

IBM TotalStorage SAN Volume Controller
Storage Software for Cisco MDS 9000



Command-Line Interface User's Guide

Version 1.1

IBM TotalStorage SAN Volume Controller
Storage Software for Cisco MDS 9000



Command-Line Interface User's Guide

Version 1.1

Note

Before using this information and the product it supports, read the information in "Notices" on page 331.

First Edition (November 2003)

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

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

Contents

About this guide v

Who should use this guide	v
Related publications	v
How to order IBM publications	vii
Related Web sites	viii
Emphasis	viii
Terminology	viii
Syntax diagrams	ix
CLI special characters	xi
Data types and value ranges	xi
CLI parameters	xvi
CLI flags	xvi
How to send your comments	xvi

Chapter 1. Preparing the SSH client system overview 1

Preparing the SSH client system to issue command-line interface (CLI) commands	2
--	---

Chapter 2. Secure Shell (SSH) configuration 3

Overview of configuring the Secure Shell (SSH)	3
--	---

Chapter 3. Cluster commands 5

addnode	6
chcluster	9
chiogrp	10
chnode	11
cleardumps	12
cpdumps	14
detectmdisk	16
dumpconfig	18
rmnode	19
startstats	23
stopcluster	24
stopstats	27

Chapter 4. Cluster diagnostic and service-aid commands 29

addnode	30
cherrstate	33
clearerrlog	34
dumperrlog	35
finderr	36
rmnode	37
setevent	39

Chapter 5. Host commands 41

addhostport	42
chhost	44
mkhost	45
rmhost	47
rmhostport	48

Chapter 6. Virtual disk commands 51

chvdisk	52
expandvdisksize	55
mkvdisk	57
mkvdiskhostmap	61
rmvdisk	63
rmvdiskhostmap	65
shrinkvdisksize	66

Chapter 7. Managed disk group commands 69

addmdisk	70
chmdiskgrp	72
mkmdiskgrp	73
rmmdisk	75
rmmdiskgrp	77

Chapter 8. Managed disk commands 79

chmdisk	80
includemdisk	81
setquorum	82

Chapter 9. FlashCopy commands 85

chfcconsistgrp	86
chfcmap	87
mkfcconsistgrp	89
mkfcmap	90
prestartfcconsistgrp	92
prestartfcmap	94
rmfcconsistgrp	96
rmfcmap	97
startfcconsistgrp	99
startfcmap	101
stopfcconsistgrp	103
stopfcmap	104

Chapter 10. Remote Copy commands 105

chpartnership	106
chrcconsistgrp	107
chrcrelationship	108
mkpartnership	110
mkrrcconsistgrp	112
mkrrcrelationship	113
rmpartnership	116
rmrrcconsistgrp	117
rmrrcrelationship	118
startrrcconsistgrp	119
startrrcrelationship	121
stoprrcconsistgrp	123
stoprrcrelationship	125
switchrrcconsistgrp	127
switchrrcrelationship	128

Chapter 11. Migration commands 129

migrateexts	130
migratevdisk	132

Chapter 12. Tracing commands. . . . 135

setdisktrace	136
settrace	137
starttrace	140
stoptrace	141

Chapter 13. Attributes of the -filtervalue argument 143

Chapter 14. Overview of the list dump commands 151

Chapter 15. Information commands 153

caterrlog	154
caterrlogbyseqnum	156
lssvcdumps (ls2145dumps)	158
lscluster	160
lsclustercandidate	163
lsconfigdumps	164
lscontroller	166
lserrlogbyfcconsistgrp	169
lserrlogbyfcmap	171
lserrlogbyhost	173
lserrlogbyiogrp	175
lserrlogbymdisk	177
lserrlogbymdiskgroup	179
lserrlogbynode	181
lserrlogbyrconsistgrp	183
lserrlogbyrrelationship	185
lserrlogbyvdisk	187
lserrlogdumps	189
lsfcconsistgrp	191
lsfcmap	193
lsfcmapcandidate	196
lsfcmapprogress	197
lsfeaturedumps	199
lsfreeextents	201
lshbaportcandidate	203
lshost	204
lshostvdiskmap	206
lsiogrp	208
lsiogrpcandidate	210
lsiostatsdumps	211
lsiotracedumps	213
lslicense	215
lsmdisk	217
lsmdiskcandidate	221
lsmdiskextent	223
lsmdiskgrp	226
lsmdiskmember	229
lsmigrate	231
lsnode	232
lsnodecandidate	235
lsnodevpd	236

lsrconsistgrp	239
lsrrelationship	242
lsrrelationshipcandidate	245
lsrrelationshipprogress	247
lsvdisk	249
lsvdiskextent	252
lsvdiskhostmap	254
lsvdiskmember	256
lsvdiskprogress	258

Chapter 16. Error log commands . . . 261

finderr	262
dumperrlog	263
clearerrlog	264
cherrstate	265
setevent	266

Chapter 17. Featurization commands 267

chlicense	268
dumpinternallog	270

Chapter 18. Service mode commands 271

cleardumps	272
dumperrlog	274
exit	275

Chapter 19. Service mode information commands 277

lssvcdumps (ls2145dumps)	278
lsclustervpd	280
lsconfigdumps	282
lserrlogdumps	283
lsfeaturedumps	284
lsiostatsdumps	285
lsiotracedumps	286
lsnodes	287
lsnodevpd	288

Chapter 20. Controller command . . . 291

chcontroller	292
------------------------	-----

Chapter 21. CLI messages 295

Appendix A. Accessibility 327

Appendix B. SAN Volume Controller CLI compatibility 329

Notices 331

Trademarks	332
----------------------	-----

Glossary 333

Index 339

About this guide

This guide introduces the IBM® TotalStorage® SAN Volume Controller Storage Software for Cisco MDS 9000 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 Storage Software for Cisco MDS 9000.

Related publications

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

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

SAN Volume Controller Storage Software for Cisco MDS 9000 library:

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

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

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

Title	Description	Order number
<i>IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000: Command-Line Interface User's Guide</i>	This guide describes the commands that you can use from the SAN Volume Controller Storage Software for Cisco MDS 9000 command-line interface (CLI).	SC26-7555
<i>IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000: Configuration Guide</i>	This guide provides guidelines for configuring your SAN Volume Controller Storage Software for Cisco MDS 9000.	SC26-7554

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

Title	Description	Order number
<i>IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000: Getting Started Guide</i>	This guide introduces the SAN Volume Controller Storage Software for Cisco MDS 9000 and lists the features you can order. It also provides guidelines for planning the installation and configuration of the SAN Volume Controller Storage Software for Cisco MDS 9000.	GA22-1055

Other IBM publications:

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

Table 2. Other IBM publications

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

Cisco Systems publications:

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

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

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

Click:

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

Ordering technical documentation from Cisco Systems:

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

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

You can order Cisco Systems documentation in the following ways:

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

Related topics:

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

How to order IBM publications

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

The IBM publications center:

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

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

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

Publications notification system:

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

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

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

Related topics:

- “Related publications” on page v

Related Web sites

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

Table 3. Web sites

Type of information	Web site
SAN Volume Controller Storage Software for Cisco MDS 9000 support	http://www.ibm.com/storage/support/2062-2300/
Technical support for IBM storage products	http://www.ibm.com/storage/support/
Cisco Systems documentation	http://www.cisco.com/univercd/home/home.htm
Cisco Systems product site	http://www.cisco.com/
Cisco Systems international site	http://www.cisco.com/public/countries_languages.shtml

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.

Terminology

This section provides a list of the abbreviations that are most commonly used for the command-line interface operations.

Table 4 shows the abbreviations that are most commonly used for the command-line interface operations.

Table 4. Abbreviated object type

Name	Object type
Host	host
Virtual disk	vdisk
Managed disk	mdisk

Table 4. Abbreviated object type (continued)

Name	Object type
Managed disk group	mdiskgrp
I/O group	io_grp
Node	node
Cluster	cluster
Controller	controller
FlashCopy mapping	fcmap
FlashCopy consistency group	fcconsistgrp
Remote Copy relationship	rrelationship
Remote Copy consistency group	rcconsistgrp
Unsupported/unknown object	unknown

Syntax diagrams

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

This topic shows you how to read the syntax diagrams that represent the SAN Volume Controller Storage Software for Cisco MDS 9000 CLI commands. In doing so, it defines the symbols that represent the CLI command elements.

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.

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



Fragment name:



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.

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.

Data types and value ranges

This topic provides information about the different data types and value ranges.

The following list defines each of the data types and value ranges.

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.

filename_arg

This is a (optionally fully qualified) file name. Maximum length is 231 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 '.'.

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/ec
- /dumps/elogs
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace
- /dumps/pl

The filename filter can be any valid filename with or without the wildcard '*'. (If using the wildcard, ensure that the `directory_or_file_filter` is surrounded by quotes.) The filename filter can be appended to the end of one of the above directory values. Maximum length is 231 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 231 characters. Valid characters consist of the following:

- a - z
- A - Z
- 0 - 9
- -
- _

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 Remote Copy 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 and is the only name that cannot be changed afterwards.

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.

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 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 figure will be rounded up to the nearest 16 MB.

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.

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

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, during 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 characters (switchname.slot.node) that is a way to physically identify the node.

sequence_number

32-bit unsigned integer, expressed in decimal.

scsi_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.

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.

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. 5 to 300 (seconds) in increments of 5.

directory_arg

Specifies a directory and or filename filter within the specified directory.

Valid directory values are:

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

The filename filter can be any valid filename with or without the wildcard '*'. (If using the wildcard, ensure that the `directory_or_file_filter` is surrounded by quotes.)

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.

CLI parameters

This topic provides information about the CLI parameters.

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

This topic provides information about issuing CLI flags.

The following flags 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.

How to send your comments

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

- e-mail

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

```
starpubs@us.ibm.com
```

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

- Mail or fax

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

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

Related topics:

- "Related publications" on page v

Chapter 1. Preparing the SSH client system overview

This topic provides an overview about how to prepare the SSH client system to enable you to issue CLI commands from the host to the cluster.

Windows operating systems::

The console workstation is a Windows 2000 system which is equipped with the PuTTY Secure Shell (SSH) client software. You can install the PuTTY SSH client software on another Windows host using the PuTTY Installation program `putty-0.53b-installer.exe` which is in the `SSHClient\PuTTY` directory of the SAN Volume Controller Console CD-ROM. Or, you can download PuTTY from the following Web site:

<http://www.chiark.greenend.org.uk/~sgtatham/putty/>

The following Web site offers SSH client alternatives for Windows:

<http://www.openssh.com/windows.html>

Cygwin software has an option to install an OpenSSH client. You can download cygwin from the following Web site:

<http://www.cygwin.com/>

AIX operating systems::

For AIX[®] 5L Power 5.1 and 5.2, you can get OpenSSH from the Bonus Packs and you will also need its prerequisite, OpenSSL, from the AIX toolbox for Linux applications for Power Systems. For AIX 4.3.3, you can get the software from the AIX toolbox for Linux applications.

You can also get the AIX installation images from IBM developer Works at the following Web site:

<http://oss.software.ibm.com/developerworks/projects/openssh>

Linux operating systems::

OpenSSH is installed by default on most Linux distributions. If it is not installed on your system, consult your installation media or visit the following Web site:

<http://www.openssh.org/portable.html>

OpenSSH is able to run on a wide variety of additional operating systems. For more information visit the following Web site:

<http://www.openssh.org/portable.html>

Preparing the SSH client system to issue command-line interface (CLI) commands

This task provides step-by-step instructions about how to prepare the SSH client system to issue CLI commands.

In order to issue CLI commands to the cluster from a host, you must prepare the Secure Shell (SSH) client on the host so that the host will be accepted by the SSH server on the cluster, and allowed to connect.

If you wish to use a host which requires a different type of SSH client, for example OpenSSH, follow the instructions for that software.

Steps:

Perform the following steps to enable your host to issue CLI commands:

For the console workstation and Windows hosts:

1. Obtain and install PuTTY from SAN Volume Controller Storage Software for Cisco MDS 9000.
2. From the Start menu, click Programs, **PuTTY**, then **PuTTY 3**.
3. From the PuTTY Configuration window, in the Portfolio section, click **Session**.
4. In the work area, place a bullet beside SSH.
5. Enter the host name or IP address in the Host Name [or IP address] field.
6. Click **Open**.

For other types of hosts, follow the instructions specific to the SSH client.

Related topics:

- Chapter 2, “Secure Shell (SSH) configuration,” on page 3

Chapter 2. Secure Shell (SSH) configuration

This topic describes how to use the secure shell client remotely from your host system.

Overview:

SSH is a client-server network application intended to ensure secure communications. The SAN Volume Controller Storage Software for Cisco MDS 9000 uses the SSH facilities within the MDS 9000 to provide these services.

When establishing a session with the SAN Volume Controller Storage Software for Cisco MDS 9000 using the CLI, any of the SSH facilities provided by the SAN Volume Controller Storage Software for Cisco MDS 9000 may be used, provided that you have a compatible SSH client.

When using the console workstation, you must enable SSH. To enable SSH on the SAN Volume Controller Storage Software for Cisco MDS 9000, at the switch, use the **switch#config-t** and **switch (config)#ssh server enable** commands.

Secure Shell (SSH) is an optional communication vehicle between the host system you are using and either:

- the SAN Volume Controller Storage Software for Cisco MDS 9000 Command-Line Interface (CLI)
- or the system on which the Console for SAN Volume Controller for Cisco MDS 9000 is installed.

In order to use the Command-Line Interface (CLI) or Console for SAN Volume Controller for Cisco MDS 9000 system, you must have an SSH Client installed on that system.

The console workstation requires the SSH client software called PuTTY. This software provides the Secure Shell (SSH) client function for users logged into the console workstation who wish to invoke the SAN Volume Controller Storage Software for Cisco MDS 9000 Command-Line Interface (CLI). You must install the PuTTY client software.

If you wish to run the SAN Volume Controller Storage Software for Cisco MDS 9000 Command-Line Interface (CLI) from a different system than the console workstation, you must install an SSH client. For your convenience, the installation program to install the PuTTY software on Windows can be found in the SSH client directory of the Console for SAN Volume Controller for Cisco MDS 9000 CD-ROM.

Overview of configuring the Secure Shell (SSH)

This topic provides an overview about configuring the SSH client system. IBM has provided the PuTTY Secure Shell client software on the console workstation. You can also install PuTTY on any Windows 2000 server where you will run the command-line interface (CLI) or where you can choose to install the Console for SAN Volume Controller for Cisco MDS 9000. If you have some other Secure Shell client software to run on another host, follow that software's documentation to perform the tasks equivalent to the following steps.

1. Install SSH client software.
2. Configure the session, if required, on the SSH client system.

Chapter 3. Cluster commands

This section lists and describes the cluster commands.

When some commands complete successfully, textual output is normally provided. However, some commands do not provide any output. The phrase `No feedback` is used to indicate that no output is provided. If the command does not complete successfully, an error is generated. For example, if the command has failed as a result of the cluster being unstable, the following output is provided:

- `CMMVC5786E` The action failed because the cluster is not in a stable state.

addnode

You can use the **addnode** command to add a new (candidate) node to an existing cluster. You can enter this command any time after a cluster has been created.

Syntax

```
▶▶ svctask — addnode — [ -panelname — switchname.slot.node ] —————▶▶
                             [ -wwnodename — wwnn_arg ]
▶ [ -name — new_name_arg ] — -iogrp — [ iogroup_name | iogroup_id ] —————▶▶
```

Parameters

- panelname** *switchname.slot.node*
Identifies the node to be added by the name as it appears in the output of the **svcinfo lsnodecandidate** command. The name is composed of *switchname.slot.node*. The *switchname* is the name of the switch with certain characters (\ / : * ? < > | .) replaced by an underscore. The *slot* is the slot of the Cisco MDS 9000 Caching Services Module of the node (1-9). The *node* is the node in the Cisco MDS 9000 Caching Services Module (1 or 2). This argument is mutually exclusive with **-wwnodename**; only one is required to uniquely identify the node.
- wwnodename** *wwnn_arg*
Identifies the node to add to the cluster by the worldwide node name (WWNN) of the node. This argument is mutually exclusive with **-panelname**; only one is required to uniquely identify the node.
- name** *new_name_arg*
Optionally specifies a name for this node.
- iogrp** *iogroup_name | iogroup_id*
Specifies the I/O group to which you want to add this node.

Description

This command adds a new node to the cluster. You can obtain a list of candidate nodes (those that are not already assigned to a cluster) by typing **svcinfo lsnodecandidate**.

For SAN Volume Controller Storage Software for Cisco MDS 9000, this command will first complete a compatibility check to ensure that the correct level of software can be run on the node before the new node is added to the cluster.

If the compatibility check fails, the following message displays:

```
CMMVC6201E The node could not be added, because incompatible
software: status code [%1].
```

Prerequisites: Before adding a node to the cluster, check the following:

- The cluster has more than one I/O group.
- The node being added to the cluster uses physical node hardware that has previously been used as a node in the cluster.

- The node being added to the cluster uses physical node hardware which has previously been used as a node in another cluster and both clusters have visibility of the same hosts.

Attention: If the conditions listed above apply, then failure to follow the procedures documented here, may result in the corruption of all data managed by the cluster.

Adding a node: If you are adding the node into the cluster for the first time, you must record the card serial number, the slot the Cisco MDS 9000 Caching Services Module is in, the node (1 or 2) on the Cisco MDS 9000 Caching Services Module, the WWNN, all WWPN's, and the I/O group to which it has been added. This can avoid a possible data corruption exposure if the node must be removed from and re-added to the cluster.

When a node is added to the cluster using the **svctask addnode** command or using the Console for SAN Volume Controller for Cisco MDS 9000, if the node has previously been a member of the cluster then either:

- the node must be added back to the same I/O group that it was previously in. The WWNN of the nodes in the cluster can be determined using the **svcinfolsnode** command. Or,
- if this information is not available, call IBM Service to add the node back into the cluster without corrupting the data.

If your attempt to add a node to the cluster results in two nodes from the same Cisco MDS 9000 Caching Services Module being put in the same I/O group, the **addnode** command will fail with the error:

CMMVC5777E The node could not be added to the I/O group, because the other node in the I/O group is in the same power domain.

Optionally, you can assign a name to the new node. You can use this name in subsequent commands to refer to the node, instead of using the node ID. If you assign a label, this label is displayed as the node name from then on. If you do not assign a label, the default label is nodeX, where X is the node ID.

Example: When 2 nodes of a 4 node cluster have been lost because of a complete failure of a Cisco MDS 9000 Caching Services Module, the two lost nodes must be added back into the cluster using the **svctask addnode** command or using the Console for SAN Volume Controller for Cisco MDS 9000.

Applications on a host system direct I/O operations to filesystems or logical volumes which are mapped by the operating system to vpaths, which are pseudo disk objects supported by the SDD driver. See the *IBM TotalStorage Subsystem Device Driver: User's Guide* for more information.

The SDD driver maintains an association between a vpath and a SAN Volume Controller Storage Software for Cisco MDS 9000 VDisk. This association uses an identifier (UID), which is unique to the VDisk and is not reused. This enables the SDD driver to unambiguously associate vpaths with VDIs.

The SDD device driver operates within a protocol stack, which also contains disk and fibre-channel device drivers that enable it to communicate with the SAN Volume Controller Storage Software for Cisco MDS 9000 using the SCSI protocol over fibre-channel as defined by the ANSI FCS standard. The addressing scheme

provided by these SCSI and fibre-channel device drivers uses a combination of a SCSI logical unit number (LUN) and the world wide name for the fibre-channel node and ports.

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

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

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5790E The node was not added to the cluster because the maximum number of nodes has been reached.
- CMMVC5791E The action failed because an entity that was specified in the command does not exist.
- CMMVC5792E The action failed because the I/O group is used for recovery.
- CMMVC5793E The node was not added to the cluster because the I/O group already contains a pair of nodes.
- CMMVC5777E The node was not added to the I/O group because the other node in the I/O group is in the same power domain.
- CMMVC6201E The node could not be added, because incompatible software: status code [%1].

Examples

An invocation example

```
svctask addnode -wwnodename 210000e08b053564 -iogrp io_grp0
```

The resulting output

```
Node, id [6], successfully added
```

chcluster

You can use the **chcluster** command to modify the attributes of an existing cluster. You can enter this command any time after a cluster has been created.

Syntax

```
▶▶—svctask— —chcluster— [—clusterip—cluster_ip_address]▶▶
```

Parameters

-clusterip *cluster_ip_address*

Changes the cluster IP address. After the cluster IP address is changed, you lose the open shell connection to the cluster. You must reconnect with the newly specified IP address.

Description

This command modifies the IP addresses that are assigned to the cluster. Any of the arguments can be used either singly or in conjunction with any of the other arguments.

If the cluster IP address is changed, the open command-line shell closes during the processing of the command. You must reconnect to the new IP address.

If you do not specify the option, the command does nothing.

Modifying an IP address: List the IP address of the cluster, by issuing the **svcinfo lscluster** command. Modify the IP address, by issuing the **svctask chcluster** command.

This command enables you to change the settings for the following:

- Cluster IP address

If you specify a new cluster IP address, the existing communication with the cluster is broken.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC6202E The cluster was not modified because the IP address is not valid.

Examples

An invocation example

```
svctask chcluster -clusterip 217.12.3.11
```

The resulting output

No feedback

chiogrp

You can use the **chiogrp** command to modify the name assigned to an I/O group.

Syntax

```
▶▶—svctask— —chiogrp— —-name—new_name_arg— —io_group_id— —io_group_name—▶▶
```

Parameters

-name *new_name_arg*

Specifies the name to assign to the I/O group.

-io_group_id | **io_group_name**

Specifies the I/O group to modify by identifying either the I/O group ID or the name already assigned to the I/O group.

Description

This command sets the name of the I/O group specified to the new name supplied.

The I/O groups already exist by default when a cluster is created, but will not contain any nodes. The first node in the cluster is always assigned to I/O group zero.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5800E The action failed because an entity that was specified in the command does not exist.
- CMMVC5792E The action failed because the I/O group is used for recovery.

Examples

An invocation example

```
svctask chiogrp -name testiogrhone io_grp0
```

The resulting output

No feedback

chnode

You can use the **chnode** command to change the name assigned to a node.

Syntax

```
svctask chnode --name new_node_name node_name  
node_id
```

Parameters

-name *new_node_name*

Specifies the name to assign to the node.

node_name | node_id

Specifies the node to be modified. The argument that follows the flag is either:

- The node name, that is, the label that you assigned when you added the node to the cluster
- The node ID that is assigned to the node (not the WWNN).

Description

This command changes the name, or label, assigned to the node in question. The name can be used from then on in subsequent command-line tools.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5798E The action failed because the node is offline.

Examples

An invocation example

```
svctask chnode -name testnodeone nodeone
```

The resulting output

No feedback

cleardumps

You can use the **cleardumps** command to clean the various dump directories on a specified node.

Syntax

```
svctask cleardumps --prefix directory_or_file_filter
node_id | node_name
```

Parameters

-prefix *directory_or_file_filter*

Specifies the directory, or files, or both to be cleaned. If a directory is specified, with no file filter, all relevant dump or log files in that directory are cleaned.

The directory arguments include:

- /dumps (cleans all files in all subdirectories)
- /dumps/configs
- /dumps/ec
- /dumps/elogs
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace
- /dumps/pl

In addition to the directory, a file filter can be specified. For example, if you specified `/dumps/elogs/*.txt` all files in the `/dumps/elogs` directory that end in `.txt` will be cleaned.

node_id | node_name

Optionally specifies the node to be cleaned. The argument that follows the flag is either:

- The node name, that is, the label that you assigned when you added the node to the cluster
- The node ID that is assigned to the node (not the WWNN).

Description

This command deletes all the files that match the `directory/file_filter` argument on the node specified. If no node is specified, the configuration node is cleaned.

You can clean all the dumps directories by specifying `/dumps` as the directory argument.

You can clean all the files in a single directory by specifying one of the directory arguments.

You can clean specific files in a given directory by specifying a directory and file name. You can use an asterisk (*) as a wildcard in part of the file name. However, you might need to quote the file name when the wildcards are used.

You can list the contents of these directories on the given node by using the `svcinfo lsxxxxdumps` commands.

Possible failures

- CMMVC6203E The action failed because the directory that was specified was not one of the following directories: `/dumps`, `/dumps/iostats`, `/dumps/iotrace`, `/dumps/feature`, `/dumps/configs`, `/dumps/elogs`, `/dumps/ec` or `/dumps/pl`.

Examples

An invocation example

```
svctask clear.dumps -prefix /dumps/configs
```

The resulting output

No feedback

cpdumps

You can use the **cpdumps** command to copy dump files from a nonconfiguration node onto the configuration node.

Note: In the rare event that the /dumps directory on the configuration node is full, the copy action will terminate when the directory is full and will provide no indicator of a failure. It is therefore advisable to clear the /dumps directory after the desired data has been migrated off the configuration node.

Syntax

```
svctask --cpdumps --prefix directory node_name  
                                file_filter node_id
```

Parameters

-prefix *directory* | *file_filter*

Specifies the directory, or files, or both to be retrieved. If a directory is specified with no file filter, all relevant dump or log files in that directory are retrieved. The directory arguments include:

- /dumps (cleans all files in all subdirectories)
- /dumps/configs
- /dumps/ec
- /dumps/elogs
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace
- /dumps/pl

In addition to the directory, you can specify a file filter. For example, if you specified /dumps/elogs/*.txt all files in the /dumps/elogs directory that end in .txt will be copied.

node_id | **node_name**

Specifies the node from which to retrieve the dumps. The argument that follows the flag is either:

- The node name, that is, the label that you assigned when you added the node to the cluster
- The node ID that is assigned to the node (not the WWNN).

If the node specified is the current configuration node, no file will be copied.

Description

This command copies any dumps that match the directory or file criteria from the given node to the current configuration node.

You can retrieve dumps that were saved to an old configuration node. When the old configuration node failed over to another node, the dumps that were on the old configuration node are not automatically copied. Because access from the IBM CLI is only provided to the configuration node, files can only be copied off the

cluster from the configuration node. This command enables you to retrieve files and place them on the configuration node so that you can then copy them off of the cluster.

You can view the contents of the directories by using the `svcinfo lsxxxxdumps` commands.

You can copy dumps off any node using the SAN Volume Controller for Cisco MDS 9000 CLI copy command. See the *Cisco MDS 9000 Family Command Reference* for more information.

Possible failures

- CMMVC6203E The action failed because the directory that was specified was not one of the following directories: `/dumps`, `/dumps/iostats`, `/dumps/iotrace`, `/dumps/feature`, `/dumps/configs`, `/dumps/elogs`, `/dumps/ec` or `/dumps/pl`.

Examples

An invocation example

```
svctask cpdumps -prefix /dumps/configs nodeone
```

The resulting output

No feedback

detectmdisk

You can use the **detectmdisk** command to manually rescan the fibre-channel network for any new managed disks that might have been added.

Syntax

►—svctask— —detectmdisk—◄

Description

This command causes the cluster to rescan the fibre-channel network and looks for managed disks that have not been seen before. This command also detects if a controller is being decommissioned.

In general, the cluster detects these disks automatically when they appear on the network. However, some fibre-channel controllers do not send the required SCSI primitives that are necessary to automatically discover the new disks.

If you have attached some new storage and the cluster has not detected it, you might need to run this command before the cluster will detect the new disks.

No parameters are required.

When back-end controllers are added to the fibre-channel SAN and are included in the same switch zone as a SAN Volume Controller Storage Software for Cisco MDS 9000 Cluster the cluster will automatically discover the back-end controller and will integrate the controller to determine what storage it is presented to the SAN Volume Controller Storage Software for Cisco MDS 9000. The SCSI LUs presented by the back-end controller will be displayed as unmanaged MDisks. If however the configuration of the back-end controller is modified after this has occurred then the SAN Volume Controller Storage Software for Cisco MDS 9000 may be unaware of these configuration changes. This task allows a user to request the SAN Volume Controller Storage Software for Cisco MDS 9000 to re-scan the fibre-channel SAN to update the list of unmanaged MDisks.

Note: The automatic discovery performed by SAN Volume Controller Storage Software for Cisco MDS 9000 does not write anything to a unmanaged MDisk. It is only when a the user instructs the SAN Volume Controller Storage Software for Cisco MDS 9000 to add a MDisk to a managed disk group or use a MDisk to create an image mode virtual disk that the storage will actually be used.

Discovering MDisks: Check to see which MDisks are available by issuing the **svctask detectmdisk** command to manually scan the fibre-channel network for any MDisks. Issue the **svcinfo lsmdiskcandidate** command to show the unmanaged MDisks. These MDisks have not been assigned to an MDisk group. Alternatively, you can issue the **svcinfo lsmdisk** command to view all of the MDisks.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svctask detectmdisk
```

The resulting output

```
No feedback
```

dumpconfig

You can use the **dumpconfig** command to dump the entire configuration of the cluster to a text file.

Syntax

```
svctask dumpconfig [-prefix filename_prefix]
```

Parameters

-prefix *filename_prefix*

Optionally specifies the file name where you want to send the dump data. If **-prefix** is not supplied, then the dump is directed to a file with a system-defined prefix of "config". The system-defined file name is created from the prefix and a timestamp. The file name takes the following format:

```
<prefix>_NNNNNN_YYMMDD_HHMMSS
```

where *NNNNNN* is the node front panel name.

Description

The dumps are written to the `/dumps/configs` directory on the node. If the **-prefix** parameter is not used then the prefix, `config`, will be added at the front of the file name.

If the **-prefix** parameter is entered then the user entered prefix will be used when naming the file.

A maximum of ten configuration dump files is kept on the cluster. When the eleventh dump is made, the oldest existing dump file is overwritten.

The contents of the `/dumps/configs` directory can be listed using the **svcinfo lsconfigdumps** command.

Possible failures

- CMMVC5983E The dump file was not created. The file system might be full.
- CMMVC5984E The dump file was not written to disk. The file system might be full.
- CMMVC6073E The maximum number of files has been exceeded.

Examples

An invocation example

```
svctask dumpconfig -prefix mydumpfile
```

The resulting output

The configuration data has been written to

rmnode

You can use the **rmnode** command to delete a node from the cluster. You can enter this command any time after a cluster has been created.

Syntax

```
► —rmnode— [node_name | node_id] ►
```

Parameters

node_name | node_id

Specifies the node to be deleted. The argument is either:

- The node name, that is, the label that you assigned when you added the node to the cluster
- The node ID that is assigned to the node (not the WWNN).

Description

This command removes a node from the cluster. This makes the node a candidate to be added back into this cluster or into another cluster. After the node is deleted, the other node in the I/O group destages the contents of its cache and goes into write-through mode until another node is added back into the I/O group.

Prerequisites:

Before you issue the **rmnode** command, perform the following tasks and read the following Attention notices to avoid losing access to data:

1. Determine which virtual disks (VDisks) are still assigned to this I/O group by issuing the following command. The command requests a filtered view of the VDIs, where the filter attribute is the I/O group.

```
svcinfolsvdisk -filtervalue IO_group_name=<name>
```

where <name> is the name of the I/O group in question.

Notes:

- a. If this is the last node in an I/O group and if there are virtual disks still assigned to the I/O group, you will *not* be able to delete the node from the cluster.
 - b. Any VDIs that are assigned to the I/O group that this node belongs to, will be assigned to the other node in the I/O group; that is, the preferred node will be changed. You cannot change this setting back.
2. Determine the hosts that the VDIs are mapped to by issuing the **svcinfolsvdiskhostmap** command.
 3. Determine if any of the VDIs assigned to this I/O group contain data that you need to maintain access to:
 - If you *do not* want to maintain access to these VDIs, go to step 5 on page 20.
 - If you *do* want to maintain access to some or all of the VDIs, back up the data or migrate the data to a different (online) I/O group.
 4. Determine if you need to shut down the node:

- If this is the last node in the cluster, you do not need to shut down the node. Go to step 5.
- If this is not the last node in the cluster, shut down the node you intend to remove using the `stopcluster` node command. This step ensures that the Subsystem Device Driver (SDD) does not rediscover the paths that are manually removed before you issue the delete node request.

If you plan to add the node back into the cluster, see 20.

5. Update the SDD configuration for each virtual path (vpath) that is presented by the VDisks that you intend to remove. Updating the SDD configuration removes the vpaths from the VDisks. Failure to update the configuration can result in data corruption. See the *IBM TotalStorage Subsystem Device Driver: User's Guide* for details about how to dynamically reconfigure SDD for the given host operating system.
6. Quiesce all I/O operations that are destined for the node you are deleting. Failure to quiesce the operations could result in failed I/O operations being reported to your host operating systems.

Attention: Removing the last node in the cluster destroys the cluster. Before you delete the last node in the cluster, ensure that you want to destroy the cluster.

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

Notes:

1. If the node you are removing is the configuration node, it might take a minute or more before you can perform the delete node request. You must wait for the configuration node failover to occur.
2. If the node you are removing is the last node in the cluster, the Console for SAN Volume Controller for Cisco MDS 9000 might seem to hang for up to 3 minutes because you have removed the last access point to the cluster.

Deleting a node from a cluster:

Notes:

1. If this is the last node in the I/O group or the last node in the cluster, you will be asked to force the deletion.
2. If this is the last node in the cluster or if it is currently assigned as the configuration node, all connections to the cluster will be lost. The user interface and any open CLI sessions will halt. A time-out might occur if a command cannot be completed before the node is deleted.

Issue the `svctask rmnode` command to delete a node from the cluster. You can enter this command any time after a cluster has been created.

Adding a node back into the cluster:

If you turn the power back on to the node that has been removed while it is still connected to the same fabric or zone, the following actions occur:

1. The node attempts to join the cluster again.
2. The cluster signals the node to remove itself from the cluster.
3. The node becomes a candidate for addition to this cluster or another cluster.

If you intend to add this node back into the cluster, ensure that you add it back to the same I/O group from which you are deleting it. Otherwise, data corruption might occur.

Before you add a node back into the cluster, you need to know the following information, which should have been recorded when the node was originally added to the cluster.

- card serial number
- slot the Cisco MDS 9000 Caching Services Module is in
- node (1 or 2) on the Cisco MDS 9000 Caching Services Module
- WWNN
- all WWPNs
- I/O group that contains the node

If you do not have access to this information, call IBM Service to add the node back into the cluster without corrupting the data.

Replacing a faulty node:

1. Verify the name or ID of the node you want to remove by issuing the **svcinfolnode** command.
2. Select one of the following actions based on why you are replacing the node:
 - a. If the node was faulty, it will be shown as offline.
 - 1) Ensure that the partner node in the I/O group is online.
 - 2) If the other node in the I/O group is offline, start the directed maintenance procedures to determine the fault.
 - b. If you have been directed here by the directed maintenance procedures because the partner node in the I/O group has subsequently failed, follow the procedure for recovering from offline VDisks after a node or an I/O group failed.
 - c. If you are replacing the node for other reasons:
 - 1) Determine which node you want to replace.
 - 2) Ensure that the partner node in the I/O group is online.

Note: If the partner node is offline and you continue, you will lose access to the VDisks that belong to this I/O group. Start the directed maintenance procedures, and fix the other node before proceeding.

3. Remove the node from the cluster by issuing the **svctask rmnode <nodename/id>** command. (First, determine if you need to shut down the node. If this is the last node in the cluster, you do not need to shut down the node. Go to step 5 on page 20. If this is *not* the last node in the cluster, shut down the node you intend to remove using the **stopcluster node** command.)
4. Verify that the node can be seen on the fabric by issuing the **svcinfolnodecandidate** command. You should see the node listed as a candidate.
5. Make a note of the WWNNs for each node.
6. If the node was repaired by replacing it with another node, the WWNN for the node will change. You can modify the WWNN of the replacement node so that it matches that of the old node. To change the WWNN and WWPN so they match the old node, see your *Cisco MDS 9000 Family Command Reference*. After a node has been replaced, the service procedures guide you through changing the WWNN. If this is the case, complete the following additional steps:

- a. At the end of the recovery process, follow the SDD procedure to discover the new paths and to check that each vpath is now presenting the correct number of paths. See the *IBM TotalStorage Subsystem Device Driver: User's Guide* sections on dynamic reconfiguration, specifically adding paths to existing vpaths.
- b. You might also need to modify the configuration of your disk controllers. If your controller uses a mapping technique to present its RAID's or partitions to the cluster, you need to modify the port groups that belong to the cluster because the WWNN or WWPNS of the node have changed.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5791E The action failed because an entity that was specified in the command does not exist.
- CMMVC5794E The action failed because the node is not a member of the cluster.
- CMMVC5795E The node was not deleted because a software upgrade is in progress.
- CMMVC5796E The action failed because the I/O group that the node belongs to is unstable.
- CMMVC5797E The node was not deleted because this is the last node in the I/O group and there are virtual disks (VDisks) associated with the I/O group.

Examples

An invocation example

```
svctask rmnode 1
```

The resulting output

No feedback

Related topics:

- "lsvdiskhostmap" on page 254
- "lsnode" on page 232

startstats

You can use the **startstats** command to start the collection of statistics for both VDisks and MDisks.

Syntax

