic creation of storage and system monitoring. The HiCommand communicates with HDS Lightning through Ethernet.

Sharing the HDS Lightning 99*xx*V between a host and the SAN Volume Controller

There are restrictions for sharing an HDS Lightning 99*xx*V between a host and a SAN Volume Controller cluster.

Sharing ports

The HDS Lightning 99*xx*V can be shared between a host and a SAN Volume Controller cluster under the following conditions:

- The same host cannot be connected to both a SAN Volume Controller cluster and an HDS Lightning at the same time because the Hitachi HiCommand Dynamic Link Manager (HDLM) and the subsystem device driver (SDD) cannot coexist.
- A controller port cannot be shared between a host and a SAN Volume Controller cluster. If a controller port is used by a SAN Volume Controller cluster, it must not be present in a switch zone that allows a host to access the port.
- Logical units (LUs) cannot be shared between a host and a SAN Volume Controller cluster.

Supported Topologies

You can connect the SAN Volume Controller cluster to the HDS Lightning under the following conditions:

- For SAN Volume Controller software version 4.2.1 and later, you can connect a maximum of 16 HDS Lightning ports to the SAN Volume Controller cluster without any special zoning requirements.

- For SAN Volume Controller software version 4.2.0, the following applies:
  - Logical Unit Size Expansion (LUSE) and Virtual LVI/LUN operations cannot be run on a disk that is managed by the SAN Volume Controller cluster. LUNs that are created using LUSE and Virtual LVI/LUN can be mapped to the cluster after they are created.
  - Only disks with open emulation can be mapped to the SAN Volume Controller cluster.
  - S/390® disks cannot be used with the SAN Volume Controller cluster.
  - Only fibre-channel connections can connect the SAN Volume Controller cluster to the HDS Lightning.

Switch zone limitations for HDS Lightning

There are limitations in switch zoning for the SAN Volume Controller and the HDS Lightning systems.

Switch zoning

The HDS Lightning systems present themselves to a SAN Volume Controller cluster as separate controllers for each port zoned to the SAN Volume Controller. For example, if one of these storage systems has 4 ports zoned to the SAN Volume Controller, each port appears as a separate controller rather than one controller with 4 WWPNs. In addition, a given logical unit (LU) must be mapped to the SAN Volume Controller through all controller ports zoned to the SAN Volume Controller using the same logical unit number (LUN).

Quorum disks on HDS Lightning 99*xx*V

HDS Lightning 99*xx*V is not an approved host for quorum disks. Therefore, configurations with only HDS Lightning are not possible.

Advanced functions for HDS Lightning

Some advanced functions of the HDS Lightning are not supported by the SAN Volume Controller.

Advanced copy functions

Advanced copy functions for HDS Lightning (for example, ShadowImage, Remote Copy, and Data Migration) are not supported for disks that are managed by the SAN Volume Controller, because the copy function does not extend to the SAN Volume Controller cache.

Logical Unit Size Expansion

The HDS Lightning 99*xx*V supports Logical Unit Expansion (LUSE). LUSE is *not* a concurrent operation. LUSE is accomplished by concatenating between 2 and 26 existing logical units (LUs) together. Before LUSE can be performed on an LU, the LU must be removed from the managed disk (MDisk) group and unmapped from the SAN Volume Controller.

Attention: LUSE destroys all data that exists on the LU, except on a Windows system.

TrueCopy

TrueCopy operations are functionally similar to Metro Mirror. TrueCopy processing is not supported when the disk controller system is used with the SAN Volume Controller. Even when an HDS Lightning 99*xx*V is shared between a host and a SAN Volume Controller, TrueCopy processing is not supported on the ports that are zoned directly with the host.

Virtual LVI/LUNs

The HDS Lightning 99*xx*V supports Virtual LVI/LUNs. Virtual LVI/LUNs is *not* a concurrent operation. Virtual LVI/LUNs allows you to divide LUNs into several smaller virtual LUNs for use by the HDS Lightning. You must first create existing LUNs into free space and then 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 are created. Therefore, LUNs that are set up using LUSE or Virtual LVI/LUNs can be used by the SAN Volume Controller after they are created.

Write protect

LUs cannot be explicitly set to write-protected. However, some of the advanced features, such as Metro Mirror, can be used to write-protect an LU as part of the function. Metro Mirror must not be used for LUs that are in use by a SAN Volume Controller.

Logical unit configuration for HDS Lightning

Logical unit (LU) configuration for HDS Lightning supports both RAID 1 and RAID 5 arrays.

The HDS Lightning system can have up to 8192 LUs defined; however, only 256 LUs 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 a LUN 0 is not configured, the HDS Lightning system presents a pseudo-LUN at LUN 0. The inquiry data for this pseudo-LUN slightly differs from the inquiry data of normal LUNs. The difference allows the SAN Volume Controller to recognize the pseudo-LUN and exclude it from I/O. The pseudo LUN can accept the report LUNs command.

The HDS Lightning system supports both open-mode attachment and S/390 attachment. The emulation mode is set when the LU is defined. All LUNs that are presented to a SAN Volume Controller must use open emulation. All LUNs with open emulation use a standard 512 byte block size.

The HDS Lightning system can only have certain sized LUs that are defined. These LUs can be expanded by merging 2 - 36 of these LUs using the Logical Unit Size Expansion (LUSE) feature. They can also be made into several, smaller virtual LUNs by using the Virtual LVI/LUN feature.

Special LUs

When an LU is mapped to a host, you have the option to make it a *command LUN*. Command LUNs support in-band configuration commands, but not I/O.

Therefore, you cannot map command LUNs to the SAN Volume Controller.

Logical unit creation and deletion on HDS Lightning

The SAN Volume Controller supports Logical Unit Size Expansion (LUSE) with certain restrictions.

The following restrictions apply:

- Before LUSE can be performed on an LU, the LU must be unmounted from a host and have no available paths. The LUSE function destroys all data that exists on the LU, except for LUs on a Windows operating 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 want to use image mode to import the data, do not use LUSE on the disk before you import the data.

Configuring settings for HDS Lightning

The Lightning configuration interface provides functions for configuration.

These options and settings can have the following scope:

- Subsystem
- Port
- Logical unit (LU)

Global settings for HDS Lightning

Global settings apply across an HDS Lightning disk controller system.

Table 41 lists the global settings for HDS Lightning.

Table 41. HDS 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
Read configuration data mode	Selected	Selected
PS off timer	Not selected	Not selected

Controller settings for HDS Lightning

Controller settings are settings that apply across the entire HDS Lightning controller.

Table 42 lists the HDS Lightning controller settings that are supported by the SAN Volume Controller.

Table 42. HDS Lightning controller settings that are supported by the SAN Volume Controller

Option	HDS Lightning default setting	SAN Volume Controller required setting
PCB mode	Standard	Standard

Port settings for HDS Lightning

Port settings are configurable at the port level.

There are no available options 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 43 lists the HDS Lightning port settings that are supported by the SAN Volume Controller.

Table 43. HDS Lightning port settings supported by the SAN Volume Controller

Option	HDS 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

Logical unit settings for HDS Lightning

Logical unit (LU) settings apply to individual LUs that are configured in the HDS Lightning controller.

HDS Lightning LUs must be configured as described in Table 44 if the LUN is associated with ports in a switch zone that is accessible to the SAN Volume Controller.

Table 44. HDS Lightning LU settings for the SAN Volume Controller

Option	HDS Lightning default setting	SAN Volume Controller required setting
Command device	Off	Off
Command security	Off	Off

Note: These settings only apply to LUs that are accessible by the SAN Volume Controller.

Configuring HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems

You can attach Hitachi Data Systems (HDS) Thunder, HDS TagmaStore Adaptable Modular Storage (AMS), and HDS TagmaStore Workgroup Modular Storage (WMS) systems to a SAN Volume Controller cluster.

Note: In Japan, the HDS Thunder 9200 is referred to as the HDS SANrise 1200. Therefore, the information in this section that refers to the HDS Thunder 9200 also applies to the HDS SANrise 1200.

Supported HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS models

You can attach certain HDS Thunder, HDS TagmaStore Adaptable Modular Storage (AMS), and HDS TagmaStore Workgroup Modular Storage (WMS) models to SAN Volume Controller clusters.

See the following Web site for the latest supported models:

www.ibm.com/storage/support/2145

Supported firmware levels for HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS

The SAN Volume Controller supports certain HDS Thunder, HDS TagmaStore Adaptable Modular Storage (AMS), and HDS TagmaStore Workgroup Modular Storage (WMS) models.

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

www.ibm.com/storage/support/2145

Concurrent maintenance on HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems

Concurrent maintenance is the capability to perform I/O operations to a system while simultaneously performing maintenance operations on it.

Important: An HDS Field Engineer must perform all maintenance operations.

The SAN Volume Controller supports concurrent hardware maintenance and firmware upgrade operations on these systems.

User interface on HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems

Ensure that you are familiar with the user interface applications that support the Hitachi Data Systems (HDS) Thunder, HDS TagmaStore Adaptable Modular Storage (AMS), and HDS TagmaStore Workgroup Modular Storage (WMS) systems.

In-band configuration

Disable the system command LUN when you use the user interface applications.

Storage Navigator Modular GUI

The Storage Navigator Modular (SNM) is the primary user interface application for configuring HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems. Use SNM to upgrade firmware, change settings, and to create and monitor storage.

SNM supports an Ethernet connection to the system. An out-of-band command-line interface is available with SNM that supports the majority of the functions that are provided in SNM.

HiCommand

HiCommand is another configuration user interface that is available for the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems. You must have access to SNM to use HiCommand to configure settings. HiCommand only allows basic creation of storage and provides some monitoring features.

HiCommand uses Ethernet to connect to the system.

Web Server

A Web server runs on each of the controllers on the system. During normal operation, the user interface only allows basic monitoring of the system and displays an error log. If you put a controller into diagnostic mode by pressing the reset button on the controller, the user interface allows firmware upgrades and system configuration resets.

Sharing the HDS Thunder, HDS TagmaStore AMS, or HDS TagmaStore WMS between a host and the SAN Volume Controller

You can share the HDS Thunder, HDS TagmaStore Adaptable Modular Storage (AMS), and HDS TagmaStore Workgroup Modular Storage (WMS) systems between a host and a SAN Volume Controller cluster, with certain restrictions.

The following restrictions apply:

- The same host cannot be connected to both a SAN Volume Controller cluster and an HDS Thunder, HDS TagmaStore AMS, or HDS TagmaStore WMS at the same time because Hitachi Dynamic Link Manager (HDLM) and the subsystem device driver (SDD) cannot coexist.
- For the HDS Thunder 9200, a target port cannot be shared between a host and a SAN Volume Controller cluster. If a target port is used by a SAN Volume Controller cluster, it cannot be present in a switch zone that allows a host to access the port.
- Logical units (LUs) cannot be shared between a host and a SAN Volume Controller cluster. The Thunder 9200 must be set into M-TID M-LUN mode and Mapping Mode must be enabled on Thunder 95*xx*. No LU can have a LUN number that is associated with a port that is zoned for host use while also having a LUN number that is associated with a port that is zoned for a SAN Volume Controller cluster.

Switch zoning limitations for HDS Thunder, HDS TagmaStore AMS, or HDS TagmaStore WMS

There are limitations in switch zoning for the SAN Volume Controller and the HDS Thunder, HDS TagmaStore AMS, or HDS TagmaStore WMS systems.

Switch zoning

The HDS Thunder, HDS TagmaStore AMS, or HDS TagmaStore WMS systems present themselves to a SAN Volume Controller cluster as separate controllers for each port zoned to the SAN Volume Controller. For example, if one of these storage systems has 4 ports zoned to the SAN Volume Controller, each port appears as a separate controller rather than one controller with 4 WWPNs. In addition, a given logical unit (LU) must be mapped to the SAN Volume Controller

through all controller ports zoned to the SAN Volume Controller using the same logical unit number (LUN).

Supported topologies

You can connect the SAN Volume Controller cluster to the HDS Thunder under the following conditions:

- For SAN Volume Controller software version 4.2.1 and later, you can connect a maximum of 16 HDS Thunder ports to the SAN Volume Controller cluster without any special zoning requirements.
- For SAN Volume Controller software version 4.2.0, the following applies:
  - Logical Unit Size Expansion (LUSE) and Virtual LVI/LUN operations cannot be run on a disk that is managed by the SAN Volume Controller cluster. LUNs that are created using LUSE and Virtual LVI/LUN can be mapped to the cluster after they are created.
  - Only disks with open emulation can be mapped to the SAN Volume Controller cluster.
  - S/390 disks cannot be used with the SAN Volume Controller cluster.
  - Only fibre-channel connections can connect the SAN Volume Controller cluster to the HDS Thunder.

Quorum disks on HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems

When a SAN Volume Controller cluster initializes, the cluster can choose managed disks (MDisks) that are presented by HDS Thunder, HDS TagmaStore Adaptable Modular Storage (AMS), and HDS TagmaStore Workgroup Modular Storage (WMS) systems as quorum disks.

You can use the set quorum disk CLI command or the SAN Volume Controller Console to select quorum disks.

Host type for HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS

When the HDS Thunder, HDS TagmaStore Adaptable Modular Storage (AMS) and HDS TagmaStore Workgroup Modular Storage (WMS), systems are attached to a SAN Volume Controller cluster, set the host mode attribute to Windows 2003 for each host group.

Advanced functions for HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS

Some advanced functions of the HDS Thunder, HDS TagmaStore Adaptable Modular Storage (AMS) and HDS TagmaStore Workgroup Modular Storage (WMS), systems are not supported by the SAN Volume Controller clusters.

Advanced copy functions

Advanced copy functions for HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems are not supported for disks that are managed by the SAN Volume Controller clusters because the copy function does not extend to the SAN Volume Controller cache. For example, ShadowImage, TrueCopy, and HiCopy are not supported.

LUN Security

LUN Security enables LUN masking by the worldwide node name (WWNN) of the initiator port. This function is not supported for logical units (LUs) that are used by SAN Volume Controller clusters.

Partitioning

Partitioning splits a RAID array into up to 128 smaller LUs, each of which serves as an independent disk like entity. The SAN Volume Controller cluster and HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems support the partitioning function.

Dynamic array expansion

The HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems allow the last LU that is defined in a RAID group to be expanded. This function is not supported when these storage systems are attached to a SAN Volume Controller cluster. Do *not* perform dynamic array expansion on LUs that are in use by a SAN Volume Controller cluster.

Note: Use in this context means that the LU has a 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 and virtual fibre-channel ports

The HDS Thunder 95*xx*V, HDS TagmaStore AMS, and HDS TagmaStore WMS systems support host storage domains (HSD) and virtual fibre-channel ports. Each fibre-channel port can support multiple HSDs. Each host in a given HSD is presented with a virtual target port and a unique set of LUNs.

The Thunder 9200 does not support HSD and virtual fibre-channel the ports.

Logical unit creation and deletion on HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems

The HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems Storage Navigator Modular Graphical User Interface (GUI) enables you to create and delete LUNs. You must avoid certain creation and deletion scenarios to prevent data corruption.

Creation and deletion scenarios

For example, the Storage Navigator Modular GUI enables you to create LUN A, delete LUN A, and then create LUN B with the same unique ID as LUN A. If a SAN Volume Controller cluster is attached, data corruption can occur because the cluster might not realize that LUN B is different than LUN A.

Attention: Before you use the Storage Navigator Modular GUI to delete a LUN, remove the LUN from the managed disk group that contains it.

