IBM TotalStorage SAN Volume Controller

Command-Line Interface User's Guide - Errata

Version 2.1.0 - Errata

SC26-7544-04

IBM TotalStorage SAN Volume Controller

Command-Line Interface User's Guide - Errata

Version 2.1.0 - Errata

SC26-7544-04

Fifth Edition (February 2005)

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

© Copyright International Business Machines Corporation 2003, 2005. All rights reserved.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

© Copyright IBM Corp. 2003, 2005

About this guide

This guide introduces the IBM® TotalStorage® SAN Volume Controller Command-Line Interface (CLI).

Who should use this guide

This guide is intended for system administrators or others who install and use the SAN Volume Controller.

Summary of Changes

This document contains terminology, maintenance, and editorial changes.

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

Summary of Changes for SC26-7544-04 SAN Volume Controller Command-Line Interface Guide

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

New information

This topic describes the changes to this guide since the previous edition, SC26-7544-03 SAN Volume Controller Command-Line Interface Guide. The following sections summarize the changes that have since been implemented from the previous version.

- The migratetoimage command was added. You can use migratetoimage to migrate an image mode vdisk to a target mdisk within the same mdisk group.
- The following messages are new to this release:
  - CMMVC3438E
  - CMMVC3439E
  - CMMVC5432E
  - CMMVC6166E
  - CMMVC6181E
  - CMMVC6182E
  - CMMVC6183E
  - CMMVC6186E

Changed information

- The filename maximum length is now defined at 169 characters.
- The filename filter maximum length was changed to allow 128 characters
- The filename prefix maximum length was changed to allow 128 characters.
- The -type parameter was added to the chhost and mkhost commands.
- The mdisk parameter was added to the mkvdisk command.

© Copyright IBM Corp. 2003, 2005

Summary of Changes for SC26-7544-03 SAN Volume Controller Command-Line Interface Guide

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

New information

This topic describes the changes to this guide since the previous edition, SC26-7544-02 SAN Volume Controller Command-Line Interface Guide. The following sections summarize the changes that have since been implemented from the previous version.

This section lists the new information added to this book.

- A new topic, PuTTy scp was added to this book.

Changed information

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

- Descriptions were changed for **rmvdisk**. **rmvdisk** is found under Virtual Disk Commands.
- Descriptions were changed for **lsvdiskextent** and **lsmdiskextent**. Both commands are found under Information Commands.
- An explanation of the **-ignore** flag was added to the **applysoftware** command section under Service Mode Information Commands.
- The Windows NT topic that was removed in the previous version has been restored.

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

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

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

- The publications that make up the library for the IBM TotalStorage SAN Volume Controller
- Other IBM publications that relate to the SAN Volume Controller

SAN Volume Controller library

The following table lists and describes the publications that make up the SAN Volume Controller library. Unless otherwise noted, these publications are available in Adobe portable document format (PDF) from the following Web site:

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

Title	Description	Order number
IBM TotalStorage SAN Volume Controller: CIM Agent Developer's Reference	This reference guide describes the objects and classes in a Common Information Model (CIM) environment.	SC26-7590
IBM TotalStorage SAN Volume Controller: Command-Line Interface User's Guide	This guide describes the commands that you can use from the SAN Volume Controller command-line interface (CLI).	SC26-7544
IBM TotalStorage SAN Volume Controller: Configuration Guide	This guide provides guidelines for configuring your SAN Volume Controller.	SC26-7543
IBM TotalStorage SAN Volume Controller: Host Attachment Guide	This guide provides guidelines for attaching the SAN Volume Controller to your host system.	SC26-7575
IBM TotalStorage SAN Volume Controller: Installation Guide	This guide includes the instructions the service representative uses to install the SAN Volume Controller.	SC26-7541
IBM TotalStorage SAN Volume Controller: Planning Guide	This guide introduces the SAN Volume Controller and lists the features you can order. It also provides guidelines for planning the installation and configuration of the SAN Volume Controller.	GA22-1052
IBM TotalStorage SAN Volume Controller: Service Guide	This guide includes the instructions the service representative uses to service the SAN Volume Controller.	SC26-7542
IBM TotalStorage SAN Volume Controller: Translated Safety Notices	This guide contains the danger and caution notices for the SAN Volume Controller. The notices are shown in English and in numerous other languages.	SC26-7577

Title	Description	Order number
IBM TotalStorage Master Console Installation and User's Guide	This guide includes the instructions on how to install and use the SAN Volume Controller Console	GC30-4090

Other IBM publications

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

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

Related Web sites

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

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

How to order IBM publications

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

The IBM publications center

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

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

Publications notification system

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

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

http://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.

Syntax diagrams

A syntax diagram uses symbols to represent the elements of a command and to specify the rules for using these elements.

This explains how to read the syntax diagrams that represent the command–line interface (CLI) commands. In doing so, it defines the symbols that represent the CLI command elements.

Element	Syntax	Description
Main path line	>>><>() () ()	>>Begins on the left with double arrowheads ()>< and ends on the right with two arrowheads facing each other (). If a diagram is longer than one line, each line to be continued ends with a single> arrowhead () and the next line begins with a single arrowhead. Read the diagrams from left–to–right, top–to–bottom, following the main path line.
Keyword		Represents the name of a command, flag, parameter, or argument. A keyword is not in italics. Spell a keyword exactly as it is shown in the syntax diagram.
Required keywords	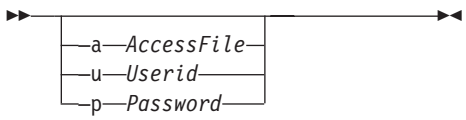	Indicate the parameters or arguments you must specify for the command. Required keywords appear on the main path line. Mutually exclusive required keywords are stacked vertically.
Optional keywords		Indicate the parameters or arguments you can choose to specify for the command. Optional keywords appear below the main path line. Mutually exclusive optional keywords are stacked vertically.
Default value	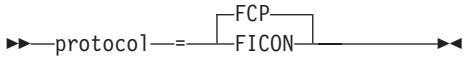	Appears above the main path line.
Repeatable keyword or value	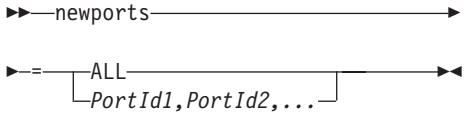	Represents a parameter or argument that you can specify more than once. A repeatable keyword or value is represented by an arrow returning to the left above the keyword or value.
Variable		Represents the value you need to supply for a parameter or argument, such as a file name, user name, or password. Variables are in italics.

Element	Syntax	Description
Space separator	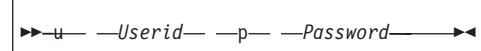	Adds a blank space on the main path line to separate keywords, parameters, arguments, or variables from each other.
Quotation mark delimiters	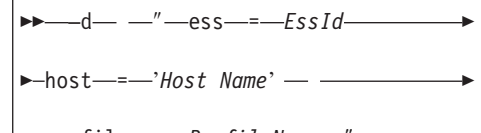	Indicates the start and end of a parameter or argument that contains multiple values. Enclose one or more name–value pairs in a set of double quotation marks for a particular parameter or argument. If the value of a parameter or name–value pair contains a blank or white space, enclose the entire value in a set of single quotation marks.
Equal–sign operator	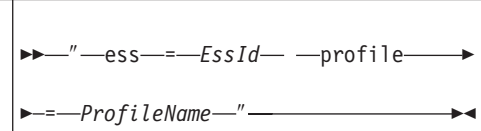	Separates a name from its value in a name–value pair.
Syntax fragment	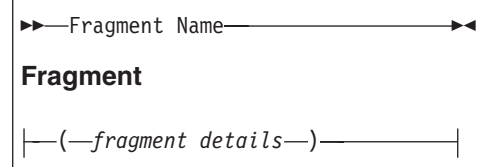	Breaks up syntax diagrams that are too long, too complex, or repetitious. The fragment name is inserted in the main diagram, and the actual fragment is shown below the main diagram.

Terminology

These are abbreviations that are most commonly used for the command-line interface operations.

The table below shows the abbreviations that are most commonly used for the command-line interface operations.

Name	Object type
Host	host
Virtual disk	vdisk
Managed disk	mdisk
Managed disk group	mdiskgrp
I/O group	iogrp
Node	node
Cluster	cluster
Controller	controller
FlashCopy mapping	fcmap
FlashCopy consistency group	fcconsistgrp

Name	Object type
IBM TotalStorage Metro Mirror (previously known as Remote Copy) relationship	rcrelationship
IBM TotalStorage Metro mirror consistency group	rcconsistgrp
Unsupported/unknown object	unknown

CLI special characters

The following special characters are used in the command-line interface (CLI) command examples.

- **-** (minus) sign. Flags are prefixed with a **-** (minus) sign. Flags define the action of a command or modify the operation of a command. You can use multiple flags, followed by parameters, when you issue a command. The - character cannot be used as the first character of an object name.
- **|** vertical bar. A vertical bar signifies that you choose only one value. For example, **[a | b]** indicates that you can choose a, b, or nothing. Similarly, **{ a | b }** indicates that you must choose either a or b.

Using wildcards in the SAN Volume Controller Command-Line Interface (CLI)

You can use wildcards in the SAN Volume Controller Command-Line Interface.

The SAN Volume Controller allows the use of the '*' as a wildcard within the arguments of certain parameters. There are some behavioral issues that must be considered when using wildcards in order to prevent unexpected results. These behavioral issues, and the ways to avoid them, are described below.