```
svctask — startstats — -interval—time_in_minutes—
```

Parameters

-interval *time_in_minutes*

Specifies the time in minutes. This is the time interval between the gathering of statistics, between 15 and 60 minutes in increments of 1.

Description

Statistics are collected at the end of each sampling period (as specified by the **-interval** parameter). These statistics are written to a file. A new file is created at the end of each sampling period. Separate files are created for managed disks and virtual disks statistics.

The files generated are written to the `/dumps/iostats` directory.

A maximum of 12 files are stored in the directory at any one time for each disk type, for example `m_xxx` and `v_xxx` files. Before the 13th file (for each type) is created, the oldest file of that type will be deleted.

These files can be listed by using the **svcinfolsiostats** command.

The naming convention for these files is:

`<disk_type>_stats_<frontpanelid>_<date>_<time>`. Where *<disk_type>* is *m* for managed disks and *v* for virtual disks, *<frontpanelid>* is the current configuration node ID, *<date>* is in the form of *yyymmdd*, and *<time>* is in the form of *hhmmss*.

An example of a managed disk file name is: `m_stats_1_030808_105224`. An example of a virtual disk file name is: `v_stats_vegas8.2.1_031123_072426`.

The statistics collected for each managed disk and each virtual disk are:

- The number of SCSI read commands processed during the sample period.
- The number of SCSI write commands processed during the sample period.
- The number of blocks of data read during the sample period.
- The number of blocks of data written during the sample period.

Possible failures

- There are no error codes.

Examples

An invocation example

```
svctask startstats -interval 25
```

The resulting output

```
No feedback
```

stopcluster

You can use the **stopcluster** command to shut down a single node or the entire cluster in a controlled manner.

Syntax

```
svctask --stopcluster [-force] [-node node_name | node_id]
```

Parameters

-node *node_name* | *node_id*

Optionally identifies the node you want to shut down. Specifies the node to be stopped. The argument that follows the flag is either:

- The node name, that is, the label that you assigned when you added the node to the cluster
- The node ID that is assigned to the node (not the WWNN).

If you supply the node ID or the name, only that node will be shut down, otherwise the entire cluster will be shutdown.

-force

The force flag is required if this is the last online node in a given I/O group.

Description

When you enter this command without any arguments, the entire cluster is shut down. All data is flushed to disk before the power is removed.

Attention: Ensure that you have stopped all FlashCopy, Remote Copy, and data migration operations before you attempt a node or cluster shutdown. You should also ensure that all asynchronous deletion operations have completed prior to a shutdown operation.

When you enter this command either a node ID or node name argument, the node in question is shut down. After the command completes, the other node in the I/O group destages the contents of its cache and goes into write-through mode until the power to the node is returned and the node rejoins the cluster.

If all input power to a SAN Volume Controller Storage Software for Cisco MDS 9000 cluster is to be removed for more than a few minutes, (for example, if the machine room power is to be shutdown for maintenance), it is important that the cluster is shutdown before the power is removed.

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

Begin the process of quiescing all I/O to the cluster by stopping the applications on your hosts that are using the VDisks provided by the cluster.

1. If you are unsure which hosts are using the VDisks provided by the cluster, follow the procedure called, Determining the hosts that a VDisk is mapped to.
2. Repeat the previous step for all VDisks.

Attention: If you are shutting down the entire cluster, you will lose access to all VDIs being provided by this cluster.

When all I/O has been stopped, issue the `svctask stopcluster` to shut down a single node or the entire cluster in a controller manner. If you specify the node ID or node name, you can shut down a single node. When you enter this command either a node ID or node name argument, the node in question is shut down. After the command completes, the other node in the I/O group destages the contents of its cache and goes into write-through mode until the power to the node is returned and the node rejoins the cluster. You must then close the SSH session if you are using ssh in interactive mode.

Attention: If this is the last node in an I/O group, you will lose all access to the virtual disks in the I/O group. Before you enter this command, ensure that this is what you want to do. You must specify the force flag.

If a node or cluster is shutdown and power remains up, the nodes will be placed in dormant state. To restore the node or nodes, issue the Cisco CLI node `svc x/y reload` command for each node and the node(s) will be restarted. If the switch that the node is in is powered off, then the node(s) in that switch will come back up when power is restored.

Shutting down a single node::

Attention: If you are shutting down a single node, and the other node in the I/O group is online, be aware that the cache on the partner node will go into write-through mode and that you are exposed to a single point of failure should the partner node fail while this node is shut down. Proceed to 2.

Attention: If you are shutting down a single node, and this is the last node in the I/O group, you will lose access to all VDIs being served by this I/O group.

Steps:

Perform the following steps to shut down a single node:

1. Begin the process of quiescing all I/O to the VDIs being served by this nodes I/O group.
 - a. Determine the VDIs in question by requesting a filtered view of VDIs where the filter attribute is the I/O group in question. This can be done using the following command:

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

where <name> is the name of the I/O group in question.
 - b. Once you have a list of VDIs, determine the hosts that these are mapped to by following the procedure called, Determining the hosts that a VDisk is mapped to.
2. When all I/O has been stopped issue the following command to shut down the node:

```
svctask stopcluster <nodename/ID>
```

where <nodename/ID> is the name or ID of the node that you want to shut down.

Note: If this is the last node in the I/O group you also need to specify the `-force` parameter. For example to force the shutdown of node1:

```
svctask stopcluster -force node1
```

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5798E The action failed because the node is offline.
- CMMVC5791E The action failed because an entity that was specified in the command does not exist.
- CMMVC5796E The action failed because the I/O group that the node belongs to is unstable.
- CMMVC5799E The shutdown was not successful because there is only one online node in the I/O group.

Examples

An invocation example

```
svctask stopcluster
```

The resulting output You will be presented with the following warning:

```
Are you sure that you want to continue with the shut down?
```

Ensure that you have stopped all FlashCopy mappings, Remote Copy relationships, data migration operations and forced deletions before continuing. Entering y to this will execute the command. No feedback is then displayed. Entering anything other than y or Y will result in the command not executing. No feedback is displayed.

stopstats

You can use the **stopstats** command to stop the collection of statistics for both VDisks and MDisks.

Syntax

```
svctask —stopstats
```

Description

This command turns off the statistics generation, until you start them again (with **svctask startstats**).

Possible failures

- There are no error codes.

Examples

An invocation example

```
svctask stopstats
```

The resulting output

No feedback

Chapter 4. Cluster diagnostic and service-aid commands

This section lists and describes the cluster diagnostic and service-aid commands.

When some commands complete successfully, textual output is normally provided. However, some commands do not provide any output. The phrase No feedback is used to indicate that no output is provided.

The SAN Volume Controller Storage Software for Cisco MDS 9000 enables service activity to be performed with a limited set of command-line tools. When you are logged in under the administrator role, all command-line activities are permitted. When you are logged in under the service role, only those commands required for service are enabled. All of these commands apply under the service role. The service commands permit problem determination and repair activities to be performed. The service commands do not permit any activity that might cause your data to become corrupted or inaccessible or that might compromise the security of the cluster.

addnode

You can use the **addnode** command to add a new (candidate) node to an existing cluster. You can enter this command any time after a cluster has been created.

Syntax

```
▶▶ svc servicetask — addnode — [ -panelname — switchname.slot.node ] —▶▶
                                   [ -wwnodename — wwnn_arg ]
▶ [ -name — new_name_arg ] — -iogrp — [ iogroup_name | iogroup_id ] —▶▶
```

Parameters

-panelname *switchname.slot.node*

Identifies the node to be added by the name as it appears in the output of the **svcinfo lsnodecandidate** command. The name is composed of *switchname.slot.node*. *Switchname* is the name of the switch with certain characters replaced by an underscore. *Slot* is the slot of the SAN Volume Controller Storage Software for Cisco MDS 9000 of the node (1-9). *Node* is the node in the Cisco MDS 9000 Caching Services Module (1 or 2). This argument is mutually exclusive with *-wwnodename*; only one is required to uniquely identify the node.

-wwnodename *wwnn_arg*

Identifies the node to add to the cluster by the worldwide node name (WWNN) of the node. This argument is mutually exclusive with *-panelname*; only one is required to uniquely identify the node.

-name *new_name_arg*

Optionally specifies a name for this node.

-iogrp *iogroup_name* | *iogroup_id*

Specifies the I/O group to which you want to add this node.

Description

This command adds a new node to the cluster. You can obtain a list of candidate nodes (those that are not already assigned to a cluster) by typing **svcinfo lsnodecandidate**.

For SAN Volume Controller Storage Software for Cisco MDS 9000, this command will first complete a compatibility check to ensure that the correct level of software can be run on the node before the new node is added to the cluster.

If the compatibility check fails, the following message displays:

```
CMMVC6201E The node could not be added, because incompatible
software: status code [%1].
```

Prerequisites: Before adding a node to the cluster, check the following:

- The cluster has more than one I/O group.
- The node being added to the cluster uses physical node hardware which has previously been used as a node in the cluster.

- The node being added to the cluster uses physical node hardware which has previously been used as a node in another cluster and both clusters have visibility of the same hosts.

Attention: If the conditions listed above apply, then failure to follow the procedures documented here, may result in the corruption of all data managed by the cluster.

Adding a node: If you are adding the node into the cluster for the first time, you must record the card serial number, the slot the Cisco MDS 9000 Caching Services Module is in, the node (1 or 2) on the Cisco MDS 9000 Caching Services Module, the WWNN, all WWPN's, and the I/O group to which it has been added. This can avoid a possible data corruption exposure if the node must be removed from and readded to the cluster.

When a node is added to the cluster using the **svctask addnode** command or using the Console for SAN Volume Controller for Cisco MDS 9000, if the node has previously been a member of the cluster then either:

- the node must be added back to the same I/O group that it was previously in. The WWNN of the nodes in the cluster can be determined using this **svcinfo lsnode** command. Or,
- if this information is not available, call IBM Service to add the node back into the cluster without corrupting the data.

When you add a node to the cluster, you must also specify the I/O group to which this node will belong. I/O groups are identifiers for a node pair. The node pairs will internally duplicate cache data for a given set of virtual disks for redundancy of the cache data. When a virtual disk is created, it is also assigned to an I/O group. All data that is destined for this virtual disk is then cached on the two nodes in the I/O group.

When you add a node, make sure that the nodes within an I/O group are attached to different uninterruptible power supplies. You can determine the uninterruptible power supply to which a node is attached from the output of the **svcinfo lsnodecandidate** (uninterruptible power supply_unique_ID). The **svcinfo lsnodes** command displays the uninterruptible power supply to which all nodes (uninterruptible power supply_unique_ID), in the cluster are attached.

If you attempt to add a node to the cluster, it results in two nodes being attached to the same uninterruptible power supply, being put in the same I/O group, the **svcservicetask addnode** command will fail with the following error:

```
CMMVC5777E The node could not be added to the I/O group, because
the other node in the I/O group is in the same power domain.
```

Optionally, you can assign a name to the new node. You can use this name in subsequent commands to refer to the node, instead of using the WWNN (node ID). If you assign a label, this label is displayed as the node name from then on. If you do not assign a label, the default label is nodeX, where X is the node ID.

Example: When 2 nodes of a 4 node cluster have been lost because of a complete failure of an uninterruptible power supply, the two lost nodes must be added back into the cluster using the **svcservicetask addnode** command or using the Console for SAN Volume Controller for Cisco MDS 9000.

Applications on a host system direct I/O operations to filesystems or logical volumes which are mapped by the operating system to vpaths, which are pseudo disk objects supported by the SDD driver. See the *IBM TotalStorage Subsystem Device Driver: User's Guide* for more information.

The SDD driver maintains an association between a vpath and a SAN Volume Controller Storage Software for Cisco MDS 9000 VDisk. This association uses an identifier (UID), which is unique to the VDisk and is not reused. This enables the SDD driver to unambiguously associate vpaths with VDIs.

The SDD device driver operates within a protocol stack, which also contains disk and fibre-channel device drivers that enable it to communicate with the SAN Volume Controller Storage Software for Cisco MDS 9000 using the SCSI protocol over fibre-channel as defined by the ANSI FCS standard. The addressing scheme provided by these SCSI and fibre-channel device drivers uses a combination of a SCSI logical unit number (LUN) and the world wide name for the fibre-channel node and ports.

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

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

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5790E The node was not added to the cluster because the maximum number of nodes has been reached.
- CMMVC5791E The action failed because an entity that was specified in the command does not exist.
- CMMVC5792E The action failed because the I/O group is used for recovery.
- CMMVC5793E The node was not added to the cluster because the I/O group already contains a pair of nodes.
- CMMVC5777E The node was not added to the I/O group because the other node in the I/O group is in the same power domain.
- CMMVC6201E The node could not be added, because incompatible software: status code [%1].

Examples

An invocation example

```
svcservicetask addnode -wwnodename 210000e08b053564 -iogrp io_grp0
```

The resulting output

```
Node, id [6], successfully added
```

cherrstate

You can use the **cherrstate** command to mark an unfixed error as fixed. You can also use it to mark a fixed error as unfixed.

Syntax

```
►—svcservicetask— —cherrstate— —-sequencenumber—sequence_number— —————►  
└─┬──────────┘  
  -unfix-
```

Parameters

-sequencenumber *sequence_number*

Specifies the error log sequence number, or numbers, to fix.

-unfix

Optionally specifies that the sequence number, or numbers, supplied should be marked as unfixed. If you supply the **-unfix** argument, the sequence numbers will be marked as unfixed. This is intended for use only when you have marked the wrong sequence number as fixed.

Description

The error log entries that the sequence number, or numbers, that you entered are marked as fixed. Use this command as a manual confirmation step that you have performed some maintenance to the cluster, fabric, or subsystems.

This step is performed as part of the directed maintenance procedures (DMPs).

Optionally, if you have marked the wrong sequence number as fixed, you can remark an entry as unfixed by specifying the **-unfix** flag.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5803E The entry in the error log was not marked because the sequence number was not found.

Examples

An invocation example

```
svcservicetask cherrstate -sequencenumber 2019
```

The resulting output

```
No feedback
```

clearerrlog

You can use the **clearerrlog** command to clear all entries from the error log including status events and any unfixed errors.

Syntax

```
▶▶—svcservicetask— —clearerrlog— —force—▶▶
```

Parameters

-force

This flag stops any confirmation requests. If the **-force** flag is not supplied, you are prompted to confirm if you are sure that you want to clear the log.

Description

This command clears all entries from the error log. The entries are cleared even if there are unfixed errors in the log. It also clears any status events that are in the log.

Attention: This command is destructive. You should only use it when you have either rebuilt the cluster, or have fixed a major problem that has caused many entries in the error log that you do not want to manually fix.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcservicetask clearerrlog -force
```

The resulting output

```
No feedback
```

dumperrlog

You can use the **dumperrlog** command to dump the contents of the error log to a text file. You can also use the command to delete unwanted error log dumps from the cluster.

Syntax

```
▶▶—svcservicetask— —dumperrlog— —[-prefix—filename_prefix—]▶▶
```

Parameters

-prefix *filename_prefix*

A file name is created from the prefix and a time stamp, and has the following format:

```
<prefix>_NNNNNN_YYMMDD_HHMMSS
```

where *NNNNNN* is the node front panel name.

Note: If the **-prefix** parameter is not supplied, the dump will be directed to a file with a system-defined prefix of "errlog".

Description

When executed with no arguments, this command dumps the cluster error log to a file using a system-supplied prefix of "errlog", which includes the node ID and time stamp. When a file name prefix is provided, the same operation is performed but the details are stored in the dumps directory a file with a name that starts with the specified prefix.

A maximum of ten error-log dump files are kept on the cluster. When the 11th dump is made, the oldest existing dump file is overwritten.

Error log dump files are written to `/dumps/elogs`. The contents of this directory can be viewed using the **svcinfo lserrlogdumps** command.

Possible failures

- CMMVC5983E The dump file was not created. The file system might be full.
- CMMVC5984E The dump file was not written to disk. The file system might be full.

Examples

An invocation example

```
svcservicetask dumperrlog -prefix testerrorlog
```

The resulting output

No feedback

finderr

You can use the **finderr** command to analyze the error log for the highest severity unfixed error.

Syntax

```
►—svcservicetask— —finderr—◄
```

Description

The command scans the error log for any unfixed errors. Given a priority ordering defined within the code, the highest priority unfixed error is returned to stdout.

You can use this command to determine the order in which to fix the logged errors.

The Web-based directed maintenance procedures (DMPs) also use this command.

Possible failures

- There are no error codes.

Examples

An invocation example

```
svcservicetask finderr
```

The resulting output

```
Highest priority unfixed error code is [1010]
```

rmnode

You can use the **rmnode** command to delete a node from the cluster. You can enter this command any time after a cluster has been created.

Syntax

```
►—svcservicetask— —rmnode— —{node_name | node_id}—►
```

Parameters

node_name | node_id

Specifies the node to be deleted. The argument that follows the flag is either:

- The node name, that is, the label that you assigned when you added the node to the cluster
- The node ID, that is, assigned to the node (not the WWNN).

Description

This command removes a node from the cluster. This makes the node a candidate to be added back into this cluster or into another cluster. After the node is deleted, the other node in the I/O group destages the contents of its cache and goes into write-through mode until another node is added back into the I/O group.

If you intend to add the node back into the cluster, you must ensure that you record the card serial number, the slot the Cisco MDS 9000 Caching Services Module is in, the node (1 or 2) on the Cisco MDS 9000 Caching Services Module, the WWNN, all WWPNs, and the I/O group to which it currently belongs. This can avoid a possible data corruption exposure if the node must be removed from and readded to the cluster. See the *IBM TotalStorage SAN Volume Controller: Configuration Guide* for more information.

If this is the last node in an I/O group and if there are virtual disks still assigned to the I/O group, you will *not* be able to delete the node from the cluster.

If this is the last node in the cluster and the I/O group has no virtual disks remaining, the cluster is deleted and all virtualization information is lost. Any data that is still required should be backed up or migrated prior to destroying the cluster.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5791E The action failed because an entity that was specified in the command does not exist.
- CMMVC5794E The action failed because the node is not a member of the cluster.
- CMMVC5795E The node was not deleted because a software upgrade is in progress.
- CMMVC5796E The action failed because the I/O group that the node belongs to is unstable.
- CMMVC5797E The node was not deleted because this is the last node in the I/O group and there are virtual disks (VDisks) associated with the I/O group.

Examples

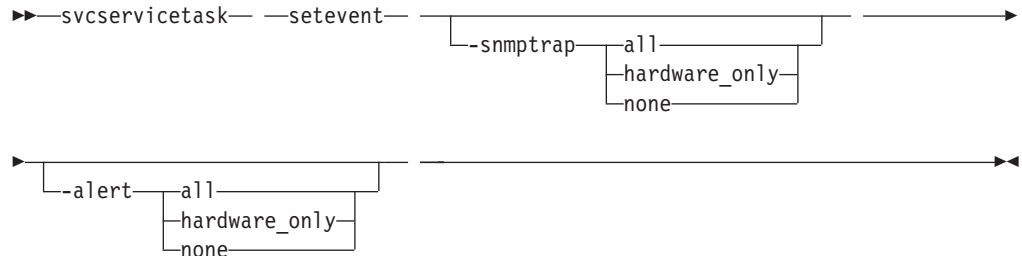
An invocation example Introduce the executed command
`svcservicetask rmnode 1`

The resulting output Introduce the resulting output
No feedback

setevent

You can use the **setevent** command to specify what happens when an error or event is logged to the error log.

Syntax



Parameters

-snmptrap *all | hardware_only | none*

Optionally specifies the SNMP trap setting, that is, when to raise a trap.

-alert *all | hardware_only | none*

Optionally specifies the alert setting. That is, when to raise an alert notification.

Description

This command modifies the settings that you want to apply to the error log. These settings define what to do when errors and events are logged. Issue the **svctask setevent** command to specify what you like to happen when an error or event is logged to the error log. You can select whether the cluster raises an SNMP trap, for entries that are added to the cluster error or event log, or both. Three levels of notification are possible:

- **None** No error or status changes will be sent.
- **Hardware_only** You will be notified of errors, but you will not be notified of status changes.
- **All** You will be notified of all errors and status changes.

If you have an SNMP manager installed or if you want to be notified by e-mail of errors or events, you should enable error notification. The notification levels for SNMP alerts can be set independently. If you choose **All** or **Hardware_only** notification, you must select a destination for the notification.

You can use this command to setup SNMP traps. For SNMP, you must supply the following information:

- When to raise a trap.

To set up e-mail and SNMP notifications, see *Cisco MDS 9000 Family Command Reference*.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcservicetask setevent -snmptrap all
```

The resulting output

No feedback

Chapter 5. Host commands

The following commands enable you to work with host options with the SAN Volume Controller Storage Software for Cisco MDS 9000.

addhostport

You can use the **addhostport** command to add WWPNs to an existing host object.

Syntax

```
svctask --addhostport --hbawwpn wwpn_list [-force]
host_name | host_id
```

Parameters

-hbawwpn *wwpn_list*

Specifies the list of ports to add to the host.

-force

Optionally forces the addition. This stops the validation of any WWPNs.

host_id | host_name

Specifies the host object to add ports to, either by ID or by name.

Description

This command adds the list of HBA WWPNs to the host object specified. Only logged-in unconfigured WWPNs can be added. For a list of candidate WWPNs, see **svcinfo lshbaportcandidate** command.

Some HBA device drivers do not log in to the fabric until they can see target LUNs. Because they do not log in, their WWPNs will not be recognized as candidate ports. You can specify the force flag with this command to stop the validation of the WWPN list.

Any virtual disks that are mapped to this host object will automatically be mapped to the new ports.

Replacing an HBA in a host: List the candidate HBA ports by issuing the **svcinfo lshbaportcandidate** command. You should see a list of the HBA ports that are available to be added to host objects. One or more of these should correspond with the one or more WWPNs that belong to the new HBA. Locate the host object that corresponds with the host in which you have replaced the HBA. The following command lists all the defined host objects:

```
svcinfo lshost
```

To list the WWPNs currently assigned to the host, issue the following:

```
svcinfo lshost <hostobjectname>
```

where *<hostobjectname>* is the name of the host object.

Add the new ports to the existing host object by issuing the following command:

```
svctask addhostport -hbawwpn <one or more existing WWPNs
separated by :> <hostobjectname/ID>
```

where *<one or more existing WWPNs separated by :>* and *<hostobjectname/id>* correspond with those listed in the previous steps.

Remove the old ports from the host object by issuing the following command:

```
svctask rmhostport -hbawpn <one or more existing WWPNs  
separated by :> <hostobjectname/ID>
```

where *<one or more existing WWPNs separated by :>* correspond with those listed in the previous step that belong to the old HBA that has been replaced. Any mappings that exist between the host object and VDisks will automatically be applied to the new WWPNs. Therefore, the host should see the VDisks as the same SCSI LUNs as before. See the *IBM TotalStorage Subsystem Device Driver: User's Guide* for additional information about dynamic reconfiguration.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5867E The action failed because the worldwide port name is already assigned or is not valid.
- CMMVC5872E The port (WWPN) was not added to the host object because an entity that was specified in the command does not exist.
- CMMVC5874E The action failed because the host does not exist.
- CMMVC5753E The object specified does not exist.

Examples

An invocation example

```
svctask addhostport -hbawpn 210100E08B251DD4 testhostlode
```

The resulting output

No feedback

chhost

You can use the **chhost** command to modify the name assigned to a host object.

Syntax

```
svctask --chhost --name new_name_arg host_name  
host_id
```

Parameters

-name *new_name_arg*

Specifies the new name to be assigned to the host object.

host_name | **host_id**

Specifies the host object to modify, either by ID or by current name.

Description

The name of the specified host object is changed to the new name. This command will not affect any of the current virtual disk-to-host mappings.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5868E The action failed because an entity that was specified in the command does not exist.
- CMMVC5869E The host object was not renamed because the host ID or name is not valid.
- CMMVC5874E The action failed because the host does not exist.

Examples

An invocation example

```
svctask chhost -name testhostlode hostone
```

The resulting output

No feedback

mkhost

You can use the **mkhost** command to create a logical host object.

Syntax

```
svctask mkhost [-name new_name_arg] [-hbawwpn wwpn_list] [-force]
```

Parameters

-name *new_name_arg*

Optionally specifies a name or label for the new object.

-hbawwpn *wwpn_list*

Specifies a list of host bus adapter (HBA) worldwide port names (WWPNs) to add to this host object.

-force

Optionally forces the creation. This argument stops any validation of the WWPNs.

Description

This command associates one or more HBA WWPNs with a logical host object. You can subsequently use this object when mapping virtual disks to hosts. This command creates a new host. The ID is returned when the command completes.

You only need to issue this command once. The cluster scans the fabric for WWPNs in the host zone. The cluster itself cannot filter into the hosts to determine which WWPNs are in which hosts. Therefore, you must use the **svctask mkhost** command to identify the hosts.

After you identify the hosts, mappings are created between hosts and virtual disks. These mappings effectively present the virtual disks to the hosts to which they are mapped. All WWPNs in the host object will be mapped to the virtual disks.

Some HBA device drivers do not log in to the fabric until they can see target logical unit numbers (LUNs). Because they do not log in, their WWPNs will not be recognized as candidate ports. You can specify the force flag with this command to stop the validation of the WWPN list.

See also the **svctask mkvdiskhostmap** and **svcinfolshbaportcandidate** commands.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5867E The action failed because the worldwide port name is already assigned or is not valid.
- CMMVC5868E The action failed because an entity that was specified in the command does not exist.
- CMMVC5729E One or more components in the list is not valid.

Examples

An invocation example

```
svctask mkhost -name hostone -hbawwpn 210100E08B251DD4 -force
```

The resulting output

```
Host id [1] successfully created.
```

rmhost

You can use the **rmhost** command to delete a host object.

Syntax

```
►►—svctask— —rmhost— —[-force]— —[host_name  
host_id]—►►
```

Parameters

-force

Optionally forces the delete. This argument deletes the host object. Any WWPNs that are still active can be added to other hosts. The active WWPNs are then listed as unconfigured WWPNs.

host_name | host_id

Specifies the host object to delete, either by ID or by name.

Description

When executed, this command deletes the logical host object. The WWPNs that were contained by this host object (if still connected and logged in to the fabric) are returned to the unconfigured state. When you issue the **svcinfo lshbaportcandidate** command, the host objects are listed as candidate ports.

If any mappings still exist between this host and virtual disks, the command will fail unless you specify the force flag. When specified, this flag causes the command to delete the mappings before deleting the host object. Thus, the command will fail if there are any VDisk mappings against this host unless the **-force** flag is specified.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5870E The host object was not deleted because an entity that was specified in the command does not exist.
- CMMVC5871E The action failed because one or more of the configured worldwide port names is in a mapping.
- CMMVC5874E The action failed because the host does not exist.

Examples

An invocation example

```
svctask rmhost testhostlode
```

The resulting output

```
No feedback
```

rmhostport

You can use the **rmhostport** command to delete WWPNs from an existing host object.

Syntax

```
svctask --rmhostport --hbawwpn wwpn_list --force  
host_name  
host_id
```

Parameters

-hbawwpn *wwpn_list*

Specifies the list of ports to delete from the host.

-force

Forces the deletion of the ports that you entered. This argument deletes the WWPNs in the list from the host specified. The ports become unconfigured WWPNs.

host_name | host_id

Specifies the host name or the host ID.

Description

This command deletes the list of HBA WWPNs from the specified host object. If these ports are still logged in to the fabric, they become unconfigured and are listed as candidate WWPNs. See also the **svcinfo lshbaportcandidate** command.

Any virtual disks that are mapped to this host object are automatically unmapped from the ports.

Replacing an HBA in a host: List the candidate HBA ports by issuing the **svcinfo lshbaportcandidate** command. You should see a list of the HBA ports that are available to be added to host objects. One or more of these should correspond with the one or more WWPNs that belong to the new HBA. Locate the host object that corresponds with the host in which you have replaced the HBA. The following command lists all the defined host objects:

```
svcinfo lshost
```

To list the WWPNs currently assigned to the host, issue the following:

```
svcinfo lshost <hostobjectname>
```

where *<hostobjectname>* is the name of the host object.

Add the new ports to the existing host object by issuing the following command:

```
svctask addhostport -hbawwpn <one or more existing WWPNS  
separated by :> <hostobjectname/ID>
```

where *<one or more existing WWPNS separated by :>* and *<hostobjectname/id>* correspond with those listed in the previous steps.

Remove the old ports from the host object by issuing the following command:

```
svctask rmhostport -hbawpn <one or more existing WWPNS  
separated by :> <hostobjectname/ID>
```

where *<one or more existing WWPNS separated by :>* correspond with those listed in the previous step that belong to the old HBA that has been replaced. Any mappings that exist between the host object and VDIsks will automatically be applied to the new WWPNS. Therefore, the host should see the VDIsks as the same SCSI LUNs as before. See the *IBM TotalStorage Subsystem Device Driver: User's Guide* for additional information about dynamic reconfiguration.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5867E The action failed because the worldwide port name is already assigned or is not valid.
- CMMVC5871E The action failed because one or more of the configured worldwide port names is in a mapping.
- CMMVC5872E The port (WWPN) was not added to the host object because an entity that was specified in the command does not exist.
- CMMVC5873E The action failed because there is no matching worldwide port name.
- CMMVC5874E The action failed because the host does not exist.

Examples

An invocation example

```
svctask rmhostport -hbawpn 210100E08B251DD4 testhostlode
```

The resulting output

```
No feedback
```

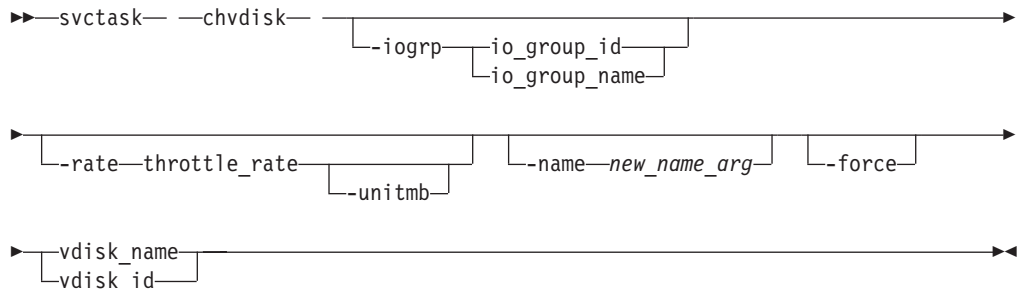
Chapter 6. Virtual disk commands

The following commands enable you to work with virtual disk options with the SAN Volume Controller Storage Software for Cisco MDS 9000.

chvdisk

You can 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.

-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.

vdisk_name | **vdisk_id**

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

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

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 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.

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 VDIs as required.

Attention:

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

Under no circumstances should VDIs be moved to an offline I/O group. You must ensure the I/O group is online before moving the VDIs 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 SDD configuration is updated to remove the vpaths in question. Failure to do this may result in data corruption. See *IBM TotalStorage 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 Remote Copy 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. *IBM TotalStorage Subsystem Device Driver: User's Guide* 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 Subsystem Device Driver: User's Guide* for details on how to dynamically reconfigure SDD for the given host operating system.

Possible failures

- 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.

Examples

An invocation example

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

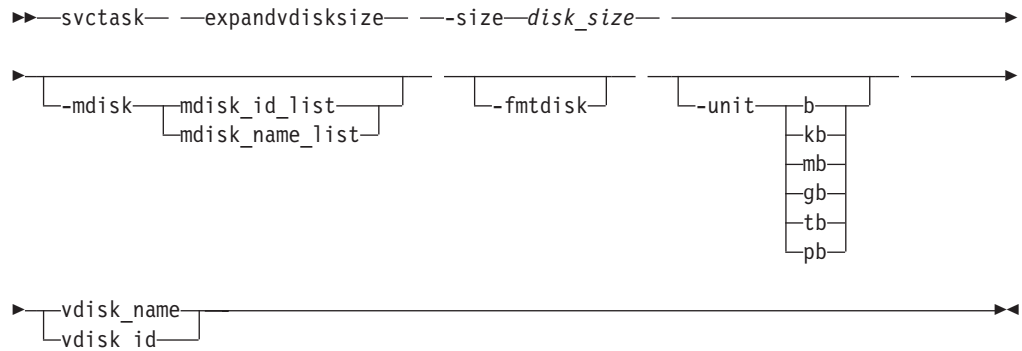
The resulting output

No feedback

expandvdisksize

You can use the **expandvdisksize** command to expand 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. The smallest granularity of bytes is 512. All capacities are rounded up to this value. However, an entire extent will be reserved even if only partial is used. The default capacity is in MB.

-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*

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.

vdisk_name | vdisk_id

Specifies the virtual disk name 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 or image mode. 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.

Examples

An invocation example

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

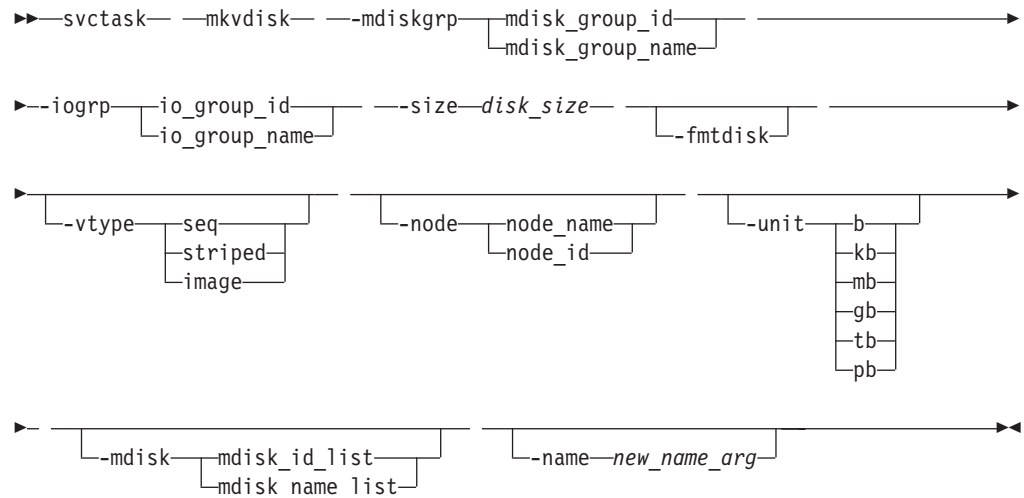
The resulting output

No feedback

mkvdisk

You can use the **mkvdisk** command to create 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.

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 are rounded up to this value. However, an entire extent will be reserved even if only partial is used. You can specify a capacity of 0. The size in bytes should be in multiples of logical block address (LBAs).

-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. You can query the status with the **svcinfor** command.

-vtype *seq* | *striped* | *image*

Optionally specifies the virtualization policy. The default virtualization type is striped. Any **-mdisk** attributes apply as if you specified a **-vtype** of striped.

-node *node_id* | *node_name*

Optionally specifies the preferred node ID or name for I/O operations to this virtual disk. You can specify the **-node** argument to specify the preferred access node. This argument is required for the subsystem device driver (SDD) and IC 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*

Optionally 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. The action of the **-mdisk** argument is determined by the virtualization type. The **-mdisk** attributes apply as if the **-vtype** is striped:

- If the **-vtype** is **seq**, a single MDisk is used.
- If the **-vtype** is **striped**, a list of MDisks is used to stripe across.

All MDisks in the list must be part of the same MDisk group.

-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, 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 **svcinfolcmdiskgrp** 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 Storage Software for Cisco MDS 9000 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 Storage Software for Cisco MDS 9000 nodes. Use the **svcinfolsiogrp** 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 remote copy 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:

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 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.

Seq (Sequential)

This policy requires the **-mdisk** flag with a single managed disk as its

argument. 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 are a special case. These 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 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.

Note: An image mode VDisk must be at least one extent in size. 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, by default 16Mb.

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.

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 MDisk has been specified and one or more of these MDisk 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 **svcinfolsfreeextents <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).

Examples

An invocation example

```
svctask mkvdisk -mdiskgrp Group0 -size 512 -unit  
kb -iogrp 0 -vtype striped -mdisk mdisk1 -node 1
```

The resulting output

```
Virtual Disk, id [1], 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

```
svctask --mkvdiskhostmap --host {host_id | host_name}
{--scsi scsi_num_arg} {vdisk_name | vdisk_id} [--force]
```

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 either 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.

vdisk_name | **vdisk_id**

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

-force

Optionally specifies the force flag to force the creation.

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.

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.

Examples

An invocation example

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

The resulting output

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

rmvdisk

You can use the **rmvdisk** command to delete a virtual disk. The command might fail if there are mappings between the virtual disk and hosts and you do not specify the force flag.

Syntax

```
svctask --rmvdisk [-force] [vdisk_id | vdisk_name]
```

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 virtual disk. The extents that made up this virtual disk are returned to the pool of free extents available on the managed disk group.

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.

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.

Examples

An invocation example

```
svctask rmvdisk -force vdisk5
```

The resulting output

No feedback

rmvdiskhostmap

You can use the **rmvdiskhostmap** command to delete 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

```
►►—svctask— —rmvdiskhostmap— —host— host_id vdisk_id◄◄  
                                  └─host_name─┘          └─vdisk_name─┘
```

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.

Examples

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. 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 Storage Software for Cisco MDS 9000 arbitrarily reduces the capacity of the VDisk by removing a partial, one or more extents from those allocated to the VDisk. You cannot control which extents are removed and so you cannot guarantee that it is unused space that is removed.

Attention: This feature should *only* be used to make a target or auxiliary VDisk the same size as the source or master VDisk when creating FlashCopy mappings or Remote Copy 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.

Examples

An invocation example

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

The resulting output

No feedback

Chapter 7. Managed disk group commands

The following commands enable you to work with managed disk group options with the SAN Volume Controller Storage Software for Cisco MDS 9000.

addmdisk

You can use the **addmdisk** command to add one or more managed disks to an existing managed disk group.

Syntax

```
svctask --addmdisk --mdisk {mdisk_id_list | mdisk_name_list}
{mdisk_group_id | mdisk_group_name}
```

Parameters

-mdisk *mdisk_id_list* | *mdisk_name_list*

Specifies one or more managed disk IDs or names to add to the group.

mdisk_group_id | **mdisk_group_name**

Specifies the ID or name of the managed disk group to add the disks to.

Description

This command adds the managed disks that you specify to the group. The disks can be specified in terms of managed disk ID or managed disk name.