Adding LUNs dynamically

To prevent the existing LUNs from rejecting I/O operations during the dynamic addition of LUNs, perform the following procedure to add LUNs:

1. Create the new LUNs using the Storage Navigator Modular GUI.
2. Quiesce all I/O operations.
3. Perform either an offline format or an online format of all new LUNs on the controller using the Storage Navigator Modular GUI. Wait for the format to complete.
4. Go into the LUN mapping function of the Storage Navigator Modular GUI. Add mapping for the new LUN to all of the controller ports that are available to the SAN Volume Controller cluster on the fabric.
5. Restart the controller. (Model 9200 only)
6. After the controller has restarted, restart I/O operations.

LUN mapping considerations

If LUN mapping is used as described in the LUN mapping topic, you must restart the controller to pick up the new LUN mapping configuration. For each managed disk group that contains an MDisk that is supported by an LU on the system, all virtual disks in those managed disk groups go offline.

Configuring settings for HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems

The Storage Navigator Modular GUI configuration interface provides functions for configuration.

These options and settings can have the following scope:

- System
- Port
- Logical unit

Global settings for the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems

Global settings apply across HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems.

Table 45 lists the global settings for these disk systems.

Table 45. HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems global settings supported by the SAN Volume Controller

Option	Default setting	SAN Volume Controller required setting
Start attribute	Dual active mode	Dual active mode
SCSI ID/Port takeover mode	Not applicable	Not applicable
Default controller	Not applicable	Not applicable
Data-share mode	Used	Used
Serial number		Same as the system default setting
Delay planned shutdown	0	0
Drive detach mode	False	False
Multipath controller (Thunder 9200 only)	False	False
PROCOM mode	False	False

Table 45. HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems global settings supported by the SAN Volume Controller (continued)

Option	Default setting	SAN Volume Controller required setting
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
Multiple stream mode (write) (Thunder 95*xx*V only)	False	False
Multiple stream mode (read) (Thunder 95*xx*V only)	False	False
RAID 3 mode (Thunder 9200 only)	False	False
Target ID (9200 only) Mapping mode on 95*xx*	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 95*xx*V)
Operation if processor failure occurs	Reset the fault	Reset the fault
Command queuing	True	True
ANSI Version	Not applicable	Not applicable
Vendor ID	HITACHI	HITACHI
Product ID (Thunder 9200)	DF500F	DF500F
Product ID (Thunder 95*xx*V)	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)	Not applicable	Not applicable
Save data pointer response (Thunder 9200 only)	Not applicable	Not applicable
Controller identifier	False	False

Table 45. HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems global settings supported by the SAN Volume Controller (continued)

Option	Default setting	SAN Volume Controller required setting
RS232C error information outflow mode	Off	Any
Execute write and verify mode	True	True

Controller settings for HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems

Controller settings apply across the entire HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems. Options are not available within the scope of a single controller.

Port settings for the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems

Port settings are configurable at the port level.

The settings listed in Table 46 apply to disk controllers that are in a switch zone that contains SAN Volume Controller nodes. If the system is shared between a SAN Volume Controller cluster and another host, you can configure with different settings than shown if both of the following conditions are true:

- 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 cluster

There are no available options with the scope of a single controller.

Table 46. HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS system port settings supported by the SAN Volume Controller

Option	Default setting	SAN Volume Controller required setting
Host connection mode 1	Standard	Standard
V*x*VM DMP mode (HDS Thunder 9200 only)	False	False
HP connection mode	False	False
Report inquiry page 83H (HDS Thunder 9200 only)	False	True
UA (06/2A00) suppress mode	False	True
HISUP mode	False	False
CCHS mode	False	False
Standard inquiry data expand (HDS Thunder 9200 only)	False	False
Host connection mode 2	False	False
Product ID DF400 mode	False	False
HBA WWN report mode (HDS Thunder 9200 only)	False	False

Table 46. HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS system port settings supported by the SAN Volume Controller (continued)

Option	Default setting	SAN Volume Controller required setting
NACA mode	False	False
SUN cluster connection mode	False	False
Persistent RSV cluster mode	False	False
ftServer connection mode 1 (HDS 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
Reset logout of third party process mode	False	False
Read Frame minimum 128 byte mode (HDS Thunder 950xxV only)	False	False
Topology	Point-to-point	Point-to-point

Logical unit settings for the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems

Logical unit (LU) settings apply to individual LUs that are configured in the HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems.

You must configure the systems LUs as described in Table 47 if the logical unit number (LUN) is associated with ports in a switch zone that is accessible to the SAN Volume Controller cluster.

Table 47. HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems 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 that are accessible by the SAN Volume Controller cluster.

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. Because 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 cluster because this can 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. If the LUN is managed by a SAN Volume Controller cluster, this scenario can cause data corruption because the cluster might not recognize that LUN B is different than LUN A.

Mapping and virtualization settings for HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems

The HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems support different modes of operation. These modes affect LUN mapping or masking and virtualization.

The 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 95*xx*. You must restart the controllers for changes to LUN mapping to take effect.

Attention: The HDS Thunder, HDS TagmaStore AMS, and HDS TagmaStore WMS systems do not provide an interface that enables a SAN Volume Controller cluster to detect and ensure that the mapping or masking and virtualization options are set properly. Therefore, you must ensure that these options are set as described in this topic.

S-TID M-LUN modes

In S-TID M-LUN mode all LUs are accessible through all ports on the system with the same LUN number on each port. You can use this mode in environments where the system is not being shared between a host and a SAN Volume Controller cluster.

M-TID M-LUN modes

If a system is shared between a host and a SAN Volume Controller cluster, you must use M-TID M-LUN mode. Configure the system so that each LU that is exported to the SAN Volume Controller cluster can be identified by a unique LUN. The LUN must be the same on all ports through which the LU can be accessed.

Example

A SAN Volume Controller cluster can access controller ports x and y. The cluster also sees an LU on port x that has LUN number p. In this situation the following conditions must be met:

- The cluster 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 cannot appear as any other LUN number on port y.

- The LU must not be mapped to any system port that is zoned for use directly by a host in a configuration where the system is shared between a host and a cluster.

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 can be LUN 0 on port 1, LUN 3 on port 2, and not visible at all on ports 3 and 4.

Important: 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 can be LUN 1 and LUN 2 on controller port 1.

Important: The SAN Volume Controller does not support this.

Configuring HDS TagmaStore USP and NSC systems

This section provides information about configuring the Hitachi Data Systems (HDS) TagmaStore Universal Storage Platform (USP) and Network Storage Controller (NSC) systems for attachment to a SAN Volume Controller. Models of the HDS USP and NSC are equivalent to HP and Sun models; therefore, the SAN Volume Controller also supports models of the HP StorageWorks XP series and the Sun StorEdge series.

The information in this section also applies to the supported models of the HP XP and the Sun StorEdge series.

Supported models of the HDS USP and NSC

The SAN Volume Controller supports models of the Hitachi Data Systems (HDS) Universal Storage Platform (USP) and Network Storage Controller (NSC) series. Models of the HDS USP and NSC are equivalent to HP and Sun models; therefore, the SAN Volume Controller also supports models of the Sun StorEdge and the HP XP series.

See the following Web site for the latest supported models:

www.ibm.com/storage/support/2145

Supported firmware levels for HDS USP and NSC

The SAN Volume Controller supports the HDS USP and NSC series of controllers.

See the following Web site for specific firmware levels and the latest supported hardware:www.ibm.com/storage/support/2145

User interface on the HDS USP and NSC

Ensure that you are familiar with the user interface application that supports the HDS USP and NSC. The HDS USP and NSC is configured, managed, and monitored by a Service Processor (SVP). The SVP is a server that is connected to the HDS USP or NSC through a private local area network (LAN).

Web server

The HDS USP and NSC use the Storage Navigator as the main configuration GUI. The Storage Navigator GUI runs on the SVP and is accessed through a Web browser.

Logical units and target ports on the HDS USP and NSC

Logical units (LUs) that are exported by the HDS USP and NSC report identification descriptors in the vital product data (VPD). The SAN Volume Controller uses the LUN associated binary type-3 IEEE Registered Extended descriptor to identify the LU.

An LU path must be defined before an LU can be accessed by a host. The LU path relates a host group to a target port and to a set of LUs. Host initiator ports are added to the host group by worldwide port name (WWPN).

The HDS USP and NSC do not use LU groups so all LUs are independent. The LU access model is active-active and does not use preferred access ports. Each LU can be accessed from any target port that is mapped to the LU. Each target port has a unique WWPN and worldwide node name (WWNN). The WWPN matches the WWNN on each port.

Note: You must wait until the LU is formatted before presenting it to the SAN Volume Controller.

Special LUs

The HDS USP and NSC can use any logical device (LDEV) as a Command Device. Command Devices are the target for HDS USP or NSC copy service functions. Therefore, do not export Command Devices to a SAN Volume Controller.

Switch zoning limitations for the HDS USP and NSC

There are limitations in switch zoning for the SAN Volume Controller and the HDS USP or NSC.

The SAN Volume Controller can be connected to the HDS USP or NSC with the following restrictions:

- If an LU is mapped to a SAN Volume Controller port as LUN *x*, the LU must appear as LUN *x* for all mappings to target ports.
- Only fibre-channel connections can be used to connect a SAN Volume Controller to the HDS USP or NSC system.
- Because the SAN Volume Controller limits the number of worldwide node names (WWNNs) for each storage system and the HDS USP and NSC present a separate WWNN for each port, the number of target ports that the SAN Volume Controller can resolve as one storage system is limited. Perform the following steps to provide connections to more target ports:
  1. Divide the set of target ports into groups of 2 to 16.
  2. Assign a discrete set of LUs to each group.

  The SAN Volume Controller can then view each group of target ports and the associated LUs as separate HDS USP or NSC systems. You can repeat this process to use all target ports.

Note: The HDS USP and NSC systems present themselves to a SAN Volume Controller cluster as separate controllers for each port zoned to the SAN Volume Controller. For example, if one of these storage systems has 4 ports zoned to the SAN Volume Controller, each port appears as a separate controller rather than one controller with 4 WWPNs. In addition, a given logical unit (LU) must be mapped to the SAN Volume Controller through all controller ports zoned to the SAN Volume Controller using the same logical unit number (LUN).

Controller splitting

You can split the HDS USP or NSC between other hosts and the SAN Volume Controller under the following conditions:

- A host cannot be simultaneously connected to both an HDS USP or NSC and a SAN Volume Controller.
- Port security must be enabled for target ports that are shared.
- An LU that is mapped to a SAN Volume Controller cannot be simultaneously mapped to another host.

Concurrent maintenance on the HDS USP and NSC

Concurrent maintenance is the capability to perform I/O operations to an HDS USP or NSC while simultaneously performing maintenance operations on it. Concurrent firmware upgrades are supported with the SAN Volume Controller.

Important: An HDS Field Engineer must perform all maintenance procedures.

Quorum disks on HDS USP and NSC

To host quorum disks on HDS USP and NSC storage systems, you must be aware of the system requirements for establishing quorum disk for these storage systems.

Note: Sun StorEdge systems are not supported to host SAN Volume Controller quorum disks.

The SAN Volume Controller cluster uses a quorum disk to store important cluster configuration data and to break a tie in the event of a SAN failure. The cluster automatically chooses three managed disks (MDisks) as quorum disk candidates. Each disk is assigned an index number: either 0, 1, or 2. Although a cluster can be configured to use up to three quorum disks, only one quorum disk is elected to resolve a tie-break situation. The purpose of the other quorum disks is to provide redundancy if a quorum disk fails before the cluster is partitioned.

Requirements for HDS TagmaStore USP, HP XP10000/12000, and NSC55

To host any of the three quorum disks on these HDS TagmaStore USP, HP XP10000/12000, or NSC55 storage systems, ensure that each of the following conditions have been met:

- Firmware version Main 50-09-72 00/00 or later is running. Contact HDS or HP support for details on installing and configuring the correct firmware version.
- **System Option 562** is enabled. Contact HDS or HP support for details on System Option 562.
- All SAN Volume Controller ports are configured in a single HDS or HP host group.

Requirements for HDS TagmaStore USPv, USP-VM, and HP XP20000/24000

To host any of the three quorum disks on these HDS TagmaStore USPv, USP-VM, or HP XP20000/24000 systems, ensure that each of the following requirements have been met:

- Firmware version Main 60-04-01-00/02 or later is running. Contact HDS or HP support for details on installing and configuring the correct firmware version.
- **Host Option 39** is enabled. Contact HDS or HP support for details on Host Option 39.

  Note: This must be applied to the HDS or HP host group that is used for SAN Volume Controller.
- All SAN Volume Controller ports are configured in a single HDS or HP host group.

After you have verified these requirements for the appropriate storage system, complete the following steps on the SAN Volume Controller command-line interface to set the quorum disks:

1. Issue the svctask chcontroller command:

   ```
   svctask chcontroller -allowquorum yes controller_id or controller_name
   ```

   where *controller_id or controller_name* is the controller that corresponds to the relevant HDS or HP storage system.
2. Repeat step 1 for each controller that is part of the relevant HDS or HP storage system.
3. Issue the svctask setquroum command:

   ```
   svctask setquroum -quorum [0|1|2] mdisk_id or mdisk_name
   ```

   where *mdisk_id or mdisk_name* is the relevant MDisk on the HDS or HP system.

Attention: Failure to meet these conditions or to follow these steps can result in data corruption.

The Support for SAN Volume Controller (2145) Web site provides current information about quorum support:

www.ibm.com/storage/support/2145

Host type for HDS USP and NSC subsystems

When the HDS USP and NSC systems are attached to a SAN Volume Controller cluster, set the host mode attribute to Windows for each host group.

Advanced functions for HDS USP and NSC

Some advanced functions of the HDS USP and NSC are not supported by the SAN Volume Controller.

Advanced system functions

The following advanced system functions for HDS USP and NSC are not supported for disks that are managed by the SAN Volume Controller:

- TrueCopy
- ShadowImage

- Extended Copy Manager
- Extended Remote Copy
- NanoCopy
- Data migration
- RapidXchange
- Multiplatform Backup Restore
- Priority Access
- HARBOR File-Level Backup/Restore
- HARBOR File Transfer
- FlashAccess

Advanced SAN Volume Controller functions

All advanced SAN Volume Controller functions are supported on logical unit (LU) that are exported by the HDS USP or NSC system.

LU Expansion

The HDS USP and NSC support Logical Unit Expansion (LUSE). LUSE is *not* a concurrent operation. LUSE allows you to create a single LU by concatenating logical devices (LDEVs). Before LUSE can be performed, the LDEVs must be unmounted from hosts and paths must be removed.

Attention:

1. LUSE destroys all data that exists on the LDEV.
2. Do not perform LUSE on any LDEV that is used to export an LU to a SAN Volume Controller.

If data exists on an LDEV and you want to use image mode migration to import the data to a SAN Volume Controller, do not perform LUSE on the disk before you import the data.

LUs that are created using LUSE can be exported to a SAN Volume Controller.

Virtual LVI/LUNs

The HDS USP and NSC support Virtual LVI/LUNs (VLL). VLL is *not* a concurrent operation. VLL allows you to create several LUs from a single LDEV. You can only create new LUs from free space on the LDEV.

Attention: Do not perform VLL on disks that are managed by the SAN Volume Controller.

LUs that are created using VLL can be exported to a SAN Volume Controller.

Configuring HP StorageWorks MA and EMA systems

This section provides information about configuring HP StorageWorks Modular Array (MA) and Enterprise Modular Array (EMA) systems for attachment to a SAN Volume Controller.

Both the HP MA and EMA use an HSG80 controller.