1. Running the command while logged onto the node.

   The shell will attempt to interpret any of the special characters if they are not escaped. Wildcards will be expanded into a list of files if any files exist that match the wildcards. If no matching files exist, the wildcard is passed to the SAN Volume Controller command untouched.

   To prevent expansion, issue the following command in one of its formats:

   `svctask cleardumps -prefix '/dumps/*.txt'` in single quotes, or

   `svctask cleardumps -prefix /dumps/\*.txt` using a backslash, or

   `svctask cleardumps -prefix "/dumps/*.txt"` in double quotes.

2. Running the command through SSH, for example from a host.

   This is slightly more complicated since the host shell will process the command line before it is passed through SSH to the shell on the cluster. This means an extra layer of protection is required around the wildcard as the host shell will strip off any protecting quotes and if the wildcard is exposed to the cluster shell, then this will result in the wildcard being expanded in the cluster shell.

   To prevent expansion, issue the following command in one of its formats:

   `svctask cleardumps "'/dumps/*.txt'"` with single quotes inside double quotes, or

   `svctask cleardumps '/dumps/\*.txt'` using a backslash inside single quotes, or

```
svctask cleardumps '"/dumps/*.txt"'
```

with double quotes inside single quotes.

Data types and value ranges

Each data type has a specified value range.

Note: If you do not specify a name when creating a new object, the cluster will assign a default name. This name is made from the object type as a prefix and the object ID as the suffix. For example, a new virtual disk (VDisk) is created with ID 5. This object will be given the default name of vdisk5. Because the system assigns these names, it will not allow you to create an object and call it vdiskx where *x* is the integer. This is because the cluster reserves these names (for example, object_type_prefix integer) for default.

Data types	Value ranges
filename_arg	This is a (optionally fully qualified) file name. Maximum length is 169 characters. Valid characters consist of the following: • . • / • - • _ • a - z • A - Z • 0 - 9 The field may not contain two consecutive '.', or start with a '.', or end with a '.'.

Data types	Value ranges
directory_or_file_filter	Specifies a directory and or filename filter within the specified directory. Valid directory values consist of the following: • /dumps • /dumps/configs • /dumps/elogs • /dumps/feature • /dumps/iostats • /dumps/iotrace • /dumps/software The filename filter can be any valid filename with or without the wildcard '*'. The filename filter can be appended to the end of one of the above directory values. Maximum length is 128 characters. Valid characters consist of the following: • * • . • / • - • _ • a - z • A - Z • 0 - 9 The field may not contain two consecutive '.', or start with a '.', or end with a '.'.
filename_prefix	This is a prefix to be used when naming a file. Maximum length is 128 characters. Valid characters consist of the following: • a - z • A - Z • 0 - 9 • - • _

Data types	Value ranges
name_arg	Names can be specified or changed using the create and modify functions. The view commands allow you to see both the name and ID of an object. A string of 1-15 characters is allowed, composed of characters A-Z, a-z, 0-9, - and _. The first character of a name_arg must not be numeric. The first character of an object name can not be a - as the CLI will interpret it as being the next parameter. When creating a name for an object, the name may not consist of the object type followed only by an integer. The exception is Metro Mirror relationships which can be named anything as long as the names are unique across the two clusters. This naming convention is used by the system to generate default names. You can not use one of the following reserved words followed by an integer: • cluster • controller • fccstgrp • fcmap • host • io_grp • mdisk • mdiskgrp • node • rccstgrp • rcmap The cluster name is set when the cluster is created.
password	This is a user defined password. A password must meet the following requirements: • may use a - z, A - Z, 0 - 9 in any sequence • may use - (dash) but not as the first character • may use _ (underscore) • may contain a maximum of 15 characters
serial_number	The format of this number conforms to IBM standard C-S 1-1121-018 1999-06 Serial Numbering for IBM products. The serial number is 7 digits, the first two of which define the manufacturing location, leaving 5 digits for the product. The standard defines a way to extend the serial number using letters in the place of numbers in the 5 digit field.

Data types	Value ranges
ip_address_arg	The decimal, dotted quad notation, standard rules.
dns_name	Dotted domain name for the subnet that the cluster is in. For example, ibm.com.
hostname	The hostname assigned to the cluster. This can be different from the cluster name and you can change the hostname at any time. A combination of the hostname and the dns_name that is used to access the cluster, for example: https://hostname.ibm.com/
capacity_value	A value with a range of 512 bytes up to 2 PetaBytes. The value can be expressed in multiples of 1 MB, ranging from 16 MB to 2 PetaBytes (PB). **Note:** The capacity can be specified as MB, KB, GB, or PB. When MB is used, the value is specified in multiples of 512 bytes. A capacity of 0 is valid for a striped/sequential vdisk. The smallest number of supported bytes is 512.
delay_arg	Unassigned integer ranging from 1 to 65535 (minutes for battery test).
node_id	Node IDs differ from other IDs as they are a unique node ID that is assigned when the node is initialized. Node IDs are expressed as 64-bit hexadecimal numbers. For example: 1A2B30C67AFFE47B Node IDs, like other IDs, cannot be modified by user commands.

Data types	Value ranges
xxx_id	All objects are referred to by unique integer IDs that are assigned by the system when the objects are created. All IDs are represented internally as 32-bit integers. Node IDs are an exception. IDs in the following ranges are used to identify the various types of objects: • node_id: 1 - 32 • mdisk_grp_id: 0 - 127 • io_grp_id: 0 - 3 (See Note.) • mdisk_id: 0 - 4095 • vdisk_id: 0 - 8191 • host_id: 0 - 127 • flash_const_grp_id: 0 - 255 • remote_const_grp_id: 0 - 255 • fcmap_id: 0 - 4095 • rcrel_id: 0 - 8191 • controller_id: 0–63 **Note:** io_group 4 exists but is only used in certain error recovery procedures. These IDs, like node IDs, cannot be modified by user commands. **Note:** IDs are assigned at run-time by the system and cannot be relied upon to be the same after, for example, the configuration restoration. Therefore, wherever possible, object names should be used in preference to IDs when working with objects.
xxx_list	A colon-delimited list of values of type *xxx*.
wwpn_arg	The Fibre Channel World Wide Port Name (wwpn). This is expressed as a 64-bit hexadecimal number, for example: 1A2B30C67AFFE47B These numbers must be composed of the characters 0 - 9, a - f, and A - F. A command will fail if you enter WWPN 0 in the command string.
panel_name	A string of up to 6 characters that correspond to the number on the printed label below the APA display on the front panel of a node in the cluster.
sequence_number	32-bit unsigned integer, expressed in decimal.
csi_num_arg	32-bit unsigned integer, expressed in decimal.
percentage_arg	8-bit unsigned integer, expressed in decimal 0 to 100.
extent_arg	32-bit unsigned integer, expressed in decimal.

Data types	Value ranges
num_extents_arg	32-bit unsigned integer, expressed in decimal.
threads_arg	8-bit unsigned integer, expressed in decimal, valid values, 1, 2, 3, or 4.
velocity_arg	The fabric speed in Giga-bits per second. Valid values are 1 or 2.
timezone_arg	The ID as detailed in the output of the **svcinfo lstimezones** command.
timeout_arg	The command timeout period. An integer from 0 to 600 (seconds).
stats_time_arg	The frequency at which statistics are gathered. 15 up to a max of 60 (minutes) in increments of 1 minute.
directory_arg	Specifies a directory and or filename filter within the specified directory. Valid directory values are: • /dumps • /dumps/configs • /dumps/elogs • /dumps/feature • /dumps/iostats • /dumps/iotrace • /home/admin/upgrade The filename filter can be any valid filename with or without the wildcard '*'. The filename filter can be appended to the end of one of the above directory values.
locale_arg	The cluster locale setting. Valid values are 0 to 9. • 0 US English (default) • 1 Chinese (simplified) • 2 Chinese (traditional) • 3 Japanese • 4 Korean • 5 French • 6 German • 7 Italian • 8 Spanish • 9 Portuguese (Brazilian)
key_arg	A user definable identifier for an SSH key. A string of up to 30 characters.
user_arg	Specifies the user, either admin or service.
copy_rate	A numeric value from 0 to 100.

CLI parameters

CLI parameters are found within the syntax diagram.

CLI parameters can be entered in any order except:

1. The first argument following the command name must be the action that is to be performed.
2. Where you are performing an action against a specific object, the object ID or name must be the last argument in the line.

CLI flags

The flags that are listed below are common to all CLI commands.

- **-?** or **-h**. Print help text. For example, issuing **svcinfo lscluster -h** will provide a list of the actions available with the **svcinfo lscluster** command.
- -nomsg. When used, this flag will prevent the display of the `successfully created` output. For example, if you issue the following:

  ```
  svctask mkmdiskgrp -ext 16
  ```

  it will display:

  MDisk Group, id [6], successfully created

  However, if the -nomsg parameter had been added, for example:

  ```
  svctask mkmdiskgrp -ext 16 -nomsg
  ```

  then the following would have been displayed:

  ```
  6
  ```

  This parameter can be entered for any command, but is only acted upon by those commands that generate the `successfully created` outputs. All other commands will ignore this parameter.

Chapter 1. Virtual disk commands

The following commands enable you to work with virtual disk options with the SAN Volume Controller.

chvdisk

Use the **chvdisk** command to modify some of the properties of a virtual disk including the name, the I/O group, and the I/O governing rates.