The managed disks must be in unmanaged mode. Disks that already belong to a group cannot be added to another group until they have been deleted from their current group. You can delete a managed disk from a group under the following circumstances:

- If the managed disk does not contain any extents in use by a virtual disk
- If you can first migrate the extents in use onto other free extents within the group.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5819E The managed disk (MDisk) was not added to the MDisk group because the MDisk is part of another MDisk group.
- CMMVC5820E The managed disk (MDisk) was not added to the MDisk group because an entity that was specified in the command does not exist.
- CMMVC5821E The managed disk (MDisk) was not added to the MDisk group because not enough MDisks were included in the list.
- CMMVC5822E The managed disk (MDisk) was not added to the MDisk group because too many MDisks were included in the list.
- 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.

Examples

An invocation example

```
svctask addmdisk -mdisk mdisk13:mdisk14 Group0
```

The resulting output
No feedback

chmdiskgrp

You can use the **chmdiskgrp** command to modify the name, or label, assigned to a managed disk group.

Syntax

```
▶▶—svctask— —chmdiskgrp— —-name—new_name_arg— —mdisk_group_id—▶▶  
                                          └──mdisk_group_name──┘
```

Parameters

-name *new_name_arg*

Specifies the new name of the managed disk group.

mdisk_group_id | **mdisk_group_name**

Specifies the ID or name of the managed disk group to modify.

Description

This command modifies the name, or label, assigned to a given managed disk group. Subsequently, you can use the new name to refer to the managed disk group.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5816E The action failed because an entity that was specified in the command does not exist.
- CMMVC5817E The managed disk (MDisk) group was not renamed because the name was not valid.

Examples

An invocation example

```
svctask chmdiskgrp -name testmdiskgrp Group0
```


The resulting output

```
No feedback
```

mkmdiskgrp

You can use the **mkmdiskgrp** command to create a new managed disk group. Managed disk groups are collections of managed disks. Each group is divided into chunks, called extents. These extents are then used to create virtual disks.

Syntax



Parameters

-name *new_name_arg*

Optionally specifies a name to assign to the new group.

-mdisk *mdisk_id_list* | *mdisk_name_list*

Optionally specifies the name of the managed disk IDs or names to add to the group. You can create an empty MDisk group by not specifying the **-mdisk** flag.

-ext *extent_size*

Specifies the size of the extents for this group in MB. The **-ext** argument can have values of: 16, 32, 64, 128, 256, or 512 (MB).

Description

This command creates a new group, assigning the name (if given). The ID of the new group is returned if the command is successful.

Optionally, you can specify a list of managed disks that will be added to this group. These managed disks cannot belong to another group, and they must have a mode of unmanaged. Use the **svcinfolsmdiskcandidate** command to get a list of suitable candidates.

Each managed disk that is a member of this group will be split into extents. The storage available on these disks will be added to a pool of extents available in this group. When a virtual disk is created from this group, free extents from the pool will be used, in accordance with the policy chosen when creating the virtual disk.

All managed disks subsequently added to this group will be split into extents of the same size as assigned to the group.

When choosing an extent size, take into account the amount of storage you want to virtualize in this group. The system maintains a mapping of extents between virtual disks and managed disks. The SAN Volume Controller can only manage a finite number of extents (4 194 304). One cluster can virtualize the following number of extents:

- 64 TB If all managed disk groups have extent sizes of 16 MB.
- 2 PB If all managed disk groups have extent sizes of 512 MB.

Note: When an image mode VDisk is created, the MDisk group will increase in capacity by the size of the image mode VDisk (not the MDisk capacity) as the image mode VDisk may be smaller than the MDisk itself. If an extent is migrated from the image mode VDisk or MDisk to elsewhere in the group then the VDisk becomes a striped VDisk (for example, no longer image mode) and at this point the available capacity may increase, as the extra capacity available on the MDisk (for example, that which was not part of the image mode VDisk) becomes available.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5815E The managed disk (MDisk) group was not created because an entity that was specified in the command does not exist.
- 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.
- 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.

Examples

An invocation example

```
svctask mkmdiskgrp -mdisk mdisk13 -ext 512
```

The resulting output

```
MDisk Group, id [1], successfully created
```

rmmdisk

You can use the **rmmdisk** command to delete a managed disk from a managed disk group. There are some constraints on this command.

Syntax

```
svctask --rmmdisk --mdisk [mdisk_id_list | mdisk_name_list] [-force]
[mdisk_group_id | mdisk_group_name]
```

Parameters

-mdisk *mdisk_id_list* | *mdisk_name_list*

Specifies one or more managed disk IDs or names to delete from the group.

-force

Optionally specifies the force flag. If you do not supply the **-force** flag and virtual disks exist that are made from extents on one or more of the managed disks specified, the command will fail. If you do supply the force flag, and virtual disks exist that are made from extents on one or more of the managed disks specified, any data on the disks will be migrated to other disks in the group if there are enough free extents in the group. This operation might take some time.

mdisk_group_id | **mdisk_group_name**

Specifies the ID or name of the managed disk group to delete the disk, or disks, from.

Description

This command attempts to remove the managed disk or disks from the group.

Deleting a managed disk from a group can only be done if the managed disk does not contain any extents in use by a virtual disk. If there are extents in use and you do not supply the force flag, the command will fail.

Attention: If this disk being removed has already been powered down, removed, or is suffering a power outage, the migrate will be pending and will not complete until the MDisk comes back online. This also means that the MDisk will not be removed from the list of MDisks contained in the group.

If the disk has been deliberately removed, the only method of removing the MDisk is to remove the entire group itself.

You should ensure that you do not destroy any controller LUNs until you have deleted them from the MDisk group they belong to.

If you do specify the force flag, an attempt will be made to migrate the extents that are in use onto other free extents within the group. If there are not enough free extents in the group, the command will fail even if the force flag is specified.

If you still want to delete the disks from the group, you have the following options:

1. Delete the virtual disk that is using the extents specified on the managed disk.
2. Add more managed disks to the group, rerun the command and specify the `-force` flag.

When data is being migrated off the managed disk, it might take some time for the command to be completed. The command itself will return with a success code, notifying you that migration is in progress. An event will be logged when the migration is complete and the disk will be deleted from the group at this time. You can also check the progress of any active migrations by running the `svcinfo lsmigrate` command.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5823E The managed disk (MDisk) was not deleted from the MDisk group because the MDisk is part of another MDisk group.
- CMMVC5824E The managed disk (MDisk) was not deleted from the MDisk group because it does not belong to the MDisk group.
- CMMVC5825E The managed disk (MDisk) was not deleted from the MDisk group because a virtual disk (VDisk) is allocated from one or more of the specified MDisks. A forced deletion is required.
- 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.
- CMMVC6006E The managed disk (MDisk) was not deleted because the resource was busy.
- CMMVC6015E A delete request is already in progress.

Examples

An invocation example

```
svctask rmdisk -mdisk mdisk12 -force Group3
```

The resulting output

```
No feedback
```

rmmdiskgrp

You can use the **rmmdiskgrp** command to delete a managed disk group. Care should be taken when using this command.

Syntax

```
svctask --rmmdiskgrp [-force] [mdisk_group_id | mdisk_group_name]
```

Parameters

-force

Optionally specifies the force flag to force the deletion. If the **-force** flag is specified, all virtual disks and virtual disk-to-host mappings are deleted. All managed disks in the group are removed and the group itself is deleted.

mdisk_group_id | mdisk_group_name

Specifies the ID or name of the managed disk group to delete.

Description

This command destroys the specified managed disk group. The force flag is required if there are virtual disks created from this group or if there are managed disks in the group. Otherwise, the command will fail.

Deleting a managed disk group is essentially like destroying a cluster or part of a cluster. The managed disk group is the central point of control of virtualization. Virtual disks are created using extents available in the group, and the mapping between virtual disk extents and managed disk extents is controlled on a group basis. Therefore, deleting a group deletes this mapping, which cannot be subsequently restored.

Attention: This command partially completes asynchronously. All virtual disks, host mappings, and copy services are deleted before the command returns. The managed disk group delete then completes asynchronously.

Attention: Before you issue the command, ensure that you really want to destroy all mapping information. Data held on virtual disks cannot be recovered after the managed disk group has been destroyed.

In detail, if you specify the force flag, the following actions take place:

1. If there are virtual disks still using extents in this group, any mappings between that disk and any host objects are deleted.
2. If there are managed disks in the group, all disks are deleted from the group. They are returned to the unmanaged state.
3. The group is deleted.

Attention: If you delete all the managed disk groups in your cluster using the force flag, you will be returned to the state you were in just after you added nodes to the cluster. All data held on virtual disks will be lost and unrecoverable.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

- CMMVC5816E The action failed because an entity that was specified in the command does not exist.
- CMMVC5818E The managed disk (MDisk) group was not deleted because there is at least one MDisk in the group.

Examples

An invocation example

```
svctask rmdiskgrp -force Group3
```

The resulting output

No feedback

Chapter 8. Managed disk commands

When the cluster detects MDisks, it will automatically add the MDisk to the clusters list of known MDisks. If you subsequently delete the RAID that corresponds to this MDisk, the cluster will only delete this MDisk from the list if the objects are:

- the MDisk has a mode of unmanaged, that is, it does not belong to a MDisk group
- and the MDisk is offline.

The following commands enable you to work with managed disk options with the SAN Volume Controller Storage Software for Cisco MDS 9000.

chmdisk

You can use the **chmdisk** command to modify the name of a managed disk.

Syntax

```
svctask --chmdisk --name new_name_arg mdisk_id  
mdisk_name
```

Parameters

-name *new_name_arg*

Specifies the new name to be applied to the managed disk.

mdisk_id_list | **mdisk_name_list**

Specifies the ID or name of the managed disk to modify.

Description

This command modifies the name, or label, assigned to a given managed disk. You can subsequently use the new name to refer to the managed disk.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5806E The action failed because an entity that was specified in the command does not exist.
- CMMVC5808E The action failed because the managed disk (MDisk) does not exist.

Examples

An invocation example

```
svctask chmdisk -name testmdisk mdisk0
```

The resulting output

```
No feedback
```

includemdisk

You can use the **includemdisk** command to include a disk that has been excluded by the cluster.

Syntax

```
svctask --includemdisk --mdisk_id | --mdisk_name
```

Parameters

mdisk_id | mdisk_name

Specifies the ID or name of the managed disk to add back into the cluster.

Description

The managed disk specified is included in the cluster.

A disk might be excluded from the cluster because of multiple I/O failures. These failures might be caused by noisy links. Once a fabric-related problem has been fixed, the excluded disk can be added back into the cluster.

Executing this command against a disk that has not been excluded will have no noticeable effect.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5806E The action failed because an entity that was specified in the command does not exist.
- CMMVC5808E The action failed because the managed disk (MDisk) does not exist.

Examples

An invocation example

```
svctask includemdisk mdisk5
```

The resulting output

```
No feedback
```

setquorum

You can use the **setquorum** command to change the managed disks that are assigned as quorum disks.

Syntax

```
svctask --setquorum --quorum {0 | 1 | 2} {mdisk_id | mdisk_name}
```

Parameters

-quorum 0 | 1 | 2

Specifies the quorum index.

mdisk_id | **mdisk_name**

Specifies the ID or name of the managed disk to assign as a quorum disk.

Description

This command sets the managed disk specified to the quorum index specified.

Quorum disks are used when the cluster loses half of the nodes that make up the cluster. The half of the cluster that contains the majority of the quorum disks continues to process the I/O transactions. The other half stops processing the I/O transactions. By establishing quorum disks, you can ensure that both halves of the cluster do not continue to operate.

The managed disk that is currently assigned the quorum index number is set to a non-quorum disk.

The cluster will automatically assign quorum indexes. You can use this command if you want to ensure a particular set of managed disks continue to be accessible if a cluster is split.

Attention: It is recommended that you set quorum disks on multiple controllers to avoid the possibility of losing all of the quorum disks with a single failure.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5806E The action failed because an entity that was specified in the command does not exist.
- CMMVC5810E The quorum index number for the managed disk (MDisk) was not set because the MDisk is offline.
- CMMVC5811E The quorum index number for the managed disk (MDisk) was not set because the quorum disk does not exist.
- CMMVC5812E The quorum index number for the managed disk (MDisk) was not set because the MDisk is in the wrong mode.
- CMMVC5814E The quorum index number for the managed disk (MDisk) was not set because the unique identifier (UID) type is not valid.

Examples

An invocation example

```
svctask setquorum -quorum 2 mdisk7
```

The resulting output

No feedback

Chapter 9. FlashCopy commands

The following commands enable you to work with FlashCopy methods and functions with the SAN Volume Controller Storage Software for Cisco MDS 9000.

chfcconsistgrp

You can use the **chfcconsistgrp** command to modify the name of an existing consistency group.

Syntax

```
▶—svctask— —chfcconsistgrp— —-name—new_name_arg— —————▶  
▶—fc_consist_group_id—|—————▶  
  └—fc_consist_group_name—┘
```

Parameters

-name *new_name_arg*

Specifies the new name to assign to the consistency group.

fc_consist_group_id | **fc_consist_group_name**

Specifies the ID or existing name of the consistency group to modify.

Description

This command changes the name of the consistency group specified.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5891E The name of the FlashCopy consistency group was not modified because the name is not valid.
- CMMVC5893E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svctask chfcconsistgrp -name testgrp1 fcconsistgrp1
```

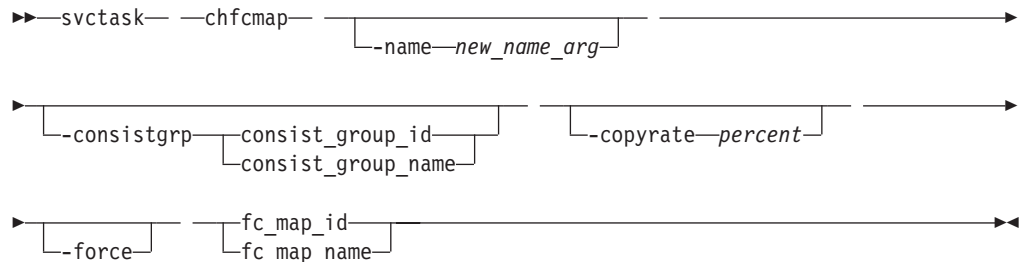
The resulting output

No feedback

chfcmap

You can use the **chfcmap** command to modify certain attributes of an existing mapping.

Syntax



Parameters

-name *new_name_arg*

Optionally specifies a new name to assign to the mapping. The `-name` argument is mutually exclusive with the other flags.

-consistgrp *consist_group_id | consist_group_name*

Optionally specifies the consistency group for which you want to modify the mapping. The `-consistgrp` and `-copyrate` arguments are not mutually exclusive with each other: that is, you can specify both arguments in one command-line invocation. The consistency group cannot be modified while a copy is active or while the target consistency group is active. This parameter is mutually exclusive with the `-name` and `-force` parameters.

-copyrate *percent*

Optionally specifies the priority of the background copy rate. The `-consistgrp` and `-copyrate` arguments are not mutually exclusive with each other: that is, you can specify both arguments in one command-line invocation. This parameter is mutually exclusive with the `-name` and `-force` parameters.

-force

If no consistency group ID is specified and the optional force flag is used, the mapping will be modified to a standalone mapping (equivalent to creating the mapping without a consistency group ID). This parameter is mutually exclusive with all other parameters.

fc_map_id | fc_map_name

Specifies the ID or name of the mapping to modify.

Description

This command modifies the specified attributes of an existing mapping. When modifying the name of a mapping, you cannot modify any of the other attributes at the same time. You can only modify the consistency group that the mapping belongs to if the mapping is inactive. A mapping is inactive if it has not been triggered or if it has been triggered, but the copy has run to completion.

If you have created several FlashCopy mappings for a group of VDisks that contain elements of data for the same application, you may find it convenient to assign these mappings to a single FlashCopy Consistency Group. Then you can

issue a single prepare or trigger command for the whole group, so that, for example, all the files for a particular database are copied at the same time.

Possible failures

- CMMVC5753E The object specified does not exist
- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5888E The action failed because an entity that was specified in the command does not exist.
- CMMVC5913E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the preparing state.
- CMMVC5914E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the prepared state.
- CMMVC5915E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the copying state.
- CMMVC5916E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the suspended state.
- CMMVC5921E The properties of the FlashCopy mapping were not modified because the consistency group is not idle.

Examples

An invocation example

```
svctask chfcmap -name testmap 1
```

The resulting output

```
No feedback
```

mkfcconsistgrp

You can use the **mkfcconsistgrp** command to create a new FlashCopy consistency group.

Syntax

```
svctask --mkfcconsistgrp --name consist_group_name
```

Parameters

-name *consist_group_name*

Specifies a name for the consistency group. If you do not specify a consistency group name, a name is automatically assigned to the consistency group. For example, if the next available consistency group ID is id=2, the consistency group name is fccstgrp2.

Description

This command creates a new consistency group. The ID of the new group is returned.

If you have created several FlashCopy mappings for a group of VDisks that contain elements of data for the same application, you may find it convenient to assign these mappings to a single FlashCopy Consistency Group. Then you can issue a single prepare or trigger command for the whole group, so that, for example, all the files for a particular database are copied at the same time.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5891E The FlashCopy consistency group was not created because the name is not valid.
- CMMVC5892E The FlashCopy consistency group was not created because it already exists.

Examples

An invocation example

```
svctask mkfcconsistgrp
```


The resulting output

```
Flash Copy Consistency Group, id [1], successfully created
```

mkfcmmap

You can use the **mkfcmmap** command to create a new FlashCopy mapping, which maps a source virtual disk to a target virtual disk ready for subsequent copying.

Syntax



Parameters

- source** *src_vdisk_id* | *src_vdisk_name*
Specifies the ID or name of the source virtual disk.
- target** *target_vdisk_id* | *target_vdisk_name*
Specifies the ID or name of the destination virtual disk.
- name** *new_name_arg*
Optionally specifies a name to assign to the new mapping.
- consistgrp** *consist_group_id* | *consist_group_name*
Optionally specifies a consistency group to add the new mapping to. If you do not specify a consistency group, the mapping is assigned to the default Consistency Group 0.
- copyrate** *percent*
Optionally specifies the priority of the background copy rate. The **-copyrate** flag specifies the background copy rate. If 0% is displayed, it means that it is in idle state.

Description

This command creates a new FlashCopy mapping logical object. This mapping persists until it is deleted. The mapping specifies the source and destination virtual disks. The destination must be identical in size to the source, or the mapping will fail. Issue the **svcinfolsvdisk -bytes** command to find the exact size of the source Vdisk that you want to create a target disk of the same size. The source and destination cannot be in an existing mapping. That is, a virtual disk can be either a source or a destination disk in **only one** mapping. A mapping is triggered at the point in time when the copy is required.

The mapping can optionally be given a name and assigned to a consistency group. These are groups of mappings that can be triggered at the same time. This enables multiple virtual disks to be copied at the same time, which creates a consistent copy of multiple disks. This is required by some database products in which the database and log files reside on different disks.

If no consistency group is defined, the mapping is assigned into the default group 0. This is a special group that cannot be started as a whole. Mappings in this group can only be started on an individual basis.

The background copy rate specifies the priority that should be given to completing the copy. If 0% is specified, the copy will not proceed in the background. The default is 50%.

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.
- CMMVC5881E The FlashCopy mapping was not created because an entity that was specified in the command does not exist.
- CMMVC5882E The FlashCopy mapping was not created because a mapping for the source or target virtual disk (VDisk) already exists.
- CMMVC5883E The FlashCopy mapping was not created because the recovery I/O group is associated with the source or target virtual disk (VDisk).
- CMMVC5884E The FlashCopy mapping was not created because the source or target virtual disk (VDisk) cannot be a member of a Remote Copy mapping.
- CMMVC5885E The FlashCopy mapping was not created because this source or target virtual disk (VDisk) cannot be a member of a FlashCopy mapping.
- CMMVC5886E The FlashCopy mapping was not created because the source or target virtual disk (VDisk) is associated with the recovery I/O group.
- CMMVC5887E The FlashCopy mapping was not created because the source or target virtual disk (VDisk) must not be in router mode.
- CMMVC5922E The FlashCopy mapping was not created because the destination virtual disk (VDisk) is too small.
- CMMVC5923E The FlashCopy mapping was not created because the I/O group is offline.
- CMMVC5924E The FlashCopy mapping was not created because the source and target virtual disks (VDisks) are different sizes.
- CMMVC5917E The FlashCopy mapping was not created because there is no memory to create the bitmap.
- CMMVC5920E The FlashCopy mapping was not created because the consistency group is not idle.

Examples

An invocation example

```
svctask mkfcmap -source 0 -target 2 -name mapone
```

The resulting output

```
FlashCopy mapping, id [1], successfully created
```

prestartfcconsistgrp

You can use the **prestartfcconsistgrp** command to prepare a FlashCopy consistency group to start. This command flushes the cache of any data destined for the source virtual disk and forces the cache into the write-through mode until the mapping is started.

Syntax

```
svctask —prestartfcconsistgrp— [fc_consist_group_id | fc_consist_group_name]
```

Parameters

fc_consist_group_id | fc_consist_group_name

Specifies the name or ID of the consistency group to prepare. Preparing Consistency Group 0 is invalid. If the mapping belongs to Consistency Group 0, you must use the **prestartfcmap** command.

Description

This command prepares a group of mappings (on a consistency group) for subsequent triggering. The preparation step ensures that any data that resides in the cache for the source virtual disk is first flushed to disk. This step ensures that the copy that is made is consistent with what the operating system thinks is on the disk.

Issue the **svctask prestartfcconsistgrp** command to prepare the FlashCopy Consistency Group before the copy process can be started (triggered). When you have assigned several mappings to a FlashCopy Consistency Group, you only have to issue a single prepare command for the whole group, to prepare all the mappings at once.

The group enters the preparing state. After the preparation is complete, they change to the prepared state. At this point, the group is ready for triggering.

Preparing and the subsequent triggering is usually performed on a consistency group basis. Only mappings belonging to Consistency Group 0 can be prepared on their own. A FlashCopy must be prepared before it can be triggered.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5888E The action failed because an entity that was specified in the command does not exist.
- CMMVC5890E The FlashCopy mapping was not prepared because preparing consistency group 0 is not a valid operation.
- CMMVC5901E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the preparing state.
- CMMVC5902E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the prepared state.
- CMMVC5903E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the copying state.

- CMMVC5904E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the suspended state.
- CMMVC5918E The FlashCopy mapping was not prepared because the I/O group is offline.
- CMMVC6031E The operation was not performed because the FlashCopy consistency group is empty.

Examples

An invocation example

```
svctask prestartfcconsistgrp 1
```

The resulting output

No feedback

prestartfcmap

You can use the **prestartfcmap** command to prepare a FlashCopy mapping to start. This command flushes the cache of any data destined for the source virtual disk and force the cache into write-through mode until the mapping is started.

Syntax

```
svctask —prestartfcmap— {fc_map_id | fc_map_name}
```

Parameters

fc_map_id | fc_map_name

Specifies the name or ID of the mapping to prepare.

Description

This command prepares a single mapping for subsequent triggering. The preparation step ensures that any data that resides in the cache for the source virtual disk is first flushed to disk. This step ensures that the copy that is made is consistent with what the operating system thinks is on the disk.

The mapping enters the preparing state. After the preparation is complete, it changes to the prepared state. At this point, the mapping is ready for triggering.

Preparing and the subsequent triggering is usually performed on a consistency group basis. Only mappings belonging to Consistency Group 0 can be prepared on their own. A FlashCopy must be prepared before it can be triggered.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5888E The action failed because an entity that was specified in the command does not exist.
- CMMVC5890E The FlashCopy mapping was not prepared because preparing consistency group 0 is not a valid operation.
- CMMVC5901E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the preparing state.
- CMMVC5902E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the prepared state.
- CMMVC5903E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the copying state.
- CMMVC5904E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the suspended state.
- CMMVC5918E The FlashCopy mapping was not prepared because the I/O group is offline.

Examples

An invocation example

```
svctask prestartfcmap 1
```

The resulting output

No feedback

rmfcconsistgrp

You can use the **rmfcconsistgrp** command to delete a FlashCopy consistency group.

Syntax

```
svctask --rmfcconsistgrp --force fc_consist_group_id fc_consist_group_name
```

Parameters

-force

Optionally specifies the force flag. If the group still contains mappings, you must specify the force flag to move all the mappings into Consistency Group 0.

fc_consist_group_id | fc_consist_group_name

Specifies the ID or name of the consistency group to delete.

Description

This command deletes the consistency group specified. If there are mappings that are members of the group, the command will fail unless you specify the force flag. If you specify the force flag, all the mappings will first be assigned to the default Consistency Group 0.

If you want to delete all the mappings in the consistency group as well, you must first delete them using the **svctask rmfcmap** command.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5893E The action failed because an entity that was specified in the command does not exist.
- CMMVC5894E The FlashCopy consistency group was not deleted because you are trying to delete consistency group 0 or the name of the consistency group is not valid.
- CMMVC5895E The FlashCopy consistency group was not deleted because it contains mappings. To delete this consistency group, a forced deletion is required.

Examples

An invocation example

```
svctask rmfcconsistgrp fcconsistgrp1
```

The resulting output

No feedback

rmfcmap

You can use the **rmfcmap** command to delete an existing mapping.

Syntax

```
svctask --rmfcmap [-force] [fc_map_id | fc_map_name]
```

Parameters

-force

Optionally specifies the force flag.

fc_map_id | fc_map_name

Specifies the ID or name of the mapping to delete. Unless you specify the force flag, you can only delete a mapping before it is triggered or after it completes.

Description

This command attempts to delete the mapping specified. If the mapping is active, the command fails unless you specified the force flag.

The **-force** flag must be used when the FlashCopy status is Stopped.

Deleting a mapping only deletes the logical *relationship* between the two virtual disks, it does not effect the virtual disks themselves. However, if you force the deletion, the data on the destination virtual disk will be inconsistent.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5889E The FlashCopy mapping was not deleted because an entity that was specified in the command does not exist.
- CMMVC5896E The FlashCopy mapping was not deleted because the mapping or consistency group is in the preparing state. The mapping or consistency group must be stopped first.
- CMMVC5897E The FlashCopy mapping was not deleted because the mapping or consistency group is in the prepared state. The mapping or consistency group must be stopped first.
- CMMVC5898E The FlashCopy mapping was not deleted because the mapping or consistency group is in the copying state. The mapping or consistency group must be stopped first.
- CMMVC5899E The FlashCopy mapping was not deleted because the mapping or consistency group is in the stopped state. To delete the mapping, a forced deletion is required.
- CMMVC5900E The FlashCopy mapping was not deleted because the mapping or consistency group is in the suspended state. The mapping or consistency group must be stopped first.

Examples

An invocation example

```
svctask rmfcmap testmap
```

The resulting output

No feedback

startfcconsistgrp

You can use the **startfcconsistgrp** command to start (trigger) a FlashCopy group of mappings. This command makes a point-in-time copy of the source virtual disk at the moment the command is executed.

Syntax

```
svctask — startfcconsistgrp — [ -prep ] [ fc_consist_group_id | fc_consist_group_name ]
```

Parameters

-prep

Optionally specifies to prepare the mapping or group prior to triggering the mapping.

fc_consist_group_id | fc_consist_group_name

Specifies the ID or name of the consistency group to trigger. Triggering Consistency Group 0 is invalid.

Description

This command triggers a group of mappings (on a consistency group basis). Triggering means to take a point-in-time copy of the source virtual disks.

The group must first be prepared for triggering. See the **svctask prestartfcconsistgrp** command to prepare the triggering. However, you can run this command with the optional **-prep** argument, which prepares the group and triggers the copy as soon as the preparation is complete. This means it is under the systems control when the trigger happens. That is, the preparation step takes some time to complete before the copy is made. If you want to control the triggering, you should use the **svctask prestartfcconsistgrp** command first.

The consistency group enters the copying state. The way the copy proceeds depends on the background copy rate attribute of the mapping. If the mapping is set to 0%, only data that is subsequently updated on the source will be copied to the destination. This operation means that the destination can only be used as a backup copy while the mapping exist in the copying state. If the copy is stopped, the destination will not be usable. If you want to end up with a duplicate copy of the source at the destination, you should set the background copy rate greater than 0. This means that the system copies all the data (even unchanged data) to the destination and eventually reaches the idle or copied state. After this data is copied, you can delete the mapping and have a usable point-in-time copy of the source at the destination.

Triggering is usually performed on a consistency group basis. Only mappings belonging to Consistency Group 0 can be triggered on their own.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5888E The action failed because an entity that was specified in the command does not exist.
- CMMVC5890E The FlashCopy mapping or consistency group was not started because starting consistency group 0 is not a valid operation.

- CMMVC5905E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the idle state. The mapping or consistency group must be prepared first.
- CMMVC5906E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the preparing state.
- CMMVC5907E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is already in the copying state.
- CMMVC5908E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the stopped state. The mapping or consistency group must be prepared first.
- CMMVC5909E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the suspended state.
- CMMVC5919E The FlashCopy mapping or consistency group was not started because the I/O group is offline.

Examples

An invocation example

```
svctask startfcconsistgrp -prep 2
```

The resulting output

No feedback

startfcmap

You can use the **startfcmap** command to start (trigger) a FlashCopy mapping. This command makes a point-in-time copy of the source virtual disk at the moment the command is executed.

Syntax

```
svctask --startfcmap [-prep] [fc_map_id | fc_map_name]
```

Parameters

-prep

Optionally specifies to prepare the mapping or group prior to triggering the mapping.

fc_map_id | fc_map_name

Specifies the ID or name of the mapping to trigger. Triggering of mappings is usually performed on a consistency group basis. If the mappings have not been prepared, then this command will fail unless **-prep** is specified which causes a prepare to happen prior to the triggering. If the mapping belongs to Consistency Group 0, then the **map_id | name** must be specified.

Description

This command triggers a single mapping. Triggering means to take a point-in-time copy of the source virtual disks.

The mapping must first be prepared for triggering. See the **svctask prestartfcmap** command to prepare the triggering. However, you can run this command with the optional **-prep** argument, which prepares the mapping and triggers the copy as soon as the preparation is complete. This means it is under the systems control when the trigger happens. That is, the preparation step takes some time to complete before the copy is made. If you want to control the triggering, you should use the **svctask prestartfcmap** command first.

The mapping enters the copying state. The way the copy proceeds depends on the background copy rate attribute of the mapping. If the mapping is set to 0%, only data that is subsequently updated on the source will be copied to the destination. This operation means that the destination can only be used as a backup copy while the mapping exist in the copying state. If the copy is stopped, the destination will not be usable. If you want to end up with a duplicate copy of the source at the destination, you should set the background copy rate greater than 0. This means that the system copies all the data (even unchanged data) to the destination and eventually reaches the idle or copied state. After this data is copied, you can delete the mapping and have a usable point-in-time copy of the source at the destination.

Triggering is usually performed on a consistency group basis. Only mappings belonging to Consistency Group 0 can be triggered on their own.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5888E The action failed because an entity that was specified in the command does not exist.

- CMMVC5890E The FlashCopy mapping or consistency group was not started because starting consistency group 0 is not a valid operation.
- CMMVC5905E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the idle state. The mapping or consistency group must be prepared first.
- CMMVC5906E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the preparing state.
- CMMVC5907E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is already in the copying state.
- CMMVC5908E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the stopped state. The mapping or consistency group must be prepared first.
- CMMVC5909E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the suspended state.
- CMMVC5919E The FlashCopy mapping or consistency group was not started because the I/O group is offline.

Examples

An invocation example

```
svctask startfcmap -prep 2
```

The resulting output

No feedback

stopfcconsistgrp

You can use the **stopfcconsistgrp** command to stop any active FlashCopy consistency group.

Syntax

```
svctask — stopfcconsistgrp — [fc_consist_group_id | fc_consist_group_name]
```

Parameters

fc_consist_group_id | fc_consist_group_name

Specifies the name or ID of the consistency group that you want to stop.

Description

This command stops a group of mappings (in a consistency group). If the copy is stopped the destination will not be usable. Before you can use the destination, you must reprepare and retrigger the group.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5888E The action failed because an entity that was specified in the command does not exist.
- CMMVC5890E The FlashCopy mapping or consistency group was not stopped because starting consistency group 0 is not a valid operation.
- CMMVC5910E The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is in the idle state.
- CMMVC5911E The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is in the preparing state.
- CMMVC5912E The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is already in the stopped state.

Examples

An invocation example

```
svctask stopfcconsistgrp testmapone
```

The resulting output

No feedback

stopfcmap

You can use the **stopfcmap** command to stop any active copying or suspended mapping.

Syntax

```
svctask --stopfcmap [fc_map_id | fc_map_name]
```

Parameters

fc_map_id | fc_map_name

Specifies the name or ID of the mapping to stop.

Description

This command stops a single mapping. If the copy is stopped the destination will not be usable. The mapping or group needs to be reprepared and retrigged.

Stopping is usually performed on a consistency group basis. Only mappings that belong to Consistency Group 0 can be stopped on their own.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5888E The action failed because an entity that was specified in the command does not exist.
- CMMVC5890E The FlashCopy mapping or consistency group was not stopped because starting consistency group 0 is not a valid operation.
- CMMVC5910E The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is in the idle state.
- CMMVC5911E The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is in the preparing state.
- CMMVC5912E The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is already in the stopped state.
- CMMVC6030E The operation was not performed because the FlashCopy mapping is part of a consistency group. The action must be performed at the consistency group level.

Examples

An invocation example

```
svctask stopfcmap testmapone
```

The resulting output

No feedback

Chapter 10. Remote Copy commands

The following commands enable you to work with the Remote Copy services provided by the SAN Volume Controller Storage Software for Cisco MDS 9000.

chpartnership

You can use the **chpartnership** command to specify the bandwidth available for background copy in a cluster partnership that has been created for Remote Copy purposes.

Syntax

```
▶▶ svctask — chpartnership — -bandwidth—bandwidth_in_mbs —————▶▶
▶ remote_cluster_id —————▶▶
  └─ remote_cluster_name ─┘
```

Parameters

-bandwidth *bandwidth_in_mbs*

Specifies the new bandwidth in MBps. This argument might be set to a value that is greater than the intercluster links can sustain. If so, the actual copy rate defaults to what is available on the link.

remote_cluster_id | **remote_cluster_name**

Specifies the cluster ID or name of the remote cluster. The intracluster bandwidth cannot be modified so if you enter the local cluster name or ID, an error will occur.

Description

This command modifies the bandwidth of the partnership between the local cluster and the remote cluster that is specified in the command. This affects the bandwidth available for background copy in Remote Copy relationships, in the direction from the local to the remote cluster. To modify the background copy bandwidth in the other direction (remote cluster → local cluster) it is necessary to issue the corresponding **chpartnership** command to the remote cluster.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5927E The action failed because the cluster ID is not valid.

Examples

An invocation example

```
svctask chpartnership -bandwidth 20 cluster1
```

The resulting output

```
No feedback
```

chrconsistgrp

You can use the **chrconsistgrp** command to modify the name of an existing Remote Copy consistency group.

Syntax

```
▶▶—svctask— —chrconsistgrp— —-name—new_name_arg— —————▶▶  
▶—rc_consist_group_name—▶▶  
└—rc_consist_group_id—┘
```

Parameters

-name *new_name_arg*

Specifies the new name to assign to the consistency group.

rc_consist_group_name | **rc_consist_group_id**

Specifies the ID or existing name of the consistency group to modify.

Description

This command changes the name of the specified consistency group.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5937E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

Change the name of the Remote Copy consistency group called "rc_testgrp" to "rctestone".