HP MA and EMA 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 that are 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 all data on 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.
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 HP MA and EMA systems

The HP MA and EMA systems provide functions that are compatible with the SAN Volume Controller.

This task assumes that the system is not in use.

Note: When you configure a SAN Volume Controller cluster to work with an HP MA or EMA, you must not exceed the limit of 96 process logins.

Perform the following procedure to enable support of an HP, MA, or EMA system.

1. Verify that the front panel of the SAN Volume Controller is clear of errors.
2. Ensure that the HP StorageWorks Operator Control Panel (OCP) on each system is clear of errors. The Operator Control Panel consists of seven green LEDs 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. Ensure that the system firmware is at a supported level. See the following Web site for the latest firmware support:

      www.ibm.com/storage/support/2145.
   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 system.
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 do not specify Host Based Logging.
8. Issue the **SHOW CONNECTIONS FULL** command to verify that you have enough spare entries for all combinations of SAN Volume Controller ports and HP MA or EMA ports.
9. Connect up to four fibre-channel cables between the fibre-channel switches and the HP MA or EMA system.
10. Ensure that the fibre-channel switches are zoned so that the SAN Volume Controller and the HP MA or EMA system are in a zone.
11. Issue the **SHOW THIS** command and **SHOW OTHER** command to verify that each connected port is running. The following is an example of the output that is displayed: `PORT_1_TOPOLOGY=FABRIC`.

12. Issue the **SHOW CONNECTIONS FULL** command to verify that the new connections have been created for each SAN Volume Controller port and HP MA or EMA port combination.
13. Verify that `No rejected hosts` is displayed at the end of the SHOW CONNECTIONS output.
14. Perform the following steps from the SAN Volume Controller command-line interface (CLI):
    a. Issue the **svctask detectmdisk** CLI command to discover the controller.
    b. Issue the **svcinfo lscontroller** CLI command to verify that the two HSG80 serial numbers appear under the `ctrl s/n`.
    c. Issue the **svcinfo lsmdisk** CLI command to verify that the additional MDisks that correspond to the UNITS shown in the HP MA or EMA system.

You can now use the SAN Volume Controller CLI commands to create an MDisk group. You can also create and map VDisks from these MDisk groups. Check the front panel of the SAN Volume Controller to ensure that there are no errors. After the host has reloaded the fibre-channel driver, you can perform I/O to the VDisks. See the *IBM System Storage SAN Volume Controller Host Attachment Guide* for detailed information.

Partitioning LUNs on HP MA and EMA systems

For SAN Volume Controller software version 4.2.1 and later, you cannot partition HSG80 LUNs. To check if any HSG80 LUNs are partitioned, use the SHOW UNITS command in the HSG80 CLI. Partition is displayed in the Used By column for the LUNs that are partitioned.

Supported models of HP MA and EMA systems

The SAN Volume Controller supports models of the HP MA and EMA systems.

See the following Web site for the latest supported models:

www.ibm.com/storage/support/2145

Attention: The SAN Volume Controller only supports configurations in which the HSG80 cache is enabled in writeback mode. Running with only a single controller results in a single point of data loss.

Supported firmware levels for HP MA and EMA systems

The HP MA and EMA systems must use a firmware level that is supported by the SAN Volume Controller.

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

www.ibm.com/storage/support/2145

Note: Concurrent upgrade of the system firmware is not supported with the SAN Volume Controller.

Concurrent maintenance on HP MA and EMA systems

Concurrent maintenance is the capability to perform I/O operations to an HP MA or EMA system while simultaneously performing maintenance operations on it.

Note: HP MA and EMA maintenance documentation uses the phrase *rolling upgrade* in place of *concurrent maintenance*. See this documentation because in some instances you must reduce the level of I/O before you can perform the maintenance procedure.

The HP MA and EMA systems allow concurrent replacement of the following components:

- Drive
- EMU
- Blower
- Dual power supply (One unit can be removed and replaced. The fan speed increases when only one power supply unit is present.)

The controller component is hot-pluggable, but concurrent maintenance of SAN Volume Controller I/O is not supported.

The HP MA and EMA systems do not allow 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 HP MA and EMA systems

The Command Console configuration and service utility is the configuration interface for the HP MA and EMA systems.

The configuration and service utility can connect to the system in the following ways:

- RS232 interface
- In-band over fibre channel
- Over TCP/IP to a proxy agent, which then communicates with the system in-band over fibre channel.

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 can therefore also access LUs that are visible to SAN Volume Controller nodes 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 able to recognize only the CCL.

Sharing the HP MA or EMA between a host and a SAN Volume Controller

There are restrictions for sharing HP MA and EMA subsystems between a host and a SAN Volume Controller cluster.

An HP MA or EMA can be shared between a host and a SAN Volume Controller cluster under the following conditions:

- A host cannot be connected to both a SAN Volume Controller cluster and an HP MA or EMA subsystem at the same time.

- Target ports cannot be shared between a host and a SAN Volume Controller cluster. Specifically, if an HSG80 port is in use by a SAN Volume Controller cluster, it cannot be present in a switch zone that enables a host to access the port.
- LUs and RAID arrays cannot be shared between a host and a SAN Volume Controller cluster.

Switch zoning limitations for HP MA and EMA systems

There are limitations in switch zoning for the SAN Volume Controller and the HP MA and EMA systems.

Attention: The HP MA and EMA systems are supported with a single HSG80 controller or dual HSG80 controllers. Because the SAN Volume Controller only supports configurations in which the HSG80 cache is enabled in write-back mode, running with a single HSG80 controller results in a single point of data loss.

Switch zoning

For SAN Volume Controller clusters that have installed software version 1.1.1, a single fibre-channel port that is attached to the system can be present in a switch zone that contains SAN Volume Controller fibre-channel ports, whether the HP MA or EMA system uses one or two HSG80 controllers. This guarantees that the nodes in the cluster can access at most one port on the HSG80 controller.

For SAN Volume Controller clusters that have software version 1.2.0 or later installed, switches can be zoned so that HSG80 controller ports are in the switch zone that contains all of the ports for each SAN Volume Controller node.

Connecting to the SAN

Multiple ports from an HSG80 controller must be physically connected to the fibre-channel SAN to enable servicing of the HP MA or EMA system. However, switch zoning must be used as described in this topic.

Note: If the HP Command Console is not able to access a fibre-channel port on each of the HSG80 controllers in a two-controller system, there is a risk of an undetected single point of failure.

Quorum disks on HP MA and EMA systems

Managed disks (MDisks) that are presented by the HP MA or EMA are chosen by the SAN Volume Controller as quorum disks.

The SAN Volume Controller uses a logical unit (LU) that is presented by an HSG80 controller as a quorum disk. The quorum disk is used even if the connection is by a single port, although this is not recommended. If you are connecting the HP MA or EMA system with a single fibre-channel port, ensure that you have another system on which to put your quorum disk. You can use the **svctask setquorum** command-line interface (CLI) command to move quorum disks to another system.

SAN Volume Controller clusters that are attached only to the HSG80 controllers are supported.

Advanced functions for HP MA and EMA

Some advanced functions of the HP MA and EMA are not supported by the SAN Volume Controller.

Advanced copy functions

Advanced copy functions for HP MA and EMA systems (for example, SnapShot and RemoteCopy) are not supported for disks that are managed by the SAN Volume Controller because the copy function does not extend to the SAN Volume Controller cache.

Partitioning

HP MA and EMA support partitioning. A partition is a logical division of a container that is represented to the host as a logical unit (LU). A container can be a RAID array or a JBOD (just a bunch of disks). All container types are candidates for partitions. Any nontransportable disk or storage set can be divided into a maximum of eight partitions.

The following restrictions apply to partitioning:

- Partitioned containers are fully supported if the HSG80 controller is connected to the SAN by a single port.
- Partitioned containers are not configured by the SAN Volume Controller if the HSG80 controller is connected to the SAN by multiple ports.
- Partitioned containers are removed from the configuration if a single port connection becomes a multiport connection.
- Partitioned containers are configured if a multiport connection becomes a single port connection.

You must partition containers such that no spare capacity exists because there is no way to detect unused partitions. With a multiport connection, subsequent attempts to use this capacity removes all partitions on the container from the configuration.

Dynamic array expansion (LU expansion)

HP MA and EMA systems do not provide dynamic array expansion.

Write protection of LUNs

Write protection of LUNs is not supported for use with the SAN Volume Controller.

SAN Volume Controller advanced functions

Virtual disks (VDisks) that are created from managed disks (MDisks) that are presented by an HSG80 controller can be used in SAN Volume Controller FlashCopy mappings, SAN Volume Controller Metro Mirror relationships, and SAN Volume Controller Global Mirror relationships.

LU creation and deletion on the HP MA and EMA

Ensure you are familiar with the HSG80 controller container types for logical unit (LU) configuration.

Table 48 lists the valid container types.

Table 48. HSG80 controller container types for LU configuration

Container	Number of Members	Maximum Size
JBOD - non transportable **Attention:** A JBOD provides no redundancy at the physical disk drive level. A single disk failure can result in the loss of an entire managed disk group and its associated virtual disks.	1	disk size minus metadata
Mirrorset	2 - 6	smallest member
RAIDset	3 - 14	1.024 terabytes
Stripeset	2 - 24	1.024 terabytes
Striped Mirrorset	2 - 48	1.024 terabytes

Note: LUs can be created and deleted on an HSG80 controller while I/O operations are performed to other LUs. You do not need to restart the HP MA or EMA subsystem.

Configuring settings for the HP MA and EMA

The HP StorageWorks configuration interface provides configuration settings and options that are supported with the SAN Volume Controller.

The settings and options can have a scope of the following:

- Subsystem (global)
- Controller
- Port
- Logical unit
- Connection

Global settings for HP MA and EMA systems

Global settings apply across HP MA and EMA systems.

The following table lists the global settings for HP MA and EMA systems:

Table 49. HP MA and EMA global settings supported by the SAN Volume Controller

Option	HSG80 controller default setting	SAN Volume Controller required setting
DRIVE_ERROR_THRESHOLD	800	Default
FAILEDSET	Not defined	n/a

Controller settings for HP MA and EMA

Controller settings apply across one HSG80 controller.

Table 50 on page 431 describes the options that can be set by HSG80 controller command-line interface (CLI) commands for each HSG80 controller.

Table 50. HSG80 controller settings that are supported by the SAN Volume Controller

Option	HSG80 controller 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

Port settings for HP MA and EMA systems

Port settings are configurable at the port level.

Restriction: Only one port per HSG80 pair can be used with the SAN Volume Controller.

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 51 lists the HSG80 controller port settings that the SAN Volume Controller supports:

Table 51. HSG80 controller port settings supported by the SAN Volume Controller

Option	HSG80 default setting	SAN Volume Controller required setting
PORT_1/2-AL-PA	71 or 72	Not applicable
PORT_1/2_TOPOLOGY	Not defined	FABRIC

Note: The HP MA and EMA systems support LUN masking that is configured with the **SET *unit number* ENABLE_ACCESS_PATH** command. When used with a SAN Volume Controller, the access path must be set to all (″SET *unit number* ENABLE_ACCESS_PATH=ALL″) and all LUN masking must be handled exclusively by the SAN Volume Controller. You can use the **SHOW CONNECTIONS FULL** command to check access rights.

LU settings for HP MA and EMA systems

Logical unit (LU) settings are configurable at the LU level.

Table 52 describes the options that must be set for each LU that is accessed by the SAN Volume Controller. LUs that are accessed by hosts can be configured differently.

Table 52. HSG80 controller LU settings supported by the SAN Volume Controller

Option	HSG80 controller default setting	SAN Volume Controller required setting
TRANSFER_RATE_REQUESTED	20MHZ	Not applicable
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	Not applicable
MAX_READ_CACHE_SIZE	32	Not applicable
MAX_WRITE_CACHE_SIZE	32	64 or higher
MAX_CACHED_TRANSFER_SIZE	32	Not applicable
PREFERRED_PATH/ NOPREFERRED_PATH	NOPREFERRED_PATH is set	Not applicable
READ_CACHE/ NOREAD_CACHE	READ_CACHE	Not applicable
READAHEAD_CACHE/ NOREADAHEAD_CACHE	READAHEAD_CACHE	Not applicable
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 must always be overridden by using ENABLE_ACCESS_PATH=ALL on all connections to the SAN Volume Controller nodes.		

Connection settings for HP MA and EMA systems

The HP MA and EMA systems provide options that are configurable at the connection level.

Table 53 on page 433 lists the default and required HSG80 controller connection settings:

Table 53. HSG80 connection default and required settings

Option	HSG80 controller default setting	HSG80 controller required setting
OPERATING_SYSTEM	Not defined	WINNT
RESERVATION_STYLE	CONNECTION_BASED	Not applicable
UNIT_OFFSET	0	0 or 199

Mapping and virtualization settings for HP MA and EMA

There are LUN mapping or masking and virtualization restrictions for HP MA and EMA subsystems that are in a SAN Volume Controller environment.

The HP StorageWorks configuration interface requires that you assign a unit number to each logical unit (LU) when it is defined. By default, the LUN is the unit number. It is possible for gaps to exist in the LUN range if the unit numbers that are used in the configuration commands are not contiguous. By default, each LUN is visible on all controller ports on both controllers.

LUN masking

The HP MA and EMA subsystems support the concept of connection names. A maximum of 96 connection names that contain the following parameters are supported:

- HOST_ID
- ADAPTER_ID
- CONTROLLER
- PORT
- REJECTED_HOST

Note: The SAN Volume Controller ports must not be in the REJECTED_HOSTS list. This list can be seen with the **SHOW CONNECTIONS FULL** command.

You cannot use LUN masking to restrict the initiator ports or the target ports that the SAN Volume Controller uses to access 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 that the SAN Volume Controller uses, but the preferred method for this is to use SAN zoning.

LU virtualization

The HP MA and EMA subsystems also provide 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 controller target ports and SAN Volume Controller initiator ports is not supported.

Configuring HP StorageWorks EVA systems

This section provides information about configuring the HP StorageWorks Enterprise Virtual Array (EVA) system for attachment to a SAN Volume Controller.

Supported models of the HP EVA

The SAN Volume Controller supports models of the HP EVA.

See the following Web site for the latest supported models:

www.ibm.com/storage/support/2145

Supported firmware levels for HP EVA

The SAN Volume Controller supports HP EVA.

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

www.ibm.com/storage/support/2145

Concurrent maintenance on the HP EVA

Concurrent maintenance is the capability to perform I/O operations to an HP EVA while simultaneously performing maintenance operations on it.

Important: All maintenance operations must be performed by an HP Field Engineer.

The SAN Volume Controller and HP EVA support concurrent hardware maintenance and firmware upgrade.

User interface on the HP EVA system

Ensure that you are familiar with the user interface that supports the HP EVA system.

Storage Management Appliance

HP EVA systems are configured, managed, and monitored through a Storage Management Appliance. The Storage Management Appliance is a PC server that runs a software agent called Command View EVA. The software agent is accessed using a user interface that is provided by a standard Web browser.

Command View EVA communicates in-band with the HSV controllers.

Sharing the HP EVA controller between a host and the SAN Volume Controller

The HP 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 HP EVA system at the same time.
- LUs and RAID arrays must not be shared between a host and a SAN Volume Controller.

Switch zoning limitations for the HP EVA system

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.

Quorum disks on HP StorageWorks EVA systems

The SAN Volume Controller cluster selects managed disks (MDisks) that are presented by HP StorageWorks EVA systems as quorum disks.