Syntax

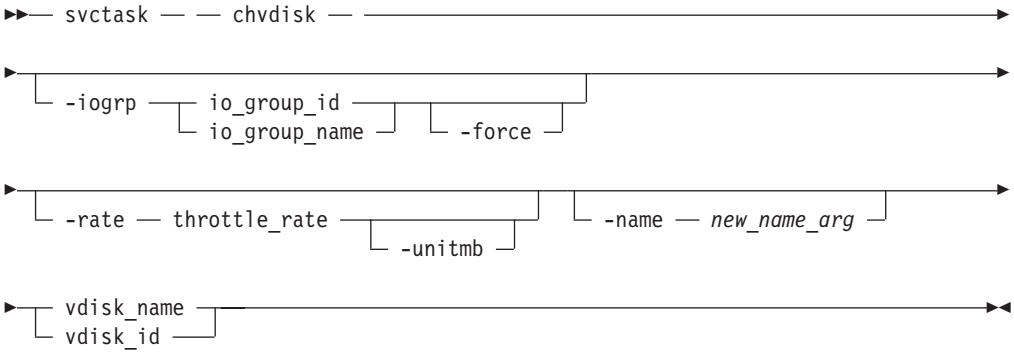

Parameters

-iogrp *io_group_id | io_group_name*
Optionally specifies a new I/O group to move the virtual disk to, either by ID or name. The -force flag can be used together with this parameter in order to force the removal of the VDisk to the I/O group.

-rate *throttle_rate [-unitmb]*
Optionally sets the I/O governing rates for the virtual disk. The default units are I/Os, but can be used in conjunction with the -unitmb argument to specify in terms of MBps.

-name *new_name_arg*
Optionally specifies a new name to assign to the virtual disk.

-force
Specifies that you want to force the VDisk to be removed from an I/O group. This parameter can only be used together with -iogrp.

CAUTION:
The -force flag can corrupt the contents of the VDisk. If the -force flag is used and if the SAN Volume Controller is unable to destage all write data from the cache, the result is that the contents of the VDisk are corrupted by the loss of the cached data.

vdisk_name | vdisk_id
Specifies the virtual disk to modify, either by ID or by name.

Note: The -iogrp, -rate, and -name parameters are mutually exclusive. Only one of these parameters can be specified per command line.

© Copyright IBM Corp. 2003, 2005

Description

This command modifies a single property of a virtual disk. You can modify one property at a time. Therefore, to change the name and modify the I/O group, you must issue the command twice.

You can specify a new name or label. You can use the new name subsequently to refer to the virtual disk.

You can change the I/O group with which this virtual disk is associated. However, to change the I/O group, you must first flush the cache within the nodes in the current I/O group to ensure that all data is written to disk. You should suspend I/O operations at the host level before performing this operation.

Attention: Do not move a VDisk to an offline I/O group. You must ensure the I/O group is online before moving the VDisks to avoid any data loss scenarios.

You can set a limit on the amount of I/O transactions that will be accepted for this virtual disk. It is set in terms of I/Os per second or MBps. By default, no I/O governing rate is set when a virtual disk is created.

Attention: All capacities, including changes should be in multiples of 512 bytes. An error will occur if you specify a capacity that is not a multiple of 512, which can only happen when byte units (-b) are used. However, an entire extent will be reserved even if only partially used. The default capacity is in MB.

When first created there is no throttling applied to a virtual disk. Using the -rate parameter can change this. To change the virtual disk back to an unthrottled state, the value 0 (zero) should be used with the -rate parameter.

You can migrate a VDisk to a new I/O group to manually balance the workload across the nodes in the cluster. You may end up with a pair of nodes that are overworked and another pair that are underworked. Follow this procedure to migrate a single VDisk to a new I/O group. Repeat for other VDisks as required.

Attention:

This is a disruptive procedure, access to the VDisk will be lost while you follow this procedure.

Under no circumstances should VDisks be moved to an offline I/O group. You must ensure the I/O group is online before moving the VDisks to avoid data loss scenarios.

Before migrating the VDisk, it is essential that for each vpath presented by the VDisk you intend to move, the subsystem device driver (SDD) configuration is updated to remove the vpaths in question. Failure to do this may result in data corruption. See *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* for details about how to dynamically reconfigure SDD for the given host operating system.

Make sure that when you migrate a VDisk to a new I/O group, you quiesce all I/O operations for the VDisk. You may need to determine the hosts that are using this VDisk. Any FlashCopy mappings or Metro Mirror relationships that use this VDisk should be stopped or deleted. Issue the following command, to check if the

VDisk is part of a relationship or mapping, issue the **svcinfo lsvdisk <vdiskname/id>** command, where <vdiskname/id> is the name or ID of the VDisk.

Look for the **FC_id** and **RC_id** fields. If these are not blank then the VDisk is part of a mapping or relationship. See "Managed disk commands" for details on how to stop or delete the mapping or relationship. Issue the following command to migrate the VDisk:

```
svctask chvdisk -iogrp <newiogrpname/id> <vdiskname/id>
```

Follow the procedure to discover the new vpaths and to check that each vpath is now presenting the correct number of paths. See the *IBM TotalStorage Multipath Subsystem Device Driver: User's Guide* for details on how to dynamically reconfigure SDD for the given host operating system.

Possible failures

- CMMVC5756E Cannot perform the request as the object is already mapped.
- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5832E The property of the virtual disk (VDisk) was not modified because an entity that was specified in the command does not exist.
- CMMVC5833E The property of the virtual disk (VDisk) was not modified because there are no nodes in the I/O group.
- CMMVC5834E The I/O group for the virtual disk (VDisk) was not modified because the group is a recovery I/O group. To modify the I/O group, use the force option.
- CMMVC5848E The action failed because the virtual disk (VDisk) does not exist or it is being deleted.
- CMMVC5853E The action failed because there was a problem with the group.
- CMMVC5856E The action failed because the virtual disk (VDisk) does not belong to the specified managed disk (MDisk) group.
- CMMVC5857E The action failed because the managed disk (MDisk) does not exist or it is not a member of the managed disk (MDisk) group.
- CMMVC5858E The action failed because the virtual disk (VDisk) is in the wrong mode, the managed disk (MDisk) is in the wrong mode, or both are in the wrong mode.
- CMMVC5860E The action failed because there were not enough extents in the managed disk (MDisk) group.
- CMMVC5861E The action failed because there were not enough extents on the managed disk (MDisk).
- CMMVC5862E The action failed because the virtual disk (VDisk) is being formatted.
- CMMVC6032E The operation was not performed because one or more of the entered parameters is invalid for this operation.

An invocation example

```
svctask chvdisk -rate 2040 -unitmb 6
```

The resulting output

```
No feedback
```

expandvdisksize

The **expandvdisksize** command expands the size of a VDisk by a given capacity.

Syntax

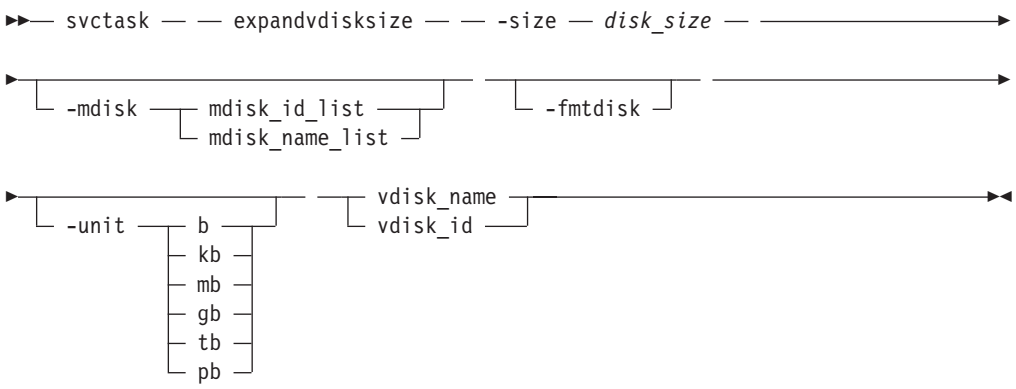

Parameters

-size *disk_size*
Specifies the capacity by which the virtual disk is to be expanded. It is used in conjunction with the units value. All capacities, including changes should be in multiples of 512 bytes. An error will occur if you specify a capacity that is not a multiple of 512, which can only happen when byte units (-b) are used. However, an entire extent will be reserved even if only partially used. The default capacity is in MB.

-mdisk *mdisk_id_list | mdisk_name_list*
Optionally specifies the list of one or more MDisks to be used as the stripe set. The extents that are used to expand the vdisk come from the specified list of MDisks. All MDisks in the list must be part of the same MDisk group.

-fmtdisk
Optionally specifies that the VDisk should be formatted before use. Using this parameter will format (all zeros) the new extents that have been added to the VDisk as a result of the expand.

-unit *b | kb | mb | gb | tb | pb*
Optionally specifies the data units to be used in conjunction with the capacity (-size).

vdisk_name | vdisk_id
Specifies the virtual disk to modify, either by ID or by name.

Description

This command expands the capacity allocated to the particular virtual disk by the given amount. The default capacity is in MB.