```
svctask chrconsistgrp -name rctestone rc_testgrp
```


The resulting output

No feedback

chrrelationship

This command enables you to modify certain attributes of an existing relationship. This command enables you to add a relationship to a consistency group or to remove a relationship from a consistency group, as well as change the name of the relationship.

Syntax



Parameters

-name *new_name_arg*

Optionally specifies a new label to assign to the relationship

-consistgrp *consist_group_id* | *consist_group_name*

Optionally specifies a new consistency group to assign the relationship to.

-force

Optionally specifies the force flag which will remove the relationship from a consistency group making the relationship a standalone relationship.

rc_rel_name | **rc_rel_id**

Specifies the ID or name of the relationship.

Note: The **-name**, **-consistgrp** and **-force** are mutually exclusive parameters. That is, only one of these parameters can be specified per command line.

Description

This command can modify the specified attributes of the relationship supplied. Only one attribute can be modified at a time. That is, all four optional flags are mutually exclusive. In addition to changing the name of a consistency group, this command can be used for the following purposes.

- **Add a relationship to a group** A standalone relationship can be added to a consistency group by specifying the **-consistgrp** parameter and the name or ID of the consistency group. The relationship and consistency group must both be connected when the command is issued, and both must have the same:
 - Master cluster
 - Auxiliary cluster
 - State (unless the group is empty)
 - Primary (unless the Group is empty)

When the first relationship is added to an empty group, the group takes on the same state and primary (copy direction) as the relationship. Subsequent relationships must have the same state and copy direction as the group in order to be added to it. A relationship may only belong to one consistency group.

- **Remove a relationship from a group** A relationship can be removed from a consistency group by simply specifying the **-force** flag and the name or ID of the

relationship. You do not have to specify or confirm the name of the consistency group, so it is recommended that you verify which group the relationship belongs to before issuing this command.

This form of the modify relationship command will succeed in the connected or disconnected states. If the clusters are disconnected, then the relationship will only be removed from the consistency group on the local cluster, at the time the command is issued. When the clusters are reconnected, the relationship will automatically be removed from the consistency group on the other cluster. Alternatively, you can issue an explicit modify (**chrrelationship**) command to remove the relationship from the group on the other cluster while it is still disconnected.

- **Move a relationship from one group to another** To move a relationship between two consistency groups you have to invoke the modify relationship command twice. First use the `-force` flag to remove the relationship from its current group, then use the `-consistgrp` parameter and the name of the new consistency group it is to be added to.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5935E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

Change the name of the relationship `rccopy1` to `testrel`
`svctask chrrelationship -name testrel rccopy1`

Add relationship `rccopy2` to group called `newgroup`.
`svctask chrrelationship -consistgrp newgroup rccopy2`

Remove relationship `rccopy3` from whichever consistency group it is a member of.
`svctask chrrelationship -force rccopy3`

The resulting output

No feedback

There is no feedback in any of these cases.

mkpartnership

You can use the **mkpartnership** command to establish a one-way Remote Copy partnership between the local cluster and a remote cluster. To establish a fully functional Remote Copy partnership, you must issue this command to both clusters. This step is a prerequisite to creating Remote Copy relationships between VDisks on the clusters.

Syntax

```
svctask --mkpartnership [-bandwidth bandwidth_in_mbs]
                        [remote_cluster_id | remote_cluster_name]
```

Parameters

-bandwidth *bandwidth_in_mbs*

Optionally specifies the bandwidth to be used by the background copy process between the clusters. It may be used to throttle the bandwidth used by Remote Copy for the initial background copy process. The bandwidth defaults to 50 MBps (megabytes per second) if you do not specify it. The bandwidth should be set to a value that is less than or equal to the bandwidth that can be sustained by the intercluster link. If the parameter is set to a higher value than the link can sustain, the background copy process will simply use the actual available bandwidth.

remote_cluster_id | remote_cluster_name

Specifies the cluster ID or name of the remote cluster. You can use the **svctask svcinfolscandidate** command to list the remote clusters that are available. If two or more remote clusters have the same name and the name is included in this command, the command will fail and ask for the ID of the cluster instead of the name.

Description

This command creates a one-way partnership between the local cluster and the remote cluster specified by the command. The equivalent **svctask mkpartnership** command must also then be issued from the other cluster to create a fully-configured two-way partnership.

Inter-cluster Remote Copy relationships can now be created between primary VDisks in the local cluster and auxiliary VDisks in the remote cluster. Intra-cluster relationships may also be created, where both VDisks reside in the local cluster.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5925E The remote cluster partnership was not created because it already exists.
- CMMVC5926E The remote cluster partnership was not created because there are too many partnerships.
- CMMVC5927E The action failed because the cluster ID is not valid.

- CMMVC5928E The action failed because the cluster name specified is a duplicate of another cluster.

Examples

An invocation example

```
svctask mkpartnership -bandwidth 20 cluster1
```


The resulting output

No feedback

mkrconsistgrp

You can use the **mkrconsistgrp** command to create a new empty Remote Copy consistency group.

Syntax



Parameters

-name *new_name_arg*

Optionally specifies a name for the new consistency group.

-cluster *cluster_id* | *cluster_name*

Optionally specifies the ID or name of the remote cluster, in which case an inter-cluster consistency group is created. If **-cluster** is not specified, an intra-cluster consistency group is created on the local cluster only.

Description

This command creates a new consistency group. The ID of the new group is returned. The name must be unique across all consistency groups known to the clusters owning this consistency group. If the consistency group involves two clusters, the clusters must be in communication throughout the create process.

The new consistency group does not contain any relationships and will be in the empty state. Remote Copy relationships can be added to the group using the **svctask chrelationship** command.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svctask mkrconsistgrp -name rc_testgrp
```

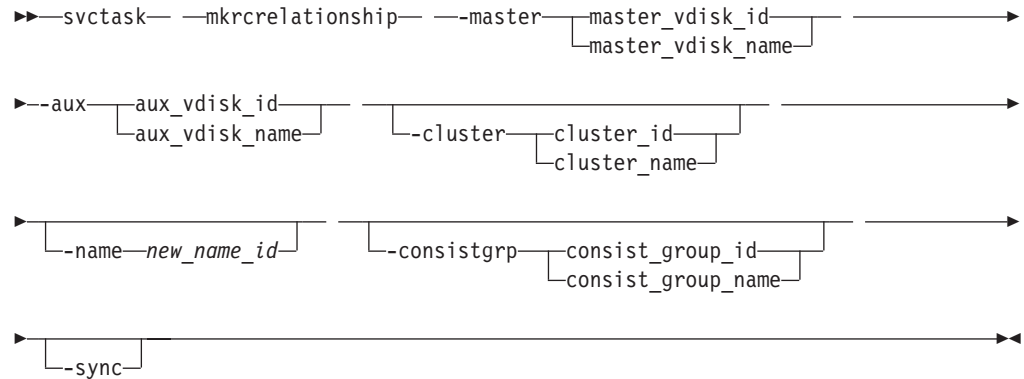
The resulting output

```
RC Consistency Group, id [255], successfully created
```

mkcrrelationship

You can use the **mkcrrelationship** command to create a new remote copy relationship with virtual disks in the same cluster (intra-cluster relationship) or in two different clusters (inter-cluster relationship).

Syntax



Parameters

- master** *master_vdisk_id* | *master_vdisk_name*
Specifies the ID or name of the master virtual disk.
- aux** *aux_vdisk_id* | *aux_vdisk_name*
Specifies the ID or name of the auxiliary virtual disk.
- cluster** *cluster_id* | *cluster_name*
Optionally specifies the ID or name of the remote cluster. This argument is required when the auxiliary virtual disk resides in the remote cluster (inter-cluster relationship). If you do not supply the **-cluster** argument, the auxiliary disk is assumed to be in the local cluster (intra-cluster relationship). To create a relationship in two different clusters, the clusters must be connected at the time the **svctask mkcrrelationship** command is received.
- name** *new_name_id*
Optionally specifies a label to assign to the relationship.
- consistgrp** *consist_group_id* | *consist_group_name*
Optionally specifies a consistency group that this relationship will join. If you do not supply the **-consistgrp** argument, the relationship will be a standalone relationship that can be started, stopped, and switched on its own.
- sync**
Optionally specifies the synchronized, or create consistency flag. Use this argument to indicate that the secondary (auxiliary) virtual disk is already synchronized with the primary (master) virtual disk. The initial background synchronization is skipped.

Description

This command creates a new Remote Copy relationship. This relationship persists until it is deleted. The auxiliary virtual disk must be identical in size to the master virtual disk or the command will fail, and if both VDisks are in the same cluster they must both be in the same I/O group. The master and auxiliary cannot be in

an existing relationship. Neither disk can be the target of a FlashCopy mapping. This command returns the new relationship (`relationship_id`) when successful.

You can optionally give the relationship a name. The name must be a unique relationship name across both clusters.

The relationship may optionally be assigned to a Remote Copy consistency group. A consistency group is used to ensure that a number of relationships are managed so that, in the event of a disconnection of the relationships, the data in all relationships within the group is in a consistent state. This can be important in say a database application where data files and log files are held on separate VDisks, and consequently are being managed by separate relationships. In the event of a disaster, the primary and secondary sites may become disconnected. If the relationships associated with the VDisks are not in a consistency group, then as the disconnection happens, and the Remote Copy relationships stop copying data from the primary to the secondary site, there is no guarantee that updates to the two separate secondary VDisks will stop in a consistent manner.

For proper database operation though it is important that updates to the log files and the database data are made in a consistent and orderly fashion. It is thus crucial in this example that the logfile VDisk and the data VDisk at the secondary site are in a consistent state. This can be achieved by putting the relationships associated with these VDisks into a consistency group. Remote Copy then ensures that updates to both VDisks at the secondary site are consistent with the updates that have been made at the primary site.

If you specify a consistency group, both the group and the relationship must have been created using the same master cluster and the same auxiliary cluster. The relationship must not be a part of another consistency group.

If the consistency group is *not empty*, then the consistency group and the relationship must be in the same state. If the consistency group is *empty*, then it will acquire the state of the first relationship that is added to it. If the state has a copy direction assigned, then the direction of the consistency group and the relationship must match that direction.

If you do not specify a consistency group, a standalone relationship is created.

If you specify the `-sync` argument, it is taken as a guarantee that the master and auxiliary virtual disks contain identical data at the point when the relationship is created. You must ensure that the auxiliary is created to match the master and that no write operation takes place to either virtual disk before you issue the `svctask mkrcrelationship` command.

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.
- CMMVC5930E The Remote Copy relationship was not created because an entity that was specified in the command does not exist.
- CMMVC5931E The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is locked.

- CMMVC5932E The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is a member of a FlashCopy mapping.
- CMMVC5933E The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is in the recovery I/O group.

Examples

An invocation example

```
svctask mkrcrelationship -master vdisk1 -aux vdisk2 -name rccopy1
```

The resulting output

```
RC Relationship, id [28], successfully created
```

rmpartnership

You can use the **rmpartnership** command to remove a Remote Copy partnership. Since the partnership exists on both clusters, it is necessary to run this command on both clusters to remove both sides of the partnership. If the command is run on only one cluster, then the Remote Copy partnership will enter a partially configured state and Remote Copy activity will cease as the relationships become disconnected.

Syntax

```
svctask --rmpartnership --remote_cluster_id | remote_cluster_name
```

Parameters

remote_cluster_id | remote_cluster_name
Specifies the cluster ID or name of the remote cluster.

Description

This command deletes the partnership between the local cluster and the remote cluster specified in the command.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5927E The action failed because the cluster ID is not valid.
- CMMVC5928E The action failed because the cluster name is a duplicate of another cluster.
- CMMVC5929E The Remote Copy partnership was not deleted because it has already been deleted.

Examples

An invocation example

```
svctask rmpartnership cluster1
```

The resulting output

No feedback

rmrcconsistgrp

You can use the **rmrcconsistgrp** command to delete an existing Remote Copy consistency group.

Syntax

```
svctask --rmrcconsistgrp [-force] rc_consist_group_id rc_consist_group_name
```

Parameters

-force

If the group contains any relationships and you do not specify the force flag, the command will fail. If one or more relationship belongs to the group and you do not specify the force flag, the deletion fails. If you specify the force flag, any relationships belonging to the group are removed from the group before it is deleted. The relationships themselves are not deleted; they become stand-alone relationships.

rc_consist_group_id | rc_consist_group_name

Specifies the ID or name of the consistency group to delete.

Description

This command deletes the specified consistency group. You can issue this command for any existing consistency group. If the consistency group is disconnected at the time that the command is issued, then the consistency group is only deleted on the cluster on which the command is being run. When the clusters reconnect, then the consistency group is automatically deleted on the other cluster. Alternatively, if the clusters are disconnected, and you still wish to remove the consistency group on both clusters, you can issue the **svctask rmrcconsistgrp** command separately on both of the clusters.

If the consistency group is not empty, the **-force** parameter is required to delete the group. This will remove the relationships from the consistency group before the group is deleted. These relationships then become standalone relationships. The state of these relationships is not changed by the action of removing them from the consistency group.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5937E The action failed because an entity that was specified in the command does not exist.
- CMMVC5938E The Remote Copy consistency group was not deleted because the consistency group contains relationships. To delete the consistency group, the force option is required.

Examples

An invocation example

```
svctask rmrcconsistgrp rctestone
```

The resulting output

```
No feedback
```

rmrcrelationship

You can use the **rmrcrelationship** command to delete an existing Remote Copy relationship.

Syntax

```
svctask --rmrcrelationship --rc_rel_id | --rc_rel_name
```

Parameters

rc_rel_id | rc_rel_name

Specifies the ID or name of the relationship. A relationship cannot be deleted if it is part of a consistency group.

Description

This command deletes the relationship that is specified.

Deleting a relationship only deletes the logical relationship between the two virtual disks, it does not affect the virtual disks themselves.

If the relationship is disconnected at the time that the command is issued, then the relationship is only deleted on the cluster on which the command is being run. When the clusters reconnect, then the relationship is automatically deleted on the other cluster. Alternatively, if the clusters are disconnected, and you still wish to remove the relationship on both clusters, you can issue the **rmrcrelationship** command independently on both of the clusters.

A relationship cannot be deleted if it is part of a consistency group. You must first remove the relationship from the consistency group using the **svctask chrrelationship -force** command.

If you delete an inconsistent relationship, the secondary virtual disk will become accessible even though it is still inconsistent. This is the one case in which Remote Copy does not inhibit access to inconsistent data.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5935E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svctask rmrcrelationship rccopy1
```


The resulting output

No feedback

startrcconsistgrp

You can use the **startrcconsistgrp** command to start the Remote Copy consistency group copy process, set the direction of copy if undefined, and optionally mark the secondary VDisks of the consistency group as clean.

Syntax



Parameters

-primary *master | aux*

This parameter specifies the copy direction by defining whether the master or auxiliary will become the primary (source). This parameter is required when the primary is undefined, for example if the consistency group is in the idling state. The primary (direction) argument specifies which disk is the primary, that is, the source disk.

-force

Optionally specifies the force parameter. This argument permits the copy operation to resume even if it might lead to a temporary loss of consistency while synchronization occurs.

-clean

Optionally specifies the clean parameter. This marks the secondary VDisk as clean for each of the relationships belonging to the group.

rc_consist_group_id | rc_consist_group_name

Specifies the ID or name of the consistency group to start.

Description

This command starts a Remote Copy consistency group.

This command can only be issued to a consistency group that is connected. For a consistency group that is idling, this command assigns a copy direction (primary and secondary roles) and begins the copy process. Otherwise this command restarts a previous copy process that was stopped either by a stop command or by some I/O error.

If the resumption of the copy process will lead to a period when the relationship is not consistent, then you must specify the force flag when restarting the relationship. This situation could arise if, say, the relationship had been stopped, and then further writes had been performed on the original primary of the relationship. The use of the force flag here is a reminder that the data on the secondary will not be useful for Disaster Recovery purposes while it is in an inconsistent state.

In the idling state, you must provide the primary argument. In other connected states, you can provide the primary argument, but it must match the existing setting.

The `-force` flag is required if consistency would be lost by starting a copy operation. This would occur, if write operations on either primary or secondary vdisks have taken place since the `ConsistentStopped` or `Idling` state was entered. If the command is issued without the `-force` flag in such circumstances, the command will fail. In general, the `-force` flag will be required if the group is in one of the following states:

- `Consistent_Stopped` but not synchronized (`sync=out_of_sync`)
- `Idling` but not synchronized

The `-force` flag is not required if the group is in one of the following states:

- `Inconsistent_Stopped`
- `Inconsistent_Copying`
- `Consistent_Synchronized`

However, the command will not fail if you do specify the `-force` flag.

The `clean` flag is used when a Remote Copy group is started and the secondary vdisks in this group are assumed to be clean. Clean in this sense means that any changes that have been made at the secondary are ignored and only changes made at the primary are considered when synchronising primary and secondary. This flag could be used in the following scenario:

1. A consistency group is created with the `synchronized` flag. (At this point it does not matter if the primary and secondary contain the same data, even though the use of the `synchronized` flag implies that this is true).
2. A `stoprconsistgrp` command is issued with the `-allow access` flag. This permits access to the secondary. Change recording begins at the primary.
3. An image of the primary is captured and loaded on to the secondary. It is permissible to allow updates to the primary during the image copy as this image need only be a fuzzy image of the primary.
4. A `starttrconsistgrp` command with `primary = master`, the `force` flag and the `clean` flag is issued. This causes the auxiliary to be marked clean, and changes on the master that have occurred since the consistency group was stopped are copied to the auxiliary.
5. Once the background copy has completed, relationships in the group will be consistent and synchronized.

Possible failures

- `CMMVC5786E` The action failed because the cluster is not in a stable state.
- `CMMVC5936E` The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svctask starttrconsistgrp rccopy1
```


The resulting output

```
No feedback
```

startrelationship

You can use the **startrelationship** command to start the Remote Copy relationship copy process, set the direction of copy if undefined, and optionally mark the secondary VDisk of the relationship as clean.

Syntax



Parameters

-primary *master* | *aux*

It specifies the copy direction by defining whether the master or auxiliary will become the primary (source). This parameter is required when the primary is undefined, for example, if the relationship is in the idling state.

-force

Optionally specifies the force parameter. This argument permits the copy operation to resume even if it might lead to a loss of consistency.

-clean

Optionally specifies the clean parameter. The clean flag marks the secondary virtual disk as clean.

rc_rel_id | **rc_rel_name**

Specifies the ID or name of the relationship that you want to start in a stand-alone relationship only.

Description

This command is used to start a standalone relationship. The command will fail if it is used to attempt to start a relationship that is part of a consistency group.

This command can only be issued to a relationship that is connected. For a relationship that is idling, this command assigns a copy direction (primary and secondary roles) and begins the copy process. Otherwise this command restarts a previous copy process that was stopped either by a stop command or by some I/O error.

If the resumption of the copy process will lead to a period when the relationship is not consistent, then you must specify the force flag when restarting the relationship. This situation could arise if, say, the relationship had been stopped, and then further writes had been performed on the original primary of the relationship. The use of the force flag here is a reminder that the data on the secondary will not be useful for Disaster Recovery purposes while it is in an inconsistent state.

In the idling state, you must provide the primary argument. In other connected states, you can provide the primary argument, but it must match the existing setting.

The `-force` flag is required if consistency would be lost by starting a copy operation. This would occur, if write operations on either primary or secondary vdisks have taken place since the `ConsistentStopped` or idling state was entered. If the command is issued without the `-force` flag in such circumstances, the command will fail. In general, the `-force` flag will be required if the relationship is in one of the following states:

- `ConsistentStopped` but not synchronized
- Idling but not synchronized

The `-force` flag is not required if the relationship is in one of the following states:

- `InconsistentStopped`
- `InconsistentCopying`
- `ConsistentSynchronized`

However, the command will not fail if you do specify the `-force` flag.

The `clean` flag is used when a Remote Copy relationship is started and the secondary VDisk in this relationship is assumed to be clean. Clean in this sense means that any changes that have been made at the secondary are ignored and only changes made at the primary are considered when synchronising primary and secondary. This flag could be used in the following scenario:

1. A relationship is created with the `synchronized` flag. (At this point it does not matter if the primary and secondary contain the same data, even though the use of the `synchronized` flag implies that this is true).
2. A `stoprelationship` command is issued with the `-allow access` flag. This permits access to the secondary. Change recording begins at the primary.
3. An image of the primary is captured and loaded on to the secondary. It is permissible to allow updates to the primary during the image copy as this image need only be a 'fuzzy' image of the primary.
4. A `startrelationship` command with `primary = master`, the `force` flag and the `clean` flag is issued. This causes the auxiliary to be marked clean, and changes on the master that have occurred since the relationship was stopped are copied to the auxiliary.
5. Once the background copy has completed, the relationship will be consistent and synchronized.

Possible failures

- `CMMVC5786E` The action failed because the cluster is not in a stable state.
- `CMMVC5936E` The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svctask startrelationship rccopy1
```

The resulting output

```
No feedback
```

stoprconsistgrp

You can use the **stoprconsistgrp** command to stop the copy process in a Remote Copy consistency group. It may also be used to enable write access to the secondary VDisks in the group if the group is in a consistent state.

Syntax

```
svctask — stoprconsistgrp — [-access] [rc_consist_group_id | rc_consist_group_name]
```

Parameters

-access

Specifies the access flag which gives the user write access to a consistent secondary. This enables write access to the secondary VDisks in the group if the group is in a consistent state.

rc_consist_group_id | rc_consist_group_name

Specifies the ID or name of the consistency group to stop.

Description

This command applies to a consistency group. You can issue this command to stop a consistency group that is copying from primary to secondary.

If the consistency group is in an inconsistent state, any copy operation stops and will not resume until you issue the **svctask startrconsistgrp** command. Write activity will no longer be copied from the primary to the secondary virtual disks belonging to the relationships in the group. For a consistency group in the ConsistentSynchronized state, this command causes a consistency freeze.

When a consistency group is in a consistent state (for example, in the ConsistentStopped, ConsistentSynchronized, or ConsistentDisconnected state) then the **-access** argument may be used with the **stoprconsistgrp** command to enable write access to the secondary virtual disks within that group.

Initial state	Final state	Notes
InconsistentStopped	InconsistentStopped	
InconsistentCopying	InconsistentStopped	
ConsistentStopped	ConsistentStopped	-access permitted
ConsistentSynchronized	ConsistentStopped	-access permitted
Idling	ConsistentStopped	-access permitted
IdlingDisconnected	unchanged	A relationship may move to stopped state when reconnected.
InconsistentDisconnected	InconsistentStopped	On the cluster issuing the svctask stoprconsistgrp command.
InconsistentDisconnected	unchanged	On the disconnected cluster.

ConsistentDisconnected	ConsistentStopped	On the cluster issuing the svctask stopprconsistgrp command, -access permitted.
ConsistentDisconnected	unchanged	On the disconnected cluster, -access permitted.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5936E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svctask stopprconsistgrp rccopy1
```

The resulting output

No feedback

stoprelationship

You can use the **stoprelationship** command to stop the copy process for a Remote Copy relationship. It may also be used to enable write access to a consistent secondary VDisk.

Syntax

```
svctask —stoprelationship— [-access] [rc_rel_id | rc_rel_name]
```

Parameters

-access

Specifies the allow access flag which gives the user write access to a consistent secondary.

rc_rel_id | rc_rel_name

Specifies the ID or name of the relationship to stop. Specify only for stand-alone relationships.

Description

This command applies to a standalone relationship. It will be rejected if it is addressed to a relationship that is part of a consistency group. You can issue this command to stop a relationship that is copying from primary to secondary.

If the relationship is in an inconsistent state, any copy operation stops and will not resume until you issue a **svctask startrelationship** command. Write activity will no longer be copied from the primary to the secondary virtual disk. For a relationship in the ConsistentSynchronized state, this command causes a consistency freeze.

When a relationship is in a consistent state (ie, in the ConsistentStopped, ConsistentSynchronized, or ConsistentDisconnected state) then the **-access** argument may be used with the **stoprelationship** command to enable write access to the secondary virtual disk.

Initial state	Final state	Notes
InconsistentStopped	InconsistentStopped	
InconsistentCopying	InconsistentStopped	
ConsistentStopped	ConsistentStopped	-access permitted
ConsistentSynchronized	ConsistentStopped	-access permitted
Idling	ConsistentStopped	-access permitted
IdlingDisconnected	unchanged	A relationship may move to stopped state when reconnected.
InconsistentDisconnected	InconsistentStopped	On the cluster issuing the svctask stoprelationship command.
InconsistentDisconnected	unchanged	On the disconnected cluster.

ConsistentDisconnected	ConsistentStopped	On the cluster issuing the svctask stopprrelationship command, -access permitted.
ConsistentDisconnected	unchanged	On the disconnected cluster, -access permitted.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5936E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svctask stopprrelationship rccopy1
```

The resulting output

No feedback

switchrcconsistgrp

You can use the **switchrcconsistgrp** command to reverse the roles of primary and secondary virtual disks in a Remote Copy consistency group when that consistency group is in a consistent state. This change will be applied to all the relationships in the consistency group.

Syntax

```
svctask — switchrcconsistgrp —-primary {master | aux}
rc_consist_group_id | rc_consist_group_name
```

Parameters

-primary *master* | *aux*

Specifies whether the master or auxiliary side of the relationships in the group will become the primary VDisks.

rc_consist_group_id | **rc_consist_group_name**

Specifies the ID or name of the consistency group to switch.

Description

This command applies to a consistency group. It is normally issued to reverse the roles of the primary and secondary in a consistency group, perhaps as part of a graceful failover. Write access to the former primary vdisks is lost and write access to the new primary vdisks is acquired. This command will only be successful when the consistency group is in a connected, consistent state, and when reversing the direction of the relationships would not lead to a loss of consistency, for example, when the consistency group is consistent and synchronized. This command will therefore only succeed when the consistency group is in one of the following states:

- ConsistentSynchronized
- ConsistentStopped and Synchronized
- Idling and Synchronized

The consistency group moves to the ConsistentSynchronized state after successful completion of this command. If you specify the **-primary** argument with the current primary, the command has no affect.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5936E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svctask switchrcconsistgrp -primary aux rccopy2
```

The resulting output

No feedback

switchrelationship

You can use the **switchrelationship** to reverse the roles of primary and secondary virtual disks in a Remote Copy relationship when that relationship is in a consistent state.

Syntax

```
svctask switchrelationship --primary master rc_rel_id  
                           aux rc_rel_name
```

Parameters

-primary *master* | *aux*

Specifies whether the master or auxiliary is to be the primary.

rc_rel_id | **rc_rel_name**

Specifies the ID or name of the relationship to switch.

Description

This command applies to a stand-alone relationship. It will be rejected if it is used to try to switch a relationship that is part of a consistency group. It is normally issued to reverse the roles of the primary and secondary virtual disk in a relationship or consistency group, perhaps as part of a graceful failover. Write access to the old primary is lost and write access to the new primary is acquired. This command will only be successful when the relationship is in a connected, consistent state, and when reversing the direction of the relationship would not lead to a loss of consistency, ie when the relationship is consistent and synchronized. This command will therefore only succeed when the relationship is in one of the following states:

- ConsistentSynchronized
- ConsistentStopped and Synchronized
- Idling and Synchronized

The relationship moves to the ConsistentSynchronized state after successful completion of this command. If you specify the **-primary** argument with the current primary, the command has no affect.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5936E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svctask switchrelationship -primary master rccopy2
```

The resulting output

No feedback

Chapter 11. Migration commands

The following commands enable you to work with migration options with the SAN Volume Controller Storage Software for Cisco MDS 9000.

migrateexts

You can use the **migrateexts** command to migrate a number of extents from a given managed disk to another.

Syntax

```
▶▶—svctask— —migrateexts— --source— {source_mdisk_id | source_mdisk_name} —————▶▶
▶▶--target— {target_mdisk_id | target_mdisk_name} --exts— number_of_extents —————▶▶
▶▶--vdisk— {vdisk_id | vdisk_name} —————▶▶
▶▶--threads— number_of_threads —————▶▶
```

Parameters

- source** *source_mdisk_id* | *source_mdisk_name*
Specifies the MDisk on which the extents currently reside.
- target** *target_mdisk_id* | *target_mdisk_name*
Specifies the MDisk onto which the extents are to be migrated.
- exts** *number_of_extents*
Specifies the number of extents to migrate.
- vdisk** *vdisk_id* | *vdisk_name*
Specifies the VDisk to which the extents belong.
- threads** *number_of_threads*
Optionally specifies the number of threads to use while migrating these extents. Valid values are 1 to 4.

Description

This command migrates a given number of extents from the source, specified in terms of the virtual disk, and the managed disk that contains some extents used to make up the virtual disk. The target is specified in terms of a managed disk (within the same managed disk group).

If a large number of extents are being migrated, you can specify the number of threads that should be started, between one and four. The progress of these can be checked by issuing the **svcinfolsmigrate** command.

The command will fail if there are insufficient free extents on the target managed disk.

The command should be used in conjunction with the **svcinfolsfreeextents** command.

Possible failures

- CMMVC5786 The action failed because the cluster is not in stable state.
- CMMVC5845 The extent was not migrated because an object that was specified in the command does not exist.
- CMMVC5849E The migration failed because some or all of the extents are already being migrated.

- CMMVC5850E The extent was not migrated because there is a problem with the source extents.
- CMMVC5851E The extent was not migrated because there is a problem with the target extents.
- CMMVC5852E The migration failed because there are too many migrations in progress.
- CMMVC5859E The migration did not complete because an error occurred while migrating the last extent on an image-mode virtual disk (VDisk).
- CMMVC5863E The migration failed because there are not enough free extents on the target managed disk (MDisk).
- CMMVC5866E The extent was not migrated because the extent contains internal data.

Examples

An invocation example

```
svctask migrateexts -vdisk vdisk4 -source mdisk4 -exts  
64 -target mdisk6 -threads 4
```

The resulting output

No feedback

migratevdisk

You can use the **migratevdisk** command to migrate an entire virtual disk from one managed disk group to another managed disk group.

Syntax

```
svctask --migratevdisk --mdiskgrp {mdisk_group_id | mdisk_group_name}
--vdisk {vdisk_id | vdisk_name}
[-threads number_of_threads]
```

Parameters

- mdiskgrp** *mdisk_group_id* | *mdisk_group_name*
Specifies the new managed disk group ID or name.
- threads** *number_of_threads*
Optionally specifies the number of threads to use while migrating these extents. You can specify 1 - 4 threads. The default of the number of threads is 1.
- vdisk** *vdisk_id* | *vdisk_name*
Specifies the virtual disk ID or name to migrate in to a new managed disk group.

Description

This command migrates the specified virtual disk into a new managed disk group. That is, all the extents that make up the virtual disk are migrated onto free extents in the new managed disk group.

The command returns a success message while the transfer completes in the background. The command returns with the `in_progress` return code. The progress of the migration can be found by issuing the **svcinfolsmigrate** command, which list the migrations in progress.

The process can be prioritized by specifying the number of threads to use while migrating. Using only one thread puts the least background load on the system.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5846E The virtual disk (VDisk) was not migrated because an entity that was specified in the command does not exist.
- CMMVC5847E The virtual disk (VDisk) was not migrated because its associated managed disk (MDisk) is already in the MDisk group.
- CMMVC5849E The migration failed because some or all of the extents are already being migrated.
- CMMVC5852E The migration failed because there are too many migrations in progress.
- CMMVC5861E The action failed because there were not enough extents on the managed disk (MDisk).

- CMMVC5863E The migration failed because there are not enough free extents on the target managed disk (MDisk).

Examples

An invocation example

```
svctask migratevdisk -vdisk 4 -mdiskgrp Group0 -threads 2
```

The resulting output

No feedback


Chapter 12. Tracing commands

The following commands enable you to work with tracing options with the SAN Volume Controller Storage Software for Cisco MDS 9000.

setdisktrace

You can use the **setdisktrace** command to set a list of disks of a given type to include in a disk trace.

Syntax



Parameters

- type** *mdisk* | *vdisk*
Specifies the object type for the disks.
- set**
Specifies the set argument. The **-set** and **-reset** arguments are mutually exclusive.
- reset**
Specifies the reset argument. The **-set** and **-reset** arguments are mutually exclusive.
- all**
Specifies that you want to trace all disks of the given type. The **-objectid** and **-all** arguments are mutually exclusive.
- objectid** *id_or_name_list*
Specifies a list of one or more disk IDs or names. The **-objectid** and **-all** arguments are mutually exclusive.

Description

This command marks a list of one, more, or all disks of the given type such that these disks will participate in the next triggered trace.

This is used in conjunction with the **svctask settrace** command which sets the options that result in a trace file being generated, and the data that is included in the trace file.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svctask setdisktrace -type mdisk -objectid  
mdisk1:mdisk3:mdisk11:mdisk10:mdisk9:mdisk5 -reset
```

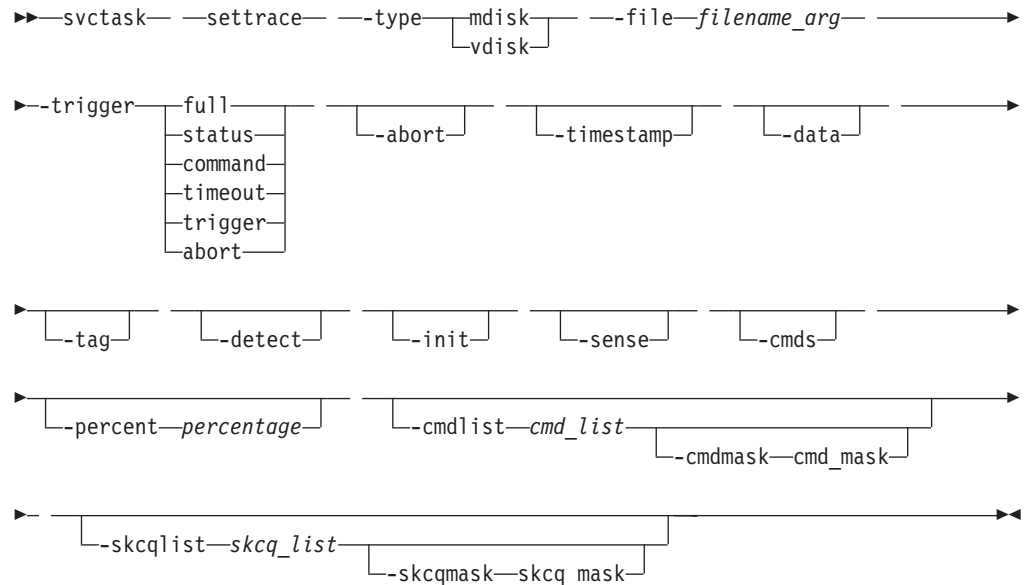
The resulting output

```
No feedback
```

settrace

You can use the **settrace** command to set options to trace certain I/O operations through the system.

Syntax



Parameters

-type *mdisk* | *vdisk*

Specifies the object type the options relate to.

-file *filename_arg*

Specifies the file name prefix for the trace file.

-trigger *full* | *status* | *command* | *timeout* | *trigger* | *abort*

Specifies the trigger option, that is, what to do when the trace is started (triggered).

The full, status, and command options are valid for both MDisks and VDIs. The timeout and trigger options are only valid for MDisks. The abort option is only valid for VDIs.

- full = when trace buffer is full, stop, that is, do not wrap
- status = when the given SCSI status (-skcqlist) is reported in sense data
- command = when the given SCSI command (-cmdlist) is sent
- timeout = when a timeout occurs
- trigger = keep running until the trigger event, that is, wrap
- abort = when an abort occurs

-abort

Optionally specifies the abort argument, which adds the abort details to the trace. This argument is only valid for VDIs.

-timestamp

Optionally specifies the time-stamp flag. Adds a time-stamp to each entry in the trace. A file name is created from the prefix plus a time-stamp. The file

name is in the form <prefix>_NN_YYMMDD_HHMMSS, where *NN* is the current configuration node ID. Files are created in the /dumps/iotrace directory.

-data

Optionally specifies the data flag, which adds the I/O data to the trace.

-tag

Optionally specifies the *ccb_tags* flag. Adds the CCB tags to the trace. This argument is for valid for MDisks.

-detect

Optionally specifies the discovery flag. Adds the MDisk discovery details to the trace for MDisks.

-init

Optionally specifies the initialization flag, which adds the MDisk initialization details to the trace for MDisks.

-sense

Optionally specifies the sense flag, which adds the SCSI sense data to the trace. This flag is only valid for VDIs.

-cmds

Optionally specifies the commands flag, which adds the commands data to the trace. This flag is only valid for VDIs only.

-percent

Optionally specifies where in the trace file the chosen trigger point should occur. That is, this flag specifies how much data to gather after the trigger point. The default is 50%, so the trigger point would be in the middle of the trace file.

-cmdlist *cmd_list*

Optionally specifies a command list, which adds only those commands to the trace file.

-cmdmask *cmd_mask*

Optionally specifies a command mask, which adds only those commands to the trace file. This can only be entered if the *-cmdlist* argument has been entered.

-skcqlist *skcq_list*

Optionally specifies a SKCQ list, which adds only those SKCQ details to the trace file.

-skcqmask *skcq_mask*

Optionally specifies a SKCQ mask, which adds only those SKCQ details to the trace file. This can only be entered if the *-skcqlist* argument has been entered.

Description

This command sets the various I/O tracing options you want for a particular disk type, that is managed disks or virtual disks. When the relevant disk type trace is subsequently triggered, the options specify the data you want to include in the trace file.

The file name specifies a file name prefix for the trace file. The system appends the node ID and a time-stamp to the file name. The node ID is the current configuration node.

A maximum of ten trace files will be kept on the cluster. When the eleventh trace is made, the oldest existing trace file is overwritten.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC6073E The maximum number of files has been exceeded.

Examples

An invocation example

```
svctask settrace -type vdisk -file tracedump -trigger abort  
-percent 100 -abort -timestamp
```

The resulting output

No feedback

starttrace

You can use the **starttrace** command to start tracing I/O operations based on the option currently set for the given object type and the list of disks to trace.

Syntax

```
svctask --starttrace --type {mdisk | vdisk}
```

Parameters

-type *mdisk* | *vdisk*
Specifies the object type to trigger.

Description

This command starts the collection of I/O tracing information. The trace file is generated according to the options that you specified in the **svctask settrace** command. The disks that will be traced are those identified in the list set by the **svctask setdisktrace** command.

The traces are written to the `/dumps/iotrace` directory. The contents of this directory can be viewed using the **svcinfolsiotracedumps** command.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5809E The tracing of I/O operations was not started because it is already in progress.
- CMMVC5986E The tracing of I/O operations was not started because the virtual disk (VDisk) or managed disk (MDisk) failed to return statistics.

Examples

An invocation example

```
svctask starttrace -type vdisk
```

The resulting output

No feedback

stoptrace

You can use the **stoptrace** command to stop the tracing for the given disk type.

Syntax

```
▶▶—svctask— —stoptrace— --type— mdisk  
vdisk————▶▶
```

Parameters

-type *mdisk* | *vdisk*

Specifies the object type to stop tracing.