Copy functions for HP StorageWorks EVA systems

Advanced copy functions for HP StorageWorks EVA systems (for example, VSnap and SnapClone) cannot be used with disks that are managed by the SAN Volume Controller cluster because the copy function does not extend to the SAN Volume Controller cache.

Logical unit configuration on the HP EVA

An EVA logical unit is referred to as a virtual disk (VDisk). An EVA system 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 can 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)

Logical unit creation and deletion on the HP 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 virtual disk (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 HP 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 HP 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 HP EVA.

Special LUs

The Console LU is a special VDisk that represents the SCSI target device. It is presented to all hosts as LUN 0.

Configuration interface for the HP EVA

The HP EVA is configured, managed, and monitored through a Storage Management Appliance. The Storage Management Appliance is a server that runs a software agent called Command View EVA. The Command View EVA is accessed using a graphical user interface that is provided by a standard Web browser.

In band

The Command View EVA system communicates in-band with the HSV controllers.

Configuration settings for HP StorageWorks EVA systems

The HP StorageWorks EVA configuration interface provides configuration settings and options that can be used with SAN Volume Controller clusters.

The settings and options can have a scope of the following:

- System (global)
- Logical unit (LU)
- Host

Global settings for HP StorageWorks EVA systems

Global settings apply across an HP StorageWorks EVA system.

Table 54 lists the system options that you can access using the Command View EVA.

Table 54. HP StorageWorks EVA global options and required settings

Option	HP EVA default setting	SAN Volume Controller required setting
Console LUN ID	0	Any
Disk replacement delay	1	Any

Logical unit options and settings for HP StorageWorks EVA systems

Logical unit (LU) settings are configurable at the LU level.

Table 55 on page 437 describes the options that must be set for each LU that is accessed by other hosts. LUs that are accessed by hosts can be configured differently.

Table 55. HP StorageWorks EVA LU options and required settings

Option	HP EVA Default Setting	SAN Volume Controller Required Setting
Capacity	None	Any
Write cache	Write-through or Write-back	Write-back
Read cache	On	On
Redundancy	Vraid0	Any
Preferred path	No preference	No preference
Write protect	Off	Off

Host options and settings for HP StorageWorks EVA systems

You must use specific settings to identify SAN Volume Controller clusters as hosts to HP StorageWorks EVA systems.

Table 56 lists the host options and settings that can be changed using the Command View EVA.

Table 56. HP EVA host options and required settings

Option	HP EVA Default Setting	SAN Volume Controller Required Setting
OS type	-	Windows
Direct eventing	Disabled	Disabled

Configuring HP StorageWorks MSA1000 and MSA1500 systems

This section provides information about configuring HP StorageWorks Modular Smart Array (MSA) 1000 and 1500 (MSA1000 and MSA1500) systems for attachment to a SAN Volume Controller.

Supported models of the HP MSA1000 and MSA1500 system

The SAN Volume Controller supports models of the HP MSA series of systems.

See the following Web site for the latest supported models:

www.ibm.com/storage/support/2145

Supported firmware levels for the HP MSA1000 and MSA1500

The HP MSA system must use a firmware level that is supported by the SAN Volume Controller.

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

www.ibm.com/storage/support/2145

User interfaces on the HP MSA1000 and MSA1500

Ensure that you are familiar with the user interface applications that are used by the HP MSA1000 and MSA1500 systems.

You can use the following configuration utilities with HP MSA1000 or MSA1500 systems in a SAN Volume Controller environment:

- The CLI through an out-of-band configuration that is accessed through a host that is connected to the serial port of the HP MSA1000 or MSA1500.
- The GUI through an in-band configuration that uses the HP Array Configuration Utility (ACU).

Notes:

1. If the HP ACU is installed in a configuration that HP does not support, some of its functionality might not be available.
2. If you use an in-band configuration, you must ensure that LUs that are used by the SAN Volume Controller cannot be accessed by a direct-attached host.

Logical unit creation, deletion, and migration for HP StorageWorks MSA systems

Before you create, delete, or migrate logical units, you must read the storage configuration guidelines that are specified in the HP StorageWorks MSA1000 or MSA1500 documentation that is provided for this system.

Creating arrays

An array is a collection of physical disks. Use the storage configuration guidelines for SAN Volume Controller clusters to create arrays on the HP StorageWorks MSA.

Creating logical drives

The following types of RAID arrays are supported:

- RAID 1+0
- RAID 1
- RAID 5
- RAID 6 (ADG)

RAID 0 is not supported because it does not provide failure protection.

All stripe sizes are supported; however, use a consistent stripe size for the HP StorageWorks MSA.

Use the following settings for logical drives:

- Set Max Boot to disabled
- Set Array Accelerator to enabled.

Note: If you are using the CLI, use the cache=enabled setting.

Presenting logical units to hosts

Set the Selective Storage Presentation (SSP), also known as ACL to enabled.

Use the following host profile settings:

```
Mode 0 = Peripheral Device LUN Addressing
Mode 1 = Asymmetric Failover
Mode 2 = Logical volumes connect as available on Backup Controller
Mode 3 = Product ID of 'MSA1000 Volume'
Mode 4 = Normal bad block handling
Mode 5 = Logout all initiators on TPRLO
Mode 6 = Fault management events not reported through Unit Attention
Mode 7 = Send FCP response info with SCSI status
Mode 8 = Do not send Unit Attention on failover
Mode 9 = SCSI inquiry revision field contains the actual version
Mode 10 = SCSI inquiry vendor field contains Compaq
Mode 11 = Power On Reset Unit Attention generated on FC Login or Logout
Mode 12 = Enforce Force Unit Access on Write
```

You can use the built-in Linux profile or Default profile to set the host profile settings. If you use the Default profile, you must issue the following Serial port CLI command to change the host profile settings:

`change mode Default` *`mode number`*

where *mode number* is the numeric value for the mode that you want to change.

See the documentation that is provided for your HP StorageWorks MSA for additional information.

Important: You must use the Serial port CLI or the SSP to recheck the connection objects after the configuration is complete.

Migrating logical units

You can use the standard migration procedure to migrate logical units from the HP StorageWorks MSA to the SAN Volume Controller cluster with the following restrictions:

- You cannot share the HP StorageWorks MSA between a host and the SAN Volume Controller cluster. You must migrate all hosts at the same time.
- The subsystem device driver (SDD) and securepath cannot coexist because they have different QLogic driver requirements.
- The QLogic driver that is supplied by HP must be removed and the driver that is supported by IBM must be installed.

Sharing the HP MSA1000 and MSA1500 between a host and the SAN Volume Controller

You must configure your environment so that only the SAN Volume Controller can access all logical units on the HP MSA1000 and MSA1500. You can zone other hosts to communicate with the HP MSA1000 and MSA1500 for in-band configuration, but nothing else.

Concurrent maintenance on the HP MSA1000 and MSA1500

Concurrent maintenance is the capability to perform I/O operations to an HP MSA1000 and MSA1500 while simultaneously performing maintenance operations on it.

You can perform nondisruptive maintenance procedures concurrently on the following components:

- HP MSA1000 or MSA1500 controller

- HP MSA1000 or MSA1500 controller cache
- Cache battery pack
- Variable speed blower
- Power supply
- Disk drive
- SFP transceiver

Quorum disks on the HP MSA

The SAN Volume Controller cannot use logical units (LUs) that are exported by the HP MSA1000 and MSA1500 as quorum disks.

Advanced functions for the HP MSA

The SAN Volume Controller Copy Service functions and RAID migration utilities are not supported for logical units (LUs) that are presented by the HP MSA.

Global settings for HP MSA systems

Global settings apply across an HP MSA system.

The following table lists the global settings for an HP MSA system:

Option	Required setting
Expand Priority	All supported **Note:** Performance impact of high priority
Rebuild Priority	All supported **Note:** Performance impact of high priority
Array Accelerator	On **Note:** Set on all logical drives that are used by the SAN Volume Controller.
Read-Write cache ratio	All supported
Name of controller	Not important

Configuring HP StorageWorks MSA2000 storage systems

This section provides information about configuring Hewlett Packard (HP) 2000 family Modular Smart Array (MSA2000) systems for attachment to a SAN Volume Controller.

HP MSA2000 supported models

SAN Volume Controller clusters can be used with MSA2000 storage systems.

See the following Web site for the latest supported models:

www.ibm.com/storage/support/2145

For SAN Volume Controller version 4.3.1.7, this is only the MSA2000fc dual controller model that is configured with each controller module attached to both fabrics. For details, refer to the *HP StorageWorks Modular Model User Guide* section on connecting two data hosts through two switches where all four ports must be used and cross-connected to both SAN fabrics.

Supported HP MSA2000 firmware levels

You must ensure that the MSA2000 firmware level can be used with the SAN Volume Controller cluster.

For the supported firmware levels and hardware, see the SAN Volume Controller (2145) Web site:

www.ibm.com/storage/support/2145

HP MSA2000 user interfaces

You can configure an MSA2000 system through the Storage Management Utility (SMU), which is a Web server on each controller, or with the command-line interface (CLI).

To access the MSA2000 system initially you can go through either a serial interface or Dynamic Host Configuration Protocol (DHCP). You can also configure user access and privileges.

MSA2000 Web graphical user interface (GUI)

The SMU is a Web-based GUI that runs on each controller that is accessible through the IP address of each controller. All management and monitoring tasks can be completed on each controller.

MSA2000 command-line interface (CLI)

The CLI is accessible through Secure Shell (SSH), Telnet, and serial port. The CLI includes all functionality that is available in the GUI.

Concurrent maintenance on MSA2000 systems

Concurrent maintenance is the capability to perform I/O operations while you simultaneously perform maintenance operations on the MSA2000 system.

Apply firmware upgrades to an MSA2000 system during a maintenance window because the MSA2000 system takes both controllers offline simultaneously multiple times during an upgrade.

Logical units and target ports on MSA2000 systems

Partitions (volumes) on MSA2000 systems are exported as logical units with a logical unit number that is assigned to that partition.

LUNs on MSA2000 systems

The controller calls a RAID array a virtual disk (VDisk). SAS and SATA disks cannot be mixed within a VDisk, and the maximum number of VDisks per controller is 16. VDisks can be divided into volumes, which are then presented to the host. There can be up to 128 volumes per controller. The capacity of a volume is between 1 MB and 16 TB.

Note: SAN Volume Controller version 4.3.1.7 has an individual 2 TB managed-disk size limit.

LUN IDs

LUNs exported by MSA2000 systems report identification descriptors 0, 3, 4, 5 in the VPD page 0x83. The LUN IDs are based on the controller MAC addresses. For example:

```
example;
# show volumes
Vdisk    Volume Name   Size  WR Policy    Class  Volume Serial Number  Cache Opt  Type

VD0                    VD0_V1              750.1GB  writeback    standard 00c0ffd76a330000a0fa124a01000000  standard  standard
VD2                    VD2_V1              750.1GB  writeback    standard 00c0ffd76a33000048fb124a01000000  standard  standard
VD_HC                  VD_CAP_V1           37.5GB   writeback    standard 00c0ffd76a3300005efc124a01000000  standard  standard
VD_1                   VD_1_V1             750.1GB  writeback    standard 00c0ffd7648f000064851d4a01000000  standard  standard
VD_3                   VD_3_V1             750.1GB  writeback    standard 00c0ffd7648f0000a6851d4a01000000  standard  standard
VD-R                   VD-R_V1             250.0GB  writeback    standard 00c0ffd7648f0000aa08234a01000000  standard  standard
VD-R                   VD-R_V2             250.0GB  writeback    standard 00c0ffd7648f0000ab08234a01000000  standard  standard
VD-R                   VD-R_V3             250.0GB  writeback    standard 00c0ffd7648f0000ab08234a02000000  standard  standard
-------------------------------------------------------------------------------------------------------------------------------
# show network-parameters
Network Parameters Controller A
-------------------------------
IP Address     : 9.71.47.27
Gateway        : 9.71.46.1
Subnet Mask    : 255.255.254.0
MAC Address    : 00:C0:FF:D7:6A:33
Addressing Mode: DHCP

Network Parameters Controller B
-------------------------------
IP Address     : 9.71.47.30
Gateway        : 9.71.46.1
Subnet Mask    : 255.255.254.0
MAC Address    : 00:C0:FF:D7:64:8F
Addressing Mode: DHCP
```

LUN creation and deletion

MSA2000 LUNs can be created, modified, or deleted either by the Storage Management Utility (SMU) or the command-line interface (CLI). LUNs can be used immediately with format to zeros as default background task.

Note: Disks appear as critical while this process is taking place.

To create a logical unit (volume from a VDisk), complete these steps:

1. In the Storage Management Utility SMU interface, go to **Manage** → **Virtual Disk Config** → **Create a VDisk**. The SMU provides a wizard to create the virtual disks.
2. You have the following options:
   - Manual
   - Virtual Disk Name
   - RAID Type

     Note: SAN Volume Controller does not support RAID 0.
   - Number of volumes
   - Expose to all hosts

     Note: The Expose to all hosts option can cause confusion in multi-cluster environments.
   - LUN assignments

You can also modify, expand, or delete a volume or VDisk using either the SMU or the CLI.

Note: Before you delete the LUN on the MSA2000 system, use the svctask rmdisk command to delete the MDisk on the SAN Volume Controller cluster.

LUN presentation

You can also use the SMU or the CLI to map and unmap MSA2000 LUNs.

To map a logical unit (volume from a VDisk), from the SMU complete these steps:

1. In the Storage Management Utility SMU interface, go to **Manage → Volume Management → VDisk or Volume → Volume Mapping**.
2. Under the section Assign Host Access Privileges, select **Map Host to Volume**.
3. For each SAN Volume Controller WWPN, select **SVC WWPN** in the HOST WWN-Name menu.
4. Enter the LUN number to present to the SAN Volume Controller. For example, use 0 for the first volume, then use 1 for the second, until all volumes are assigned.
5. Select read-write for the Port 0 Access and Port 1 Access.
6. Click **Map it**. The resulting mapping is displayed in the Current Host-Volume Relationships section.

   Important: Use this section to verify that the LUN ID is consistent and all SAN Volume Controller WWPNs have been mapped.

Because there are 8 nodes in the following example, 32 WWPNs show in the show volume-maps output (four ports per node).

```
example shown for an 8-node cluster, that is, 32 WWPNs;
# show volume-maps
Volume [SN 00c0ffd76a330000a0fa124a01000000, Name (VD0_V1)] mapping view:
CH       ID LUN Access Host-Port-Identifier                    Nickname
-------------------------------------------------------------------------------
0,1       0   0 rw     50050768012FFFFF
0,1       0   0 rw     5005076801105CEE
0,1       0   0 rw     500507680110008A
0,1       0   0 rw     50050768011FFFFF
0,1       0   0 rw     50050768013FFFFF
0,1       0   0 rw     50050768014FFFFF
0,1       0   0 rw     500507680140008A
0,1       0   0 rw     500507680130008A
0,1       0   0 rw     500507680120008A
0,1       0   0 rw     5005076801405CEE
0,1       0   0 rw     5005076801205CEE
0,1       0   0 rw     5005076801305CEE
0,1       0   0 rw     500507680110596B
0,1       0   0 rw     5005076801305FB8
0,1       0   0 rw     5005076801205FB8
0,1       0   0 rw     5005076801405FB8
0,1       0   0 rw     5005076801105FB8
0,1       0   0 rw     500507680120596B
0,1       0   0 rw     500507680140596B
0,1       0   0 rw     500507680130596B
0,1       0   0 rw     5005076801400009
0,1       0   0 rw     5005076801300009
0,1       0   0 rw     5005076801100009
0,1       0   0 rw     5005076801200009
0,1       0   0 rw     50050768014FFFFE
0,1       0   0 rw     50050768013FFFFE
0,1       0   0 rw     50050768012FFFFE
0,1       0   0 rw     50050768011FFFFE
0,1       0   0 rw     5005076801200001
0,1       0   0 rw     5005076801400001
0,1       0   0 rw     5005076801300001
0,1       0   0 rw     5005076801100001
```

Note: LUNs from controller module A and controller module B can have the same LUN IDs (0). Controller module A and Controller module B appear on the SAN Volume Controller cluster as separate controllers. Managed disks (MDisks) on the cluster should be in separate MDisk groups so that each controller module has its own separate MDisk group for its presented MDisks.