When a VDisk is expanded, the policy may change. Its mode will become striped even if it was previously sequential. See the **svctask mkvdisk** command for details of the virtualization policies. This command will complete asynchronously if the **fmtdisk** argument is selected.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5808E The action failed because the managed disk (MDisk) does not exist.
- CMMVC5835E The virtual disk (VDisk) was not expanded because an object that was specified in the command does not exist.
- CMMVC5837E The action failed because the virtual disk (VDisk) is part of a FlashCopy mapping.
- CMMVC5838E The action failed because the virtual disk (VDisk) is part of a Remote Copy mapping.
- CMMVC5848E The action failed because the virtual disk (VDisk) does not exist or it is being deleted.
- CMMVC5853E The action failed because there was a problem with the group.
- CMMVC5856E The action failed because the virtual disk (VDisk) does not belong to the specified managed disk (MDisk) group.
- CMMVC5857E The action failed because the managed disk (MDisk) does not exist or it is not a member of the managed disk (MDisk) group.
- CMMVC5858E The action failed because the virtual disk (VDisk) is in the wrong mode, the managed disk (MDisk) is in the wrong mode, or both are in the wrong mode.
- CMMVC5860E The action failed because there were not enough extents in the managed disk (MDisk) group.
- CMMVC5861E The action failed because there were not enough extents on the managed disk (MDisk).
- CMMVC5862E The action failed because the virtual disk (VDisk) is being formatted.
- CMMVC5998W The virtualized storage capacity exceeds the amount that you are licensed to use.

An invocation example

```
svctask expandvdisksize -size 2048 -unit b -mdisk
mdisk0:mdisk1 -fmtdisk vdisk1
```

The resulting output

```
No feedback
```

mkvdisk

The **mkvdisk** command creates sequential, striped or image mode virtual disk objects. Once mapped to a host object, these objects are seen as disk drives with which the host can perform I/O operations.

Note: The first syntax diagram depicts the creation of a sequential or striped mode virtual disk. The second syntax diagram depicts the creation of an image mode virtual disk.

Syntax

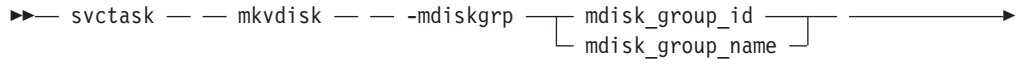

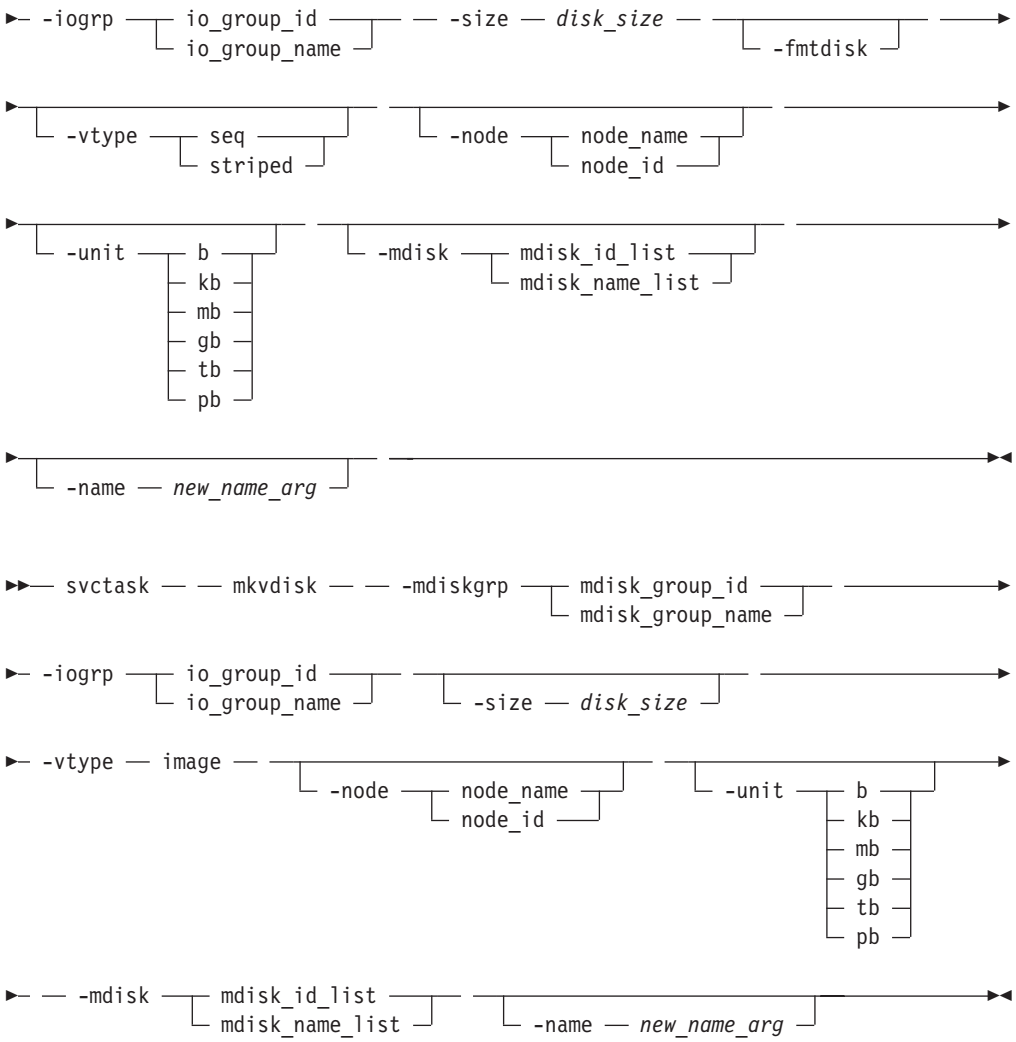

Parameters

-mdiskgrp *mdisk_group_id | mdisk_group_name*
Specifies the managed disk group to use when creating this virtual disk.

-iogrp *io_group_id | io_group_name*
Specifies the I/O group (node pair) with which to associate this virtual disk.

-size *disk_size*
Specifies the capacity of the virtual disk, which is used in conjunction with the units value. The smallest granularity of bytes is 512. All capacities, including changes should be in multiples of 512 bytes. An error will occur if you specify a capacity that is not a multiple of 512, which can only happen when byte units (-b) are used. However, an entire extent will be reserved even if only partially used. The default capacity is in MB. You can specify a capacity of 0. The size in bytes should be in multiples of logical block address (LBAs). When creating an image mode disk, if this parameter is not specified, then the entire managed disk capacity will be used.

-fmtdisk
Optionally specifies that the virtual disk should be formatted before use. The -fmtdisk argument formats (sets to all zeros) the extents that make up this VDisk after it is created. If this parameter is used, the command completes

asynchronously and you can query the status with the **svcinfo lsvdiskprogress** command. This flag may not be used when creating an image mode VDisk.

-vtype *seq* | *striped* | *image*
Optionally specifies the virtualization policy. The default virtualization type is striped. Refer to the notes below for more information.

-node *node_id* | *node_name*
Optionally specifies the preferred node ID or name for I/O operations to this virtual disk. You can use the -node argument to specify the preferred access node. This argument is required for the subsystem device driver (SDD) and the SAN Volume Controller will choose a default if you do not supply this argument.

-unit *b* | *kb* | *mb* | *gb* | *tb* | *pb*
Optionally specifies the data units to be used in conjunction with the capacity (-size).

-mdisk *mdisk_id_list* | *mdisk_name_list*
Specifies a list of one or more managed disks. This argument is used in conjunction with -vtype and has different uses depending upon the policy chosen. Refer to the notes below for more information.

-name *new_name_arg*
Optionally specifies a name to assign to the new virtual disk.

Description

This command creates a new virtual disk object. You can use the command to create a variety of types of virtual disk objects and, as such, it is one of the most complex commands.

You will need to decide which managed disk group will provide the storage for the VDisk. Use the **svcinfo lsmdiskgrp** command to list the available managed disk groups and the amount of free storage in each group.

Decide which I/O group the VDisk should be assigned to. This determines which SAN Volume Controller nodes in the cluster process the I/O requests from the host systems. If you have more than one I/O group then make sure you distribute the VDisks between the I/O groups so that the I/O workload is shared evenly between all SAN Volume Controller nodes. Use the **svcinfo lsiogrp** command to show the I/O groups and the number of virtual disks assigned to each I/O group.

Note: It is normal for clusters with more than one I/O group to have MDisk groups that have VDisks in different I/O groups. FlashCopy can be used to make copies of VDisks regardless of whether the source and destination VDisk are in the same I/O group. If however you plan to use intra-cluster Metro Mirror then make sure that both the master and auxiliary VDisk are in the same I/O group.

The virtualization policy controls the type of virtual disk to create. These policies include striped and seq and image:

Striped
This is the default policy. If the -vtype is not specified, then this policy is used in its default form. That is, all managed disks in the managed disk group will be used to create the virtual disk. The striping is at an extent level, in a circular fashion, one extent from each managed disk in the

group is used. For example, a managed disk group with 10 managed disks uses one extent from each managed disk, then it uses the 11th extent from the first managed disk, and so on.

If the -mdisk argument is also specified, you can supply a list of managed disks to use as the stripe set. This can be two or more managed disks from the same managed disk group. The same circular algorithm is used across the striped set. However, a single managed disk can be specified more than once in the list. For example, if you enter `-m 0:1:2:1` from the extents will be from the following maintenance disks: 0, 1, 2, 1, 0, 1, 2, and so forth. All MDisks specified in the -mdisk argument must be in the managed mode.