Description

This command stops the tracing of I/O operations for the given object type. When you issue the **svctask stoptrace** command, you might not get a trace file if the trigger options have not been met.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svctask stoptrace -type mdisk
```

The resulting output

No feedback

Chapter 13. Attributes of the `-filtervalue` argument

The list report style can be used to obtain two types of reports:

- A list containing concise information about all the objects of a given type. Using this style, each entry in the list corresponds to a single object. A concise view gives summary information about all the objects in table format.
- A list containing detailed information about a single object of a given type. Using this style, each entry in the list corresponds to a single attribute. A detailed view gives information about more attributes for one object in a list format.

You can use the `-filtervalue` argument to filter a view based on some 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. When using a wildcard you must ensure that the `filtervalue` is surrounded by quotes. For example, to list all the VDisks with names ending with the characters `E3`:

```
svcinfo lsvdisk -filtervalue 'vdisk_name=*E3'
```

Table 5. 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.
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: <ul style="list-style-type: none"> • 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.

Table 5. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
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: <ul style="list-style-type: none"> • online • degraded • excluded • offline
	mode	=	No	The mode of the MDisk. The valid input options for MDisk mode are: <ul style="list-style-type: none"> • 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.	
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 VDIs in the group.
	status	=	No	The status of the MDisk group. The valid input options are: <ul style="list-style-type: none"> • online • degraded • offline
	extent_size	=, <, <=, >, >=	No	The extent size. (MB)

Table 5. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • seq • striped • image
	FC_name	=	Yes	The FlashCopy mapping name.
	FC_id	=, <, <=, >, >=	No	The FlashCopy mapping ID.
	RC_name	=	Yes	The Remote Copy relationship name.
RC_id	=, <, <=, >, >=	No	The Remote Copy 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 5. 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • idle_or_copied • preparing • prepared • copying • stopped • suspended

Table 5. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
rrelation-ship	RC_rel_id or id	=, <, <=, >, >=	No	The Remote Copy relationship ID.
	RC_rel_name or name	=	Yes	The Remote Copy 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: <ul style="list-style-type: none"> • master • aux
	consistency_group_id	=, <, <=, >, >=	No	The Remote Copy consistency group ID.
	consistency_group_name	=	Yes	The Remote Copy consistency group name.
	state	=	Yes	The relationship state. The valid input values are: <ul style="list-style-type: none"> • inconsistent_stopped • inconsistent_copying • consistent_stopped • consistent_synchronised • idling • idling_disconnected • inconsistent_disconnected • consistent_disconnected
	progress	=, <, <=, >, >=	No	The progress of the initial background copy (synchronization) for the relationship.

Table 5. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
rconsist-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: <ul style="list-style-type: none"> • master • aux
	state	=	No	The consistency group state. The valid input values are: <ul style="list-style-type: none"> • inconsistent_stopped • inconsistent_copying • consistent_stopped • consistent_synchronised • idling • idling_disconnected • inconsistent_disconnected • consistent_disconnected • empty
	relationship_count	=, <, <=, >, >=	No	The relationship count.

Chapter 14. 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 Storage Software for Cisco MDS 9000 are contained in the following directory structure:

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

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.

The file names used for storing I/O statistics dumps are `m_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 VDIs. 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: Application abend dumps and trace files are contained in the `/dumps`, `/dumps/ec` and `/dumps/pl` directories. Application abend dumps in the `/dumps/ec` directory will be named `ecdumps.NNNNNN.YYMMDD.HHMMSS.tar`, where `NNNNNN` is the front panel name. Application abend dumps in the `/dumps/pl` directory will be similarly named `pldumps.NNNNNN.YYMMDD.HHMMSS.tar`. In addition there may be some trace files written to the `/dumps` directory. These will be named `NNNNNN.trc` or `NNNNN_XXX.trc` where `XXX` designates a particular type of trace file.

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

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.

Chapter 15. Information commands


The following commands enables you to work with displaying specific types of information with the SAN Volume Controller Storage Software for Cisco MDS 9000.

Note: Since an object may be created, deleted, and recreated, and have different IDs, it is preferable to use object names whenever possible.

caterlog

You can use the **caterlog** command to return the contents of the cluster error and event log.

Syntax



Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view providing general information about objects of a particular type), and each item of data (in a detailed style view providing much more information about a specific object of a particular type). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim delimiter

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

-config

Optionally specifies to list the configuration events.

-unfixed

Optionally specifies to list the unfixed errors.

-first *number_of_entries_to_return*

Optionally specifies to display the first *n* number of entries in the log, where *n* is the number entered by the user as an argument to the **-first** flag.

Description

This command returns a list of the specified error log entries. When no flags are passed, all error log entries are listed.

The list can be filtered to only include configuration events, or unfixed errors by specifying the **-config** or **-unfixed** arguments.

Using the **-first** parameter will result in the first *x* number of records being displayed, where *x* is the number entered as an argument for the **-first** parameter.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcinfo caterrlog -delim :
```

The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code
0:cluster:no:no:6:node1:100:100:030407052547:030407052547:1:00981001
0:fc_card:no:no:1:node1:101:101:030407052547:030407052547:1:00073001
1:node:no:no:1:node1:102:102:030407052547:030407052547:1:00074001
0:cluster:no:no:6:node1:103:100:030407052547:030407052547:1:00981001
1:fc_card:no:no:1:node1:104:104:030407052632:030407052632:1:00073003
0:node:no:no:6:node1:105:105:030407082202:030407082717:2:00980500
2:remote:no:no:6:n/a:106:106:030407090117:030407090117:1:00985002
1:node:no:no:5:node1:0:0:030407052546:030407052546:1:00990383
0:cluster:no:no:5:node1:0:0:030407080630:030407080630:1:00990117
0:mdisk_grp:no:no:5:node1:0:0:030407081610:030407081610:1:00990148
128:mdisk_grp:no:no:5:node1:0:0:030407081610:030407081610:1:00990173
1:mdisk_grp:no:no:5:node1:0:0:030407081619:030407081619:1:00990148
0:vdisk:no:no:5:node1:0:0:030407081836:030407081836:1:00990169
1:vdisk:no:no:5:node1:0:0:030407081843:030407081843:1:00990169
0:vdisk:no:no:5:node1:0:0:030407081854:030407081854:1:00990169
0:vdisk:no:no:5:node1:0:0:030407082015:030407082015:1:00990169
0:vdisk:no:no:5:node1:0:0:030407082145:030407082145:1:00990169
0:vdisk:no:no:5:node1:0:0:030407082148:030407082148:1:00990169
0:vdisk:no:no:5:node1:0:0:030407082158:030407082158:1:00990169
1:vdisk:no:no:5:node1:0:0:030407082213:030407082213:1:00990169
0:host:no:no:5:node1:0:0:030407082441:030407082441:1:00990106
1:host:no:no:5:node1:0:0:030407082457:030407082457:1:00990106
2:host:no:no:5:node1:0:0:030407082523:030407082523:1:00990106
0:flash:no:no:5:node1:0:0:030407082704:030407082704:1:00990184
1:node:no:no:5:node1:0:0:030407082716:030407082716:1:00990501
1:node:no:no:5:node1:0:0:030407082722:030407082722:1:00990501
1:fc_const_grp:no:no:5:node1:0:0:030407083141:030407083141:1:00990204
2:fc_const_grp:no:no:5:node1:0:0:030407083143:030407083143:1:00990204
3:fc_const_grp:no:no:5:node1:0:0:030407083145:030407083145:1:00990204
0:flash:no:no:5:node1:0:0:030407083318:030407083318:1:00990185
0:flash:no:no:5:node1:0:0:030407083355:030407083355:1:00990185
0:flash:no:no:5:node1:0:0:030407085753:030407085753:1:00990185
1:remote:no:no:5:node1:0:0:030407085932:030407085932:1:00990225
2:vdisk:no:no:5:node1:0:0:030407085959:030407085959:1:00990169
3:vdisk:no:no:5:node1:0:0:030407090004:030407090004:1:00990169
4:vdisk:no:no:5:node1:0:0:030407090013:030407090013:1:00990169
2:remote:no:no:5:node1:0:0:030407090106:030407090106:1:00990225
255:rc_const_grp:no:no:5:node1:0:0:030407090323:030407090323:1:00990240
254:rc_const_grp:no:no:5:node1:0:0:030407090327:030407090327:1:00990240
253:rc_const_grp:no:no:5:node1:0:0:030407090333:030407090333:1:00990240
2:remote:no:no:5:node1:0:0:030407090442:030407090442:1:00990226
1:vdisk:no:no:5:node1:0:0:030407090820:030407090820:1:00990182
3:vdisk:no:no:5:node1:0:0:030407090825:030407090825:1:00990182
```

caterlogbyseqnum

You can use the **caterlogbyseqnum** command to display all the errors with a sequence number, or root cause number, as specified by the user.

Syntax

```
▶ svcinfo --caterlogbyseqnum --num sequence_number --root root_cause_number  
└──nohdr┘ └──delim delimiter┘
```

Parameters

-num *sequence_number*

Specifies the sequence number to view.

-root *root_cause_number*

Specifies the root sequence number. All errors marked with this root cause will be displayed.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command will return a single error log entry, as specified by the sequence number passed with the **-num** argument.

If **-root** argument is used, the log will be searched for all entries marked with a root cause sequence number as specified. This returns a list of all entries marked with this root cause.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcinfo caterrlogbyseqnum -num 100 -delim :
```

The resulting output


```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:  
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:  
error_code  
0:cluster:no:no:6:node1:100:100:030407052547:030407052547:1:00981001
```

lssvcdumps (ls2145dumps)

You can use the **lssvcdumps** command to obtain a list of dumps from the /dumps directory.

Note: For compatibility, **ls2145dumps** can be used as a synonym for **lssvcdumps**.

Syntax



Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim delimiter

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

node_id | node_name

Specifies the node ID or name to list the available dumps of the given type. If you do not specify a node, the dumps available on the configuration node are listed.

Description

This command returns a list of node assert dumps and associated output files. These dumps are created as a result of the assertion of a node. If you do not specify a node, the dumps available on the configuration node will be listed. The command will display files from the /dumps, /dumps/ec, and /dumps/pl directories. Files will be listed by date regardless of directory. Files with names starting with ecdump will be from the /dumps/ec directory. Files with names starting with pldump will be from the /dumps/pl directory. Other files (files with a .trc extension) will be from /dumps.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcinfo lssvcdumps -delim :
```

The resulting output

```
id:svc_filename  
0:pldump.vegas3.9.2.031122.191834.tar  
1:pldump.vegas3.9.2.031123.031812.tar  
2:pldump.vegas3.9.2.031123.070501.tar  
3:vegas3.9.2.trc  
4:vegas3.9.2_icsmi.trc  
5:vegas3.9.2_icsp.trc.old  
6:vegas3.9.2_icsp.trc
```

Related topics

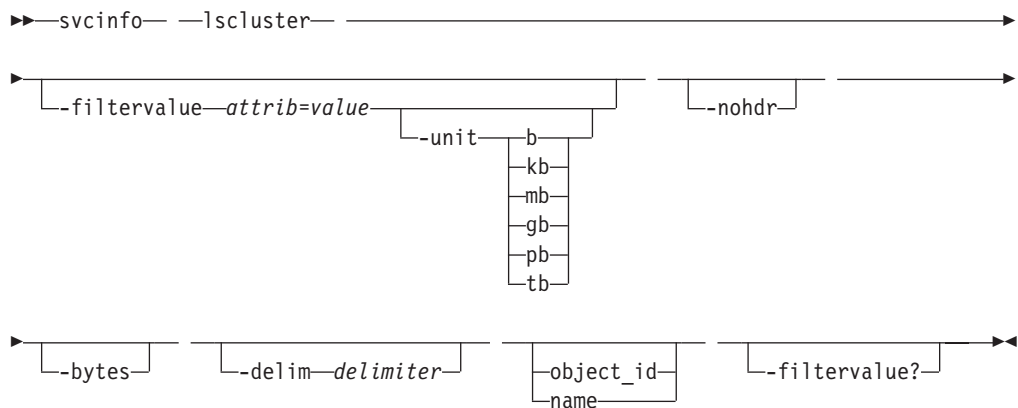
- Chapter 14, “Overview of the list dump commands,” on page 151

lscluster

The list report style can be used to obtain two styles of report.

1. A list containing concise information about all clusters. (Each entry in the list corresponds to a single cluster.)
2. The detailed information about a single, user specified, cluster.

Syntax



Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-bytes

Optionally used to display all capacities as bytes.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

-filtervalue *attribute=value*

Optionally specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

object_id | name

Optionally specifies the name or ID of an object. If not supplied, the concise view of all objects, or all of those objects matching the filtering requirements in

-filtervalue (if specified), of the given type are returned. If supplied, the detailed view of the specific object is returned, and any filtervalue entry (if entered) is ignored.

-filtervalue?

Display a list of valid filter attributes. The valid filters for the **svcinfolcluster** command are:

- cluster_name
- cluster_unique_id
- id
- name

Description

This command will return a concise list or a detailed view, of clusters.

The following list provides possible values that are applicable to the attributes that are displayed as data in the output views:

location	local, remote
statistics status	on, off
SNMP setting	none, all, hardware_only
E-mail setting	none, all, hardware_only

The location, partnership and bandwidth fields are relevant to Remote Copy configurations where the SAN fabrics of two clusters are linked together. Information about the remote cluster will be reported by the **lscluster** command if the **mkpartnership** command has been issued from the local cluster to the remote cluster, ie. if the partnership has been at least partially established from the local cluster.

You can issue the **svcinfolcluster** command to display a concise view of the cluster.

```
svcinfolcluster -delim : 10030a007e5
```

where *10030a007e5* is the name of the cluster. The output from this command will include the following for each cluster on the fabric:

- cluster ID
- cluster name
- cluster IP address

For the remote cluster, these fields indicate the following:

location: remote

partnership: partially_configured (mkpartnership command has only been issued from the local cluster to the remote cluster)

fully_configured (mkpartnership command has been issued in both directions)

bandwidth: MB/sec (the bandwidth available on the inter-cluster link for background copy)

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

A concise invocation example

```
svcinfolcluster -delim :
```

The concise resulting output

```
id:name:location:partnership:bandwidth:cluster_IP_address: 0000020062813ABA:  
clusterA:local:::9.20.247.210: 0000020062006746:clusterB:remote:  
fully_configured:50:9.20.247.211:
```

Examples

A detailed invocation example

```
svcinfolcluster -delim : cluster1
```

The detailed resulting output

```
id:00000200EC2144FE  
name:cluster1  
location:local  
partnership:  
bandwidth:  
cluster_IP_address:9.20.165.16  
total_mdisk_capacity:71.8GB  
space_in_mdisk_grps:48.5GB  
space_allocated_to_vdisks:0  
total_free_space:71.8GB  
statistics_status:off  
statistics_frequency:15  
required_memory:4096  
SNMP_setting:all  
email_setting:all  
code_level:1.3(1) (IBM build 0.3.3)
```

Examples

A concise invocation example for a Remote Copy configuration, where clusterA has issued the mkpartnership to clusterB, and the intercluster bandwidth is set to 50 MB/s.

```
svcinfolcluster -delim :
```

The concise resulting output

```
id:name:location:partnership:bandwidth:  
cluster_IP_address  
0000020062813ABA:clusterA:local:::9.20.247.210  
0000020062006746:clusterB:remote:  
fully_configured:50:9.20.247.211
```

Related topics

- Chapter 13, “Attributes of the -filtervalue argument,” on page 143

lsclustercandidate

You can use the **lsclustercandidate** command to list the clusters that are available for setting up a two-cluster partnership. This is a prerequisite for creating Remote Copy relationships.

Syntax

```
►► svcinfo — lsclustercandidate — [nohdr] [delim delimiter] ►►
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-delim delimiter

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command returns a list of clusters that are available as candidate partner clusters to form a Remote Copy Partnership between two clusters.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svcinfo lsclustercandidate
```


The resulting output

```
id                configured  cluster_name
0000010034E0F430 no          ldcluster26
```

lsconfigdumps

You can use the **lsconfigdumps** command to return a list of the config dumps on the node in service mode.

Syntax



Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

node_id | node_name

Specifies the node ID or name to list the available dumps of the given type. If you do not specify a node, the dumps available on the configuration node are listed.

Description

This command returns a list of configuration dumps. These dumps are created as a result of issuing the **svctask dumpconfig** command. A configuration dump describes the configuration of the cluster. If you do not specify a node, the dumps available on the configuration node will be listed. The command displays the files from the `/dumps/configs` directory.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcinfo lsconfigdumps
```


The resulting output

id	config_filename
0	config_myswitch1.3.2_030403_101205

Related topics


- Chapter 14, “Overview of the list dump commands,” on page 151

lscontroller

The list report style can be used to obtain two styles of report.

1. A list containing concise information about controllers. (Each entry in the list corresponds to a single controller.)
2. The detailed information about a single, user specified, controller.

Syntax



Parameters

-filtervalue *attribute=value*

Optionally specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-bytes

Optionally used to display all capacities as bytes.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

object_id | name

Optionally specifies the name or ID of a controller. If not supplied, the concise view of all objects, or all of those objects matching the filtering requirements in

-filtervalue (if specified), of the given type are returned. If supplied, the detailed view of the specific object is returned, and any filtervalue entry (if entered) is ignored.

-filtervalue?

Display a list of valid filter attributes. The valid filters for the **svcinfo lscontroller** command is:

- controller_id
- id

Description

This command will return a concise list or a detailed view, of controllers visible to the cluster.

The following list provides possible values that are applicable to the attributes that are displayed as data in the output views:

degraded no, yes

Determining a storage controllers name from its SAN Volume Controller Storage Software for Cisco MDS 9000 name: List the storage controllers by issuing the **svcinfo lscontroller** command. Remember the controller name or ID for the controller you want to determine. For the controller in question, issue the **svcinfo lscontroller <controllername/id>** command, where *<controllername/id>* is the controller name or ID. Remember the WWNN for the controller. Make a written record of it. The WWNN can be used to determine the actual storage controller by launching the native controller user interface or using the command line tools it provides to verify the actual controller that has this WWNN.

Determining the relationship between MDisks and RAID arrays or LUNs: Each MDisk corresponds with a single RAID array, or a single partition on a given RAID array. Each RAID controller will define a LUN number for this disk. The LUN number and controller name or ID are needed to be able to determine the relationship between mdisks and RAID arrays or partitions.

Show the detailed view of the given MDisk *<mdiskname>*, by issuing the **svcinfo lsmdisk <mdiskname>** command, where *<mdiskname>* is the name of the MDisk.

Note: Remember the controller name or controller ID and controller LUN number. Show the detailed view of the controller determined in by issuing the **svcinfo lscontroller <controllername>** command, where *<controllername>* is the name of the controller.

Note: Remember the vendor ID, product ID, and WWNN. Use these to determine what is being presented to the MDisk.

From the native user interface for the given controller, list the LUNs it is presenting and match the LUN number. This will tell you the exact RAID array or partition that corresponds with the MDisk.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

A concise invocation example

```
svcinfo lscontroller -delim :
```

The concise resulting output

```
id:controller_name:ctrl_s/n:vendor_id:product_id_low:product_id_high
7:controller7:3EK0J5Y8:SEAGATE :ST373405:FC
8:controller8:3EK0J6CR:SEAGATE :ST373405:FC
9:controller9:3EK0J4YN:SEAGATE :ST373405:FC
10:controller10:3EK0GKGH:SEAGATE :ST373405:FC
11:controller11:3EK0J85C:SEAGATE :ST373405:FC
12:controller12:3EK0JBR2:SEAGATE :ST373405:FC
13:controller13:3EKYNJF8:SEAGATE :ST373405:FC
14:controller14:3EK0HVTM:SEAGATE :ST373405:FC
```

Examples

A detailed invocation example

```
svcinfo lscontroller -delim = 7
```

The detailed resulting output

```
id=7
controller_name=controller7
WWNN=20000004CF2412AC
mdisk_link_count=1
max_mdisk_link_count=1
degraded=no
vendor_id=SEAGATE
product_id_low=ST373405
product_id_high=FC
product_revision=0003
ctrl_s/n=3EK0J5Y8
WWPN=22000004CF2412AC
path_count=1
max_path_count=1
WWPN=21000004CF2412AC
path_count=0
max_path_count=0
```


Related topics

- Chapter 13, “Attributes of the -filtervalue argument,” on page 143

lserrlogbyfcconsistgrp

You can use the **lserrlogbyfcconsistgrp** command to display the error log by FlashCopy consistency group.

Syntax



Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim delimiter

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

-count number

Optionally specifies to list the last number entries in the log. The **-count** argument specifies the maximum number of errors or events to list.

-config

Optionally specifies to list configuration events. When **-config** argument is used, the command will act as described above, but only list configuration events.

-unfixed

Optionally specifies to list unfixed errors. When the **-unfixed** argument is used, the command will act as above, but will only list unfixed errors.

fcconsistgrp_id | fcconsistgrp_name

Optionally specifies the object ID used to filter the log.

Description

When executed this command will display a list of the errors and events in the log related to FlashCopy consistency groups. The list can be filtered further by specifying a specific object ID or name. This will return only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events or the unfixed errors for the given object type or object ID. Similarly the last x entries against a given object type or object ID can be listed.

Note: Although there is an object type of unknown displayed in the error log, there is no command with which to filter these object types.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcinfo lserrlogbyfcconsistgrp -delim :
```


The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:  
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code  
3:fc_const_grp:no:no:5:node1:0:0:030407083145:030407083145:1:00990204  
2:fc_const_grp:no:no:5:node1:0:0:030407083143:030407083143:1:00990204  
1:fc_const_grp:no:no:5:node1:0:0:030407083141:030407083141:1:00990204
```

lserrlogbyfcmmap

This command displays the error log by FlashCopy mapping.

Syntax



Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

-count *number*

Optionally specifies to only list the last *number* entries in the log. `-count` specifies the maximum number of errors or events to list.

-config

Optionally specifies to only list configuration events. When `-config` is passed, the command will act as described above, but only list configuration events.

-unfixed

Optionally specifies to only list unfixed errors. When `-unfixed` is passed, the command will act as above, but will only list unfixed errors.

fcmmap_id | fcmmap_name

Optionally specifies the object id used to filter the log.

Description

When executed this command will display a list of the errors and events in the log related to flash copy mappings. The list can be filtered further by specifying a specific object ID or name. This will return only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events or the unfixed errors for the given object type or object ID. Similarly the last *x* entries against a given object type or object ID can be listed.

Note: Although there is an object type of unknown displayed in the error log, there is no command with which to filter these object types.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcinfolerrlogbyfcmap -delim :
```


The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:  
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code  
0:flash:no:no:5:node1:0:0:030407085753:030407085753:1:00990185  
0:flash:no:no:5:node1:0:0:030407083355:030407083355:1:00990185  
0:flash:no:no:5:node1:0:0:030407083318:030407083318:1:00990185  
0:flash:no:no:5:node1:0:0:030407082704:030407082704:1:00990184
```

Iserrlogbyhost

You can use the **iserrlogbyhost** command to display the error log by host.

Syntax



Parameters

-count *number*

Optionally specifies to list the last *number* entries in the log. The **-count** argument specifies the maximum number of errors or events to list.

-config

Optionally specifies to list configuration events. When the **-config** argument is used, the command will act as described above, but only list configuration events.

-unfixed

Optionally specifies to list unfixed errors. When the **-unfixed** argument is used, the command will act as above, but will only list unfixed errors.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

host_id | **host_name**

Optionally specifies the object ID used to filter the log.

Description

This command displays a list of the errors and events in the log related to hosts. The list can be filtered further by specifying a specific object ID or name. This will return only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events or the

unfixed errors for the given object type, or object ID. Similarly, the last x entries against a given object type or object ID can be listed.

Note: Although there is an object type of unknown displayed in the error log, there is no command with which to filter these object types.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcinfo lserrlogbyhost -delim :
```


The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:  
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code  
2:host:no:no:5:node1:0:0:030407082523:030407082523:1:00990106  
1:host:no:no:5:node1:0:0:030407082457:030407082457:1:00990106  
0:host:no:no:5:node1:0:0:030407082441:030407082441:1:00990106
```

Iserrlogbyiogrp

You can use the **Iserrlogbyiogrp** command to display the error log by I/O group.

Syntax



Parameters

-count *number*

Optionally specifies to list the last *number* entries in the log. The **-count** argument specifies the maximum number of errors or events to list.

-config

Optionally specifies to list configuration events. When the **-config** argument is used, the command will act as described above, but only list configuration events.

-unfixed

Optionally specifies to only list unfixed errors. When the **-unfixed** argument is used, the command will act as above, but will only list unfixed errors.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

iogroup_id | iogroup_name

Optionally specifies the object ID used to filter the log.

Description

This command displays a list of the errors and events in the log related to I/O groups. The list can be filtered further by specifying a specific object ID or name. This will return only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events

or the unfixed errors for the given object type or object ID. Similarly the last x entries against a given object type or object ID can be listed.

Note: Although there is an object type of unknown displayed in the error log, there is no command with which to filter these object types.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcinfo lserrlogbyiogrp -delim :
```


The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:  
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code  
1:io_grp:no:no:1:node1:109:109:030407094417:030407094417:1:00000001
```

lserrlogbymdisk

You can use the **lserrlogbymdisk** command to display the error log by MDisk.

Syntax



Parameters

-count *number*

Optionally specifies to only list the last number entries in the log. The `-count` argument specifies the maximum number of errors or events to list.

-config

Optionally specifies to only list configuration events. When the `-config` argument is used, the command will act as described above, but only list configuration events.

-unfixed

Optionally specifies to only list unfixed errors. When the `-unfixed` argument is used, the command will act as above, but will only list unfixed errors.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

mdisk_id | mdisk_name

Specifies the object ID used to filter the log.

Description

This command displays a list of the errors and events in the log related to MDisks. The list can be filtered further by specifying a specific object ID or name. This will return only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events or the

unfixed errors for the given object type, or object ID. Similarly, the last x entries against a given object type or object ID can be listed.

Note: Although there is an object type of unknown displayed in the error log, there is no command with which to filter these object types.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcinfo lserrlogbydisk -delim :
```


The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:  
sequence_number:root_sequence_number:first_timestamp:  
last_timestamp:number_of_errors:error_code  
11:mdisk:no:no:3:node1:108:108:030407092947:030407092947:1:00000016  
11:mdisk:no:no:2:node1:107:107:030407092947:030407092947:1:00000016
```

lserrlogbydiskgroup

You can use the **lserrlogbydiskgroup** command to display the error log by MDisk group.

Syntax



Parameters

-count *number*

Optionally specifies to list the last number entries in the log. The **-count** argument specifies the maximum number of errors or events to list.

-config

Optionally specifies to list configuration events. When the **-config** argument is used, the command will act as described above, but only list configuration events.

-unfixed

Optionally specifies to list unfixed errors. When the **-unfixed** argument is used, the command will act as above, but will only list unfixed errors.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

mdisk_group_id | mdisk_group_name

Optionally specifies the object ID used to filter the log.

Description

This command displays a list of the errors and events in the log related to MDisk groups. The list can be filtered further by specifying a specific object ID or name. This will return only the errors and events that have been logged against the

specified object. The list can also be filtered to show only the configuration events or the unfixed errors for the given object type, or object ID. Similarly, the last x entries against a given object type or object ID can be listed.

Note: Although there is an object type of unknown displayed in the error log, there is no command with which to filter these object types.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcinfo lserrlogbydiskgrp -delim :
```


The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:  
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code  
1:mdisk_grp:no:no:5:node1:0:0:030407081619:030407081619:1:00990148  
128:mdisk_grp:no:no:5:node1:0:0:030407081610:030407081610:1:00990173  
0:mdisk_grp:no:no:5:node1:0:0:030407081610:030407081610:1:00990148
```

Iserrlogbynode

You can use the **iserrlogbynode** command to display the error log by node.

Syntax



Parameters

-count *number*

Optionally specifies to list the last *number* entries in the log. The **-count** argument specifies the maximum number of errors or events to list.

-config

Optionally specifies to list configuration events. When the **-config** argument is used, the command will act as described above, but only list configuration events.

-unfixed

Optionally specifies to only list unfixed errors. When the **-unfixed** argument is used, the command will act as above, but will only list unfixed errors.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

node_id | **node_name**

Optionally specifies the object ID used to filter the log.

Description

This command displays a list of the errors and events in the log related to nodes. The list can be filtered further by specifying a specific object ID or name. This will return only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events or the

unfixed errors for the given object type or object ID. Similarly, the last x entries against a given object type or object ID can be listed.

Note: Although there is an object type of unknown displayed in the error log, there is no command with which to filter these object types.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcinfo lserrlogbynode -delim :
```


The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:  
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code  
1:node:no:no:5:node1:0:0:030407082722:030407082722:1:00990501  
1:node:no:no:5:node1:0:0:030407082716:030407082716:1:00990501  
1:node:no:no:5:node1:0:0:030407052546:030407052546:1:00990383  
0:node:no:no:6:node1:105:105:030407082202:030407082717:2:00980500  
1:node:no:no:1:node1:102:102:030407052547:030407052547:1:00074001
```

lserrlogbyrconsistgrp

You can use the **lserrlogbyrconsistgrp** command to display the error log by Remote Copy consistency group.

Syntax



Parameters

-count *number*

Optionally specifies to only list the last number entries in the log. -count specifies the maximum number of errors or events to list.

-config

Optionally specifies to only list configuration events. When -config is passed, the command will act as described above, but only list configuration events.

-unfixed

Optionally specifies to only list unfixed errors. When -unfixed is passed, the command will act as above, but will only list unfixed errors.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the -nohdr parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the -nohdr option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the -delim parameter will override this behavior. Valid input for the -delim parameter is a one byte character. If, for example, you entered -delim : a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

rconsistgrp_id | rconsistgrp_name

Optionally specifies the object id used to filter the log.

Description

When executed this command will display a list of the errors and events in the log related to remote copy consistency groups. The list can be filtered further by specifying a specific object ID or name. This will return only the errors and events that have been logged against the specified object. The list can also be filtered to

show only the configuration events or the unfixed errors for the given object type or object ID. Similarly the last x entries against a given object type or object ID can be listed.

Note: Although there is an object type of unknown displayed in the error log, there is no command with which to filter these object types.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcinfo lserrlogbyrcconsistgrp -delim :
```


The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:  
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code  
253:rc_const_grp:no:no:5:node1:0:0:030407090333:030407090333:1:00990240  
254:rc_const_grp:no:no:5:node1:0:0:030407090327:030407090327:1:00990240  
255:rc_const_grp:no:no:5:node1:0:0:030407090323:030407090323:1:00990240
```

lserrlogbyrcrelationship

You can use the **lserrlogbyrcrelationship** command to display the error log by Remote Copy relationship.

Syntax



Parameters

-count *number*

Optionally specifies to only list the last number entries in the log. The `-count` argument specifies the maximum number of errors or events to list.

-config

Optionally specifies to only list configuration events. When `-config` is passed, the command will act as described above, but only list configuration events.

-unfixed

Optionally specifies to only list unfixed errors. When `-unfixed` is passed, the command will act as above, but will only list unfixed errors.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

rrelationship_id | rrelationship_name

Optionally specifies the object id used to filter the log.

Description

When executed this command will display a list of the errors and events in the log related to remote copy relationships. The list can be filtered further by specifying a specific object ID or name. This will return only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events or the unfixed errors for the given object type or object ID. Similarly the last x entries against a given object type or object ID can be listed.

Note: Although there is an object type of unknown displayed in the error log, there is no command with which to filter these object types.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcinfo lserrlogbyrcrelationship -delim :
```


The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:  
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code  
2:remote:no:no:5:node1:0:0:030407090442:030407090442:1:00990226  
2:remote:no:no:5:node1:0:0:030407090106:030407090106:1:00990225  
1:remote:no:no:5:node1:0:0:030407085932:030407085932:1:00990225  
2:remote:no:no:6:n/a:106:106:030407090117:030407090117:1:00985002
```

Iserrlogbyvdisk

You can use the **Iserrlogbyvdisk** command to display the error log by VDisk.

Syntax



Parameters

-count *number*

Optionally specifies to list the last *number* entries in the log. The **-count** argument specifies the maximum number of errors or events to list.

-config

Optionally specifies to list configuration events. When the **-config** argument is used, the command will act as described above, but only list configuration events.

-unfixed

Optionally specifies to only list unfixed errors. When the **-unfixed** argument is used, the command will act as above, but will only list unfixed errors.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

vdisk_id | **vdisk_name**

Optionally specifies the object ID used to filter the log.

Description

This command displays a list of the errors and events in the log related to VDIs. The list can be filtered further by specifying a specific object ID or name. This will return only the errors and events that have been logged against the specified object. The list can also be filtered to show only the configuration events or the

unfixed errors for the given object type, or object ID. Similarly, the last x entries against a given object type or object ID can be listed.

Note: Although there is an object type of unknown displayed in the error log, there is no command with which to filter these object types.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcinfolerrlogbyvdisk -delim :
```


The resulting output

```
id:type:fixed:SNMP_trap_raised:error_type:node_name:sequence_number:
root_sequence_number:first_timestamp:last_timestamp:number_of_errors:error_code
3:vdisk:no:no:5:node1:0:0:030407090825:030407090825:1:00990182
1:vdisk:no:no:5:node1:0:0:030407090820:030407090820:1:00990182
4:vdisk:no:no:5:node1:0:0:030407090013:030407090013:1:00990169
3:vdisk:no:no:5:node1:0:0:030407090004:030407090004:1:00990169
2:vdisk:no:no:5:node1:0:0:030407085959:030407085959:1:00990169
1:vdisk:no:no:5:node1:0:0:030407082213:030407082213:1:00990169
0:vdisk:no:no:5:node1:0:0:030407082158:030407082158:1:00990169
0:vdisk:no:no:5:node1:0:0:030407082148:030407082148:1:00990169
0:vdisk:no:no:5:node1:0:0:030407082145:030407082145:1:00990169
0:vdisk:no:no:5:node1:0:0:030407082015:030407082015:1:00990169
0:vdisk:no:no:5:node1:0:0:030407081854:030407081854:1:00990169
1:vdisk:no:no:5:node1:0:0:030407081843:030407081843:1:00990169
0:vdisk:no:no:5:node1:0:0:030407081836:030407081836:1:00990169
```

lserrlogdumps

You can use the **lserrlogdumps** command to return a list of dumps in the `/dumps/elogs` directory.

Syntax



Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-delim delimiter

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

node_id | node_name

Specifies the node ID or name to list the available dumps of the given type. If you do not specify a node, the dumps available on the configuration node are listed.

Description

This command returns a list of error log dumps. These dumps are created as a result of issuing the **svctask dumperrlog** command. An error log dump describes the contents of the error log at the time that the command was executed. If you do not specify a node, the dumps available on the configuration node will be listed. The command will display files from the `/dumps/elogs` directory.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- There are no error codes associated with the **svcservicemodeinfo lserrlogdumps** command.

Examples

An invocation example

```
svcinfo lserrlogdumps
```

The resulting output

id	filename
0	errlog_myswitch1.3.2_030327_154511
1	aaa.txt_myswitch1.3.2_030327_154527
2	aaa.txt_myswitch1.3.2_030327_154559
3	errlog_myswitch1.3.2_030403_110628

Related topics

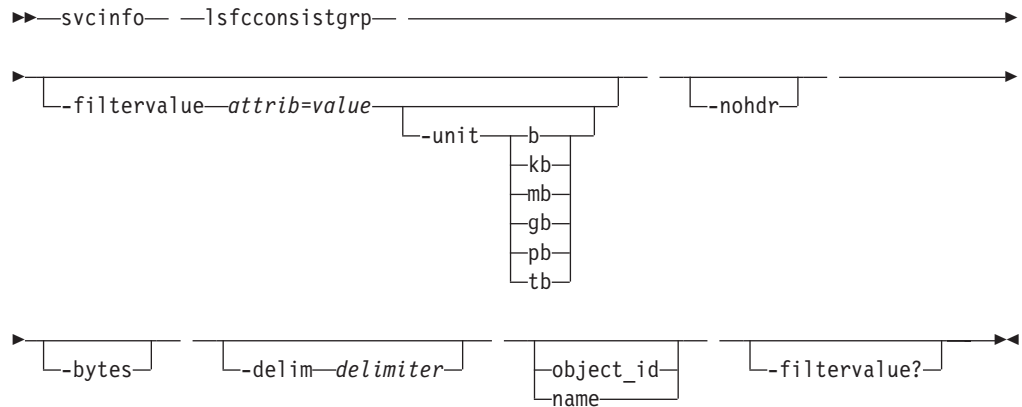
- Chapter 14, “Overview of the list dump commands,” on page 151

lsfcconsistgrp

The list report style can be used to obtain two styles of report.

1. A list containing concise information about all the FlashCopy consistency groups on a cluster. (Each entry in the list corresponds to a single FlashCopy consistency group.)
2. The detailed information about a single FlashCopy consistency group.

Syntax



Parameters

-filtervalue *attribute=value*

Optionally specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-bytes

Optionally used to display all capacities as bytes.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

object_id | name

Optionally specifies the name or ID of an object. If not supplied, the concise

view of all objects, or all of those objects matching the filtering requirements in `-filtervalue` (if specified), of the given type are returned. If supplied, the detailed view of the specific object is returned, and any `filtervalue` entry (if entered) is ignored.

-filtervalue?

Display a list of valid filter attributes. The valid filters for the `svcinfolsfcconsistgrp` command are:

- name
- FC_group_id
- status
- id

Description

This command will return a concise list or a detailed view, of FlashCopy consistency groups visible to the cluster.

The following list provides possible values that are applicable to the attributes that are displayed as data in the output views:

status idle_or_copied, preparing, prepared, copying, stopped, suspended

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

A concise invocation example

```
svcinfolsfcconsistgrp -delim :
```

The concise resulting output

```
id:name:status
1:ffccg0:idle_or_copied
2:ffccg1:idle_or_copied
3:ffccg2:idle_or_copied
```

A detailed invocation example

```
svcinfolsfcconsistgrp -delim : 1
```

The detailed resulting output

```
id:1
name:ffccg0
status:idle_or_copied
```

Related topics

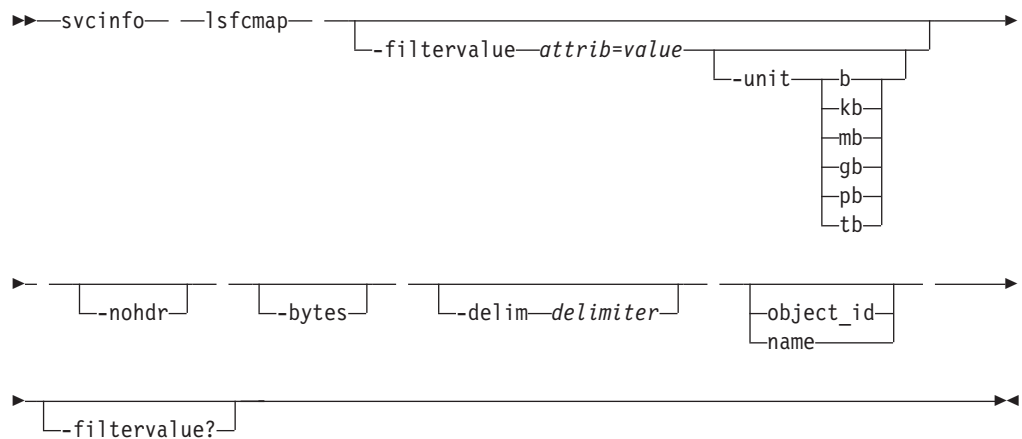
- Chapter 13, “Attributes of the `-filtervalue` argument,” on page 143

lsfcmap

The list report style can be used to obtain two styles of report.

1. A list containing concise information about all FlashCopy mappings visible to the cluster. (Each entry in the list corresponds to a single FlashCopy mapping.)
2. The detailed information about a single FlashCopyMapping.

Syntax



Parameters

-filtervalue *attribute=value*

Optionally specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-bytes

Optionally used to display all capacities as bytes.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

object_id | name

Optionally specifies the name or ID of an object. If not supplied, the concise

view of all objects, or all of those objects matching the filtering requirements in `-filtervalue` (if specified), of the given type are returned. If supplied, the detailed view of the specific object is returned, and any `filtervalue` entry (if entered) is ignored.

-filtervalue?

Display a list of valid filter attributes. The valid filters for the `svcinfolsfcmmap` command are:

- `FC_mapping_name`
- `FC_id`
- `source_vdisk_id`
- `source_vdisk_name`
- `target_vdisk_id`
- `target_vdisk_name`
- `group_name`
- `group_id`
- `status copy_rate`
- `name`
- `id`

Description

This command will return a concise list or a detailed view, of FlashCopy mappings visible to the cluster.

The following list provides possible values that are applicable to the attributes that are displayed as data in the output views:

`status` `idle_or_copied`, `preparing`, `prepared`, `copying`, `stopped`, `suspended`

Possible failures

- `CMMVC5786E` The action failed because the cluster is not in a stable state.
- `CMMVC5804E` The action failed because an entity that was specified in the command does not exist.

Examples

A concise and detailed invocation example

```
svcinfolsfcmmap -delim :
```

The concise and detailed resulting output

```
id:name:source_vdisk_id:source_vdisk_name:target_vdisk_id:
target_vdisk_name:group_id:group_name:status:progress:copy_rate
0:ffcmmap1:0:vdisk0:1:vvdisktwo:::idle_or_copied::75
```

A detailed and detailed invocation example

```
svcinfolsfcmmap -delim : 0
```

The detailed and detailed resulting output

```
id:0
name:ffcmmap1
source_vdisk_id:0
source_vdisk_name:vdisk0
target_vdisk_id:1
```

```
target_vdisk_name:vvdisktwo
group_id:
group_name:
status:idle_or_copied
progress:
copy_rate:75
```

Related topics

- Chapter 13, “Attributes of the -filtervalue argument,” on page 143

lsfcmapcandidate

You can use the **lsfcmapcandidate** command to list all the VDisks that can be source or destinations for FlashCopy, for example, those that are not already in a mapping.

Syntax

```
▶▶ svcinfo — lsfcmapcandidate — [ -nohdr ] [ -delim delimiter ] ▶▶
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command returns a list of VDisks that are not in a FlashCopy mapping. Only the VDisk IDs are returned.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svcinfo lsfcmapcandidate
```

The resulting output

```
id  
2  
3  
4
```

lsfcmapprogress

You can use the **lsfcmapprogress** command to return the progress of the background copy of a FlashCopy mapping.

Syntax

```
▶—svcinfo— —lsfcmapprogress— —nohdr— —delim—delimiter—▶  
▶—fcmap_id—  
  |  
  └—fcmap_name—▶
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

fcmap_id | **fcmap_name**

Specifies the specific object ID or name of the given type.