Special LUNs

Volumes can have a LUN ID from 0 to 126 on each controller. LUN 0 on the MSA2000 is visible from both controllers, but can only be used to access storage from the preferred controller. LUN 0 on the other controller does not present storage.

Target ports on MSA2000 systems

The MSA2000 system has two dual-active controllers with two ports each. You must set these as point-to-point using the SMU interface.

In the Storage Management Utility SMU interface, go to **Manage** → **General Config** → **Host Port Configuration**. Select Advanced Options and specify point to point for Change Host Topology.

Each WWPN is identified with the pattern 2P:7N:CC:CC:CC:MM:MM:MM where *P* is the port number on the controller and *N* is the address of the controller port (0 or 8), *CC:CC:CC* represents the Organizationally Unique Identifier (OUI), and *MM:MM:MM* is unique for the particular controller.

```
example;
# show port-wwn
CTRL CH WWPN
-------------------------
A    0  207000C0FFD75198
A    1  217000C0FFD75198
B    0  207800C0FFD75198
B    1  217800C0FFD75198
```

LU access model

The MSA2000 is a dual-active system. Each LUN has an owning controller, and I/O is serviced only by ports on that controller. Failover automatically takes place when one controller fails (shuts down). There is no way for SAN Volume Controller to force failover.

LU grouping

The MSA2000 system does not support LU grouping.

LU preferred access port

The MSA system has two ports per controller. The I/O is through port 0, and port 1 is linked to port 0 of the other controller during a failure or a code upgrade.

Detecting ownership

The LUN is reported only by the target ports of the owning controller.

Failover

The only way to cause failover of LUs from one controller to the other is to shut down one of the controllers. The MSA2000 system cannot normally present all the system LUNs through both controllers. Therefore, it requires a four-port connection to two SAN fabrics. Failover for MS2000 systems involves the surviving controller taking its ports offline, then returning with one of its ports, emulating the WWPNs of the failed controller.

Note: This behavior also means that half of the operational paths from the surviving controller are taken away when failover takes place, which allows the port from the controller that is shutting down to be emulated.

Switch zoning for MSA2000 storage systems

Switch zoning configurations for the MSA2000 system include considerations for fabric zoning, target port sharing, host splitting, and controller splitting.

Fabric zoning

Each SAN Volume Controller switch zone must include at least one target port from each controller to have no single point of failure. This means, for example, that the zone on the first fabric has Port 0 MSA Controller A with Port 1 of MSA Controller B and the SAN Volume Controller ports. The zone on the second fabric has Port 0 MSA Controller B and Port 1 MSA Controller A and the SAN Volume Controller ports. For more information about the fibre-channel dual-fabric setup, refer to the relevant MSA documentation.

Target port sharing

Target ports may not be shared between SAN Volume Controller and other hosts.

Host splitting

A single host must not be connected to SAN Volume Controller and an MSA2000 system simultaneously.

Controller splitting

MSA2000 system LUNs must be mapped only to the SAN Volume Controller cluster. The four target ports are all required for dual SAN-fabric connections and cannot be shared.

Configuration settings for MSA2000 systems

The MSA2000 System Storage Management Utility (SMU) provides configuration settings and options that can be used with SAN Volume Controller clusters.

Target port options

Table 57 on page 446 describes the port settings that are supported by the SAN Volume Controller.

Table 57. MSA2000 system port settings for use with the SAN Volume Controller

Option	Values (any limits on the possible values)	Description
Host Port Configuration	2 Gbps or 4 Gbps	Set according to the fabric speed.
Internal Host Port Interconnect	Straight-through	Set to Straight-through for a point-to-point fibre-channel connection.
Host Port Configuration	Point-to-Point	Set to Point-to-point for use with SAN Volume Controller.

LU options and settings

The MSA volumes can be created after you create a virtual disk (VDisk) (RAID 0 is not supported) or later added to the VDisk. LUNs can be configured in 16K, 32K, and 64K (default) chunks by using the advanced option. Table 58 describes the preferred options available when you create a logical unit (LU).

Table 58. Preferred options for logical units (LU)

Option	Value	Description
Expose to All Hosts	Yes	After the mapping of the volume to the SAN Volume Controller completes, this is modified to All other hosts (none no access). This can be done under the Assign Host Access Privileges frame.
Automatically assign LUNs	Yes	This forces option expose to all hosts and is necessary for consistent LUN numbering.
write-policy	write-back	
optimization	any	
read-ahead-size	default	
independent	disable	This setting controls cache mirroring. Because SAN Volume Controller requires mirroring, the independent=disable option must be used.

Host options and settings for MSA2000 systems

There is no specific host option to present the MSA2000 systems to SAN Volume Controller clusters. Use Microsoft Windows 2003 (Microsoft Windows 2003) as the host setting for SAN Volume Controller.

Quorum disks on MSA2000 systems

The SAN Volume Controller cluster requires managed disks (MDisks) that are quorum disks for system metadata storage. The MSA2000 system failover method is not compatible with the requirements for these disks. Quorum disks must be on another separate and suitable managed controller.

Copy functions for MSA2000 systems

The MSA2000 system provides optional copy and replicate features, called *clone* and *snapshot*. However, these functions must not be used with SAN Volume Controller.

Configuring NEC iStorage systems

This section provides information about configuring NEC iStorage systems for attachment to a SAN Volume Controller.

Supported firmware levels for the NEC iStorage

The NEC iStorage system must use a firmware level that is supported by the SAN Volume Controller.

See the following Web site for specific firmware levels and the latest supported hardware:www.ibm.com/storage/support/2145

Logical unit creation and deletion for NEC iStorage systems

You can create or delete logical units for NEC iStorage systems. See the storage configuration guidelines that are specified in the NEC iStorage documentation that is provided for this system.

Platform type for NEC iStorage

You must set all logical units that the SAN Volume Controller accesses to platform type AX (AIX).

Access control methods for NEC iStorage

You can use access control to restrict access from hosts and SAN Volume Controller clusters. You do not need to use access control to allow a SAN Volume Controller cluster to use all of the defined logical units on the system.

The following table lists the access control methods that are available:

Method	Description
Port Mode	Allows access to logical units that you want to define on a per storage controller port basis. SAN Volume Controller visibility (through switch zoning, physical cabling, etc.) must allow the SAN Volume Controller cluster to have the same access from all nodes and the accessible controller ports have been assigned the same set of logical units with the same logical unit number. This method of access control is not recommended for SAN Volume Controller connection.
WWN Mode	Allows access to logical units using the WWPN of each of the ports of an accessing host device. All WWPNs of all the SAN Volume Controller nodes in the same cluster must be added to the list of linked paths in the controller configuration. This becomes the list of host (SAN Volume Controller) ports for an LD Set or group of logical units. This method of access control allows sharing because different logical units can be accessed by other hosts.

Setting cache allocations for NEC iStorage

Cache allocations can be set manually; however, changes to the default settings can adversely effect performance and cause you to lose access to the system.

Snapshot Volume and Link Volume for NEC iStorage

You cannot use Copy Services logical volumes with logical units that are assigned to the SAN Volume Controller.

Configuring NetApp FAS systems

This section provides information about configuring the Network Appliance (NetApp) Fibre-attached Storage (FAS) systems for attachment to a SAN Volume Controller. Models of the NetApp FAS system are equivalent to the IBM System Storage N5000 series and the IBM System Storage N7000 series; therefore, the SAN Volume Controller also supports models of the IBM N5000 series and the IBM N7000 series.

Attention: You must configure NetApp FAS systems in single-image mode. SAN Volume Controller does not support NetApp FAS systems that are in multiple-image mode.

The information in this section also applies to the supported models of the IBM N5000 series and the IBM N7000 series.

Supported models of the NetApp FAS system

The SAN Volume Controller supports models of the NetApp FAS200, FAS900, FAS3000 and FAS6000 series of systems.

See the following Web site for the latest supported models:

www.ibm.com/storage/support/2145

Supported firmware levels for the NetApp FAS

The NetApp FAS must use a firmware level that is supported by the SAN Volume Controller.

See the following Web site for specific firmware levels and the latest supported hardware:www.ibm.com/storage/support/2145

User interfaces on the NetApp FAS

Ensure that you are familiar with the user interface applications that support the NetApp FAS.

See the documentation that is provided with your NetApp FAS system for more information about the Web server and CLI.

Web server

You can manage, configure, and monitor the NetApp FAS through the FileView GUI.

CLI

You can access the command-line interface through a direct connection to the filer serial console port or by using the filer IP address to establish a telnet session.

Logical units and target ports on NetApp FAS systems

For the NetApp FAS systems, a logical unit (LU) is a subdirectory in an internal file system.

LUs that are exported by the NetApp FAS system report identification descriptors in the vital product data (VPD). The SAN Volume Controller cluster uses the

LUN-associated binary type-3 IEEE Registered Extended descriptor to identify the LU. For a NetApp LUN that is mapped to the SAN Volume Controller cluster, set the LUN Protocol Type to Linux.

The NetApp FAS system does not use LU groups so that all LUs are independent. The LU access model is active-active. Each LU has a preferred filer, but can be accessed from either filer. The preferred filer contains the preferred access ports for the LU. The SAN Volume Controller cluster detects and uses this preference.

The NetApp FAS reports a different worldwide port name (WWPN) for each port and a single worldwide node name (WWNN).

Creating logical units on the NetApp FAS

To create logical units, you must identify a volume from which to create the logical unit and specify the amount of space to use.

Perform the following steps to create logical units:

1. Log on to the NetApp FAS.
2. Go to **Filer View** and authenticate.
3. Click **Volumes** and identify a volume to use to create an LU. A list of volumes is displayed.
4. Identify a volume that has sufficient free space for the LUN size that you want to use.
5. Click **LUNs** on the left panel.
6. Click **Add** in the list.
7. Enter the following:
   a. In the **Path** field, enter /vol/*volx/lun_name* where *volx* is the name of the volume identified above and *lun_name* is a generic name.
   b. In the **LUN Type** field, enter Image.
   c. Leave the **Description** field blank.
   d. In the **Size** field, enter a LUN Size.
   e. In the **Units** field, enter the LUN Size in units.
   f. Select the **Space Reserved** box.

      Note: If the Space Reserved box is not selected and the file system is full, the LUN goes offline. The managed disk group also goes offline and you cannot access the virtual disks.

   g. Click **Add**.

      Note: To check the LUN settings, go to the Manage LUNs section and click the LUN you want to view. Ensure that the Space Reserved setting is set.

Deleting logical units on the NetApp FAS

You can delete logical units.

Perform the following steps to delete logical units:

1. Log on to the NetApp FAS.
2. Go to **Filer View** and authenticate.
3. Click **LUNs** on the left panel.

4. Click **Manage**. A list of LUNs is displayed.
5. Click the LUN that you want to delete.
6. Click **Delete**.
7. Confirm the LUN that you want to delete.

Creating host objects for the NetApp FAS

You can create host objects.

Perform the following steps to create host objects:

1. Log on to the NetApp FAS.
2. Go to **Filer View** and authenticate.
3. Click **LUNs** on the left panel.
4. Click **Initiator Groups**.
5. Click **Add** in the list.
6. Enter the following:
   a. In the **Group Name** field, enter the name of the initiator group or host.
   b. In the **Type** list, select FCP.
   c. In the **Operating System** field, select Linux.
   d. In the **Initiators** field, enter the list of WWPNs of all the ports of the nodes in the cluster that are associated with the host.

      Note: Delete the WWPNs that are displayed in the list and manually enter the list of SAN Volume Controller node ports. You must enter the ports of all nodes in the SAN Volume Controller cluster.
7. Click **Add**.

Presenting LUNs to hosts for NetApp FAS

You can present LUNs to hosts.

Perform the following steps to present LUNs to hosts:

1. Log on to the NetApp FAS.
2. Go to **Filer View** and authenticate.
3. Click **LUNs** on the left panel.
4. Click **Manage**. A list of LUNs is displayed.
5. Click the LUN that you want to map.
6. Click **Map LUN**.
7. Click **Add Groups to Map**.
8. Select the name of the host or initiator group from the list and click **Add**.

   Notes:

   a. You can leave the LUN ID section blank. A LUN ID is assigned based on the information the controllers are currently presenting.
   b. If you are re-mapping the LUN from one host to another, you can also select the **Unmap** box.
9. Click **Apply**.

Switch zoning limitations for NetApp FAS systems

There are limitations in switch zoning for SAN Volume Controller clusters and NetApp FAS systems.

Fabric zoning

The SAN Volume Controller switch zone must include at least one target port from each filer to avoid a single point of failure.

Target port sharing

Target ports can be shared between the SAN Volume Controller cluster and other hosts. However, you must define separate initiator groups (igroups) for the SAN Volume Controller initiator ports and the host ports.

Host splitting

A single host cannot be connected to both the SAN Volume Controller cluster and the NetApp FAS to avoid the possibility of an interaction between multipathing drivers.

Controller splitting

You can connect other hosts directly to both the NetApp FAS and the SAN Volume Controller cluster under the following conditions:

- Target ports are dedicated to each host or are in a different igroup than the SAN Volume Controller cluster
- LUNs that are in the SAN Volume Controller cluster igroup are not included in any other igroup

Concurrent maintenance on the NetApp FAS

Concurrent maintenance is the capability to perform I/O operations to a NetApp FAS while simultaneously performing maintenance operations on it.

The SAN Volume Controller supports concurrent maintenance on the NetApp FAS.

Quorum disks on the NetApp FAS

The SAN Volume Controller can use logical units (LUs) that are exported by the NetApp FAS as quorum disks.

Advanced functions for the NetApp FAS

The SAN Volume Controller copy and migration functions are supported for logical units (LUs) that are presented by the NetApp FAS.

Configuring Pillar Axiom systems

This section provides information about configuring Pillar Axiom systems for attachment to a SAN Volume Controller cluster.

Supported models of Pillar Axiom systems

SAN Volume Controller clusters can be used with some models of the Pillar Axiom series of systems.

See the following Web site for the latest models that can be used with SAN Volume Controller clusters:

www.ibm.com/storage/support/2145

Supported firmware levels of Pillar Axiom systems

You must ensure that the firmware level of the Pillar Axiom system can be used with the SAN Volume Controller cluster.

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

www.ibm.com/storage/support/2145

Concurrent maintenance on Pillar Axiom systems

Concurrent maintenance is the capability to perform I/O operations to a Pillar Axiom system while simultaneously performing maintenance operations on it.

Because some maintenance operations restart the Pillar Axiom system, you cannot perform hardware maintenance or firmware upgrades while the system is attached to a SAN Volume Controller cluster.

Pillar Axiom user interfaces

Ensure that you are familiar with the user interface applications that Pillar Axiom systems use. For more information, see the documentation that is included with the Pillar Axiom system.

AxiomONE Storage Services Manager

The AxiomONE Storage Services Manager is a browser-based GUI that allows you to configure, manage, and troubleshoot Pillar Axiom systems.