A capacity of 0 is allowed.

Seq (Sequential)
This policy requires the -mdisk flag with a single managed disk as its argument. This MDisk must be in the managed mode.

It will create the virtual disk only using extents from the given managed disk (assuming there are enough free extents on the managed disk).

Image Image mode virtual disks can be used when a managed disk already has data on it, perhaps from a previrtualized subsystem. When an image mode virtual disk is created, it directly corresponds to the (previously unmanaged) managed disk it was created from, therefore, virtual disk logical block address (LBA) x equals managed disk LBA x. This command can be used to bring a nonvirtualized disk under control of the cluster. The data can then be migrated from the single managed disk at which time the virtual disk is no longer an image mode virtual disk.

You may add image mode VDisks to an already populated mdiskgrp with other types of VDisks, such as a striped or sequential VDisk.

Note: An image mode VDisk must be at least 512 bytes (capacity can not be 0). That is, the minimum size that can be specified for an image mode VDisk should be the same as the MDisk group extent size that it will be added to, with the minimum being 16Mb.

The -mdisk flag must be used to specify an MDisk that has a mode of unmanaged. The -fmtdisk flag can not be used when creating an image mode VDisk.

The command returns the IDs of the newly created VDisk.

Attention: Do not create a VDisk in an offline I/O group. You must ensure the I/O group is online before creating a VDisk to avoid any data loss scenarios. This applies in particular to recreating VDisks that are assigned the same object ID.

CAUTION:
To create an image mode disk you must already have a quorum disk in the cluster, since an image mode disk cannot be used to hold quorum data. See "Creating a quorum disk" in the *IBM TotalStorage SAN Volume Controller: Configuration Guide* **for more details.**

Possible failures

Note: If you receive an error for this command that indicates that the licensed virtualization capacity has been exceeded, then the command was still effective. However, the return code will indicate the license violation.

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5807E The action failed because the managed disk (MDisk) cannot be changed to the specified mode.
- CMMVC5808E The action failed because the managed disk (MDisk) does not exist.
- CMMVC5826E The virtual disk (VDisk) was not created because an entity that was specified in the command does not exist.
- CMMVC5827E The command failed as a result of an inconsistency between two or more of the entered parameters.
- CMMVC5828E The virtual disk (VDisk) was not created because the I/O group contains no nodes.
- CMMVC5829E The image-mode virtual disk (VDisk) was not created because the number of managed disks (MDisks) specified is greater than one.
- CMMVC5830E The image-mode virtual disk (VDisk) was not created because no managed disk (MDisk) was specified in the command.
- CMMVC5831E The virtual disk (VDisk) was not created because the preferred node for I/O operations is not part of the I/O group.
- CMMVC5857E The action failed because the managed disk (MDisk) does not exist or it is not a member of the managed disk (MDisk) group.
- CMMVC5858E The action failed because the virtual disk (VDisk) is in the wrong mode, the managed disk (MDisk) is in the wrong mode, or both are in the wrong mode.
- CMMVC5860E The action failed because there were not enough extents in the managed disk (MDisk) group.

  Note: This error is also returned if a stripe set of MDisks has been specified and one or more of these MDisks does not contain enough free extents to complete the creation of the VDisk. In this case, the MDisk group will be reporting that it has enough free capacity to create the VDisk. You can check the free capacity on each MDisk by issuing the **svcinfo lsfreeextents <mdiskname/ID>** command. Alternatively, do not specify a stripe set and let the system choose the free extents automatically.
- CMMVC5861E The action failed because there were not enough extents on the managed disk (MDisk).

An invocation example

```
svctask mkvdisk -mdiskgrp Group0 -size 0
-iogrp 0 -vtype striped -mdisk mdisk1 -node 1
```

The resulting output

```
Virtual Disk, id [1], successfully created
```

An invocation example for creating an image mode VDisk

```
svctask mkvdisk -mdiskgrp Group0
-iogrp 0 -vtype image -mdisk mdisk2 -node 1
```

The resulting output

```
Virtual Disk, id [2], successfully created
```

mkvdiskhostmap

You can use the **mkvdiskhostmap** command to create a new mapping between a virtual disk and a host. That is, the virtual disk is made accessible for I/O operations to the specified host.

Syntax

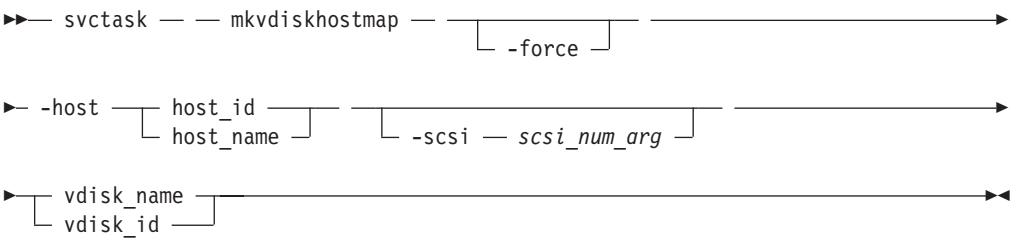

Parameters

-host *host_id* | *host_name*
Specifies the host to map the virtual disk to, either by ID or by name.

-scsi *scsi_num_arg*
Optionally specifies the SCSI LUN ID to assign to this virtual disk on the given host. The scsi_num argument contains the SCSI LUN ID that will be assigned to the VDisk on the given host. You need to check your host system for the next available SCSI LUN ID on the given HBA. This is an optional flag, if it is not supplied, the next available SCSI LUN ID will be provided to the host.

-force
Optionally specifies the force flag to force the creation.

vdisk_name | vdisk_id
Specifies the name of the virtual disk to map, either by ID or by name.

Description

This command creates a new mapping between the virtual disk and the specified host. The virtual disk is presented to the host as if the disk is directly attached to the host. It is only after this command is executed that the host can perform I/O transactions to the virtual disk.

Optionally, you can assign a SCSI LUN ID to the mapping. When the HBA in the host scans for devices attached to it, it discovers all virtual disks that are mapped to its fibre-channel ports. When the devices are found, each one is allocated an identifier (SCSI LUN ID). For example, the first disk found is usually SCSI LUN 1, and so on. You can control the order in which the HBA discovers virtual disks by assigning the SCSI LUN ID as required. If you do not specify a SCSI LUN ID, the cluster automatically assigns the next available SCSI LUN ID, given any mappings that already exist with that host.

Some HBA device drivers will stop when they find a gap in the SCSI LUN IDs. For example:

- Virtual Disk 1 is mapped to Host 1 with SCSI LUN ID 1
- Virtual Disk 2 is mapped to Host 1 with SCSI LUN ID 2
- Virtual Disk 3 is mapped to Host 1 with SCSI LUN ID 4

When the device driver scans the HBA, it must stop after discovering virtual disks 1 and 2, because there is no SCSI LUN mapped with ID 3. You should, therefore, ensure that the SCSI LUN ID allocation is contiguous.

You can create multiple Vdisk assignments. Normally, multiple vdisk to host assignments should not be used, as corruption is likely to occur if more than one host can access a disk. However, in *certain* multiple path environments, such as that found in IBM's SAN File System (SFS), a vdisk **must** be mapped to more than one host. To do this you must use the command line interface and use the -force flag. For example:

svctask mkvdiskhostmap -host host1 -force 4

svctask mkvdiskhostmap -host host2 -force 4

The above will create two host to vdisk mappings for vdisk 4, that map to host1 and host2. Note that omitting the -force flag will cause the mapping to fail if that vdisk is already mapped to a host.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5842E The action failed because an entity that was specified in the command does not exist.
- CMMVC5843E The virtual disk (VDisk)-to-host mapping was not created because the VDisk does not have a capacity greater than zero bytes.
- CMMVC5844E The virtual disk (VDisk)-to-host mapping was not created because the SCSI logical unit number (LUN) ID is not valid.
- CMMVC5862E The action failed because the virtual disk (VDisk) is being formatted.
- CMMVC5874E The action failed because the host does not exist.
- CMMVC5875E The action failed because the virtual disk (VDisk) does not exist.
- CMMVC5876E The virtual disk (VDisk)-to-host mapping was not created because the maximum number of mappings has been reached.
- CMMVC5877E The virtual disk (VDisk)-to-host mapping was not created because the maximum number of SCSI LUNs has been allocated.
- CMMVC5878E The virtual disk (VDisk)-to-host mapping was not created because this VDisk is already mapped to this host.
- CMMVC5879E The virtual disk (VDisk)-to-host mapping was not created because this VDisk is already mapped to this host with this SCSI LUN.
- CMMVC5880E The virtual disk (VDisk)-to-host mapping was not created because the VDisk has a capacity of zero bytes.
- CMMVC6071E This action will result in the creation of multiple mappings. Use the -force flag if you are sure that this is what you wish to do.

An invocation example

```
svctask mkvdiskhostmap -host host1 -scsi 1 5
```

The resulting output

```
Virtual Disk to Host map, id [1], successfully created
```

rmvdisk

The **rmvdisk** command deletes a virtual disk. The command will fail if there are mappings between the virtual disk and hosts and you do not specify the force flag.