Description

This command returns the percentage progress of the background copy of a FlashCopy mapping.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.
- CMMVC5805E The progress information was not returned because the FlashCopy statistics are not ready yet.

Examples

An invocation example

```
svcinfo lsfcmapprogress 0
```

The resulting output

```
id      progress
0
```

lsfeaturedumps

You can use the **lsfeaturedumps** command to return a list of dumps in `/dumps/feature`.

Syntax

```
▶▶ —lsfeaturedumps— [ _nohdr_ ] [ _delim—delimiter_ ] ▶▶
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command returns a list of featurization dumps. These dumps are created as a result of issuing the **svctask dumpinternallogs** command. A featurization dump file describes the contents of the featurization log at the time that the command was executed. If you do not specify a node, the dumps available on the configuration node will be listed. The command will display files from the `/dumps/feature` directory.

Issue the **svcinfo lsfeaturedumps** command to return a list of dumps in the `/dumps/feature` destination directory. The feature log is maintained by the cluster. The feature log records events that are generated when license parameters are entered or when the current license settings have been breached.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- There are no error codes associated with the **svcserviceinfo lsfeaturedumps** command.

Examples

An invocation example

```
svcinfo lsfeaturedumps
```

The resulting output

```
id          feature_filename
0           feature.txt
```


Related topics

- Chapter 14, “Overview of the list dump commands,” on page 151

lsfreeextents

You can use the **lsfreeextents** command to list the number of free extents available on a specified MDisk.

Syntax



Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

mdisk_id | **mdisk_name**

Specifies the ID or the name of the MDisk for which you want to know the number of free extents.

Description

This command returns a count of the number of free extents on the specified MDisk.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svcinfo lsfreeextents 2
```

The resulting output

```
id 2  
number_of_extents 4372
```

Ishbaportcandidate

You can use the **Ishbaportcandidate** command to list all of the unconfigured, logged-in host bus adapter (HBA) ports.

Syntax

```
►►—svcinfo— —Ishbaportcandidate— —nohdr— —delim—delimiter—►►
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command returns a list of unconfigured, logged-in HBA ports.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svcinfo Ishbaportcandidate
```

The resulting output

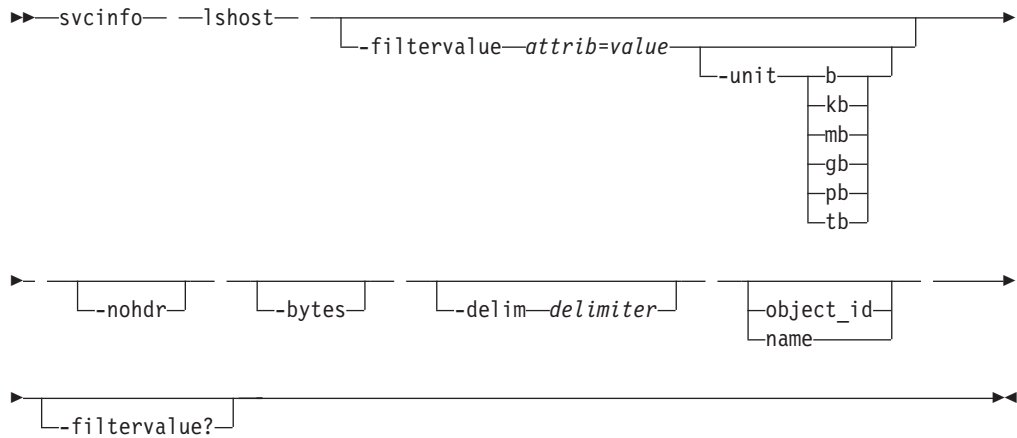
```
id  
210100E08B2520D4
```

lshost

The list report style can be used to obtain two styles of report.

1. A list containing concise information about all the hosts visible to the cluster. (Each entry in the list corresponds to a single host.)
2. The detailed information about a single host.

Syntax



Parameters

-filtervalue *attribute=value*

Optionally specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-bytes

Optionally used to display all capacities as bytes.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

object_id | name

Optionally specifies the name or ID of an object. If not supplied, the concise

view of all objects, or all of those objects matching the filtering requirements in -filtervalue (if specified), of the given type are returned. If supplied, the detailed view of the specific object is returned, and any filtervalue entry (if entered) is ignored.

-filtervalue?

Display a list of valid filter attributes. The valid filters for the **svcinfolshost** command are:

- host_name
- host_id
- port_count
- name
- id

Description

This command will return a concise list or a detailed view, of hosts visible to the cluster.

The following list provides possible values that are applicable to the attributes that are displayed as data in the output views:

status offline, online, degraded, degraded (offline)

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

A concise invocation example

```
svcinfolshost -delim :
```

The concise resulting output

```
id:name:port_count
0:hhost1port:1
1:hhost3ports:3
2:hhost:1
```

A detailed invocation example

```
svcinfolshost -delim : 1
```

The detailed resulting output

```
id:1
name:hhost3ports
port_count:3
WWPN:00000000000000AB
port_logged_in_count:0
WWPN:00000000000000AC
port_logged_in_count:0
WWPN:00000000000000AD
port_logged_in_count:0
```


Related topics

- Chapter 13, “Attributes of the -filtervalue argument,” on page 143

Ishostvdiskmap

You can use the **Ishostvdiskmap** command to obtain a list of the virtual disks that are mapped (visible) to a given host.

Syntax



Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim delimiter

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

host_id | host_name

Optionally specifies the host in terms of its ID or name. The SAN Volume Controller Storage Software for Cisco MDS 9000 returns a list of all the virtual disks that are mapped to the specified host and the SCSI ID by which they are mapped. If neither a host ID or name are entered, the command will return a list of all the host to VDisk mappings that can be seen.

Description

This command returns a list of virtual disk IDs and names. These are the virtual disks that have been mapped to the specified host, that is, they are visible to the specified host. The SCSI LUN ID is also shown. This SCSI LUN ID is the ID by which the virtual disk is being presented to the host.

Determining the VDisk name from the vpath number on the host: Each VDisk exported by the SAN Volume Controller Storage Software for Cisco MDS 9000 is assigned a unique vpath number. This number uniquely identifies the VDisk and can be used to determine which VDisk corresponds to the volume that the hosts sees. This procedure can only be performed using the command line interface.

For the volume in question, find the vpath serial number by issuing the **datapath query device** command. Find the host object defined to the SAN Volume Controller Storage Software for Cisco MDS 9000 that corresponds with the host you are working with.

1. The WWPNs are an attribute of the HBA. You can find these by looking at the device definitions stored by your operating system. For example, on AIX they will be in the ODM, in Windows they will be in the Device Manager details for the given HBA.
2. Verify which host object defined to the SAN Volume Controller Storage Software for Cisco MDS 9000 that these ports belong to. The ports are stored as part of the detailed view, so you will need to list each host in turn by issuing the following:

```
svcinfolshost <name/id>
```

where *<name/id>* is the name or ID of the host. Check for matching WWPNs.

Note: You should name your hosts accordingly, for example, if the actual host is called *orange* you should also name the host object defined to the SAN Volume Controller Storage Software for Cisco MDS 9000 as *orange*.

Now that you have the *<host name>* as defined to the SAN Volume Controller Storage Software for Cisco MDS 9000 and the *<vpath serial number>*, issue the following command:

```
svcinfolshostvdiskmap <hostname>
```

where *<hostname>* is the name of the host. A list is displayed. Look for the VDisk UID that matches the *<vpath serial number>* and remember the VDisk name or ID.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svcinfolshostvdiskmap -delim : 2
```

The resulting output

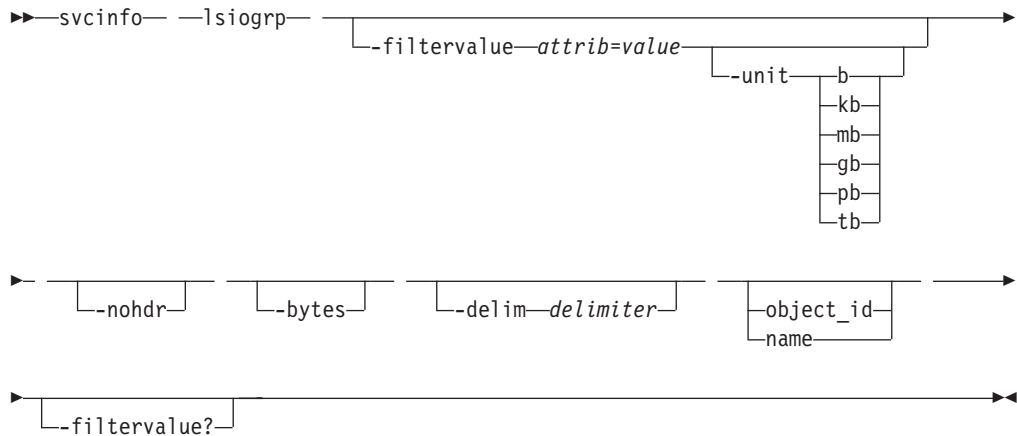
```
id:name:SCSI_id:vdisk_id:vdisk_name:wwpn:vdisk_UID
2:host2:0:10:vdisk10:0000000000000000ACA:6005076801958001500000000000000A
2:host2:1:11:vdisk11:0000000000000000ACA:6005076801958001500000000000000B
2:host2:2:12:vdisk12:0000000000000000ACA:6005076801958001500000000000000C
2:host2:3:13:vdisk13:0000000000000000ACA:6005076801958001500000000000000D
2:host2:4:14:vdisk14:0000000000000000ACA:6005076801958001500000000000000E
```

Isiogrp

The list report style can be used to obtain two styles of report.

1. A list containing concise information about all the IO groups visible to the cluster. (Each entry in the list corresponds to a single I/O group.)
2. The detailed information about a single I/O group.

Syntax



Parameters

-filtervalue *attribute=value*

Optionally specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-bytes

Optionally used to display all capacities as bytes.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

object_id | name

Optionally specifies the name or ID of an object. If not supplied, the concise

view of all objects, or all of those objects matching the filtering requirements in `-filtervalue` (if specified), of the given type are returned. If supplied, the detailed view of the specific object is returned, and any `filtervalue` entry (if entered) is ignored.

-filtervalue?

Display a list of valid filter attributes. The valid filters for the `svcinfolsiogrp` command are:

- `HWS_name`
- `HWS_unique_id`
- `node_count`
- `name`
- `id`

Description

This command will return a concise list or a detailed view, of I/O groups visible to the cluster.

Possible failures

- `CMMVC5786E` The action failed because the cluster is not in a stable state.
- `CMMVC5804E` The action failed because an entity that was specified in the command does not exist.

Examples

A concise and detailed invocation example

```
svcinfolsiogrp -delim :
```

The concise and detailed resulting output

```
id:name:node_count:vdisk_count
0:io_grp0:1:0
1:io_grp1:0:0
2:io_grp2:0:0
3:io_grp3:0:0
4:recovery_io_grp:0:0
```

A detailed and detailed invocation example

```
svcinfolsiogrp -delim : 2
```

The detailed and detailed resulting output

```
id:2
name:io_grp2
node_count:0
vdisk_count:0
```

Related topics

- Chapter 13, “Attributes of the `-filtervalue` argument,” on page 143

lsgiogrpcandidate

You can use the **lsgiogrpcandidate** command to list the I/O groups that can have nodes added to them.

Syntax

```
▶▶ svcinfo — lsgiogrpcandidate — [ -nohdr ] [ -delim delimiter ] ▶▶
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command returns a list of I/O groups to which nodes can be added. Only the I/O group IDs are returned.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svcinfo lsgiogrpcandidate -delim :
```


The resulting output

```
id:  
0:  
1:  
2:  
3:  
4:
```

Isiostatsdumps

You can use the **Isiostatsdumps** command to return a list of dumps in /dumps/iostats directory.

Syntax



Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim delimiter

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

node_id | node_name

Specifies the node ID or name to list the available dumps of the given type. If you do not specify a node, the dumps available on the configuration node are listed.

Description

This command returns a list of I/O statistics dumps. These dumps are created as a result of issuing the **svctask startstats** command. If you do not specify a node, the dumps available on the configuration node will be listed. The command will display files from the /dumps/iostats directory.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcinfo Isiostatsdumps
```

The resulting output

id	iostat_filename
0	v_stats_vegas8.2.1_031123_072426
1	m_stats_vegas8.2.1_031123_072425

Related topics

- Chapter 14, “Overview of the list dump commands,” on page 151

Isiotracedumps

You can use the **Isiotracedumps** command to returns a list of files in the `/dumps/iotrace` directory.

Syntax

```
▶▶ svcinfo — lsiotracedumps — [ -nohdr ] [ -delim delimiter ] ▶▶  
▶▶ [ node_id | node_name ] ▶▶
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

node_id | node_name

Specifies the node ID or name to list the available dumps of the given type. If you do not specify a node, the dumps available on the configuration node are listed.

Description

This command returns a list of I/O trace dumps. These dumps are created as a result of issuing the **svctask settrace** command. If you do not specify a node, the dumps available on the configuration node will be listed. The command will display files from the `/dumps/iotrace` directory.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- There are no error codes associated with the **svcservicemodeinfo lsiotracedumps** command.

Examples

An invocation example

```
svcinfo lsiotracedumps
```

The resulting output

id	iotrace_filename
0	c1_myswitch1.3.2_030405_092155
1	c2_myswitch1.3.2_030405_092156
2	c3_myswitch1.3.2_030405_092158
3	c4_myswitch1.3.2_030405_092159
4	c5_myswitch1.3.2_030405_092201

Related topics

- Chapter 14, “Overview of the list dump commands,” on page 151

Islicense

You can use the **Islicense** command to return the current license (featurization) settings for the cluster.

Syntax

```
►►—svcinfo— —Islicense— —[nohdr]— —[delim—delimiter]— ►►
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command returns the licensed features of the cluster. That is, the copy services status and the capacity of virtual storage licensed for use by this cluster.

You can issue the **svcinfo Islicense** command to return the current license (featurization) settings for the cluster. You can issue the **svctask chlicense** command to change the licensed settings of the cluster. Because the feature settings are entered when the cluster is first created, you need to update the settings only if you have changed your license. You can change the following values:

- FlashCopy: disabled or enabled
- Remote Copy: disabled or enabled
- Virtualization limit: number, in gigabytes (1073741824 bytes)

The output displayed lists the feature functions in a list and displays whether they are enabled or disabled.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svcinfo lslicense
```

The resulting output


```
feature_flash on  
feature_remote on  
feature_num_gb 32
```

lsmdisk

The list report style can be used to obtain two styles of report.

1. A list containing concise information about all the managed disks visible to the cluster. (Each entry in the list corresponds to a single managed disk.)
2. The detailed information about a single managed disk.

Syntax



Parameters

-filtervalue *attribute=value*

Optionally specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-bytes

Optionally used to display all capacities as bytes.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

object_id | name

Optionally specifies the name or ID of an object. If not supplied, the concise

view of all objects, or all of those objects matching the filtering requirements in -filtervalue (if specified), of the given type are returned. If supplied, the detailed view of the specific object is returned, and any filtervalue entry (if entered) is ignored.

-filtervalue?

Display a list of valid filter attributes. The valid filters for the **svcinfo lsmdisk** command are:

- name
- id
- status
- mode
- mdisk_grp_id
- mdisk_grp_name
- capacity
- controller_name

Description

This command will return a concise list or a detailed view, of MDisks visible to the cluster.

The following list provides possible values that are applicable to the attributes that are displayed as data in the output views:

status	offline, excluded, degraded, online
mode	unmanaged, managed, image
quorum index	0/1/2 valid Quorum index

When back-end controllers are added to the fibre-channel SAN and are included in the same switch zone as a SAN Volume Controller Storage Software for Cisco MDS 9000 Cluster the cluster will automatically discover the back-end controller and will integrate the controller to determine what storage it is presented to the SAN Volume Controller Storage Software for Cisco MDS 9000. The SCSI LUs presented by the back-end controller will be displayed as unmanaged MDisks. If however the configuration of the back-end controller is modified after this has occurred then the SAN Volume Controller Storage Software for Cisco MDS 9000 may be unaware of these configuration changes. This task allows a user to request the SAN Volume Controller Storage Software for Cisco MDS 9000 to re-scan the fibre-channel SAN to update the list of unmanaged Mdisks.

Note: The automatic discovery performed by SAN Volume Controller Storage Software for Cisco MDS 9000 does not write anything to a unmanaged MDisk. It is only when a the user instructs the SAN Volume Controller Storage Software for Cisco MDS 9000 to add a MDisk to a managed disk group or use a Mdisk to create an image mode virtual disk that the storage will actually be used.

Discovering MDisks: Check to see which MDisks are available by issuing the **svctask detectmdisk** command to manually scan the fibre-channel network for any MDisks. Issue the **svcinfo lsmdiskcandidate** command to show the unmanaged MDisks. These MDisks have not been assigned to an MDisk group. Alternatively, you can issue the **svcinfo lsmdisk** command to view all of the MDisks.

lsmdiskcandidate

You can use the **lsmdiskcandidate** command to list all of the unmanaged MDisks.

Syntax

```
➤➤—svcinfo—lsmdiskcandidate—┌-nohdr┐┌-delim—delimiter┐➤➤
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command returns a list of MDisks that are unmanaged. Only the MDisk IDs are returned.

When back-end controllers are added to the fibre-channel SAN and are included in the same switch zone as a SAN Volume Controller Storage Software for Cisco MDS 9000 Cluster the cluster will automatically discover the back-end controller and will integrate the controller to determine what storage it is presented to the SAN Volume Controller Storage Software for Cisco MDS 9000. The SCSI LUs presented by the back-end controller will be displayed as unmanaged MDisks. If however the configuration of the back-end controller is modified after this has occurred then the SAN Volume Controller Storage Software for Cisco MDS 9000 may be unaware of these configuration changes. This task allows a user to request the SAN Volume Controller Storage Software for Cisco MDS 9000 to re-scan the fibre-channel SAN to update the list of unmanaged Mdisks.

Note: The automatic discovery performed by SAN Volume Controller Storage Software for Cisco MDS 9000 does not write anything to a unmanaged MDisk. It is only when a the user instructs the SAN Volume Controller Storage Software for Cisco MDS 9000 to add a MDisk to a managed disk group or use a Mdisk to create an image mode virtual disk that the storage will actually be used.

Discovering MDisks: Check to see which MDisks are available by issuing the `svctask detectmdisk` command to manually scan the fibre-channel network for any MDisks. Issue the `svcinfolismdiskcandidate` command to show the unmanaged MDisks. These MDisks have not been assigned to an MDisk group. Alternatively, you can issue the `svcinfolismdisk` command to view all of the MDisks.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svcinfolismdiskcandidate
```


The resulting output

```
id  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14
```

lsmdiskextent

You can use the **lsmdiskextent** command to return the extent allocation between managed disks and virtual disks.

Syntax



Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

mdisk_name | **mdisk_id**

Specifies the specific object ID or name of the given type.

Description

The command returns a list, in which each entry contains a VDisk ID and the number of extents. These VDIsks are using extents on the specified MDisk. The number of extents being used on each MDisk is also shown.

Every VDisk is constructed from one or more mdisks. At times you may need to determine the relationship between the two objects. The following procedure allows you to determine the relationships.

Determining the relationship between VDIsks and MDIsks: For a given VDisk `<vdiskname/id>`, issue the following command:

```
svcinfo lsvdiskmember <vdiskname/id>
```

where `<vdiskname/id>` is the name or ID of the VDisk. This will return a list of IDs that correspond to the MDIsks that make up the VDisk.

Determining the relationship between VDIs and MDIs and the number of extents provided by each MDI: If you wish more details, you can also determine the number of extents that are being provided by each MDI. This procedure can only be performed using the command line interface. For a given VDI `<vdiskname/id>`, issue the following command:

```
svcinfo lsvdiskextent <vdiskname/id>
```

where `<vdiskname/id>` is the name or ID of the VDI. This will return a table of MDI IDs and the corresponding number of extents each MDI is providing as storage for the given VDI.

Determining the relationship between MDIs and VDIs: For a given MDI `<mdiskname/id>`, issue the following command:

```
svcinfo lsmdiskmember <mdiskname/id>
```

where `<mdiskname/id>` is the name or ID of the MDI. This will return a list of IDs that correspond to the VDIs that are using this MDI.

Determining the relationship between MDIs and VDIs and the number of extents used by each VDI: If you wish more details, you can also determine the number of extents that this MDI is providing for each VDI. This procedure can only be performed using the command line interface. For a given MDI `<mdiskname/id>`, issue the following command:

```
svcinfo lsmdiskextent <mdiskname/id>
```

where `<mdiskname/id>` is the name or ID of the MDI. This returns a table of VDI IDs and the corresponding number of extents being used by each VDI.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.
- CMMVC5854E The extent information was not returned because the extent is not used or does not exist.
- CMMVC5855E The extent information was not returned because the managed disk (MDI) is not used by any virtual disk (VDI).
- CMMVC5864E The extent information was not returned because the source extent is not used.
- CMMVC5865E The extent information was not returned because the extent is out of range for the managed disk (MDI) or virtual disk (VDI).
- CMMVC6005E The view request failed as the specified object is not a member of an appropriate group.

Examples

An invocation example

```
svcinfo lsmdiskextent 2
```

The resulting output

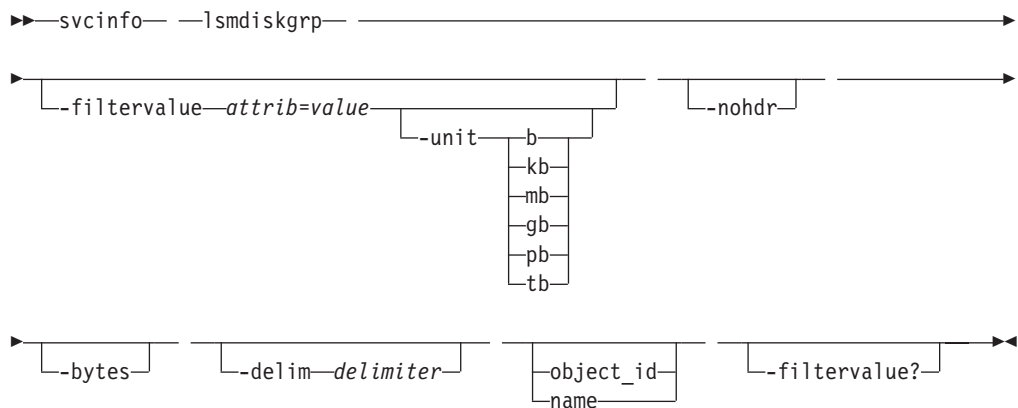
id	number_of_extents
1	1
2	1

lsmdiskgrp

The list report style can be used to obtain two styles of report.

1. A list containing concise information about all the managed disk groups in a cluster. (Each entry in the list corresponds to a single managed disk group.)
2. The detailed information about a single managed disk group.

Syntax



Parameters

-filtervalue *attribute=value*

Optionally specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-bytes

Optionally used to display all capacities as bytes.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

object_id | name

Optionally specifies the name or ID of an object. If not supplied, the concise view of all objects, or all of those objects matching the filtering requirements in

-filtervalue (if specified), of the given type are returned. If supplied, the detailed view of the specific object is returned, and any filtervalue entry (if entered) is ignored.

-filtervalue?

Display a list of valid filter attributes. The valid filters for the **svcinfo lsmdiskgrp** command are:

- name
- storage_pool_id
- mdisk_count
- vdisk_count
- extent_size
- status
- id

Description

This command will return a concise list or a detailed view, of MDisk groups visible to the cluster.

The following list provides possible values that are applicable to the attributes that are displayed as data in the output views:

status online, degraded, offline

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

A concise invocation example

```
svcinfo lsmdiskgrp -delim :
```

The concise resulting output

```
id:name:status:mdisk_count:vdisk_count:capacity:extent_size:free_capacity  
0:mdiskgrp0:online:5:0:341.8GB:16:341.8GB  
1:mdiskgrp1:online:0:0:0:16:0
```

A detailed invocation example

```
svcinfo lsmdiskgrp -delim : 0
```

The detailed resulting output

```
id:0  
name:mdiskgrp0  
status:online  
mdisk_count:5  
vdisk_count:0  
capacity:341.8GB  
extent_size:16  
free_capacity:341.8GB
```


Related topics

- Chapter 13, “Attributes of the -filtervalue argument,” on page 143

lsmdiskmember

You can use the **lsmdiskmember** command to return a list of VDIsks that are using extents on the specified MDisk.

Syntax



Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

mdisk_id | mdisk_name

Specifies the ID or name of the MDisk for which the user wants a list of VDIsks using extents of that MDisk.

Description

This command returns a list of virtual disks that use extents on the specified managed disk. These are the virtual disks that are using extents on the managed disk specified by the ID. The list returned is the members of the respective object and is independent of the state of the individual members, that is, if they are in offline state, they are still returned.

Every VDisk is constructed from one or more mdisks. At times you may need to determine the relationship between the two objects. The following procedure allows you to determine the relationships.

Determining the relationship between VDIsks and MDIsks: For a given VDisk <vdiskname/id>, issue the following command:

```
svcinfo lsvdiskmember <vdiskname/id>
```

where *<vdiskname/id>* is the name or ID of the VDisk. This will return a list of IDs that correspond to the MDisks that make up the VDisk.

Determining the relationship between VDIsks and MDisks and the number of extents provided by each MDisk: If you wish more details, you can also determine the number of extents that make are being provided by each MDisk. This procedure can only be performed using the command line interface. For a given VDisk *<vdiskname/id>*, issue the following command:

```
svcinfo lsvdiskextent <vdiskname/id>
```

where *<vdiskname/id>* is the name or ID of the VDisk. This will return a table of MDisk IDs and the corresponding number of extents each MDisk is providing as storage for the given VDisk.

Determining the relationship between MDisks and VDIsks: For a given MDisk *<mdiskname/id>*, issue the following command:

```
svcinfo lsmdiskmember <mdiskname/id>
```

where *<mdiskname/id>* is the name or ID of the MDisk. This will return a list of IDs that correspond to the VDIsks that are using this MDisk.

Determining the relationship between MDisks and VDIsks and the number of extents used by each VDisk: If you wish more details, you can also determine the number of extents that this MDisk is providing for each VDisk. This procedure can only be performed using the command line interface. For a given MDisk *<mdiskname/id>*, issue the following command:

```
svcinfo lsmdiskextent <mdiskname/id>
```

where *<mdiskname/id>* is the name or ID of the MDisk. This returns a table of VDisk IDs and the corresponding number of extents being used by each VDisk.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svcinfo lsmdiskmember 1
```

The resulting output

```
id  
0
```

lsmigrate

You can use the **lsmigrate** command to show the progress of all the data migration operations currently in progress.

Syntax

```
►► svcinfo — lsmigrate — [ -nohdr ] [ -delim delimiter ] ►►
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command displays information of all migrations currently in progress.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svcinfo lsmigrate -delim :
```

The resulting output

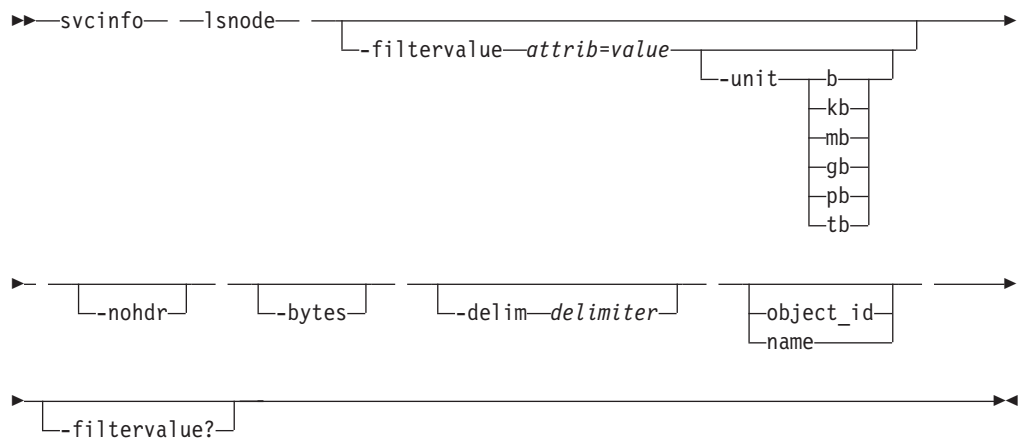
```
migrate_type:progress:migrate_source_vdisk_index:  
migrate_target_mdisk_grp:max_thread_count  
3:53:0:1:2
```

lsnode

The list report style can be used to obtain two styles of report.

1. A list containing concise information about all the nodes on a cluster. (Each entry in the list corresponds to a single node.)
2. The detailed information about a single node.

Syntax



Parameters

-filtervalue *attribute=value*

Optionally specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-bytes

Optionally used to display all capacities as bytes.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

object_id | name

Optionally specifies the name or ID of an object. If not supplied, the concise

view of all objects, or all of those objects matching the filtering requirements in -filtervalue (if specified), of the given type are returned. If supplied, the detailed view of the specific object is returned, and any filtervalue entry (if entered) is ignored.

-filtervalue?

Display a list of valid filter attributes. The valid filters for the **svcinfolnode** command are:

- node_name
- id
- status
- IO_group_name
- IO_group_id
- name

Description

This command will return a concise list or a detailed view, of nodes visible to the cluster.

The following list provides possible values that are applicable to the attributes that are displayed as data in the output views:

status	offline, pending, online, adding, deleting
config node	no, yes
port status	not present, online, offline

Attention: `svcinfolnode <nodename/id>` will report only one WWPN in SAN Volume Controller Storage Software for Cisco MDS 9000. To see all the WWPNs associated with a node, Cisco switch CLI commands will have to be used. See the *Cisco MDS 9000 Family Command Reference*. If the node is in adding state, the WWPN will be displayed as 0000000000000000. Once the node has successfully become a member of the cluster, the state will change to online and the WWPN will be displayed correctly.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

A concise invocation example

```
svcinfolnode -delim :
```

The concise resulting output

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

A detailed invocation example

```
svcinfolnode -delim = 1
```

The detailed resulting output

```
id=1
name=node1
card_serial_number=JAB071007HB
WWNN=2548000530003920
status=online
IO_group_id=0
IO_group_name=io_grp0
partner_node_id=2
partner_node_name=node2
config_node=yes
card_unique_id=0100071030374842
port_id=1E39003005000220
port_status=active
```

Related topics

- Chapter 13, “Attributes of the -filtervalue argument,” on page 143

lsnodecandidate

You can use the **lsnodecandidate** command to list all of the nodes that are not assigned to a cluster.

Syntax

```
▶▶—svcinfol—lsnodecandidate—┬──nohdr┬──┬──delim—delimiter──▶▶
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command returns a list of nodes that are not assigned to a cluster.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svcinfol lsnodecandidate -delim :
```


The resulting output

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

lsnodevpd

You can use the **lsnodevpd** command to return the vital product data (VPD) for the given node.

Syntax



Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim delimiter

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

node_id | node_name

Specifies the node to view in terms of its node ID or name.