Pillar Data Systems CLI

The Pillar Data Systems command-line interface (CLI) communicates with the system through an XML-based application programming interface (API) over a TCP/IP network. The Pillar Data Systems CLI is installed through the AxiomOne Storage Service Manager. You can use the Pillar Data Systems CLI to issue all commands, run scripts, request input files to run commands, and run commands through a command prompt. The Pillar Data Systems CLI can run on all operating systems that can be used with Pillar Axiom systems.

AxiomONE CLI

The AxiomONE CLI is installed through the AxiomONE Storage Service Manager. You can use the AxiomONE CLI to perform administrative tasks. The AxiomONE CLI can run on a subset of operating systems that can be used with Pillar Axiom systems.

Logical units and target ports on Pillar Axiom systems

For Pillar Axiom systems, logical units are enumerated devices that have the same characteristics as LUNs.

LUNs

You can use the AxiomONE Storage Services Manager to create and delete LUNs.

Important:

1. When you create a LUN, it is not formatted and therefore can still contain sensitive data from previous usage.
2. You cannot map more than 256 Pillar Axiom LUNs to a SAN Volume Controller cluster.

You can create LUNs in a specific volume group or in a generic volume group. A single LUN can use the entire capacity of a disk group. However, for SAN Volume Controller clusters, LUNs cannot exceed 2 TB. When LUNs are exactly 2 TB, a warning is issued in the SAN Volume Controller cluster error log.

The amount of capacity that the LUN uses is determined by the capacity of the LUN and the level of redundancy. You can define one of three levels of redundancy:

- Standard, which stores only the original data
- Double, which stores the original data and one copy
- Triple, which stores the original data and two copies

For all levels of redundancy, data is striped across multiple RAID-5 groups.

LUNs that are exported by the Pillar Axiom system report identification descriptors in the vital product data (VPD). The SAN Volume Controller cluster uses the LUN-associated binary type-2 IEEE Registered Extended descriptor to identify the LUN. The following format is used:

CCCCCCLLLLMMMMMM

where CCCCCC is the IEEE company ID (0x00b08), LLLL is a number that increments each time a LUN is created (0000–0xFFFD) and MMMMMM is the system serial number.

You can find the identifier in the AxiomONE Storage Services Manager. From the AxiomONE Storage Services Manager, click **Storage → LUNs → Identity**. The identifier is listed in the LUID column. To verify that the identifier matches the UID that the SAN Volume Controller cluster lists, issue the **svcinfo lsmdisk** *mdisk_id or mdisk_name* from the command-line interface and check the value in the `UID` column.

Moving LUNs

If you want to migrate more than 256 LUNs on an existing Pillar Axiom system to the SAN Volume Controller cluster, you must use the SAN Volume Controller cluster migration function. The Pillar Axiom system allows up to 256 LUNs per host and the SAN Volume Controller cluster must be configured as a single host. Because the SAN Volume Controller cluster is not limited to 256 virtual disks, you can migrate your existing Pillar Axiom system set up to the SAN Volume Controller cluster. You must then virtualize groups of LUNs and then migrate the group to larger managed mode disks.

Target ports

Pillar Axiom systems with one pair of controllers report a different worldwide port name (WWPN) for each port and a single worldwide node name (WWNN). Systems with more than one pair of controllers report a unique WWNN for each controller pair.

LUN groups are not used so that all LUNs are independent. The LUN access model is active-active/asymmetric with one controller having ownership of the LUN. All I/O operations to the LUN on this controller is optimized for performance. You can use the **svcinfo lsmdisk** *mdisk_id or mdisk_name* CLI command to determine the assigned controller for a LUN.

To balance I/O load across the controllers, I/O operations can be performed through any port. However, performance is higher on the ports of the controller that own the LUNs. By default, the LUNs that are mapped to the SAN Volume Controller cluster are accessed through the ports of the controller that owns the LUNs.

Switch zoning limitations for Pillar Axiom systems

There are limitations in switch zoning for SAN Volume Controller clusters and Pillar Axiom systems.

Fabric zoning

The SAN Volume Controller switch zone must include at least one target port from each Pillar Axiom controller to avoid a single point of failure.

Target port sharing

Target ports can be shared between the SAN Volume Controller cluster and other hosts.

Host splitting

A single host cannot be connected to both the SAN Volume Controller cluster and the Pillar Axiom system to avoid the possibility of an interaction between multipathing drivers.

Controller splitting

Pillar Axiom system LUNs that are mapped to the SAN Volume Controller cluster cannot be mapped to other hosts. Pillar Axiom system LUNs that are *not* mapped to the SAN Volume Controller cluster can be mapped to other hosts.

Configuration settings for Pillar Axiom systems

The AxiomONE Storage Services Manager provides configuration settings and options that can be used with SAN Volume Controller clusters.

The settings and options can have a scope of the following:

- System (global)
- Logical unit (LU)
- Host

Global settings for Pillar Axiom systems

Global settings apply across a Pillar Axiom system.

Table 59 lists the system options that you can access using the AxiomONE Storage Services Manager.

Table 59. Pillar Axiom global options and required settings

Option	Pillar Axiom default setting	SAN Volume Controller required setting
Enable Automatic Failback of NAS Control Units	Y	N/A
Link Aggregation	N	N/A
DHCP/Static	-	Any
Call-home	-	Any

Logical unit options and settings for Pillar Axiom systems

Logical unit (LU) settings are configurable at the LU level.

Table 60 lists the options that must be set for each LU that is accessed by other hosts. LUs that are accessed by hosts can be configured differently. You can use the AxiomONE Storage Services Manager to change these settings.

Table 60. Pillar Axiom LU options and required settings

Option	Pillar Axiom Default Setting	SAN Volume Controller Required Setting
LUN Access	All hosts	Select hosts
Protocol	FC	FC
LUN Assignment	Auto	Any **Attention:** Do not change the LUN assignment after the LUNs are mapped to the SAN Volume Controller cluster.
Select Port Mask	All On	All On
Quality of Service	Various	No preference. See the note below.
Note: If you do not know the Quality of Service setting, you can use the following: • Priority vs other Volumes = Medium • Data is typically accessed = Mixed • I/O Bias = Mixed		

Host options and settings for Pillar Axiom systems

You must use specific settings to identify SAN Volume Controller clusters as hosts to Pillar Axiom systems.

Table 61 on page 456 lists the host options and settings that can be changed using the AxiomONE Storage Services Manager.

Table 61. Pillar Axiom host options and required settings

Option	Pillar Axiom default setting	SAN Volume Controller required setting
Load balancing	Static	Static
HP-UX	N	N

Quorum disks on Pillar Axiom systems

The SAN Volume Controller cluster selects managed disks (MDisks) that are presented by Pillar Axiom systems as quorum disks.

Copy functions for Pillar Axiom systems

Advanced copy functions for Pillar Axiom systems (for example, Snap FS, Snap LUN, Volume Backup, Volume Copy, and Remote Copy) cannot be used with disks that are managed by the SAN Volume Controller cluster.

Configuring Texas Memory Systems RamSan Solid State Storage systems

This section provides information about configuring Texas Memory Systems (TMS) RamSan systems for attachment to a SAN Volume Controller.

TMS RamSan Solid State Storage supported models

SAN Volume Controller clusters can be used with the RamSan Solid State Storage systems.

For the latest RamSan models that can be used with SAN Volume Controller clusters, see the SAN Volume Controller (2145) Web site:

www.ibm.com/storage/support/2145

Supported TMS RamSan firmware levels

You must ensure that the RamSan firmware level can be used with the SAN Volume Controller cluster.

For the supported firmware levels and hardware, see the SAN Volume Controller (2145) Web site:

www.ibm.com/storage/support/2145

Concurrent maintenance on RamSan systems

Concurrent maintenance is the capability to perform I/O operations while you simultaneously perform maintenance operations on the RamSan system.

Apply firmware upgrades to a RamSan system during a maintenance window. A power cycle of the RamSan system is required for the upgraded firmware to take effect.

RamSan user interfaces

You can configure a RamSan system through a Web GUI based on Java™ or a command-line interface (CLI). You can also perform some system-critical operations by using the front panel on the RamSan system.

RamSan Web GUI

The Web GUI is an applet based on Java that is accessible through the IP address of the RamSan system. All configuration and monitoring steps are available through this interface. Be default, the Web GUI uses SSL encryption to communicate with the RamSan system.

RamSan CLI

The command-line interface (CLI) is accessible through SSH, Telnet, and RS-232 port. The CLI includes all functionality that is available in the GUI with the exception of statistics monitoring. The CLI includes a diagnostics interface, however, for internal hardware checks.

Logical units and target ports on RamSan systems

Partitions on RamSan systems are exported as logical units (LUs) with a logical unit number (LUN) that is assigned to the partition.

LUNs on RamSan systems

RamSan systems are shipped with a particular capacity of user space, which depends on the model. Capacities on one model can range from 1 - 2 TB. A partition with this capacity is known as a *logical unit*.

RamSan systems can export up to 1024 LUNs to the SAN Volume Controller through various exported FC ports. The maximum logical-unit size is the full, usable capacity of the RamSan system.

LUN IDs

RamSan systems identify exported LUs through identification descriptors 0, 1, and 2. The EUI-64 identifier for the LU is in the `CCCCCCLLLLMMMMMM` notation where `CCCCCC` is the Texas Memory Systems IEEE Company ID of 0020C2h, `LLLL` is the logical unit handle, and `MMMMMM` is the serial number of the chassis. The EUI-64 identifier is available on the detailed view of each logical unit in the GUI.

LUN creation and deletion

RamSan LUNs are created, modified, or deleted either by using a wizard tutorial in the GUI or by entering a CLI command. LUNs are not formatted to all zeros upon creation.

To create a logical unit, highlight **Logical Units** and select **Create toolbar**. To modify, resize, or delete an LU, select the appropriate toolbar button when the specific logical unit is highlighted in the navigation tree.

Note: Delete the MDisk on the SAN Volume Controller cluster before you delete the LUN on the RamSan system.

LUN presentation

LUNs are exported through the available FC ports of RamSan systems by access policies. Access policies are associations of the logical unit, port, and host. A RamSan system requires that one of the three items is unique across all available access policies. LUNs that are to be presented to SAN Volume Controller must be

presented to all node ports in the cluster through at least two ports on the RamSan system. Present each LU to the SAN Volume Controller on the same LUN through all target ports.

To apply access policies to a logical unit, highlight the specific logical unit in the GUI and click the **Access** toolbar button.

Special LUNs

The RamSan system has no special considerations for logical unit numbering. LUN 0 can be exported where necessary. In one RamSan model, a licensed Turbo feature is available to create a logical unit up to half the size of the cache to keep locked in the DRAM cache for the highest performance. No identification difference exists with a Turbo or locked LUN as opposed to any other LUN ID.

Target ports on RamSan systems

A RamSan system is capable of housing 4 dual-ported FC cards. Each worldwide port name (WWPN) is identified with the pattern `2P:0N:00:20:C2:MM:MM:MM` where `P` is the port number on the controller and `N` is the address of the controller. The `MMMMMM` represents the chassis serial number.

The controller address is as follows:

```
04:  FC77-1
08:  FC77-2
0C:  FC77-3
10:  FC77-4
```

Port 2B has a WWPN of 21:08:00:20:C2:07:83:32 for a system with serial number G-8332. The same system has a worldwide node name (WWNN) of 10:00:00:20:C2:07:83:32 for all ports.

LU access model

On a RamSan system, all controllers are Active/Active on a nonblocking crossbar. To avoid an outage from controller failure, configure multipathing across FC controller cards in all conditions. Because all RamSan systems are equal in priority, there is no benefit to using an exclusive set for a specific LU.

LU grouping

The RamSan system does not use LU grouping.

LU preferred access port

There are no preferred access ports for the RamSan system because all ports are Active/Active across all controllers.

Detecting ownership

Ownership is not relevant to the RamSan system.

Switch zoning for RamSan storage systems

Switch zoning configurations for the RamSan system include considerations for fabric zoning, target port sharing, host splitting, and controller splitting.

Fabric zoning

To enable multipathing, ensure that you have multiple zones or multiple RamSan and SAN Volume Controller ports for each zone when you are zoning a RamSan system to the SAN Volume Controller back-end ports.

Target port sharing

The RamSan system can support LUN masking to enable multiple servers to access separate LUNs through a common controller port. There are no issues with mixing workloads or server types in this setup. LUN Masking is a licensed feature of the RamSan system.

Host splitting

There are no issues with host splitting on a RamSan system.

Controller splitting

RamSan system LUNs that are mapped to the SAN Volume Controller cluster cannot be mapped to other hosts. LUNs that are not presented to the SAN Volume Controller can be mapped to other hosts.

Configuration settings for RamSan systems

The RamSan GUI provides configuration settings and options that can be used with SAN Volume Controller clusters.

LU options and settings

When you create a logical unit (LU), the options in Table 62 are available on RamSan systems.

Table 62. RamSan LU options

Option	Data type	Range	Default	SAN Volume Controller setting	Comments
Name	String	1 character - 32 characters	Logical unit number	Any	This is only for management reference.
Number	Integer	0 - 1023	Next available LUN	0 - 254	Some hosts have known limitations of 254 as their highest LUN ability. One logical unit can appear at multiple LUNs. For example, the same data could appear at LUN 1, LUN 7, and LUN 124.

Table 62. RamSan LU options (continued)

Option	Data type	Range	Default	SAN Volume Controller setting	Comments
Size	Integer	1 MB - maximum capacity	Maximum available capacity	Any	MB and GB are BASE2 offerings.
Backup mode	Option list	Writeback caching or writethrough caching	Writeback caching	Writeback caching	Use writeback caching in production. Use writethrough caching strictly for diagnostics.
Device ID	Integer	Blank, 0 - 32768	Blank	Blank	Specific only to OpenVMS.
Report corrected media errors	Checkbox	Checked or Unchecked	Checked	Checked	Notifies the host if ECC was used to correct the requested data.
Report uncorrected media errors	Checkbox	Checked or Unchecked	Checked	Checked	Always report uncorrected media errors.

Host options and settings for RamSan systems

No host options are required to present the RamSan systems to SAN Volume Controller clusters.

Quorum disks on RamSan systems

The SAN Volume Controller cluster selects managed disks (MDisks) that are presented by RamSan systems as quorum disks. To maintain availability with the cluster, ensure that each quorum disk resides on a separate disk system.

Clearing SCSI reservations and registrations

You must not use the RamSan CLI to clear SCSI reservations and registrations on volumes that are managed by the SAN Volume Controller. This option is not available on the GUI.

Copy functions for RamSan systems

The RamSan system does not provide copy, replicate, or SnapShot features. The RamSan system also does not provide thin provisioning.

Configuring Xiotech Emprise systems

This section provides information about configuring Xiotech Emprise systems for attachment to a SAN Volume Controller cluster.

Supported Xiotech Emprise models

SAN Volume Controller clusters can be used with the Xiotech Emprise storage system.

See the SAN Volume Controller (2145) Web site for the latest Xiotech Emprise models that can be used with SAN Volume Controller clusters:

www.ibm.com/storage/support/2145

Supported Xiotech Emprise firmware levels

You must ensure that the Xiotech Emprise firmware level can be used with the SAN Volume Controller cluster.

See the SAN Volume Controller (2145) Web site for the supported firmware levels and hardware:

www.ibm.com/storage/support/2145

Concurrent maintenance on Xiotech Emprise systems

Concurrent maintenance is the capability to perform I/O operations on a Xiotech Emprise system while simultaneously performing maintenance operations on the system.