Syntax

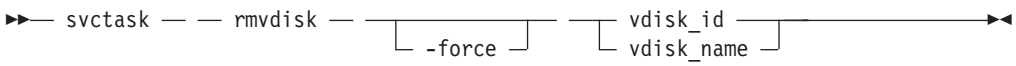

Parameters

-force
Optionally forces the deletion. This argument deletes any host-to-VDisk mappings and any FlashCopy mappings that exist for this VDisk.

vdisk_id | vdisk_name
Specifies the name of the virtual disk to delete, either by ID or by name.

Description

This command deletes an existing managed mode virtual disk or an existing image mode virtual disk. The extents that made up this virtual disk are returned to the pool of free extents that are available on the managed disk group, if the vdisk is in managed mode.

Attention: Any data that was on the virtual disk is lost. Care should be taken when executing this command to ensure that the virtual disk (and any data that resides on it) is no longer required.

If any mappings still exist between this virtual disk and hosts, the deletion will fail unless the force flag is specified. When the force flag is specified, any mappings that remain are deleted and then the virtual disk is deleted.

Deleting a managed mode virtual disk

When you use this command to delete a managed mode virtual disk, all the data on the virtual disk is deleted. The extents that make up the virtual disk are returned to the pool of free extents that are available in the managed disk group.

If FlashCopy mappings or host mappings exist for the virtual disk, then the deletion will fail. You can use the force flag to force the deletion. If you use the force flag, mappings are deleted, and then the virtual disk is deleted.

If the virtual disk is in the process of migrating to an image mode virtual disk (using the **svctask migratetoimage** command), then the deletion will fail unless the force flag is used. In this case, the migration is halted and then the virtual disk is deleted. Care should be taken when executing this command to ensure that the virtual disk (and any data that resides on it) is no longer required.

Deleting an image mode virtual disk

When you use this command to delete an image mode virtual disk, the data on the controller logical unit will be consistent with the data that had been available on the image mode virtual disk before it was deleted. In other words, the fast write

data will be moved to the controller logical unit. If the force flag is used, then the data is not moved to the controller logical unit.

If there are any virtual medium errors on the virtual disk, then the command will fail. You can force the deletion with the force flag, however this option can cause data integrity issues.

Note: A virtual medium error occurs when you copy from one disk (the source) to another (the target). If you read the source, you should detect that there is a medium error. At that moment you have to have two identical copies of data and you then have to simulate a medium error on the target disk. You can simulate that medium error on the target disk by creating a virtual medium error on the target disk.

If FlashCopy mappings or host mappings exist for the virtual disk, then the deletion will fail. You can use the force flag to force the deletion. If you use the force flag, mappings are deleted, and then the virtual disk is deleted. If there is any un-destaged data in the fast write cache for this virtual disk then the deletion of the virtual disk will fail. When the force flag is specified, any un-destaged data in the fast write cache will be deleted. Deleting an image mode virtual disk causes the managed disk associated with the virtual disk to be ejected from the managed disk group. The mode of the managed disk will be returned to "unmanaged."

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5807E The action failed because the managed disk (MDisk) cannot be changed to the specified mode.
- CMMVC5840E The virtual disk (VDisk) was not deleted because it is mapped to a host or because it is part of a FlashCopy or Remote Copy mapping.
- CMMVC5841E The virtual disk (VDisk) was not deleted because it does not exist.
- CMMVC5848E The action failed because the virtual disk (VDisk) does not exist or it is being deleted.
- CMMVC5858E The action failed because the virtual disk (VDisk) is in the wrong mode, the managed disk (MDisk) is in the wrong mode, or both are in the wrong mode.
- CMMVC5862E The action failed because the virtual disk (VDisk) is being formatted.

An invocation example

```
svctask rmvdisk -force vdisk5
```

The resulting output

```
No feedback
```

rmvdiskhostmap

The **rmvdiskhostmap** command deletes an existing virtual disk-to-host mapping. That is, the virtual disk will no longer be accessible for I/O transaction on the given host.

Syntax

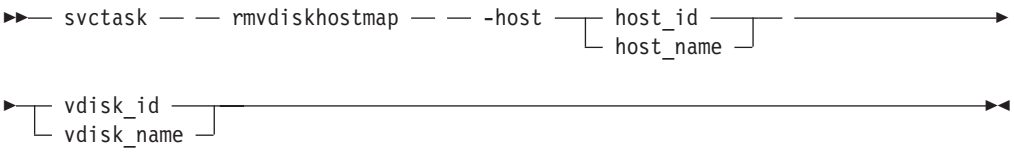

Parameters

-host *host_id* | *host_name*
Specifies the host to remove from the map with the virtual disk, either by ID or by name.

vdisk_id | vdisk_name
Specifies the name of the virtual disk from the map, either by ID or by name.

Description

This command deletes an existing mapping between the given virtual disk and the host. This effectively stops the virtual disk from being available for I/O transactions on the given host.

Care should be taken when executing this command because to the host it seems as if the virtual disk has been deleted or is offline.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5842E The action failed because an entity that was specified in the command does not exist.
- CMMVC5874E The action failed because the host does not exist.
- CMMVC5875E The action failed because the virtual disk (VDisk) does not exist.

An invocation example

```
svctask rmvdiskhostmap -host host1 vdisk8
```

The resulting output

```
No feedback
```

shrinkvdisksize

You can use the **shrinkvdisksize** command to shrink a VDisk by a given capacity.

Syntax

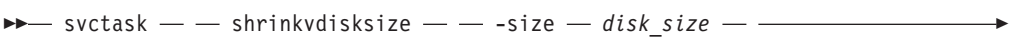

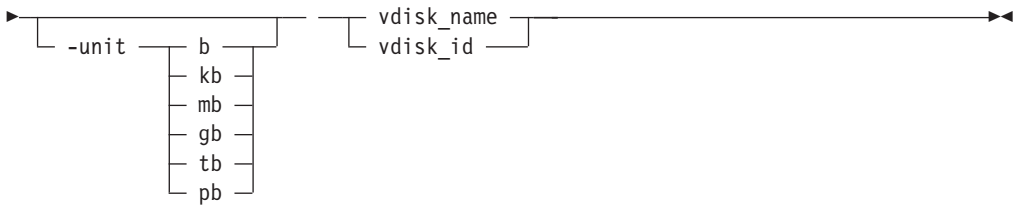

Parameters

-size *disk_size*
The command line will reduce the size by the specified amount.

-unit *b | kb | mb | gb | tb | pb*
Optionally specifies the data units to be used in conjunction with the capacity (-size).

vdisk_name | vdisk_id
Specifies the virtual disk name to modify, either by ID or by name.

Description

This command reduces the capacity allocated to the particular virtual disk by the given amount. All capacities, including changes should be in multiples of 512 bytes. An error will occur if you specify a capacity that is not a multiple of 512, which can only happen when byte units (-b) are used. However, an entire extent will be reserved even if only partially used. The default capacity is in MB.

VDisks can be reduced in size should it be required. However, if the VDisk contains data that is being used, **under no circumstances should you attempt to shrink a VDisk without first backing up your data**. The SAN Volume Controller arbitrarily reduces the capacity of the VDisk by removing a partial, one or more extents from those allocated to the VDisk. You cannot control which extents are removed and so you cannot assume that it is unused space that is removed.

Attention: This feature should *only* be used to make a target or auxiliary VDisk the same size as the source or master VDisk when creating FlashCopy mappings or Metro Mirror relationships. You should also ensure that the target VDisk is not mapped to any hosts prior to performing this operation.

Attention: If the virtual disk contains data, you should not shrink the disk.

Note: Some operating systems or file systems use what they consider to be the outer edge of the disk for performance reasons. This command is provided to shrink FlashCopy target virtual disks to the same capacity as the source.

Validate that the VDisk is not mapped to any host objects. If the VDisk is mapped, data is displayed. You can determine the exact capacity of the source or master VDisk by issuing the **svcinfo lsvdisk -bytes <vdiskname>** command. Shrink the VDisk by the required amount by issuing the **svctask shrinkvdisksize -size <capacitytoshrinkby> -unit <unitsforreduction> <vdiskname/ID>** command.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5836E The virtual disk (VDisk) was not shrunk because it is locked.

- CMMVC5837E The action failed because the virtual disk (VDisk) is part of a FlashCopy mapping.
- CMMVC5838E The action failed because the virtual disk (VDisk) is part of a Remote Copy mapping.
- CMMVC5839E The virtual disk (VDisk) was not shrunk because an entity that was specified in the command does not exist.
- CMMVC5848E The action failed because the virtual disk (VDisk) does not exist or it is being deleted.
- CMMVC5862E The action failed because the virtual disk (VDisk) is being formatted.
- CMMVC6010E Unable to complete the command as there are insufficient free extents.

An invocation example

```
svctask shrinkvdisksize -size 2048 -unit b vdisk1
```

The resulting output

```
No feedback
```

Chapter 2. Attributes of the -filtervalue argument

The **-filtervalue** argument is used to filter a view based on specific attribute values that relate to each object type. You can combine multiple filters to create specific searches, for example, -filtervalue name=fred:status=online. The help (-filtervalue?) specifies the attributes available for each object type.

The -filtervalue argument, if used, must have attrib=value entered. The -filtervalue? and -filtervalue arguments are mutually exclusive.