Description

This command returns the VPD for the specified node. Each field is reported on a new line. All fields are strings. The VPD is split into sections. Each section has a section heading. Following the heading is the number of fields in that section. Each section is separated by an empty line. For example:

```
section name:3 fields
field1:value
field2:value
field3:value
```

```
new section:x fields
...
```

Some sections contain information about multiple objects of that type. Each object within the section is separated by an empty line. For example:


```
section name:4 fields
object1 field1:value
object1 field2:value

object2 field1:value
object2 field2:value

new section: x fields
...
```

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svcinfo lsnodevpd -delim : 1
```

The resulting output

```
id:1

system board: 4 fields
Card serial number:JAB071105kv
number_of_processors:2
number_of_scsi/ide_devices:1
Switch IP address:9.43.224.172

processor: 8 fields
processor_location:Processor 1
manufacturer:BROADCOM
speed:750
status:Enabled
:
processor_location:Processor 2
manufacturer:BROADCOM
speed:750
status:Enabled

memory module: 1 fields
size (MB):4096

device: 6 fields
bus:ide0
device:0
model:IC25N020ATCS05-0
revision:CS20A61A
serial_number: CLP225F2G2RLAA
approx_capacity:18

software: 4 fields
code_level:1.3(1) (IBM build 0.3.3)
node_name:node1
WWNN:0x225d000b4679ec82
id:0x1

front panel assembly: 1 fields
front_panel_id:vegas3.9.2
```


```
UPS: 2 fields  
Card Serial number:JAB071105kv  
Card unique id:100071130356b76
```

lsrcconsistgrp

The list report style can be used to obtain two styles of report.

1. A list containing concise information about all the Remote Copy consistency groups visible to the cluster. (Each entry in the list corresponds to a single Remote Copy consistency group.)
2. The detailed information about a single Remote Copy consistency group.

Syntax



Parameters

-filtervalue *attribute=value*

Optionally specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-bytes

Optionally used to display all capacities as bytes.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

object_id | name

Optionally specifies the name or ID of an object. If not supplied, the concise view of all objects, or all of those objects matching the filtering requirements in `-filtervalue` (if specified), of the given type are returned. If supplied, the detailed view of the specific object is returned, and any `filtervalue` entry (if entered) is ignored.

-filtervalue?

Display a list of valid filter attributes. The valid filters for the **svcinfolsrcconsistgrp** command are:

- group_id
- name
- master_cluster_id
- master_cluster_name
- aux_cluster_id
- aux_cluster_name
- primary
- state
- relationship_count
- id

Description

This command will return a concise list or a detailed view, of Remote Copy consistency groups visible to the cluster.

The following list provides possible values that are applicable to the attributes that are displayed as data in the output views:

primary	n/a, master, aux
state	inconsistent_stopped, inconsistent_copying, consistent_stopped, consistent_synchronized, idling, idling_disconnected, inconsistent_disconnected, consistent_disconnected, empty
freeze_time	The time in YY/MM/DD/HH/MM format.
status	online, primary_offline, secondary_offline
sync	in_sync, out_of_sync

Note: The names of the Remote Copy relationships and consistency groups may be blank if the relationship or consistency groups are inter-cluster and the cluster partnership is disconnected.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

A concise invocation example

```
svcinfolsrcconsistgrp -delim :
```

The concise resulting output

```
id:name:master_cluster_id:master_cluster_name:  
aux_cluster_id:aux_cluster_name:primary:state:relationship_count  
253:bbcc:0000010030A007E5:kkk:0000010030A007E5:  
kkk::empty:0  
254:rcstgrp0:0000010030A007E5:kkk:0000010030A007E5:  
kkk:master:inconsistent_stopped:1  
255:bbb:0000010030A007E5:kkk:  
0000010030A007E5:kkk::empty:0
```

A detailed invocation example

```
svcinfolsrconsgtgrp -delim : 254
```

The detailed resulting output

```
id:254
name:rccstgrp0
master_cluster_id:0000010030A007E5
master_cluster_name:kkk
aux_cluster_id:0000010030A007E5
aux_cluster_name:kkk
primary:master
state:inconsistent_stopped
relationship_count:1
freeze_time:
status:online
sync:
RC_rel_id:2
RC_rel_name:aaa
```

Related topics

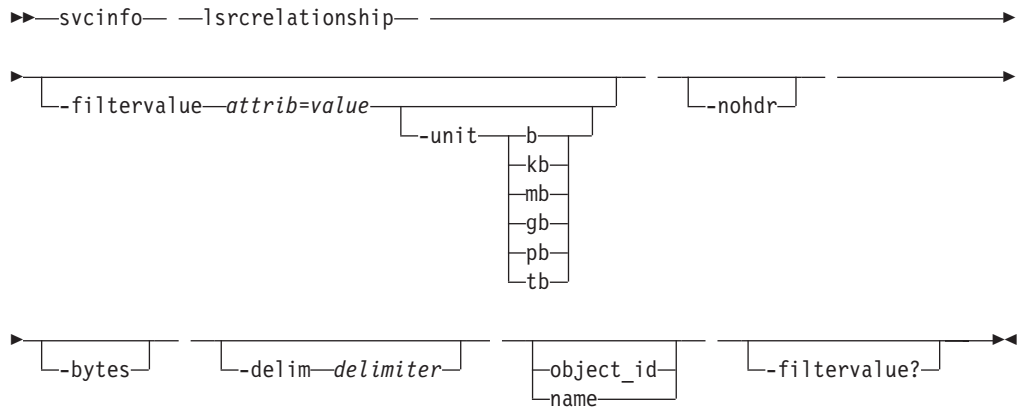
- Chapter 13, “Attributes of the -filtervalue argument,” on page 143

lsrcrelationship

The list report style can be used to obtain two styles of report.

1. A list containing concise information about all the Remote Copy relationships visible to the cluster. (Each entry in the list corresponds to a single Remote Copy relationship.)
2. The detailed information about a single Remote Copy relationship.

Syntax



Parameters

-filtervalue *attribute=value*

Optionally specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-bytes

Optionally used to display all capacities as bytes.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

object_id | name

Optionally specifies the name or ID of an object. If not supplied, the concise

view of all objects, or all of those objects matching the filtering requirements in `-filtervalue` (if specified), of the given type are returned. If supplied, the detailed view of the specific object is returned, and any `filtervalue` entry (if entered) is ignored.

-filtervalue?

Display a list of valid filter attributes. The valid filters for the `svcinfolsrcrelationship` command are:

- RC_rel_id
- RC_rel_name
- master_cluster_id
- master_cluster_name
- master_vdisk_id
- master_vdisk_name
- aux_cluster_id
- aux_cluster_name
- aux_vdisk_id
- aux_vdisk_name
- primary
- consistency_group_id
- consistency_group_name
- state
- progress

Description

This command will return a concise list or a detailed view, of Remote Copy relationships visible to the cluster.

The following list provides possible values that are applicable to the attributes that are displayed as data in the output views:

primary	n/a, master, aux
state	inconsistent_stopped, inconsistent_copying, consistent_stopped, consistent_synchronized, idling, idling_disconnected, inconsistent_disconnected, consistent_disconnected
progress	0-100, or n/a
freeze time	The time in YY/MM/DD/HH/MM format.
status	online, primary_offline, secondary_offline
sync	in_sync, out_of_sync

Note: The names of the Remote Copy relationships and consistency groups may be blank if the relationship or consistency groups are inter-cluster and the cluster partnership is disconnected.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

A concise and detailed invocation example

```
svcinfolsrrelationship -delim :
```

The concise and detailed resulting output

```
id:name:master_cluster_id:master_cluster_name:
master_vdisk_id:master_vdisk_name:aux_cluster_id:
aux_cluster_name:aux_vdisk_id:aux_vdisk_name:primary:
consistency_group_id:consistency_group_name:state:
bg_copy_priority:progress
2:aaa:0000010030A007E5:kkk:2:vdisk2:
0000010030A007E5:kkk:3:vdisk3:master:::
inconsistent_stopped:50:0
```

A detailed invocation example

```
svcinfolsrrelationship -delim : 2
```

The detailed resulting output

```
id:2
name:aaa
master_cluster_id:0000010030A007E5
master_cluster_name:kkk
master_vdisk_id:2
master_vdisk_name:vdisk2
aux_cluster_id:0000010030A007E5
aux_cluster_name:kkk
aux_vdisk_id:3
aux_vdisk_name:vdisk3
primary:master
consistency_group_id:
consistency_group_name:
state:inconsistent_stopped
bg_copy_priority:50
progress:0
freeze_time:
status:online
sync:
```


Related topics

- Chapter 13, “Attributes of the -filtervalue argument,” on page 143

lsrrelationshipcandidate

You can use the **lsrrelationshipcandidate** to list VDisks that are eligible to form Remote Copy relationships. You can list suitable VDisks on the local or remote cluster.

Syntax



Parameters

-master *master_vdisk_id* | *master_vdisk_name*

You can use this parameter to specify a particular vdisk that you want to use as the master vdisk. The command will then look for candidates that match the size of this vdisk. If you are requesting candidate vdisks on the local cluster, this command will also match the `io_group`.

-aux *aux_cluster_id* | *aux_cluster_name*

Specifies a remote cluster to find vdisk candidates for an inter-cluster relationship. If you do not specify this parameter, the candidates on the local cluster are returned.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command returns a list of VDisks that can be the master or auxiliary disk for a Remote Copy relationship. VDisk IDs and names are returned.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svcinfo lsrelationshipcandidate -delim :
```

The resulting output

```
id:vdisk_name  
0:vdisk0  
4:vdisk4
```

lsrcrelationshipprogress

You can use the **lsrcrelationshipprogress** command to return the progress of the background copy of a Remote Copy relationship. When the initial background copy process for a relationship has completed, null will be displayed for the progress of that relationship.

Syntax

```
svcinfo --lsrcrelationshipprogress [-nohdr]
                                     [rcrelationship_id]
                                     [-delim delimiter] [rcrelationship_name]
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

rcrelationship_id | **rcrelationship_name**

Specifies the specific object ID or name of the given type.

Description

This command returns the percentage progress of the background copy of a Remote Copy relationship.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svcinfo lsrcrelationshipprogress -delim : 0
```

The resulting output

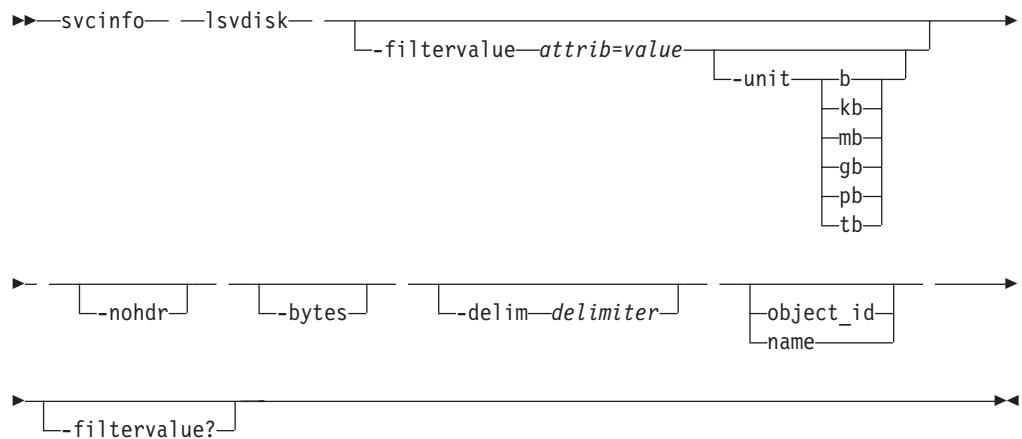
```
id:progress  
0:58
```

lsvdisk

The list report style can be used to obtain two styles of report.

1. A list containing concise information about all the virtual disks visible to the cluster. (Each entry in the list corresponds to a single virtual disk.)
2. The detailed information about a single virtual disk.

Syntax



Parameters

-filtervalue *attribute=value*

Optionally specifies a list of one or more filters. Only objects with a value that matches the filter attribute value are returned. If a capacity is specified, the units must also be included.

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-bytes

Optionally used to display all capacities as bytes.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

object_id | name

Optionally specifies the name or ID of an object. If not supplied, the concise

view of all objects, or all of those objects matching the filtering requirements in -filtervalue (if specified), of the given type are returned. If supplied, the detailed view of the specific object is returned, and any filtervalue entry (if entered) is ignored.

-filtervalue?

Display a list of valid filter attributes. The valid filters for the **svcinfolsvdisk** command are:

- vdisk_name
- vdisk_id
- IO_group_id
- IO_group_name
- status
- mdisk_grp_name
- mdisk_grp_id
- capacity
- type
- FC_id
- FC_name
- RC_id
- RC_name
- name
- id

Description

This command will return a concise list or a detailed view, of VDisks visible to the cluster.

The following list provides possible values that are applicable to the attributes that are displayed as data in the output views:

status	offline, online
capacity	Rounded to GB (2 decimal places) or MB if less than 1GB.
type	sequential, striped, image
formatted	no, yes
mdisk id	Not used for striped
mdisk name	Not used for striped

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

A concise invocation example

```
svcinfolsvdisk -delim :
```

The concise resulting output

```
id:name:IO_group_id:IO_group_name:status:mdisk_grp_id:mdisk_grp_name
:capacity:type:FC_id:FC_name:RC_id:RC_name
0:vdisk0:0:io_grp0:degraded:0:mdiskgrp0:16.0MB:seq:::
1:vvdisktwo:0:io_grp0:degraded:0:mdiskgrp0:16.0MB:seq:::
```

A detailed invocation example

```
svcinfolsvdisk -delim : 1
```

The detailed resulting output

```
id:1
name:vvdisktwo
IO_group_id:0
IO_group_name:io_grp0
status:degraded
mdisk_grp_id:0
mdisk_grp_name:mdiskgrp0
capacity:16.0MB
type:seq
formatted:no
mdisk_id:2
mdisk_name:mdisk2
FC_id:0
FC_name:ffcmap1
RC_id:
RC_name:
throttling:0
preferred_node_id:1
```

Related topics

- Chapter 13, “Attributes of the -filtervalue argument,” on page 143

lsvdiskextent

You can use the **lsvdiskextent** command to list the number of extents per MDisk making up the VDisk. The number of extents being provided by each MDisk is displayed.

Syntax

```
▶▶ svcinfo — lsvdiskextent — [ -nohdr ] [ -delim delimiter ] ▶▶
▶ [ vdisk_name | vdisk_id ] ▶▶
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

vdisk_name | vdisk_id

Specifies one or more virtual disk IDs or names.

Description

The command returns a list, in which each entry contains a MDisk ID and the number of extents. These MDisks are using extents on the specified VDisk. The number of extents being used on each VDisk is also shown.

Every VDisk is constructed from one or more mdisks. At times you may need to determine the relationship between the two objects. The following procedure allows you to determine the relationships.

Determining the relationship between VDIsks and MDisks: For a given VDisk `<vdiskname/id>`, issue the following command:

```
svcinfo lsvdiskmember <vdiskname/id>
```

where `<vdiskname/id>` is the name or ID of the VDisk. This will return a list of IDs that correspond to the MDisks that make up the VDisk.

Determining the relationship between VDIs and MDIs and the number of extents provided by each MDI: If you wish more details, you can also determine the number of extents that are being provided by each MDI. This procedure can only be performed using the command line interface. For a given VDI <vdiskname/id>, issue the following command:

```
svcinfo lsvdiskextent <vdiskname/id>
```

where <vdiskname/id> is the name or ID of the VDI. This will return a table of MDI IDs and the corresponding number of extents each MDI is providing as storage for the given VDI.

Determining the relationship between MDIs and VDIs: For a given MDI <mdiskname/id>, issue the following command:

```
svcinfo lsmdiskmember <mdiskname/id>
```

where <mdiskname/id> is the name or ID of the MDI. This will return a list of IDs that correspond to the VDIs that are using this MDI.

Determining the relationship between MDIs and VDIs and the number of extents used by each VDI: If you wish more details, you can also determine the number of extents that this MDI is providing for each VDI. This procedure can only be performed using the command line interface. For a given MDI <mdiskname/id>, issue the following command:

```
svcinfo lsmdiskextent <mdiskname/id>
```

where <mdiskname/id> is the name or ID of the MDI. This returns a table of VDI IDs and the corresponding number of extents being used by each VDI.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.
- CMMVC5854E The extent information was not returned because the extent is not used or does not exist.
- CMMVC5855E The extent information was not returned because the managed disk (MDI) is not used by any virtual disk (VDI).
- CMMVC5864E The extent information was not returned because the source extent is not used.
- CMMVC5865E The extent information was not returned because the extent is out of range for the managed disk (MDI) or virtual disk (VDI).

Examples

An invocation example

```
svcinfo lsvdiskextent -delim : vdisk0
```


The resulting output

```
id:extent offset  
0:0
```

lsvdiskhostmap

You can use the **lsvdiskhostmap** command to list the VDisk to host mapping.

Syntax



Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim delimiter

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

vdisk_id | vdisk_name

Specifies the virtual disk in terms of its ID or name. The SAN Volume Controller Storage Software for Cisco MDS 9000 returns a list of all the hosts to which this virtual disk is mapped and the SCSI ID by which the virtual disk is mapped.

Description

This command returns a list of host IDs and names. These hosts have the specified virtual disk mapped to them, that is, the virtual disk is visible to these hosts. The SCSI LUN ID is also shown. The SCSI LUN ID is the ID by which the virtual disk is being presented to the host.

Determining the host that a VDisk is mapped to: List the hosts that this VDisk is mapped, by issuing the following command:

```
svcinfo lsvdiskhostmap <vdiskname/id>
```

where *<vdiskname/id>* is the name or ID of the VDisk. A list is displayed. Look for the host name or ID to determine which host this VDisk is mapped to. If no data is returned, the VDisk is not mapped to any hosts.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

The concise invocation example

```
svcinfolsvdiskhostmap -delim : 3
```

The concise resulting output

```
id:name:SCSI_id:host_id:host_name:wwpn:vdisk_UID  
3:vdisk3:0:2:host2:0000000000100ABC:60050768018A00015000000000000003
```

lsvdiskmember

You can use the **lsvdiskmember** command to return a list of MDisks that are members of the specified VDisk.

Syntax

```
▶▶ svcinfo — lsvdiskmember — [ -nohdr ] [ -delim delimiter ] —▶▶
[ vdisk_id | vdisk_name ] —▶▶
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

vdisk_id | vdisk_name

Specifies the name or ID of the VDisk for which the user requires the list of member MDisks.

Description

This command returns a list of managed disks, which provide extents that make up the virtual disk specified by the ID.

Every VDisk is constructed from one or more mdisks. At times you may need to determine the relationship between the two objects. The following procedure allows you to determine the relationships.

If the command used is the **svcinfo lsmdiskmember** command, the concise view will return a list of virtual disks. These are the virtual disks that are using extents on the managed disk specified by the ID. The list returned is the members of the respective object and is independent of the state of the individual members, that is, if they are in offline state, they are still returned.

Determining the relationship between VDisks and MDisks: For a given VDisk <vdiskname/id>, issue the following command:

```
svcinfo lsvdiskmember <vdiskname/id>
```

where *<vdiskname/id>* is the name or ID of the VDisk. This will return a list of IDs that correspond to the MDisks that make up the VDisk.

Determining the relationship between VDIsks and MDisks and the number of extents provided by each MDisk: If you wish more details, you can also determine the number of extents that make are being provided by each MDisk. This procedure can only be performed using the command line interface. For a given VDisk *<vdiskname/id>*, issue the following command:

```
svcinfo lsvdiskextent <vdiskname/id>
```

where *<vdiskname/id>* is the name or ID of the VDisk. This will return a table of MDisk IDs and the corresponding number of extents each MDisk is providing as storage for the given VDisk.

Determining the relationship between MDisks and VDIsks: For a given MDisk *<mdiskname/id>*, issue the following command:

```
svcinfo lsmdiskmember <mdiskname/id>
```

where *<mdiskname/id>* is the name or ID of the MDisk. This will return a list of IDs that correspond to the VDIsks that are using this MDisk.

Determining the relationship between MDisks and VDIsks and the number of extents used by each VDisk: If you wish more details, you can also determine the number of extents that this MDisk is providing for each VDisk. This procedure can only be performed using the command line interface. For a given MDisk *<mdiskname/id>*, issue the following command:

```
svcinfo lsmdiskextent <mdiskname/id>
```

where *<mdiskname/id>* is the name or ID of the MDisk. This returns a table of VDisk IDs and the corresponding number of extents being used by each VDisk.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svcinfo lsvdiskmember 1
```

The resulting output

```
id  
2
```

lsvdiskprogress

You can use the **lsvdiskprogress** command to return the progress of the format of a new virtual disk.

Syntax

```
▶▶ svcinfo — lsvdiskprogress — [ -nohdr ] [ -delim delimiter ] ▶▶
▶ [ vdisk_id | vdisk_name ] ▶▶
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

vdisk_id | **vdisk_name**

Specifies the specific object ID or name of the given type.

Description

This command returns the progress of the format of a new virtual disk as a percentage completed.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.
- CMMVC5805E The progress information was not returned because the FlashCopy statistics are not ready yet.

Examples

An invocation example

```
svcinfo lsvdiskprogress -delim : 0
```

The resulting output

id:progress
0:58

Chapter 16. Error log commands

The following commands enable you to work with error logs with the SAN Volume Controller Storage Software for Cisco MDS 9000.

finderr

You can use the **finderr** command to analyze the error log for the highest severity unfixed error.

Syntax

▶—svctask— —finderr—▶

Description

The command scans the error log for any unfixed errors. Given a priority ordering defined within the code, the highest priority unfixed error is returned to stdout.

You can use this command to determine the order in which to fix the logged errors.

The Web-based directed maintenance procedures (DMPs) also use this command.

Possible failures

- There are no error codes.

Examples

An invocation example

```
svctask finderr
```

The resulting output

```
Highest priority unfixed error code is [1010]
```

dumperrlog

You can use the **dumperrlog** command to dump the contents of the error log to a text file.

Syntax

```
svctask dumperrlog [-prefix filename_prefix]
```

Parameters

-prefix *filename_prefix*

A file name is created from the prefix and a time stamp, and has the following format:

```
<prefix>_NNNNNN_YYMMDD_HHMMSS
```

where *NNNNNN* is the node front panel name.

Note: If the **-prefix** parameter is not supplied, the dump will be directed to a file with a system-defined prefix of "errlog".

Description

When executed with no arguments, the command dumps the cluster error log to a file using a system-supplied prefix of "errlog", which includes the node ID and time stamp. When a file name prefix is provided, the same operation is performed but the details are stored in the dumps directory in a file with a name that starts with the specified prefix.

A maximum of ten error-log dump files are kept on the cluster. When the 11th dump is made, the oldest existing dump file is overwritten.

Error log dump files are written to `/dumps/elogs`. The contents of this directory can be viewed using the **svcinfolerrlogdumps** command.

Possible failures

- CMMVC5983E The dump file was not created. The file system might be full.
- CMMVC5984E The dump file was not written to disk. The file system might be full.
- CMMVC6073E The maximum number of files has been exceeded.
- There are no error codes associated with the **svcservicemodetask dumperrlog** command.

Examples

An invocation example

```
svctask dumperrlog -prefix testerrorlog
```

The resulting output

No feedback

clearerrlog

You can use the **clearerrlog** command to clear all entries from the error log including status events and any unfixed errors.

Syntax

```
svctask --clearerrlog [-force]
```

Parameters

-force

This flag stops any confirmation requests. If the **-force** flag is not supplied, you are prompted to confirm if you are sure that you want to clear the log.

Description

This command clears all entries from the error log. The entries are cleared even if there are unfixed errors in the log. It also clears any status events that are in the log.

Attention: This command is destructive. You should only use it when you have either rebuilt the cluster, or have fixed a major problem that has caused many entries in the error log that you do not want to manually fix.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svctask clearerrlog -force
```

The resulting output

```
No feedback
```

cherrstate

You can use the **cherrstate** command to mark an unfixed error as fixed. You can also use it to mark a fixed error as unfixed.

Syntax

```
svctask —cherrstate— --sequencenumber—sequence_number— [-unfix]
```

Parameters

-sequencenumber *sequence_number*

Specifies the error log sequence number, or numbers, to fix.

-unfix

Optionally specifies that the sequence number, or numbers, supplied should be marked as unfixed. If you supply the **-unfix** argument, the sequence numbers will be marked as unfixed. This is intended for use only when you have marked the wrong sequence number as fixed.

Description

The error log entries that the sequence number, or numbers, that you entered are marked as fixed. Use this command as a manual confirmation step that you have performed some maintenance to the cluster, fabric, or subsystems.

This step is performed as part of the directed maintenance procedures (DMPs).

Optionally, if you have marked the wrong sequence number as fixed, you can remark an entry as unfixed by specifying the **-unfix** flag.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5803E The entry in the error log was not marked because the sequence number was not found.

Examples

An invocation example

```
svctask cherrstate -sequencenumber 2019
```

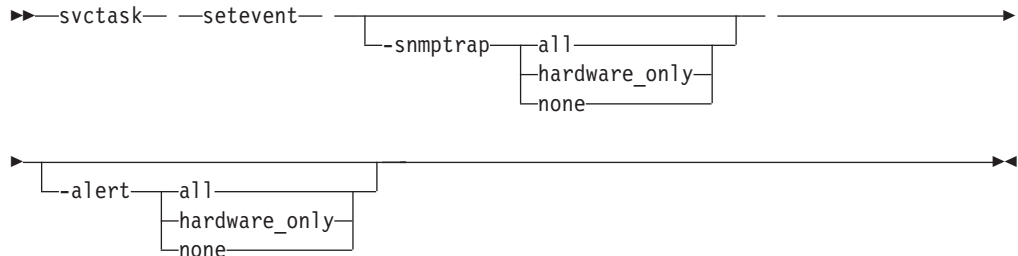
The resulting output

```
No feedback
```

setevent

You can use the **setevent** command to specify what happens when an error or event is logged to the error log.

Syntax



Parameters

-snmptrap *all | hardware_only | none*

Optionally specifies the SNMP trap setting, that is, when to raise a trap.

-alert *all | hardware_only | none*

Optionally specifies the alert setting. That is, when to raise an alert notification.

Note: These arguments are not mutually exclusive.

Description

This command modifies the settings that you want to apply to the error log. These settings define what to do when errors and events are logged. You can set the following values for the **-snmptrap** and **-alert** arguments:

all Sends an SNMP trap for all errors and state changes that are logged.

hardware_only

Sends an SNMP trap for all errors, but not for object state changes.

none Does not send any SNMP traps or errors. This is the default setting for a new cluster.

You can use this command to setup SNMP traps. For SNMP, you must supply the following information:

- When to raise a trap.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svctask setevent -snmptrap all
```

The resulting output

No feedback


Chapter 17. Featurization commands

The following commands enable you to work with featurization functions with the SAN Volume Controller Storage Software for Cisco MDS 9000.

chlicense

You can use the **chlicense** command to change the licensed settings of the cluster.

Syntax



Parameters

-flash *on | off*

Optionally specifies if flash copy has been licensed for this cluster. Default is off.

-remote *on | off*

Optionally specifies if remote copy has been licensed for this cluster. Default is off.

-size *capacity*

Optionally specifies the amount of virtualization licensed for this cluster. Default is 0 GB.

Note: All three arguments are mutually exclusive.

Description

This command will change the license settings for the cluster. Any change made is logged as an event in the featurization log.

The current feature settings for the cluster are displayed in the viewing feature log panel. These settings show if you are licensed to use FlashCopy or Remote Copy. They also show the amount of storage you are licensed to virtualize. Normally, the feature log contains entries because feature options must be set as part of the Web-based create cluster process.

Note: Dumping an empty feature log produces a file that contains headers, 256 lines of formatted zeros, and a couple of lines that includes a checksum.

By default the copy services are disabled, this will not stop you from creating and using copy services. However errors will be placed in the featurization log that state you are using an unlicensed feature. The command line tool return code will also notify you that you are using an unlicensed feature.

The total virtualized capacity can also be modified with this command. This is the number of gigabytes (GB) of virtual disk capacity that can be configured by the cluster.

When you reach 90% capacity, any attempt to create or extend virtual disks will result in a warning message from the command line tool. Again, this will not stop you from creating and expanding virtual disks, but errors will be placed in the featurization log when your usage reaches 100%.

Any error placed in the featurization log results in a generic featurization error being placed in the normal cluster error log. These occur when you issue a command that breaks the license agreement. The return code to the command will also notify you that you are breaking or exceeding the license settings, that is a return code of 1 is returned with any commands that exceed, or break the current license settings. Featurization events are only logged in the featurization log.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.

Examples

An invocation example

```
svctask chlicense -flash on
```

The resulting output

No feedback

dumpinternallog

You can use the **dumpinternallog** command to dump the contents of the featurization error and event log to a file on the current configuration node.

Syntax

```
▶▶—svctask— —dumpinternallog—————▶▶
```

Description

This command will dump the contents of the internal featurization error and event log to a file on the current configuration node.

This file is always called `feature.txt` and will be created (or overwritten) in the `/dumps/feature` directory on the configuration node.

This file may be requested by IBM service personnel.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5983E The dump file was not created. The file system might be full.

Examples

An invocation example

```
svctask dumpinternallog
```

The resulting output

```
No feedback
```

Chapter 18. Service mode commands

These commands can only be run on a node that is in service mode. Attempting to run them on a working configuration node will result in the message:

CMMVC5997E This command can only be run on a node that is in service mode.

Attempting to run any of the other **svctask** and **svcservicetask** commands on a node that is in service mode will result in the following message:

CMMVC5998E This command can not be run on a node that is in service mode.

cleardumps

You can use the **cleardumps** command to clean the various dumps directories on the node that is in service mode.

Syntax

```
▶▶—svcservicemodetask— —cleardumps— —-prefix—directory_or_file_filter—▶▶
```

Parameters

-prefix *directory_or_file_filter*

Specifies the directory and or files to be cleaned. If a directory is specified, with no file filter, all relevant dump/log files in this directory are cleaned. The allowable directory arguments are:

- /dumps (cleans all files in all subdirectories, that is, those listed below)
- /dumps/configs
- /dumps/elogs
- /dumps/ec
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace
- /dumps/pl

In addition to the directory, a file filter can be specified. For example `/dumps/elogs/*.txt` will clean all files in the `/dumps/elogs` directory that end in `.txt`.

Description

This command will delete all the files matching the `directory/file_filter` argument on the node. The node must be in service mode.

This command can be used to clean all the dumps directories by specifying `/dumps` as the directory argument.

The command can also be used to clean all the files in a single directory by specifying one of the directory arguments listed in the note above.

The command can also be used to clean specific files in a given directory by specifying a directory and filename. The wildcard `*` can be used as part of the filename although you may need to quote the filename when wildcards are used.

The contents of these directories on the given node can be listed by using the `svcservicemodeinfo lsxxxdumps` command.

Possible failures

- There are no error codes.

Examples

An invocation example

```
svcservicemodetask cleardumps -prefix /dumps/configs
```

The resulting output
No feedback

dumperrlog

You can use the **dumperrlog** command to dump the contents of the error log to a text file from a node in service mode.

Syntax

```
▶▶—svcservicemodetask— —dumperrlog— —-prefix—filename_prefix————▶▶
```

Parameters

-prefix *filename_prefix*

If **-prefix** is not supplied, then the dump will be directed to a file with a system defined name. A filename is created from the prefix and a timestamp, and will take the form <prefix>_NN_YYMMDD_HHMMSS, where *NN* is the current configuration node ID.

Description

This command dumps the cluster error log to a file using a system supplied name (which will include the node ID and time-stamp). When a filename prefix is used, the same operation is performed but the details are stored in a file whose name starts with the user defined prefix, in the dumps directory.

A maximum of ten error log dump files will be kept on the cluster. When the eleventh dump is made, the oldest existing dump file will be overwritten.

Possible failures

- There are no error codes.

Examples

An invocation example

```
svcservicemodetask dumperrlog -prefix testerrorlog
```

The resulting output

```
No feedback
```

exit

You can use the **exit** command to exit service mode and restart the node.

Syntax

```
▶▶—svcservicemodetask— —exit—————▶▶
```

Description

This command will cause the node to be restarted. The node will come up in standard operating mode and will attempt to rejoin the cluster.

At some point during the execution of this command the SSH and Web server connection that was being used by the user to access the node will be terminated as a result of the restart.

Possible failures

- There are no error codes.

Examples

An invocation example

```
svcservicemodetask exit
```

The resulting output

No feedback

Chapter 19. Service mode information commands

These commands can only be run on a node that is in service mode. Attempting to run them on a working configuration node will result in the message:

CMMVC5997E This command can only be run on a node that is in service mode.

Attempting to run any of the other svcinfo commands on a node that is in service mode will result in the following message:

CMMVC5998E This command can not be run on a node that is in service mode.

lssvcdumps (ls2145dumps)

You can use the **lssvcdumps** command to return a list of the files that exist of the type specified on the given node. The **lssvcdumps** command is the preferred format if you are not in service mode.

Note: For compatibility, **ls2145dumps** can be used as a synonym for **lssvcdumps**.

Syntax

```
▶▶svcserviceinfo — lssvcdumps — [ -nohdr ] [ -delim delimiter ] ▶▶
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command displays files from the `/dumps`, `/dumps/ec`, and `/dumps/pl` directories. Files are listed by date regardless of directory. Files with names starting with `ecdump` are from the `/dumps/ec` directory. Files with names starting with `pldump` are from the `/dumps/pl` directory. Other files (files with a `.trc` extension) are from `/dumps`.

Possible failures

- There are no error codes.

Examples

An invocation example

```
svcserviceinfo lssvcdumps
```

The resulting output

```
id: svc_filename
0: pldump.vegas3.9.2.031122.191834.tar
1: pldump.vegas3.9.2.031123.031812.tar
2: pldump.vegas3.9.2.031123.070501.tar
```

```
3: vegas3.9.2.trc
4: vegas3.9.2_icsmi.trc
5: vegas3.9.2_icsp.trc.old
6: vegas3.9.2_icsp.trc
```

lsclustervpd

You can use the **lsclustervpd** command to return the vital product data (VPD) for the cluster to which the node belonged.

Syntax

```
svcserviceinfo lsclustervpd [-nohdr] [-delim delimiter]
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim delimiter

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command displays the VPD of the cluster to which the node belonged.

Possible failures

- There are no error codes.

Examples

An invocation example

```
svcserviceinfo lsclustervpd
```

The resulting output

```
id 00000200EC2144FE
name cluster1
location local
partnership
bandwidth
cluster_IP_address 9.20.165.16
total_mdisk_capacity 71.8GB
space_in_mdisk_grps 48.5GB
space_allocated_to_vdisks 0
```

```
total_free_space 71.8GB
statistics_status off
statistics_frequency 15
required_memory 4096
SNMP_setting all
email_setting all
code_level 1.3(1) (IBM build 0.3.3)
```

lsconfigdumps

You can use the **lsconfigdumps** command to return a list of the files that exist of the type specified on the given node.

Syntax

```
svcserviceinfo lsconfigdumps [-nohdr] [-delim delimiter]
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim delimiter

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command displays a list of dumps in the `/dumps/configs` directory.

Possible failures

- There are no error codes.

Examples

An invocation example

```
svcserviceinfo lsconfigdumps
```

The resulting output

```
id          config_filename
0          config_myswitch1.3.2_030403_101205
```

lserrlogdumps

You can use the **lserrlogdumps** command to return a list of the files that exist of the type specified on the given node.

Syntax

```
svcservicemodeinfo lserrlogdumps [-nohdr] [-delim delimiter]
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim delimiter

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command displays a list of dumps in the `/dumps/elog`s directory.

Possible failures

- There are no error codes.

Examples

An invocation example

```
svcservicemodeinfo lserrlogdumps
```

The resulting output

id	filename
0	errlog_lynn02_030327_154511
1	aaa.txt_lynn02_030327_154527
2	aaa.txt_lynn02_030327_154559
3	errlog_lynn02_030403_110628

lsfeaturedumps

You can use the **lsfeaturedumps** command to return a list of the files that exist of the type specified on the given node.

Syntax

```
svcserviceinfo lsfeaturedumps [-nohdr] [-delim delimiter]
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command displays a list of dumps in the `/dumps/feature` directory.

Possible failures

- There are no error codes.

Examples

An invocation example

```
svcserviceinfo lsfeaturedumps
```

The resulting output

```
id          feature_filename
0           feature.txt
```

Isiostatsdumps

You can use the **Isiostatsdumps** command to return a list of the files that exist of the type specified on the given node.

Syntax

```
svcservicemodeinfo -lsiostatsdumps [-nohdr] [-delim delimiter]
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command displays a list of dumps in the `/dumps/iostats` directory.

Possible failures

- There are no error codes.

Examples

An invocation example

```
svcservicemodeinfo lsiostatsdumps
```

The resulting output

```
id          iostat_filename
0           v_stats_vegas8.2.1_031123_072426
1           m_stats_vegas8.2.1_031123_072425
```

lsiotracedumps

You can use the **lsiotracedumps** command to return a list of the files that exist of the type specified on the given node.

Syntax

```
svcserviceinfo --lsiotracedumps --nohdr  
--delim delimiter
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim delimiter

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command displays a list of dumps in the `/dumps/iotrace` directory.

Possible failures

- There are no error codes.

Examples

An invocation example

```
svcserviceinfo lsiotracedumps
```

The resulting output

```
id          iotrace_filename  
0           c1_myswitch1.3.2_030405_092155  
1           c2_myswitch1.3.2_030405_092156  
2           c3_myswitch1.3.2_030405_092158  
3           c4_myswitch1.3.2_030405_092159  
4           c5_myswitch1.3.2_030405_092201
```

lsnodes

You can use the **lsnodes** command to get an annotated list of the nodes in the cluster in service mode.

Syntax

```
▶▶—svcserviceinfo—lsnodes—┌─nohdr─┐┌─delim—delimiter─┐▶▶
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the `-nohdr` parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the `-nohdr` option was used or not.

-delim delimiter

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the `-delim` parameter will override this behavior. Valid input for the `-delim` parameter is a one byte character. If, for example, you entered `-delim :` a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command returns an annotated list of nodes in the cluster.

Possible failures

- There are no error codes.

Examples

An invocation example

```
svcserviceinfo lsnodes -delim :
```

The resulting output

```
id:WWNN:front_panel_id:node_name:cluster:fabric
1:2548000530003920:vegas1.1.2:node1:yes:yes
2:2537000530003920:vegas1.8.2:node2:yes:yes
```

lsnodevpd

You can use the **lsnodevpd** command to return the vital product data (VPD) for the given node.

Syntax

```
svcserviceinfo --lsnodevpd [-nohdr] [-delim delimiter]
```

Parameters

-nohdr

By default, headings are displayed for each column of data (in a concise style view), and each item of data (in a detailed style view). Using the **-nohdr** parameter will suppress the display of these headings.

Note: If there is no data to be displayed (for example, an empty view has been returned), then headings are not displayed irrespective of whether the **-nohdr** option was used or not.

-delim *delimiter*

By default in a concise view, all columns of data are space separated. The width of each column is set to the maximum possible width of each item of data. In a detailed view, each item of data has its own row, and if the headers are displayed the data is separated from the header by a space. Using the **-delim** parameter will override this behavior. Valid input for the **-delim** parameter is a one byte character. If, for example, you entered **-delim :** a colon will be used to separate all items of data in a concise view (for example, the spacing of columns does not occur) and in a detailed view the data is separated from its header by a colon.

Description

This command returns the VPD for the specified node. Each field is reported on a new line. All fields are strings.

The VPD is split into sections. Each section has a section heading. Following the heading is the number of fields in that section. Each section is separated by an empty line.

Possible failures

- There are no error codes.

Examples

An invocation example

```
svcserviceinfo lsnodevpd id 1
```

The resulting output

```
id 1

system board: 4 fields
Card serial number JAB071105kn
number_of_processors 2
number_of_scsi/ide_devices 1
```

Switch IP address 9.43.224.172

processor: 8 fields
processor_location Processor 1
manufacturer BROADCOM
speed 750
status Enabled

processor_location Processor 2
manufacturer BROADCOM
speed 750
status Enabled

memory module: 1 fields
size (MB) 4096
device: 6 fields
bus ide0
device 0
model IC25N020ATCS05-0
revision CS20A61A
serial_number CLP225F2G2RTDA
approx_capacity 18

software: 4 fields
code_level 1.3(1) (IBM build 0.3.3)
node_name node2
WWNN 0x2252000b4679ec82
id 0x1

front panel assembly: 1 fields
front_panel_id vegas3.7.2

UPS: 2 fields
Card Serial number JAB071105kn
Card unique id 100071130356b6e

Chapter 20. Controller command

The following command enables you to work with controllers in the SAN Volume Controller Storage Software for Cisco MDS 9000.

chcontroller

You can use the **chcontroller** command to modify the name of a controller.