Concurrent maintenance cannot be supported during I/O operations. Because some maintenance operations, such as firmware updates, restart Xiotech Emprise systems, consult the appropriate maintenance manual at the Xiotech Web site before you perform maintenance:

www.xiotech.com

Xiotech Emprise user interfaces

Ensure that you are familiar with the Xiotech Emprise user interface applications. For more information about the user interface applications, see the documentation that is included with the Xiotech Emprise system.

Xiotech Emprise Storage Management GUI

The Xiotech Emprise Storage Management GUI is a Java-based interface that you can use to configure, manage, and troubleshoot Xiotech Emprise storage systems. The Xiotech Emprise Storage Management GUI is designed and supported on Microsoft Windows systems and has the following minimum requirements:

Internet Explorer v6.02800.1106, SP1, Q903235 or higher (JavaScript™ enabled; XML/XSL rendering enabled)

Xiotech Emprise CLI

The Xiotech Emprise command-line interface (CLI) communicates with the system through a serial port that is connected to a computer that runs a terminal emulation program, such as Microsoft HyperTerminal or PuTTY. The Xiotech Emprise CLI is primarily used to configure the network adapter TCP/IP settings.

A null modem cable is required. Configure the serial port on the computer as follows:

- 115200 baud
- 8 data bits
- No parity
- 1-stop bit
- No flow control

Logical units and target ports on Xiotech Emprise systems

On Xiotech Emprise systems, logical units (LUs) are enumerated devices that have the same characteristics as logical unit numbers (LUNs).

LUNs

An Xiotech Emprise logical unit is referred to as a *volume*.

A single Xiotech Emprise volume can potentially consume the entire capacity that is allocated for SAN Volume Controller managed disk groups, but it cannot exceed the SAN Volume Controller 2 TB LUN size limit. Any LUN that is 2 TB or larger is truncated to 2 TB, and a warning message is generated for each path to the LUN.

LUN IDs

LUNs that are exported by Xiotech Emprise systems are guaranteed to be unique. They are created with a combination of serial numbers and counters along with a standard IEEE registered extended format.

LUN creation and deletion

Xiotech Emprise LUNs are created and deleted by using either the Xiotech Emprise Storage Management GUI or CLI. LUNs are formatted to all zeros at creation.

When a new LUN is created, the Xiotech Emprise system begins a background zeroing process. If a read operation comes in to an area that has not been processed yet, the system returns zeros as a read response. This is the normal procedure. If a previous LUN with data was in that storage area, it is zeroed out. If a non-zeroed-out area gets read, the system returns zeros if it has not been written to yet.

LUN presentation on Xiotech Emprise systems

Xiotech Emprise LUNs are presented to the SAN Volume Controller interface using the following rules:

- LUNs can be presented to one or more selected hosts.
- Configuration is easier if you create one host name for the SAN Volume Controller.
- No individual LUN volume on the Xiotech Emprise system can exceed 2 TB in size.
- For the managed reliability features to be effective on the Xiotech Emprise system, use either RAID 1 or RAID 5 when you create volumes.
- The write-back and write-through cache options are available depending on the performance requirements on each individual LUN. Generally, write-back caching provides the best performance.
- Although either Linux or Windows can be used, Linux is recommended for volumes that are intended for use on the SAN Volume Controller.

To present Xiotech Emprise LUNs to the SAN Volume Controller, follow these steps:

1. On the Xiotech Emprise system, create a single host name for the SAN Volume Controller and assign all SAN Volume Controller host bus adapter (HBA) ports to that host name as shown in Table 63.

Table 63. Host information for Xiotech Emprise

Name	Operating system type	HBA ports	Mapping
SVC_Cluster	Linux	500507680130535F 5005076801305555 500507680140535F 5005076801405555	Volume01 (lun:1) Volume02 (lun:2)

2. When you create new volumes that are intended for use on the SAN Volume Controller, assign them to the host name that is used to represent the SAN Volume Controller.

Special LUNs

The Xiotech Emprise storage system does not use a special LUN. Storage can be presented by using any valid LUN, including 0.

Target ports on Xiotech Emprise systems

Each Xiotech Emprise system has two physical fibre-channel ports. They are, by default, intended to provide failover or multipath capability. The worldwide node name (WWNN) and worldwide port name (WWPN) are typically similar, such as in the following example:

```
WWNN 20:00:00:14:c3:67:3f:c4
WWPN 20:00:00:14:c3:67:3f:c4
WWPN 20:00:00:14:c3:67:3f:c5
```

LU access model

The Xiotech Emprise system has no specific ownership of any LUN by any module. Because data is striped over all disks in a DataPac, performance is generally unaffected by the choice of a target port.

LU grouping

The Xiotech Emprise system does not use LU grouping; all LUNs are independent entities.

LU preferred access port

There are no preferred access ports for the Xiotech Emprise system.

Detecting ownership

Ownership is not relevant to the Xiotech Emprise system.

Switch zoning limitations for Xiotech Emprise storage systems

Limitations exist in switch zoning for SAN Volume Controller clusters and the Xiotech Emprise storage system.

Fabric zoning

To avoid a single point of failure, the SAN Volume Controller switch zone must include both target ports from each Xiotech Emprise controller.

Target port sharing

Target ports can be shared between the SAN Volume Controller cluster and other hosts.

Host splitting

To avoid the possibility of an interaction between multipathing drivers, a single host cannot be connected to both the SAN Volume Controller cluster and the Xiotech Emprise system.

Controller splitting

Xiotech Emprise system logical unit numbers (LUNs) that are mapped to the SAN Volume Controller cluster cannot be mapped to other hosts. Xiotech Emprise system LUNs that are *not* mapped to the SAN Volume Controller cluster can be mapped to other hosts.

Configuration settings for Xiotech Emprise systems

The Xiotech Emprise Storage Management GUI provides configuration settings and options that can be used with SAN Volume Controller clusters.

The only specific setting is the host operating system type: Windows or Linux. For SAN Volume Controller clusters, use Linux.

LU options and settings

Logical unit (LU) settings for the Xiotech Emprise system are configurable at the LU level.

Table 64 lists the options that must be set for each LU that is accessed by other hosts. LUs that are accessed by hosts can be configured differently. You can use the Xiotech Emprise Storage Management GUI or CLI to change these settings.

Table 64. Xiotech Emprise LU settings

Option	Data type	Range	Default	SAN Volume Controller setting	Comments
Capacity	Int	1 GB to 2 TB	No	Any	SAN Volume Controller supports up to 2 TB.

Host options and settings for Xiotech Emprise

You must use specific settings to identify SAN Volume Controller clusters as hosts to the Xiotech Emprise storage system.

A Xiotech Emprise host is a single WWPN; however, multiple WWPNs can be included in a single host definition on the Xiotech Emprise system.

A Xiotech Emprise host also can consist of more than one WWPN. The recommended method is to make each SAN Volume Controller node a Xiotech Emprise host and to make a Xiotech Emprise cluster that corresponds to all the nodes in the SAN Volume Controller cluster. To do this, include all of the SAN Volume Controller WWPNs under the same Xiotech Emprise host name.

Quorum disks on Xiotech Emprise systems

The SAN Volume Controller cluster selects managed disks (MDisks) that are presented by Xiotech Emprise systems as quorum disks. The clearing of Small Computer System Interface (SCSI) reservations and registrations is not supported by the Xiotech Emprise system.

Copy functions for Xiotech Emprise systems

Advanced copy functions for Xiotech Emprise systems such as SnapShot and remote mirroring cannot be used with disks that are managed by the SAN Volume Controller cluster. Thin provisioning is not supported for use with SAN Volume Controller.

Configuring IBM XIV Storage System models

This section provides information about configuring IBM XIV Storage System models for attachment to a SAN Volume Controller cluster.

Supported IBM XIV Storage System models

SAN Volume Controller support for IBM XIV Storage System systems is specific to certain models.

The supported IBM XIV Storage System models are:

- IBM XIV Storage System Model A14

Note: For Model A14, partially populated racks are supported.

See the following Web site for the latest IBM XIV Storage System models that can be used with SAN Volume Controller clusters:

www.ibm.com/storage/support/2145

Supported IBM XIV firmware levels

You must ensure that SAN Volume Controller supports your IBM XIV Storage System firmware level.

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

www.ibm.com/storage/support/2145

Concurrent maintenance on IBM XIV Storage System models

Concurrent maintenance is the capability to perform I/O operations on an IBM XIV Storage System models while simultaneously performing maintenance operations on the system.

Some maintenance operations require a complete restart of IBM XIV Storage System systems. Such procedures are not supported when the system is attached to the SAN Volume Controller.

Firmware upgrades of IBM XIV Storage System systems are not supported while the system is attached to the SAN Volume Controller.

All other concurrent maintenance procedures are supported.

IBM XIV user interfaces

Ensure that you are familiar with the IBM XIV Storage System user interface applications. For more information, see the documentation that is included with your IBM XIV Storage System system.

IBM XIV Storage Management GUI

The IBM XIV Storage System Storage Management GUI is a Java-based GUI that you can use to configure, manage, and troubleshoot IBM XIV Storage System systems. The IBM XIV Storage System Storage Management GUI can run on all operating systems that can be used with IBM XIV Storage System systems.

IBM XIV CLI

The IBM XIV Storage System command-line interface (XCLI) communicates with the systems through an XML-based API over a TCP/IP network. You can use the XCLI to issue all commands, run scripts, request input files to run commands, and run commands through a command prompt. The XCLI can run on all operating systems that can be used with IBM XIV Storage System systems.

Logical units and target ports on IBM XIV Storage System models

On IBM XIV Storage System, logical units (LUs) are enumerated devices that have the same characteristics as LUNs.

LUNs

An IBM XIV Storage System Logical Unit is referred to as a *volume*. IBM XIV Storage System and volumes are enumerated devices that all share identical characteristics.

A single IBM XIV Storage System volume can potentially consume the entire capacity that is allocated for SAN Volume Controller managed disk (MDisk) groups, and it can also exceed the SAN Volume Controller 2 TB LUN size limit. Any LUN that is 2 TB or larger is truncated to 2 TB, and a warning message is generated for each path to the LUN.

IBM XIV Storage System volumes consume chunks of 17,179,869,184 bytes (17 GB), although you can create volumes with an arbitrary block count.

LUN IDs

LUNs that are exported by IBM XIV Storage System models report Identification Descriptors 0, 1, and 2 in VPD page 0x83. SAN Volume Controller uses the EUI-64 compliant type 2 descriptor *CCCCCCMMMMMMLLLL*, where *CCCCCC* is the IEEE company ID, *MMMMMM* is the System Serial Number transcribed to hexadecimal (**10142->0x010142**, for example) and *LLLL* is **0000-0xFFFF**, which increments each time a LUN is created. You can identify the *LLLL* value by using the IBM XIV Storage System GUI or CLI to display the volume serial number.

LUN creation and deletion

IBM XIV Storage System LUNs are created and deleted using the IBM XIV Storage System GUI or CLI. LUNs are formatted to all zeros upon creation, but to avoid a significant formatting delay, zeros are not written.

Special LUNs

IBM XIV Storage System systems do not use a special LUN; storage can be presented using any valid LUN, including **0**.

LU access model

IBM XIV Storage System systems have no specific ownership of any LUN by any module. Because data is striped over all disks in the system, performance is generally unaffected by the choice of a target port.

LU grouping

IBM XIV Storage System models do not use LU grouping; all LUNs are independent entities. To protect a single IBM XIV Storage System volume from accidental deletion, you can create a consistency group containing all LUNs that are mapped to a single SAN Volume Controller cluster.

LU preferred access port

There are no preferred access ports for IBM XIV Storage System models.

Detecting ownership

Ownership is not relevant to IBM XIV Storage System models.

LUN presentation on XIV Nextra™ systems

XIV Nextra LUNs are presented to the SAN Volume Controller interface using the following rules:

- LUNs can be presented to one or more selected hosts.
- XIV Nextra maps consist of sets of LUN pairs and linked hosts.
- A volume can only appear once in a map.
- A LUN can only appear once in a map.
- A host can only be linked to one map.

To present XIV Nextra LUNs to the SAN Volume Controller, perform the following steps:

1. Create a map with all of the volumes that you intend to manage with the SAN Volume Controller cluster.
2. Link the WWPN for all node ports in the SAN Volume Controller cluster into the map. Each SAN Volume Controller node port WWPN is recognized as a separate host by XIV Nextra systems.

LUN presentation on IBM XIV Type Number 2810 systems

IBM XIV Storage System Type Number 2810 LUNs are presented to the SAN Volume Controller interface using the following rules:

- LUNs can be presented to one or more selected hosts or clusters.
- Clusters are collections of hosts.

To present IBM XIV Storage System Type Number 2810 LUNs to the SAN Volume Controller, perform the following steps:

1. Use the IBM XIV Storage System GUI to create an IBM XIV Storage System cluster for the SAN Volume Controller cluster.
2. Create a host for each node in the SAN Volume Controller.
3. Add a port to each host that you created in step 2. You must add a port for each port on the corresponding node.
4. Map volumes to the cluster that you created in step 1.

Target ports on XIV Nextra systems

XIV Nextra systems are single-rack systems. All XIV Nextra WWNNs include zeros as the last two hexadecimal digits. In the following example, `WWNN 2000001738279E00` is IEEE extended; the WWNNs that start with the number 1 are IEEE 48 bit:

```
WWNN 2000001738279E00
WWPN 1000001738279E13
WWPN 1000001738279E10
WWPN 1000001738279E11
WWPN 1000001738279E12
```

Target ports on IBM XIV Type Number 2810 systems

IBM XIV Storage System Type Number 2810 systems are multi-rack systems, but only single racks are supported. All IBM XIV Storage System Type Number 2810 WWNNs include zeros as the last four hexadecimal digits. For example:

```
WWNN 5001738000030000
WWPN 5001738000030153
WWPN 5001738000030121
```

Switch zoning limitations for IBM XIV systems

There are limitations in switch zoning for SAN Volume Controller clusters and IBM XIV Storage System systems.

Fabric zoning

To avoid a single point of failure, the SAN Volume Controller switch zone must include at least one target port from each IBM XIV Storage System controller.

Target port sharing

Target ports can be shared between the SAN Volume Controller cluster and other hosts.

Host splitting

To avoid the possibility of an interaction between multipathing drivers, a single host cannot be connected to both the SAN Volume Controller cluster and the IBM XIV Storage System system.

Controller splitting

IBM XIV Storage System system LUNs that are mapped to the SAN Volume Controller cluster cannot be mapped to other hosts. IBM XIV Storage System system LUNs that are *not* mapped to the SAN Volume Controller cluster can be mapped to other hosts.

Configuration settings for IBM XIV systems

The IBM XIV Storage System Storage Management GUI provides configuration settings and options that can be used with SAN Volume Controller clusters.

Logical unit options and settings for IBM XIV systems

Logical unit (LU) settings for IBM XIV Storage System systems are configurable at the LU level.

Table 65 lists the options that must be set for each LU that is accessed by other hosts. LUs that are accessed by hosts can be configured differently. You can use the IBM XIV Storage System and XIV Nextra Storage Management GUI or CLI to change these settings.

Table 65. IBM XIV options and required settings

Option	Data Type	Range	IBM XIV Storage System and XIV Nextra default setting	SAN Volume Controller setting
Capacity	int	17,179,869,184 bytes (17 GB), up to the total system capacity ORBlock count	None	Any
Notes: • SAN Volume Controller supports up to 2 TB. • LUNs are allocated in 17-GB chunks. • Using a block count results in LUNs that are arbitrarily sized, but that still consume multiples of 17 GB.				

Host options and settings for IBM XIV systems

You must use specific settings to identify SAN Volume Controller clusters as hosts to IBM XIV Storage System systems.