Note: The qualifiers < and > should be enclosed in quotes. For example,

-filtervalue vdisk_count ″<″4 or port_count ″>″1

It is also valid to include the whole expression in quotes. For example,

-filtervalue ″vdisk_count<4″

When an attribute requires the -unit argument, it is specified after the attribute. For example,

-filtervalue capacity=24 -unit mb

The valid input options for the -unit parameter are:

- b (bytes)
- mb (Megabytes)
- gb (Gigabytes)
- tb (Terabytes)
- pb (Petabytes)

The * character may be used as a wildcard at the beginning or end of a text string, but not both.

Table 1. Valid filter attributes

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
cluster	cluster_name or name	=	Yes	The cluster name.
	cluster_unique_id or id	=, <, <=, >, >=	No	The cluster ID.

© Copyright IBM Corp. 2003, 2005

Table 1. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
node	node_name or name	=	Yes	The node name.
	id	=, <, <=, >, >=	No	The node ID.
	status	=	No	The status of the node. The valid input options for node status are: • adding • deleting • online • offline • pending
	IO_group_name	=	Yes	The I/O group name.
	IO_group_id	=, <, <=, >, >=	No	The I/O group ID.
io_grp	HWS_name or name	=	Yes	The I/O group name
	HWS_unique_id or id	=, <, <=, >, >=	No	The I/O group ID.
	node_count	=, <, <=, >, >=	No	The number of nodes in the I/O group.
controller	controller_id or id	=, <, <=, >, >=	No	The controller ID.
mdisk	name	=	Yes	The name of the MDisk.
	id	=, <, <=, >, >=	No	The ID of the MDisk.
	controller_name	=	Yes	The name of the controller the MDisk belongs to.
	status	=	No	The status of the MDisk. The valid input options for MDisk status are: • online • degraded • excluded • offline
	mode	=	No	The mode of the MDisk. The valid input options for MDisk mode are: • unmanaged • managed • image
	mdisk_grp_name	=	Yes	The MDisk group name.
	mdisk_grp_id	=, <, <=, >, >=	No	The MDisk group ID.
	capacity	=, <, <=, >, >=	No	The capacity. Requires the -unit argument.

Table 1. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
mdiskgrp	name	=	Yes	The MDisk group name.
	storage_pool_id or id	=, <, <=, >, >=	No	The MDisk group ID.
	mdisk_count	=, <, <=, >, >=	No	The number of MDisks in the group.
	vdisk_count	=, <, <=, >, >=	No	The number of VDisks in the group.
	status	=	No	The status of the MDisk group. The valid input options are: • online • degraded • offline
	extent_size	=, <, <=, >, >=	No	The extent size. (MB)
vdisk	vdisk_name or name	=	Yes	The name of the VDisk.
	vdisk_id or id	=, <, <=, >, >=	No	The ID of the VDisk.
	IO_group_name	=	Yes	The name of the I/O group.
	IO_group_id	=, <, <=, >, >=	No	The ID of the I/O group.
	status	=	No	The status of the VDisk. The valid input options for VDisk status are: • online • degraded • offline
	mdisk_grp_name	=	Yes	The MDisk group name.
	mdisk_grp_id	=, <, <=, >, >=	No	The MDisk group ID.
	capacity	=, <, <=, >, >=	No	The capacity. Requires the -unit argument.
	type	=	No	The VDisk type. The valid value options are: • seq • striped • image
	FC_name	=	Yes	The FlashCopy mapping name.
	FC_id	=, <, <=, >, >=	No	The FlashCopy mapping ID.
	RC_name	=	Yes	The Metro Mirror relationship name.
	RC_id	=, <, <=, >, >=	No	The Metro Mirror relationship ID.
host	host_name or name	=	Yes	The host name.
	host_id or id	=, <, <=, >, >=	No	The host ID.
	port_count	=, <, <=, >, >=	No	The number of ports.

Table 1. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
fcmap	FC_mapping_name or name	=	Yes	The FlashCopy mapping name.
	FC_id or id	=, <, <=, >, >=	No	The FlashCopy mapping ID.
	source_vdisk_name	=	Yes	The source VDisk name.
	source_vdisk_id	=, <, <=, >, >=	No	The source VDisk ID.
	target_vdisk_name	=	Yes	The target VDisk name.
	target_vdisk_id	=, <, <=, >, >=	No	The target VDisk ID.
	group_name	=	Yes	The consistency group name.
	group_id	=, <, <=, >, >=	No	The consistency group ID.
	status	=	No	The mapping status. The valid input options for fcmap status are: • idle_copied • preparing • copying • stopped • suspended
	copy_rate	=, <, <=, >, >=	No	The background copy rate.
fcconsist-grp	name	=	Yes	The consistency group name.
	FC_group_id or id	=, <, <=, >, >=	No	The consistency group ID.
	status	=	No	The consistency group status. The valid value options are: • idle_or_copied • preparing • prepared • copying • stopped • suspended

Table 1. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
rcrelation-ship	RC_rel_id or id	=, <, <=, >, >=	No	The Metro Mirror relationship ID.
	RC_rel_name or name	=	Yes	The Metro Mirror relationship name.
	master_cluster_id	=, <, <=, >, >=	No	The master cluster ID.
	master_cluster_name	=	Yes	The master cluster name.
	master_vdisk_id	=, <, <=, >, >=	No	The master VDisk ID.
	master_vdisk_name	=	Yes	The master VDisk name.
	aux_cluster_id	=, <, <=, >, >=	No	The aux cluster ID.
	aux_cluster_name	=	Yes	The aux cluster name.
	aux_vdisk_id	=, <, <=, >, >=	No	The aux VDisk ID.
	aux_vdisk_name	=	Yes	The aux VDisk name.
	primary	=	No	The relationship primary. The valid input values are: • master • aux
	consistency_group_id	=, <, <=, >, >=	No	The Metro Mirror consistency group ID.
	consistency_group_name	=	Yes	The Metro Mirror consistency group name.
	state	=	Yes	The relationship state. The valid input values are: • inconsistent_stopped • inconsistent_copying • consistent_stopped • consistent_synchronized • idling • idling_disconnected • inconsistent_disconnected • consistent_disconnected
	progress	=, <, <=, >, >=	No	The progress of the initial background copy (synchronization) for the relationship.

Table 1. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
rcconsist-grp	group_id or id	=, <, <=, >, >=	No	The consistency group ID.
	name	=	Yes	The consistency group name.
	master_cluster_id	=, <, <=, >, >=	No	The master cluster ID.
	master_cluster_name	=	Yes	The master cluster name.
	aux_cluster_id	=, <, <=, >, >=	No	The aux cluster ID.
	aux_cluster_name	=	Yes	The aux cluster name.
	primary	=	No	The consistency group primary. The valid input values are: • master • aux
	state	=	No	The consistency group state. The valid input values are: • inconsistent_stopped • inconsistent_copying • consistent_stopped • consistent_synchronized • idling • idling_disconnected • inconsistent_disconnected • consistent_disconnected • empty
	relationship_count	=, <, <=, >, >=	No	The relationship count.

Chapter 3. Overview of the list dump commands

You can use the list dumps command to return a list of dumps in the appropriate directory.

The dumps in the SAN Volume Controller are contained in the following directory structure:

- /dumps
- /dumps/configs
- /dumps/elogs
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace

Software upgrade packages are contained in the /home/admin/upgrade directory. These directories exist on every node in the cluster.

Configuration dump: Dumps contained in the /dumps/configs directory are dumps of the cluster configuration data. A configuration dump is created by using the **svctask dumpconfig** command. This will dump the configuration of the cluster, including all object details, to the /dumps/configs directory. If no filename prefix is supplied, the default config_ will be used. The full, default file name, will be config_NNNNNN_YYMMDD_HHMMSS (where *NNNNNN* is the node front panel name). If the command is used with the -prefix option, then the value entered for the -prefix will be used instead of config. The command to list all dumps in the /dumps/configs directory is **svcinfo lsconfigdumps**.

Error or event dump: Dumps contained in the /dumps/elogs directory are dumps of the contents of the error and event log at the time that the dump was taken. An error or event log dump is created by using the **svctask dumperrlog** command. This will dump the contents of the error or event log to the /dumps/elogs directory. If no filename prefix is supplied, the default errlog_ will be used. The full, default file name, will be errlog_NNNNNN_YYMMDD_HHMMSS (where *NNNNNN* is the node front panel name). If the command is used with the -prefix option, then the value entered for the -prefix will be used instead of errlog. The command to list all dumps in the /dumps/elogs directory is **svcinfo lserrlogdumps**.

Featurization log dump: Dumps contained in the /dumps/feature directory are dumps of the featurization log. A featurization log dump is created by using the **svctask dumpinternallog** command. This will dump the contents of the featurization log to the /dumps/feature directory to a file called feature.txt. Only one of these files exists, so every time the **svctask dumpinternallog** command is run, this file is overwritten. The command to list all dumps in the /dumps/feature directory is **svcinfo lsfeaturedumps**.

I/O statistics dump: Dumps contained in the /dumps/iostats directory are dumps of the I/O statistics for disks on the cluster. An I/O statistics dump is created by using the **svctask startstats** command. As part of this command, you can specify a time interval at which you want the statistics to be written to the file (the default is 15 minutes). Every time the time interval is encountered, the I/O statistics that have been collected this far are written to a file in the /dumps/iostats directory.