An XIV Nextra host is a single WWPN, so one XIV Nextra host must be defined for each SAN Volume Controller node port in the cluster. An XIV Nextra host is considered to be a single SCSI initiator. Up to 256 XIV Nextra hosts can be presented to each port. Each SAN Volume Controller host object that is associated with the XIV Nextra system must be associated with the same XIV Nextra LUN map because each LU can only be in a single map.

An IBM XIV Storage System Type Number 2810 host can consist of more than one WWPN. Configure each SAN Volume Controller node as an IBM XIV Storage System Type Number 2810 host, and create a cluster of IBM XIV Storage System systems that corresponds to each of the SAN Volume Controller nodes in the SAN Volume Controller cluster.

Table 66 lists the host options and settings that can be changed using the IBM XIV Storage System and XIV Nextra Storage Management GUI.

Table 66. IBM XIV Type Number 2810 and XIV Nextra host options and required settings

Option	Data type	Range	IBM XIV Storage System Type Number 2810 and XIV Nextra default setting	SAN Volume Controller required setting	Notes
Type	Enum	FC/iSCSI	Not applicable	FC	The Type must be FC for SAN Volume Controller.
XIV Nextra map_set_special_type CLI command or IBM XIV Storage System Type Number 2810 special_type_set CLI command	Enum	default/hpux	default	default	This command is used by hpux hosts only. Do not use the command for SAN Volume Controller LUNs.
WWPN	int64	Any	Not applicable	Node port	For XIV Nextra, one host for each node port WWPN in the cluster must be defined. For IBM XIV Storage System, Type Number 2810 one host port for each node port WWPN in the cluster must be defined.
LUN Map	String	Any	Not applicable	Any	For XIV Nextra, each SAN Volume Controller host in the XIV Nextra system must be associated with the same LUN map because each LU can be only in a single map.

Quorum disks on IBM XIV systems

The SAN Volume Controller cluster selects managed disks (MDisks) that are presented by IBM XIV Storage System systems as quorum disks.

Clearing SCSI reservations and registrations

You must not use the **vol_clear_keys** command to clear SCSI reservations and registrations on volumes that are managed by SAN Volume Controller.

Copy functions for IBM XIV Storage System models

Advanced copy functions for IBM XIV Storage System models such as taking a snapshot and remote mirroring cannot be used with disks that are managed by the SAN Volume Controller cluster. Thin provisioning is not supported for use with SAN Volume Controller.

Chapter 12. IBM System Storage support for Microsoft Volume Shadow Copy Service and Virtual Disk Service for Windows

The SAN Volume Controller provides support for the Microsoft Volume Shadow Copy Service and Virtual Disk 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 Microsoft Virtual Disk Service provides a single vendor and technology-neutral interface for managing block storage virtualization, whether done by operating system software, RAID storage hardware, or other storage virtualization engines.

The following components are used to provide support for the service:

- SAN Volume Controller
- The cluster CIM server
- IBM System Storage hardware provider, known as the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software
- Microsoft Volume Shadow Copy Service

The IBM System Storage 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 System Storage hardware provider that a copy is needed.
3. The SAN Volume Controller prepares the volumes for a snapshot.
4. The Volume Shadow Copy Service quiesces 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 creates the shadow copy using the FlashCopy Copy Service.
6. The Volume Shadow Copy Service notifies the writing applications that I/O operations can resume, and 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 on the SAN Volume Controller.

Installation overview

The steps for implementing the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software 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 also have experience with or knowledge of administering a SAN Volume Controller.

Complete the following tasks:

1. Verify that the system requirements are met.
2. Install the SAN Volume Controller Console if it is not already installed.
3. Install the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software.
4. Verify the installation.
5. Create a free pool of volumes and a reserved pool of volumes on the SAN Volume Controller.
6. Optionally, you reconfigure the services to change the configuration that you established during the installation.

System requirements for the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software

Ensure that your system satisfies the following requirements before you install the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software on a Microsoft Windows Server 2003 or 2008 operating system.

The following software is required:

- SAN Volume Controller must have licenses for FlashCopy enabled.
- IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software version 4.0 or later.
- Windows Server 2003 R2 or later or Windows Server 2008 operating system. The following editions of Windows Server 2003 and Windows Server 2008 are supported:
  - Standard Server Edition 32-bit version
  - Enterprise Edition, 32-bit version
  - Standard Server Edition 64-bit version
  - Enterprise Edition, 64-bit version

Installing the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software

This section includes the steps to install the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software on a Windows server.

You must satisfy all of the prerequisites that are listed in the system requirements section before starting the installation.

Perform the following steps to install the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software on the Windows server:

1. Log on to Windows as an administrator.
2. Download the IBM VSS Host Installation Package file from the following Web site:

   www.ibm.com/storage/support/2145

3. Double click on the name of the file that you downloaded in step 2 on page 474 to start the installation process. The Welcome panel is displayed.
4. Click **Next** to continue. The License Agreement panel is displayed. You can click **Cancel** at any time to exit the installation. To move back to previous screens while using the wizard, click **Back**.
5. 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. The Choose Destination Location panel is displayed.
6. 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**. The Ready to Install the Program panel is displayed.
7. Click **Install** to begin the installation. To exit the wizard and end the installation, click **Cancel**. The Setup Status panel is displayed.

   The program setup verifies your configuration.

   The Select CIM Server panel is displayed.
8. Select the required CIM server, or select **Enter the CIM Server address manually**, and click **Next**. The Enter CIM Server Details panel is displayed.
9. Enter the following information in the fields:
   - In the **CIM Server Address** field, type the name of the server where the SAN Volume Controller is installed. For example, enter `https://server.company.com:5989`.
   - In the **CIM User** field, type the user name for the SAN Volume Controller that the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software will use to gain access to the CIM server.
   - In the **CIM Password** field, type the password for the user name that the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software will use to gain access to the CIM server and click **Next**.

   Notes:

   a. If these settings change after installation, you can use the **ibmvcfg.exe** tool to update Microsoft Volume Shadow Copy and Virtual Disk Services software with the new settings.
   b. If you do not have the CIM agent server, port, or user information, contact your CIM agent administrator.

   The InstallShield Wizard Complete panel is displayed.
10. Click **Finish**. If necessary, the InstallShield Wizard prompts you to restart the system.
11. Make the IBM Hardware Provider for VSS-VDS aware of the SAN Volume Controller, as follows:
    a. Open a command prompt.
    b. Change directories to the hardware provider directory; the default directory is **C:\Program Files\IBM\Hardware Provider for VSS-VDS**.
    c. Use the ibmvcfg command to set the IP address for the SAN Volume Controller cluster, as follows:

       ```
       ibmvcfg set targetSVC ip_address
       ```

       The *ip_address* value must be the SAN Volume Controller cluster IP address.

Creating the free and reserved pools of volumes

The IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software maintains a free and a reserved pool of volumes. Because these objects do not exist on the SAN Volume Controller, the free and reserved pool of volumes are implemented as virtual host systems. You must define these two virtual host systems on the SAN Volume Controller.

When a shadow copy is created, the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software 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 virtual disks (VDisks) mapped to the free pool. The VDisks must be the same size as the source VDisks.

Use the SAN Volume Controller Console or the SAN Volume Controller command-line interface (CLI) to perform the following steps:

1. Create a host for the free pool of VDisks.
   - You can use the default name VSS_FREE or specify a different name.
   - Associate the host with the worldwide port name (WWPN) 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 WWPN 5000000000000001 (14 zeroes).
3. Map the logical units (VDisks) to the free pool of volumes.

   Restriction: The VDisks cannot be mapped to any other hosts.
   - If you already have VDisks created for the free pool of volumes, you must assign the VDisks to the free pool.
4. Create VDisk-to-host mappings between the VDisks selected in step 3 and the VSS_FREE host to add the VDisks to the free pool. Alternatively, you can use the **ibmvcfg add** command to add VDisks to the free pool.
5. Verify that the VDisks have been mapped.

If you do not use the default WWPNs 5000000000000000 and 5000000000000001, you must configure the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software with the WWPNs.

Verifying the installation

This task verifies that the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software is correctly installed on the Windows server.

Perform the following steps to verify the installation:

1. Click **Start** → **All Programs** → **Administrative Tools** → **Services** from the Windows server task bar. The **Services** panel is displayed.
2. Ensure that the service named IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software appears and that **Status** is set to Started and **Startup Type** is set to Automatic.
3. Open a command prompt window and issue the following command:

```
vssadmin list providers
```

4. Ensure that the service named IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software is listed as a provider.
5. Use the **ibmvcfg listvols** command to test the connection to the IBM System Storage Productivity Center or master console.

If you are able to successfully perform all of these verification tasks, the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software was successfully installed on the Windows server.

Changing the configuration parameters

You can change the parameters that you defined when you installed the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software. You must use the ibmvcfg.exe utility to change the parameters.

Table 67 describes the configuration commands that are provided by the ibmvcfg.exe utility.

Table 67. Configuration commands

Command	Description	Example
ibmvcfg showcfg	Lists the current settings.	`showcfg`
ibmvcfg set username <username>	Sets the user name to access the CIM server.	`set username johnny`
ibmvcfg set password <password>	Sets the password of the user name that will access the CIM server.	`set password mypassword`
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`
ibmvcfg set backgroundCopy	Sets the background copy rate for FlashCopy.	`set backgroundCopy 80`
ibmvcfg set incrementalFC	Specifies if incremental FlashCopy has to be used on SAN Volume Controller for the shadow copy.	`ibmvcfg set incrementalFC yes`
ibmvcfg set usingSSL	Specifying the Secure Sockets Layer (SSL) protocol is required to use a CIM server.	`ibmvcfg set usingSSL yes`
ibmvcfg set cimomHost <server name>	Sets the CIM server for the cluster.	`ibmvcfg set cimomHost cimomserver`

Table 67. Configuration commands (continued)

Command	Description	Example
ibmvcfg set namespace <namespace>	Specifies the namespace value that master console is using.	`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, removing, or listing volumes and FlashCopy relationships

You can use the ibmvcfg.exe. utility to perform the pool management tasks of adding, removing, or listing volumes and FlashCopy relationships.

The IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software 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 68 describes the ibmvcfg.exe commands for adding or removing volumes from the free pool of volumes and listing or deleting FlashCopy relationships.

Table 68. Pool management commands

Command	Description	Example
ibmvcfg list all -l	Lists all VDisks, including information about volume ID, UUID, volume name, size, operational state, health status, type of volume, VDisks to host mappings, and host name. Use the **l** parameter for output in verbose-list column format.	`ibmvcfg list all` `ibmvcfg list all -l`
ibmvcfg list free -l	Lists the volumes that are currently in the free pool. Use the **l** parameter for output in verbose-list column format.	`ibmvcfg list free` `ibmvcfg list free -l`
ibmvcfg list reserved -l	Lists the volumes that are currently in the reserved pool. Use the **l** parameter for output in verbose-list column format.	`ibmvcfg list reserved` `ibmvcfg list reserved -l`

Table 68. Pool management commands (continued)

Command	Description	Example
ibmvcfg list assigned -l	Lists the volumes that are currently in the assigned pool or host. Use the **l** parameter for output in verbose-list column format.	`ibmvcfg list assigned` `ibmvcfg list assigned -l`
ibmvcfg list unassigned -l	Lists the volumes that are currently in the unassigned pool or host. Use the **l** parameter for output in verbose-list column format.	`ibmvcfg list unassigned` `ibmvcfg list unassigned -l`
ibmvcfg list infc -l	Lists all the FlashCopy relationships on the SAN Volume Controller. This command lists both incremental and nonincremental FlashCopy relationships.	`ibmvcfg list infc` `ibmvcfg list infc -l`
ibmvcfg add	Adds one or more volumes to the free pool of volumes.	`ibmvcfg add 13036511188` `ibmvcfg add vdisk18`
ibmvcfg rem	Removes one or more volumes from the free pool of volumes.	`ibmvcfg rem 13036511188` `ibmvcfg rem vdisk18`
ibmvcfg del	Deletes one or more FlashCopy relationships. Use the serial number of the FlashCopy target to delete the relationship.	`ibmvcfg del vdisk18`

Error codes for IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software

The IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software logs error messages in the Windows Event Viewer and in private log files.

You can view error messages by going to the following locations on the Windows server where the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software 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 System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software is installed.

Table 69 lists the errors messages that are reported by the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software.

Table 69. Error messages for the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software

Code	Message	Symbolic name
1000	JVM Creation failed.	ERR_JVM

Table 69. Error messages for the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software (continued)

Code	Message	Symbolic name
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
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

Table 69. Error messages for the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software (continued)

Code	Message	Symbolic name
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 System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software

You must use Windows to uninstall the IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software from the Windows server.

Perform the following steps to uninstall the software:

1. Log on to the Windows server as the local administrator.
2. Click **Start** → **Control Panel** from the task bar. The Control Panel window is displayed.
3. Double-click **Add or Remove Programs**. The Add or Remove Programs window is displayed.
4. Select **IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software** and click **Remove**.
5. Click **Yes** when you are prompted to verify that you want to completely remove the program and all of its components.
6. Click **Finish**.

The IBM System Storage Support for Microsoft Volume Shadow Copy Service and Virtual Disk Service software is no longer installed on the Windows server.

Appendix. Accessibility

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

Features

These are the major accessibility features in the SAN Volume Controller Console:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. The following screen reader has been tested: Window-Eyes v6.1.
- You can operate all features using the keyboard instead of the mouse.
- When setting or changing an IP address on the SAN Volume Controller front panel, you can disable the fast increase and decrease address scrolling speed function of the up and down buttons to two seconds. This feature is documented in the topic that discusses initiating cluster creation from the front panel, which is located in the IBM System Storage SAN Volume Controller Information Center and the *IBM System Storage SAN Volume Controller Software Installation and Configuration Guide*.

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 find the HTML version of the IBM System Storage SAN Volume Controller information at the following Web site:

http://publib.boulder.ibm.com/infocenter/svcic/v3r1m0/index.jsp

You can access this information using screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. JAWS version 10 has been tested.

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.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
3-2-12, Roppongi, Minato-ku, Tokyo 106-8711 Japan

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.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Almaden Research
650 Harry Road
Bldg 80, D3-304, Department 277
San Jose, CA 95120-6099
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

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 may 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.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application

programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided ″AS IS″, without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at Copyright and trademark information at www.ibm.com/legal/copytrade.shtml.

Adobe and the Adobe logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel, Intel logo, Intel Xeon, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo 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.

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

Index

Numerics

A

B

C

D

E

F

G

H

I

K

L

M

N

O

P

Q

R

S

T

U

V

W

X

Z

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

IBM System Storage SAN Volume Controller
Software Installation and Configuration Guide
Version 5.1.0

Publication No. SC23-6628-05

We appreciate your comments about this publication. Please comment on specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. The comments you send should pertain to only the information in this manual or product and the way in which the information is presented.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer.

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. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Comments:

Thank you for your support.

Send your comments to the address on the reverse side of this form.

If you would like a response from IBM, please fill in the following information:

Name	Address
Company or Organization	
Phone No.	E-mail address

Readers' Comments — We'd Like to Hear from You
SC23-6628-05

Cut or Fold
Along Line

Fold and Tape | **Please do not staple** | Fold and Tape

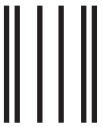

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

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

Cut or Fold
Along Line

Part Number: 31P1344

Printed in USA

(1P) P/N: 31P1344

SC23-6628-05

Spine information:

IBM System Storage SAN Volume Controller

SAN Volume Controller Software Installation and Configuration Guide

Version 5.1.0