© Copyright IBM Corp. 2003, 2005

The file names used for storing I/O statistics dumps are m_stats_NNNNNN_YYMMDD_HHMMSS, Nm_stats_NNNNNN_YYMMDD_HHMMSS, or v_stats_NNNNNN_YYMMDD_HHMMSS (where *NNNNNN* is the node front panel name), depending on if the statistics are for MDisks or VDisks. The command to list all dumps in the /dumps/iostats directory is **svcinfo lsiostatsdumps**.

I/O trace dump: Dumps contained in the /dumps/iotrace directory are dumps of I/O trace data. The type of data that is traced depends on the options specified by the **svctask settrace** command. The collection of the I/O trace data is started by using the **svctask starttrace** command. The I/O trace data collection is stopped when the **svctask stoptrace** command is used. It is when the trace is stopped that the data is written to the file. The file name will be <prefix>_NNNNNN_YYMMDD_HHMMSS (where *NNNNNN* is the node front panel name, and <prefix> is the value entered by the user for the -filename parameter in the **svctask settrace** command.) The command to list all dumps in the /dumps/iotrace directory is **svcinfo lsiotracedumps**.

Application abends dump: Dumps contained in the /dumps directory are dumps resulting from application abends. Such dumps will be written to the /dumps directory. The default file names are dump.NNNNNN.YYMMDD.HHMMSS (where *NNNNNN* is the node front panel name). In addition to the dump file, it is possible that there may be some trace files written to this directory. These will be named NNNNNN.trc.

The command to list all dumps in the /dumps directory is **svcinfo ls2145dumps**.

The final option available in the list dumps command series is the **svcinfo lssoftwaredumps** command. This command will list the contents of the /home/admin/upgrade directory. Any files in this directory will have been copied there at the time that you wanted to perform a software upgrade.

All of the list dumps commands can accept a node identifier as input. If this identifier is not specified then the list of files on the current configuration node will be displayed. If the node identifier is specified, then the list of files on that node will be displayed.

Because files can only be copied off of the current configuration node (using secure copy), you can issue the **svctask cpdumps** command to copy the files off of a non-configuration node to the current configuration node.

Accessibility

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

Features

These are the major accessibility features in the SAN Volume Controller master console:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen. The following screen readers have been tested: JAWS v4.5 and IBM Home Page Reader v3.0.
- You can operate all features using the keyboard instead of the mouse.

Navigating by keyboard

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

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

Accessing the publications

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

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

Related reference

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

© Copyright IBM Corp. 2003, 2005

Notices

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

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

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

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

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

© Copyright IBM Corp. 2003, 2005

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

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

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

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

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

Trademarks

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

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

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

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

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

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

Glossary

This is the glossary for the SAN Volume Controller.

A

application server
A host that is attached to the storage area network (SAN) and that runs applications.

C

cluster
In SAN Volume Controller, a pair of nodes that provides a single configuration and service interface.

configuration node
A node that acts as the focal point for configuration commands and manages the data that describes the cluster configuration.

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

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

copied
In a FlashCopy® relationship, a state that indicates that a copy has been started after the copy relationship was created. The copy process is complete and the target disk has no further dependence on the source disk.

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

D

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

degraded
Pertaining to a valid configuration that has suffered a failure but continues to be supported and legal. Typically, a repair action can be performed on a degraded configuration to restore it to a valid configuration.

dependent write operations
A set of write operations that must be applied in the correct order to maintain cross-volume consistency.

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

© Copyright IBM Corp. 2003, 2005

directed maintenance procedures
The set of maintenance procedures that can be run for a cluster. These procedures are documented in the service guide.

disconnected
In a Global Mirror relationship, pertains to two clusters when they cannot communicate.

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

disk zone
A zone defined in the storage area network (SAN) fabric in which the SAN Volume Controller can detect and address the logical units that the disk controllers present.

E

error code
A value that identifies an error condition.

ESS See *IBM® TotalStorage® Enterprise Storage Server®*.

exclude
To remove a managed disk (MDisk) from a cluster because of certain error conditions.

excluded
In SAN Volume Controller, the status of a managed disk that the cluster has removed from use after repeated access errors.

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

F

failover
In SAN Volume Controller, the function that occurs when one redundant part of the system takes over the workload of another part of the system that has failed.

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

FC See *fibre channel*.

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

FlashCopy mapping
A relationship between two virtual disks.

FlashCopy relationship
See *FlashCopy mapping*.

H

HBA See *host bus adapter*.

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

host An open-systems computer that is connected to the SAN Volume Controller through a fibre-channel interface.

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

host zone
A zone defined in the storage area network (SAN) fabric in which the hosts can address the SAN Volume Controllers.

I

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

IBM TotalStorage Enterprise Storage Server (ESS)
An IBM product that provides an intelligent disk-storage subsystem across an enterprise.

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

illegal configuration
A configuration that will not operate and will generate an error code to indicate the cause of the problem.

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

image VDisk
A virtual disk (VDisk) in which there is a direct block-for-block translation from the managed disk (MDisk) to the VDisk.

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

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

integrity
The ability of a system to either return only correct data or respond that it cannot return correct data.

Internet Protocol (IP)
In the Internet suite of protocols, a connectionless protocol that routes data

through a network or interconnected networks and acts as an intermediary between the higher protocol layers and the physical network.

I/O See *input/output.*

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

I/O throttling rate
The maximum rate at which an I/O transaction is accepted for this virtual disk (VDisk).

IP See *Internet Protocol*.

L

LBA See *logical block address*.

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

local/remote fabric interconnect
The storage area network (SAN) components that are used to connect the local and remote fabrics together.

logical block address (LBA)
The block number on a disk.

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

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

LU See *logical unit*.

LUN See *logical unit number*.

M

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

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

mapping
See *FlashCopy mapping*.

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

MDisk
See *managed disk.*

migration
See *data migration*.

N

node One SAN Volume Controller. Each node provides virtualization, cache, and Copy Services to the storage area network (SAN).

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

O

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

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

P

partnership
In Global Mirror, the relationship between two clusters. In a cluster partnership, one cluster is defined as the local cluster and the other cluster as the remote cluster.

paused
In SAN Volume Controller, the process by which the cache component quiesces all ongoing I/O activity below the cache layer.

pend To cause to wait for an event.

port The physical entity within a host, SAN Volume Controller, or disk controller system that performs the data communication (transmitting and receiving) over the fibre channel.

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

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

Q

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

R

RAID See *redundant array of independent disks*.

RAID 1
SNIA dictionary definition: A form of storage array in which two or more identical copies of data are maintained on separate media. IBM definition: A form of storage array in which two or more identical copies of data are maintained on separate media. Also known as mirrorset. HP definition: See *mirrorset*.

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

RAID 5

- SNIA definition: A form of parity RAID in which the disks operate independently, the data strip size is no smaller than the exported block size, and parity check data is distributed across the array's disks. (S)
- IBM definition: See above.
- HP definition: A specially developed RAID storageset that stripes data and parity across three or more members in a disk array. A RAIDset combines the best characteristics of RAID level 3 and RAID level 5. A RAIDset is the best choice for most applications with small to medium I/O requests, unless the application is write intensive. A RAIDset is sometimes called parity RAID. RAID level 3/5 storagesets are referred to as RAIDsets.

RAID 10
A type of RAID that optimizes high performance while maintaining fault tolerance for up to two failed disk drives by striping volume data across several disk drives and mirroring the first set of disk drives on an identical set.

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

rejected
A status condition that describes a node that the cluster software has removed from the working set of nodes in the cluster.

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

Global Mirror
In SAN Volume Controller, a copy service that enables host data on a particular source virtual disk (VDisk) to be copied to the target VDisk designated in the relationship.

S

SAN See *storage area network*.

SAN Volume Controller fibre-channel port fan in
The number of hosts that can see any one SAN Volume Controller port.

SCSI See *Small Computer Systems Interface*.

sequential VDisk
A virtual disk that uses extents from a single managed disk.

Small Computer System Interface (SCSI)
A standard hardware interface that enables a variety of peripheral devices to communicate with one another.

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

Secure Shell
A program to log in to another computer over a network, to execute commands in a remote machine, and to move files from one machine to another.

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

SNMP
See *Simple Network Management Protocol*.

stand-alone relationship
In FlashCopy and Global Mirror, relationships that do not belong to a consistency group and that have a null consistency group attribute.

SSH See *Secure Shell*.

stop A configuration command that is used to stop the activity for all copy relationships in a consistency group.

stopped
The status of a pair of virtual disks (VDisks) that have a copy relationship that the user has temporarily broken because of a problem.

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

superuser authority
The level of access required to add users.

suspended
The status of a pair of virtual disks (VDisks) that have a copy relationship that has been temporarily broken because of a problem.

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

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

T

trigger
To initiate or reinitiate copying between a pair of virtual disks (VDisks) that have a copy relationship.

U

unconfigured mode
A mode in which I/O operations cannot be performed. See also *image mode* and *managed space mode*.

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

V

valid configuration
A configuration that is supported.

VDisk See *virtual disk*.

virtual disk (VDisk)
In SAN Volume Controller, a device that host systems attached to the storage area network (SAN) recognize as a Small Computer System Interface (SCSI) disk.

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

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

vital product data (VPD)
Information that uniquely defines system, hardware, software, and microcode elements of a processing system.

W

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

WWNN
See *worldwide node name*.

WWPN
See *worldwide port name*.

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

Printed in USA

SC26-7544-04

Spine information:

IBM TotalStorage SAN Volume Controller

Command-Line Interface User's Guide - Errata

Version 2.1.0 - Errata