Syntax

```
svctask chcontroller --name new_name controller_id controller_name
```

Parameters

-name *new_name*

Specifies the new name to be assigned to the controller.

controller_id | **controller_name**

Specifies the controller to be modified. Use either name or id.

Description

This command changes the name of the controller specified by `controller_id` or `controller_name` to that specified by `-name`.

You can add a new disk controller system to your SAN at any time. Follow the switch zoning guidelines in the section about switch zoning. Also, ensure the controller is setup correctly for use with the SAN Volume Controller Storage Software for Cisco MDS 9000 .

You should create one or more arrays on the new controller. It is recommend that you use, RAID-5, RAID-1 or RAID-0+1 (sometimes called RAID-10) for maximum redundancy and reliability. Generally 5+P arrays are recommend. If your controller provides array partitioning we recommend that you create a single partition from the entire capacity available in the array, remember the LUN number that you assign to each partition as you will need this later. You should also follow the mapping guidelines (if your disk controller system requires LUN mapping) to map the partitions or arrays to the SAN Volume Controller Storage Software for Cisco MDS 9000 ports.

Adding a new disk controller system to a running configuration: Ensure that the cluster has detected the new storage (MDisks) by issuing the **svctask detectmdisk** command. The controller itself will have automatically been assigned a default name. If you are unsure which controller is presenting the MDisks, list the controllers by issuing the **svcinfolcontroller** command. You should see a new controller listed (the one with the highest numbered default name). Remember the controller name and follow the instructions in the section about determining a disk controller system name.

You should give this controller a name that you can easily use to identify it. Issue the following command:

```
svctask chcontroller -name <newname> <oldname>
```

List the unmanaged MDisks by issuing the following command:

```
svcinfolsmdisk -filtervalue mode=unmanaged:controller_name=<new_name>
```


These MDisks should correspond with the RAID arrays or partitions you have created. Remember the field controller LUN number. This corresponds with the LUN number you assigned to each of the arrays or partitions.

It is recommended that you create a new managed disk group and add only the RAID arrays that belong to the new controller to this MDisk group. You should also avoid mixing RAID types, so for each set of RAID array types (for example, RAID-5, RAID-1) you should create a new MDisk group. Give this MDisk group an appropriate name, so if your controller is called FAST650-fred, and the MDisk group contains RAID-5 arrays, call it something like F600-fred-R5). Issue the following command:

```
svctask mkmdiskgrp -ext 16 -name <mdisk_grp_name>  
-mdisk <colon separated list of RAID-x mdisks returned  
in step 4.
```

Note: This will create a new MDisk group with an extent size of 16MB.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5816E The action failed because an entity that was specified in the command does not exist.

Examples

An invocation example

```
svctask chcontroller -name newtwo 2
```

The resulting output

No feedback

Chapter 21. CLI messages

The command line tools will return a return value on completion. If the command completes normally and without error, then the return code will be 0. If the command fails then the return code will be 1 and an Error Code will be printed on stderr. If the command succeeds, but the cluster is operating close to its licensed virtualization limit, then the return code may still be 1, and a warning Error Code will be printed to stderr.

When a create command is issued, the message ID that has been assigned to the new object is returned as part of the success message sent to stdout. If the -quiet command is used, only the message ID is sent to stdout.

CMMVC5700E The parameter list is not valid.

Explanation: The parameter list that you specified is not valid.

Action: Specify the correct parameter list and issue the command again.

CMMVC5701E No object ID was specified.

Explanation: You did not specify an object ID.

Action: Specify an object ID and issue the command again.

CMMVC5702E [%1] is below the minimum level.

Explanation: [%1] is below the minimum level.

Action: Specify the correct level and issue the command again.

CMMVC5703E [%1] is above the maximum level.

Explanation: [%1] is above the maximum level.

Action: Specify the correct level and issue the command again.

CMMVC5704E [%1] is not divisible by the permitted step level.

Explanation: [%1] is not divisible by the permitted step level.

Action: Not applicable.

CMMVC5705E A required parameter is missing.

Explanation: A required parameter is missing.

Action: Specify the required parameter and issue the command again.

CMMVC5706E An invalid argument has been entered for the [%1] parameter.

Explanation: [%1] is not a valid argument for the specified parameter.

Action: Specify the correct argument and issue the command again.

CMMVC5707E Required parameters are missing.

Explanation: There are required parameters that are missing.

Action: Specify the required parameters and issue the command again.

CMMVC5708E The %1 parameter is missing its associated arguments.

Explanation: The [%1] parameter is missing its associated arguments.

Action: Specify the associated arguments and issue the command again.

CMMVC5709E [%1] is not a supported parameter.

Explanation: [%1] is not a supported parameter.

Action: Specify the correct parameter and issue the command again.

CMMVC5710E No self describing structure for identifier parameter [%1].

Explanation: There is no self describing structure for the identifier parameter [%1].

Action: Not applicable.

CMMVC5711E [%1] is not valid data.

Explanation: [%1] is not valid data.

Action: Not applicable.

CMMVC5712E Required data is missing.

Explanation: There is required data that is missing.

Action: Specify the missing data and issue the command again.

CMMVC5713E Some parameters are mutually exclusive.

Explanation: There are some parameters which are mutually exclusive.

Action: Not applicable.

CMMVC5714E There are no items in the parameter list.

Explanation: There are no items in the parameter list.

Action: Specify an item in the parameter list and issue the command again.

CMMVC5715E There is no parameter list.

Explanation: There is no parameter list.

Action: Not applicable.

CMMVC5716E Nonnumeric data was entered for a numeric field ([%1]). Enter a numeric value.

Explanation: Nonnumeric data was specified for a numeric field.

Action: Specify a numeric value in the numeric field and issue the command again.

CMMVC5717E No match was found for the specified unit.

Explanation: No match was found for the specified unit.

Action: Specify the correct unit and issue the command again.

CMMVC5718E An unexpected return code was received.

Explanation: An unexpected return code was received.

Action: Not applicable.

CMMVC5719E A value of %2 requires the parameter %1 to be specified.

Explanation: A value of %2 requires the parameter %1 to be specified.

Action: Specify the required parameter and issue the command again.

CMMVC5720E [%1] is not a valid argument for the -o parameter.

Explanation: [%1] is not a valid argument for the -o parameter.

Action: Specify the correct argument and issue the command again.

CMMVC5721E [%1] is not a valid time-stamp format. The valid format is MMDDHHmmYY.

Explanation: [%1] is not a valid time-stamp format. The valid format is MMDDHHmmYY.

Action: Follow the correct time-stamp format and issue the command again.

CMMVC5722E [%1] is not a valid month.

Explanation: [%1] is not a valid month.

Action: Specify the correct month (MM) and issue the command again.

CMMVC5723E [%1] is not a valid day.

Explanation: [%1] is not a valid day.

Action: Specify the correct day (DD) and issue the command again.

CMMVC5724E [%1] is not a valid hour.

Explanation: [%1] is not a valid hour.

Action: Specify the correct hour (HH) and issue the command again.

CMMVC5725E [%1] is not a valid minute.

Explanation: [%1] is not a valid minute.

Action: Specify the correct minute (mm) and issue the command again.

CMMVC5726E [%1] are not valid seconds.

Explanation: [%1] are not valid seconds.

Action: Specify the correct seconds (ss) and issue the command again.

CMMVC5727E [%1] is not a valid filter.

Explanation: [%1] is not a valid filter.

Action: Not applicable.

CMMVC5728E [%1] should be in the format
minute:hour:day:month:weekday.

Explanation: [%1] should be in the format
minute:hour:day:month:weekday.

Action: Follow the correct format and issue the
command again.

CMMVC5729E One or more components in the list
is not valid.

Explanation: You specified one or more components
that are not valid.

Action: Specify the correct component and issue the
command again.

CMMVC5730E %1 is only valid when %2 has a
value of %3.

Explanation: %1 is only valid when %2 has a value of
%3.

Action: Not applicable.

CMMVC5731E %1 can only be entered when %2 has
been entered.

Explanation: %1 can only be entered when %2 has
been entered.

Action: Not applicable.

CMMVC5732E The shared-memory interface is not
available.

Explanation: The shared memory interface (SMI) is
not available.

Action: Not applicable.

CMMVC5733E Enter at least one parameter.

Explanation: At least one parameter needs to be
specified.

Action: Specify the correct parameter and issue the
command again.

CMMVC5734E A combination of values was entered
that is not valid.

Explanation: You specified a combination of values
that are not valid.

Action: Specify the correct combination of values and
issue the command again.

CMMVC5735E The name entered is not valid. Enter
an alphanumeric string that does not
start with a numeric.

Explanation: You specified a name that is not valid.

Action: Specify an alphanumeric string that does not
start with a numeric.

CMMVC5736E -c is not a valid unit.

Explanation: You specified a parameter that is not a
valid unit.

Action: Specify the correct parameter and issue the
command again.

CMMVC5737E The parameter %1 has been entered
multiple times. Enter the parameter
once.

Explanation: You specified the same parameter more
than once.

Action: Delete any duplicate parameters and issue the
command again.

CMMVC5738E The name contains too many
characters. Enter an alphanumeric string
of 1 - 15 characters that is composed of
any of the following characters: A - Z, a
- z, 0 - 9, -, or _. The first character
cannot be numeric.

Explanation: You specified an argument that contains
too many characters.

Action: Specify the correct argument and issue the
command again.

CMMVC5739E The argument %1 does not contain
enough letters.

Explanation: You specified an argument that does not
contain enough characters.

Action: Specify the correct argument and issue the
command again.

CMMVC5740E The filter flag %1 is not valid.

Explanation: The filter flag %1 is not valid.

Action: Specify the correct flag and issue the
command again.

CMMVC5741E The filter value %1 is not valid.

Explanation: The filter value %1 is not valid.

Action: Specify the correct value and issue the
command again.

CMMVC5742E A specified parameter is out of its valid range.

Explanation: You specified a parameter that is out of its valid range.

Action: Specify the correct parameter and issue the command again.

CMMVC5743E A specified parameter does not comply with the step value.

Explanation: You specified a parameter that does not comply with the step value.

Action: Specify the correct parameter and issue the command again.

CMMVC5744E Too many objects were specified in the command.

Explanation: There were too many objects specified in the command.

Action: Specify the correct object and issue the command again.

CMMVC5745E Not enough objects were specified in the command.

Explanation: There were not enough objects specified in the command.

Action: Specify the correct object and issue the command again.

CMMVC5746E The operation that was requested is not valid for the this object.

Explanation: The requested operation is not valid for this object.

Action: Specify a valid operation and issue the command again.

CMMVC5747E The operation that was requested is not valid.

Explanation: The operation that was requested is not valid.

Action: Specify the correct operation and issue the command again.

CMMVC5748E The operation that was requested is not valid.

Explanation: The operation that was requested is not valid.

Action: Specify the correct operation and issue the command again.

CMMVC5749E The dump file name already exists.

Explanation: The dump file name that you specified already exists.

Action: Specify a different dump file name and issue the command again.

CMMVC5750E The dump file was not created. The file system might be full.

Explanation: The dump file was not created. The file system might be full.

Action: Not applicable.

CMMVC5751E The dump file could not be written to disk.

Explanation: The dump file could not be written to disk.

Action: Not applicable.

CMMVC5752E The operation failed because the object contains child objects. Delete the child objects and then resubmit the request.

Explanation: The operation failed because the specified object contains child objects.

Action: Delete the child objects and issue the command again.

CMMVC5753E The specified object does not exist.

Explanation: The specified object does not exist.

Action: Specify the correct object and issue the command again.

CMMVC5754E The specified object does not exist, or the name does not meet the naming rules.

Explanation: The specified object does not exist, or the name of the object does not meet the naming requirements.

Action: Specify the correct object name and issue the command again.

CMMVC5755E The sizes of the specified objects do not match.

Explanation: The sizes of the specified objects do not match.

Action: Not applicable.

CMMVC5756E The operation failed because the object is already mapped.

Explanation: The operation failed because the specified object is already mapped.

Action: Specify a different object and issue the command again.

CMMVC5757E The defaults for the self describing structure were not found.

Explanation: The defaults for the self describing structure were not found.

Action: Not applicable.

CMMVC5758E The object filename already exists.

Explanation: The object filename already exists.

Action: Specify a different object filename and issue the command again.

CMMVC5759E Memory cannot be allocated.

Explanation: The memory cannot be allocated.

Action: Not applicable.

CMMVC5760E The node could not be added to the cluster.

Explanation: The node could not be added to the cluster.

Action: Not applicable.

CMMVC5761E The node could not be deleted from the cluster.

Explanation: The node could not be deleted from the cluster.

Action: Not applicable.

CMMVC5762E The operation failed because the timeout period expired.

Explanation: The operation failed because the timeout period expired.

Action: Issue the command again.

CMMVC5763E The node failed to go online.

Explanation: The node failed to go online.

Action: Not applicable.

CMMVC5764E The specified mode change is not valid.

Explanation: The specified mode change is not valid.

Action: Specify a different mode and issue the command again.

CMMVC5765E The select object is no longer a candidate. A change occurred during the request.

Explanation: The specified object is no longer a candidate. A change occurred during the request.

Action: Specify a different object and issue the command again.

CMMVC5767E One or more of the specified parameters is not valid.

Explanation: One or more of the specified parameters is not valid.

Action: Specify the correct parameter and issue the command again.

CMMVC5769E The operation requires that all nodes be online. One or more nodes are not online.

Explanation: The operation requires that all nodes be online. One or more nodes are not online.

Action: Check that each node is online and issue the command again.

CMMVC5771E The operation failed, probably, because the object contains child objects. To complete the operation, specify the force flag.

Explanation: The operation failed, probably, because the object contains child objects.

Action: Specify the -force flag to complete the operation and then issue the command again.

CMMVC5772E The operation failed because a software upgrade is in progress.

Explanation: The operation failed because a software upgrade is in progress.

Action: Wait for the software upgrade to complete and then issue the command again.

CMMVC5773E The operation failed because the selected object is in the wrong mode.

Explanation: The operation failed because the selected object is in the wrong mode.

Action: Specify the correct mode and issue the command again.

CMMVC5775E The directory attribute is not valid.

Explanation: The directory attribute is not valid.

Action: Specify a different directory and issue the command again.

CMMVC5776E The directory listing could not be retrieved.

Explanation: The directory listing could not be retrieved.

Action: Specify a different directory listing and issue the command again.

CMMVC5777E The node was not added to the I/O group because the other node in the I/O Group is in the same power domain.

Explanation: The node was not added to the I/O group because the other node in the I/O group is on the same Cisco MDS 9000 Caching Services Module.

Action: Specify a different node from another I/O group and issue the command again.

CMMVC5778E The cluster was not created because one already exists.

Explanation: The cluster was not created because one already exists.

Action: Not applicable.

CMMVC5781E The cluster ID is not valid.

Explanation: The cluster ID is not valid.

Action: Specify a different cluster ID and issue the command again.

CMMVC5782E The object is offline.

Explanation: The object is offline.

Action: Specify an object that is online and issue the command again.

CMMVC5784E The cluster name is not unique. Specify the cluster using the cluster ID.

Explanation: The cluster name is not unique.

Action: Specify the cluster using the cluster ID and issue the command again.

CMMVC5785E The filename contains an illegal character.

Explanation: The filename contains an illegal character.

Action: Specify a valid filename and issue the command again.

CMMVC5786E The action failed because the cluster is not in a stable state.

Explanation: The action failed because the cluster is not in a stable state.

Action: Not applicable.

CMMVC5787E The cluster was not created because a cluster already exists.

Explanation: The cluster was not created because a cluster already exists.

Action: Not applicable.

CMMVC5790E The node was not added to the cluster because the maximum number of nodes has been reached.

Explanation: The node was not added to the cluster because the maximum number of nodes has been reached.

Action: Not applicable.

CMMVC5791E The action failed because an entity that was specified in the command does not exist.

Explanation: An entity that was specified in the command does not exist, therefore the action failed.

Action: Specify the correct entity and issue the command again.

CMMVC5792E The action failed because the I/O group is used for recovery.

Explanation: The action failed because the I/O group is used for recovery.

Action: Not applicable.

CMMVC5793E The node was not added to the cluster because the I/O group already contains a pair of nodes.

Explanation: The node was not added to the cluster because the I/O group already contains a pair of nodes.

Action: Not applicable.

CMMVC5794E The action failed because the node is not a member of the cluster.

Explanation: The node is not a member of the cluster, therefore the action failed.

Action: Specify a node that is contained in the cluster and issue the command again.

CMMVC5795E The node was not deleted because a software upgrade is in progress.

Explanation: The node was not deleted because a software upgrade is in progress.

Action: Wait for the software upgrade to complete and then issue the command again.

CMMVC5796E The action failed because the I/O group that the node belongs to is unstable.

Explanation: The I/O group that the node belongs too is unstable, therefore the action failed.

Action: Not applicable.

CMMVC5797E The node was not deleted because this is the last node in the I/O group and there are virtual disks (VDisks) associated with the I/O group.

Explanation: The specified node is the last node in the I/O group and there are VDisks associated with the I/O group, therefore the node could not be deleted.

Action: Not applicable.

CMMVC5798E The action failed because the node is offline.

Explanation: The action failed because the node is offline.

Action: Specify a node that is online and issue the command again.

CMMVC5799E The shut down was not successful because there is only one online node in the I/O group.

Explanation: There is only one online node is the I/O group, therefore the shut down operation was not successful.

Action: Not applicable.

CMMVC5800E The action failed because an entity that was specified in the command does not exist.

Explanation: The entity that was specified in the command does not exist, therefore the action failed.

Action: Specify a different entity and issue the command again.

CMMVC5801E The upgrade of the cluster software could not proceed because every node in the cluster must be online. Either delete the node that is offline or bring the node online and resubmit the command

Explanation: The upgrade of the cluster software could not proceed because every node in the cluster must be online.

Action: Either delete the node that is offline or bring the node online and issue the command again.

CMMVC5802E The upgrade of the cluster software could not proceed because there is an I/O group in the cluster that contains only one node. The software upgrade requires that each node in an I/O group be shut down and restarted. If there is only one node in an I/O group, I/O operations could be lost if I/O operations are not stopped before beginning the software upgrade. To upgrade the cluster, the force option is required.

Explanation: The upgrade of the cluster software could not proceed because there is an I/O group in the cluster that contains only one node. The software upgrade requires that each node in an I/O group be shut down and restarted. If there is only one node in an I/O group, I/O operations could be lost if I/O operations are not stopped before beginning the software upgrade. To upgrade the cluster, the force option is required.

Action: Either upgrade the cluster using the -force option or specify a different node and issue the command again.

CMMVC5803E The entry in the error log was not marked because the sequence number was not found.

Explanation: The entry in the error log was not marked because the sequence number was not found.

Action: Not applicable.

CMMVC5804E The action failed because an entity that was specified in the command does not exist.

Explanation: The entity that was specified in the command does not exist, therefore the action failed.

Action: Specify a different entity and issue the command again.

CMMVC5805E The progress information was not returned because the FlashCopy statistics are not ready yet.

Explanation: The progress information was not returned because the FlashCopy statistics are not ready yet.

Action: Not applicable.

CMMVC5806E The action failed because an entity that was specified in the command does not exist.

Explanation: The entity that was specified in the command does not exist, therefore the action failed.

Action: Specify a different entity and issue the command again.

CMMVC5807E The action failed because the managed disk (MDisk) cannot be changed to the specified mode.

Explanation: The action failed because the managed disk (MDisk) cannot be changed to the specified mode.

Action: Not applicable.

CMMVC5808E The action failed because the managed disk (MDisk) does not exist.

Explanation: The action failed because the managed disk (MDisk) does not exist.

Action: Specify a different MDisk and issue the command again.

CMMVC5809E The tracing of I/O operations was not started because it is already in progress.

Explanation: The tracing of I/O operations was not started because it is already in progress.

Action: Not applicable.

CMMVC5810E The quorum index number for the managed disk (MDisk) was not set because the MDisk is offline.

Explanation: The quorum index number for the managed disk (MDisk) was not set because the MDisk is offline.

Action: Either change the status of the MDisk to online or specify a different MDisk and issue the command again.

CMMVC5811E The quorum index number for the managed disk (MDisk) was not set because the quorum disk does not exist.

Explanation: The quorum index number for the managed disk (MDisk) was not set because the quorum disk does not exist.

Action: Specify a different quorum disk and issue the command again.

CMMVC5812E The quorum index number for the managed disk (MDisk) was not set because the MDisk is in the wrong mode. Select an MDisk that has a mode of managed.

Explanation: The quorum index number for the managed disk (MDisk) was not set because the MDisk is not in the managed mode.

Action:

- Change the mode of the MDisk and issue the command again
 - Select an MDisk with a mode of managed and issue the command again
-

CMMVC5813E The quorum index number for the managed disk (MDisk) was not set because the MDisk has a sector size that is not valid.

Explanation: The parameter list that you specified is not valid.

Action: Specify a different sector size for the MDisk and issue the command again.

CMMVC5814E The quorum index number for the managed disk (MDisk) was not set because the unique identifier (UID) type is not valid.

Explanation: The quorum index number for the managed disk (MDisk) was not set because the unique identifier (UID) type is not valid.

Action: Specify a different unique identifier (UID) and issue the command again.

CMMVC5815E The managed disk (MDisk) group was not created because an entity that was specified in the command does not exist.

Explanation: The managed disk (MDisk) group was not created because an entity that was specified in the command does not exist.

Action: Specify a different entity and issue the command again.

CMMVC5816E The action failed because an entity that was specified in the command does not exist.

Explanation: The action failed because an entity that was specified in the command does not exist.

Action: Specify a different entity and issue the command again.

CMMVC5817E The managed disk (MDisk) group was not renamed because the name was not valid.

Explanation: The managed disk (MDisk) group was not renamed because the name was not valid.

Action: Specify a different MDisk group name and issue the command again.

CMMVC5818E The managed disk (MDisk) group was not deleted because there is at least one MDisk in the group.

Explanation: The managed disk (MDisk) group was not deleted because there is at least one MDisk in the group.

Action: Not applicable.

CMMVC5819E The managed disk (MDisk) was not added to the MDisk group because the MDisk is part of another MDisk group.

Explanation: The managed disk (MDisk) was not added to the MDisk group because the MDisk is part of another MDisk group.

Action: Not applicable.

CMMVC5820E The managed disk (MDisk) was not added to the MDisk group because an entity that was specified in the command does not exist.

Explanation: The managed disk (MDisk) was not added to the MDisk group because an entity that was specified in the command does not exist.

Action: Specify a different entity and issue the command again.

CMMVC5821E The managed disk (MDisk) was not added to the MDisk group because not enough MDisks were included in the list.

Explanation: The managed disk (MDisk) was not added to the MDisk group because not enough MDisks were included in the list.

Action: Include more MDisks in the list and issue the command again.

CMMVC5822E The managed disk (MDisk) was not added to the MDisk group because too many MDisks were included in the list.

Explanation: The managed disk (MDisk) was not added to the MDisk group because too many MDisks were included in the list.

Action: Delete the extra MDisks in the list and issue the command again.

CMMVC5823E The managed disk (MDisk) was not deleted from the MDisk group because the MDisk is part of another MDisk group.

Explanation: The managed disk (MDisk) was not deleted from the MDisk group because the MDisk is part of another MDisk group.

Action: Not applicable.

CMMVC5824E The managed disk (MDisk) was not deleted from the MDisk group because it does not belong to the MDisk group.

Explanation: The managed disk (MDisk) was not deleted from the MDisk group because it does not belong to the MDisk group.

Action: Not applicable.

CMMVC5825E The managed disk (MDisk) was not deleted from the MDisk group because a virtual disk (VDisk) is allocated from one or more of the specified MDisks. A forced deletion is required.

Explanation: The managed disk (MDisk) was not deleted from the MDisk group because a virtual disk (VDisk) is allocated from one or more of the specified MDisks.

Action: Specify the -force option and issue the command again.

CMMVC5826E The virtual disk (VDisk) was not created because an entity that was specified in the command does not exist.

Explanation: The virtual disk (VDisk) was not created because an entity that was specified in the command does not exist.

Action: Specify a different entity and issue the command again.

CMMVC5827E The command failed as a result of an inconsistency between two or more of the entered parameters.

Explanation: The command failed as a result of an inconsistency between two or more of the entered parameters.

Action: Specify one parameter and issue the command again.

CMMVC5828E The virtual disk (VDisk) was not created because the I/O group contains no nodes.

Explanation: The virtual disk (VDisk) was not created because the I/O group contains no nodes.

Action: Not applicable.

CMMVC5829E The image-mode virtual disk (VDisk) was not created because the number of managed disks (MDisks) specified is greater than one.

Explanation: The image-mode virtual disk (VDisk) was not created because the number of managed disks (MDisks) specified is greater than one.

Action: Specify a different MDisk and issue the command again.

CMMVC5830E The image-mode virtual disk (VDisk) was not created because no managed disk (MDisk) was specified in the command.

Explanation: The image-mode virtual disk (VDisk) was not created because no managed disk (MDisk) was specified in the command.)

Action: Specify a MDisk and issue the command again.

CMMVC5831E The virtual disk (VDisk) was not created because the preferred node for I/O operations is not part of the I/O group.

Explanation: The virtual disk (VDisk) was not created because the preferred node for I/O operations is not part of the I/O group.

Action: Specify a different node and issue the command again.

CMMVC5832E The property of the virtual disk (VDisk) was not modified because an entity that was specified in the command does not exist.

Explanation: The property of the virtual disk (VDisk) was not modified because an entity that was specified

in the command does not exist.

Action: Specify a different entity and issue the command again.

CMMVC5833E The property of the virtual disk (VDisk) was not modified because there are no nodes in the I/O group.

Explanation: The property of the virtual disk (VDisk) was not modified because there are no nodes in the I/O group.

Action: Not applicable.

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.

Explanation: 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.

Action: Specify the -force option and issue the command again.

CMMVC5835E The virtual disk (VDisk) was not expanded because an entity that was specified in the command does not exist.

Explanation: The virtual disk (VDisk) was not expanded because an entity that was specified in the command does not exist.

Action: Specify a different entity and issue the command again.

CMMVC5836E The virtual disk (VDisk) was not shrunk because it is locked.

Explanation: The virtual disk (VDisk) was not shrunk because it is locked.

Action: Unlock the VDisk and issue the command again.

CMMVC5837E The action failed because the virtual disk (VDisk) is part of a FlashCopy mapping.

Explanation: The action failed because the virtual disk (VDisk) is part of a FlashCopy mapping.

Action: Specify a different VDisk that is not part of a FlashCopy mapping and issue the command again.

CMMVC5838E The action failed because the virtual disk (VDisk) is part of a Remote Copy mapping.

Explanation: The action failed because the virtual disk (VDisk) is part of a Remote Copy mapping.

Action: Specify a different VDisk that is not part of a Remote Copy mapping and issue the command again.

CMMVC5839E The virtual disk (VDisk) was not shrunk because an entity that was specified in the command does not exist.

Explanation: The virtual disk (VDisk) was not shrunk because an entity that was specified in the command does not exist.

Action: Specify a different entity and issue the command again.

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.

Explanation: 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.

Action: Specify a different VDisk and issue the command again.

CMMVC5841E The virtual disk (VDisk) was not deleted because it does not exist.

Explanation: The virtual disk (VDisk) was not deleted because it does not exist.

Action: Specify a different VDisk and issue the command again.

CMMVC5842E The action failed because an entity that was specified in the command does not exist.

Explanation: The action failed because an entity that was specified in the command does not exist.

Action: Specify a different entity and issue the command again.

CMMVC5843E The virtual disk (VDisk)-to-host mapping was not created because the VDisk does not have a capacity greater than zero bytes.

Explanation: The virtual disk (VDisk)-to-host mapping was not created because the VDisk does not have a capacity greater than zero bytes.

Action: Specify a VDisk in which its capacity is

greater than zero bytes and issue the command again.

CMMVC5844E The virtual disk (VDisk)-to-host mapping was not created because the SCSI logical unit number (LUN) ID is not valid.

Explanation: The virtual disk (VDisk)-to-host mapping was not created because the SCSI logical unit number (LUN) ID is not valid.

Action: Specify the correct SCSI logical unit number (LUN) ID and issue the command again.

CMMVC5845E The extent was not migrated because an entity that was specified in the command does not exist.

Explanation: The extent was not migrated because an entity that was specified in the command does not exist.

Action: Specify a different entity and issue the command again.

CMMVC5846E The virtual disk (VDisk) was not migrated because an entity that was specified in the command does not exist.

Explanation: The virtual disk (VDisk) was not migrated because an entity that was specified in the command does not exist.

Action: Specify a different entity and issue the command again.

CMMVC5847E The virtual disk (VDisk) was not migrated because its associated managed disk (MDisk) is already in the MDisk group.

Explanation: The virtual disk (VDisk) was not migrated because its associated managed disk (MDisk) is already in the MDisk group.

Action: Not applicable.

CMMVC5848E The action failed because the virtual disk (VDisk) does not exist or it is being deleted.

Explanation: The action failed because the virtual disk (VDisk) does not exist or it is being deleted.

Action: Specify a different VDisk and issue the command again.

CMMVC5849E The migration failed because some or all of the extents are already being migrated.

Explanation: The migration failed because some or all of the extents are already being migrated.

Action: Not applicable.

CMMVC5850E The extent was not migrated because there is a problem with the source extents.

Explanation: The extent was not migrated because there is a problem with the source extents.

Action: Not applicable.

CMMVC5851E The extent was not migrated because there is a problem with the target extents.

Explanation: The extent was not migrated because there is a problem with the target extents.

Action: Not applicable.

CMMVC5852E The migration failed because there are too many migrations in progress.

Explanation: The migration failed because there are too many migrations in progress.

Action: Wait for the migration process to complete and issue the command again.

CMMVC5853E The action failed because there was a problem with the MDisk group.

Explanation: An attempt was made to work on a VDisk which is using an MDisk group with one of the following problems:

- The target and source MDisk groups have different extent sizes (group migrate).
- The target and source MDisk groups are the same (group migrate).
- The target and source MDisk groups are different (extents migrate).
- Target invalid group (Group Migrate).
- Source invalid group (Group Migrate).

Action: Ensure that none of the above conditions exist before issuing the command again.

CMMVC5854E The extent information was not returned because the extent is not used or does not exist.

Explanation: The extent information was not returned because the extent is not used or does not exist.

Action: Specify the correct extent and issue the command again.

CMMVC5855E The extent information was not returned because the managed disk (MDisk) is not used by any virtual disk (VDisk).

Explanation: The extent information was not returned because the managed disk (MDisk) is not used by any virtual disk (VDisk).

Action: Specify the correct MDisk and issue the command again.

CMMVC5856E The action failed because the virtual disk (VDisk) does not belong to the specified managed disk (MDisk) group.

Explanation: The action failed because the virtual disk (VDisk) does not belong to the specified managed disk (MDisk) group.

Action: Specify a different VDisk and issue the command again.

CMMVC5857E The action failed because the managed disk (MDisk) does not exist or it is not a member of the managed disk (MDisk) group.

Explanation: The action failed because the managed disk (MDisk) does not exist or it is not a member of the managed disk (MDisk) group.

Action: Specify a different MDisk and issue the command again.

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.

Explanation: 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.

Action: Check that the VDisk and MDisk are in the correct mode and issue the command again.

CMMVC5859E The migration did not complete because an error occurred while migrating the last extent on an image-mode virtual disk (VDisk).

Explanation: The migration did not complete because an error occurred while migrating the last extent on an image-mode virtual disk (VDisk).

Action: Not applicable.

CMMVC5860E The action failed because there were not enough extents in the managed disk (MDisk) group.

Explanation: 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.

Action: In this case, the MDisk group reports that it has enough free capacity to create the VDisk. You can check the free capacity on each MDisk by issuing the `svcinfolsfreextents <mdiskname/ID>`. 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).

Explanation: The action failed because there were not enough extents on the managed disk (MDisk).

Action: Specify another extent and issue the command again.

CMMVC5862E The action failed because the virtual disk (VDisk) is being formatted.

Explanation: The action failed because the virtual disk (VDisk) is being formatted.

Action: Wait for the VDisk to be successfully formatted and then issue the command again.

CMMVC5863E The migration failed because there are not enough free extents on the target managed disk (MDisk).

Explanation: The migration failed because there are not enough free extents on the target managed disk (MDisk).

Action: Specify another free extent and issue the command again.

CMMVC5864E The extent information was not returned because the source extent is not used.

Explanation: The extent information was not returned because the source extent is not used.

Action: Specify a different source extent and issue the command again.

CMMVC5865E The extent information was not returned because the extent is out of range for the managed disk (MDisk) or virtual disk (VDisk).

Explanation: The extent information was not returned because the extent is out of range for the managed disk

(MDisk) or virtual disk (VDisk).

Action: Specify a different extent which is in range for the MDisk or VDisk and issue the command again.

CMMVC5866E The extent was not migrated because the extent contains internal data.

Explanation: The extent was not migrated because the extent contains internal data.

Action: Not applicable.

CMMVC5867E The action failed because the worldwide port name is already assigned or is not valid.

Explanation: The action failed because the worldwide port name is already assigned or is not valid.

Action: Specify a different worldwide port name and issue the command again.

CMMVC5868E The action failed because an entity that was specified in the command does not exist.

Explanation: The action failed because an entity that was specified in the command does not exist.

Action: Specify a different entity and issue the command again.

CMMVC5869E The host object was not renamed because the host ID or name is not valid.

Explanation: The host object was not renamed because the host ID or name is not valid.

Action: Specify a different host ID or name and issue the command again.

CMMVC5870E The host object was not deleted because an entity that was specified in the command does not exist.

Explanation: The host object was not deleted because an entity that was specified in the command does not exist.

Action: Specify the correct entity and issue the command again.

CMMVC5871E The action failed because one or more of the configured worldwide port names is in a mapping.

Explanation: The action failed because one or more of the configured worldwide port names is in a mapping.

Action: Specify a worldwide port name that is not in a mapping and issue the command again.

CMMVC5872E The port (WWPN) was not added to the host object because an entity that was specified in the command does not exist.

Explanation: The port (WWPN) was not added to the host object because an entity that was specified in the command does not exist.

Action: Specify the correct entity and issue the command again.

CMMVC5873E The action failed because there is no matching worldwide port name.

Explanation: The action failed because there is no matching worldwide port name.

Action: Not applicable.

CMMVC5874E The action failed because the host does not exist.

Explanation: The action failed because the host does not exist.

Action: Specify a different host and issue the command again.

CMMVC5875E The action failed because the virtual disk (VDisk) does not exist.

Explanation: The action failed because the virtual disk (VDisk) does not exist.

Action: Specify a different VDisk and issue the command again.

CMMVC5876E The virtual disk (VDisk)-to-host mapping was not created because the maximum number of mappings has been reached.

Explanation: The virtual disk (VDisk)-to-host mapping was not created because the maximum number of mappings has been reached.

Action: Not applicable.

CMMVC5877E The virtual disk (VDisk)-to-host mapping was not created because the maximum number of SCSI LUNs has been allocated.

Explanation: The virtual disk (VDisk)-to-host mapping was not created because the maximum number of SCSI LUNs has been allocated.

Action: Not applicable.

CMMVC5878E The virtual disk (VDisk)-to-host mapping was not created because this VDisk is already mapped to this host.

Explanation: The virtual disk (VDisk)-to-host mapping was not created because this VDisk is already mapped to this host.

Action: Specify a different VDisk and issue the command again.

CMMVC5879E The virtual disk-to-host mapping was not created because this SCSI LUN is already assigned to another mapping.

Explanation: The virtual disk-to-host mapping was not created because this SCSI LUN is already assigned to another mapping.

Action: Specify a different SCSI LUN and issue the command again.

CMMVC5880E The virtual disk (VDisk)-to-host mapping was not created because the VDisk has a capacity of zero bytes.

Explanation: The virtual disk (VDisk)-to-host mapping was not created because the VDisk has a capacity of zero bytes.

Action: Specify a different VDisk and issue the command again.

CMMVC5881E The FlashCopy mapping was not created because an entity that was specified in the command does not exist.

Explanation: The FlashCopy mapping was not created because an entity that was specified in the command does not exist.

Action: Specify a different entity and issue the command again.

CMMVC5882E The FlashCopy mapping was not created because a mapping for the source or target virtual disk (VDisk) already exists.

Explanation: The FlashCopy mapping was not created because a mapping for the source or target virtual disk (VDisk) already exists.

Action: Specify a different source or target VDisk and issue the command again.

CMMVC5883E The FlashCopy mapping was not created because the recovery I/O group is associated with the source or target virtual disk (VDisk).

Explanation: The FlashCopy mapping was not created because the recovery I/O group is associated with the source or target virtual disk (VDisk).

Action: Specify a different recovery I/O group and issue the command again.

CMMVC5884E The FlashCopy mapping was not created because the source or target virtual disk (VDisk) cannot be a member of a Remote Copy mapping.

Explanation: The FlashCopy mapping was not created because the source or target virtual disk (VDisk) cannot be a member of a Remote Copy mapping.

Action: Specify a different source or target VDisk and issue the command again.

CMMVC5885E The FlashCopy mapping was not created because this source or target virtual disk (VDisk) cannot be a member of a FlashCopy mapping.

Explanation: The FlashCopy mapping was not created because this source or target virtual disk (VDisk) cannot be a member of a FlashCopy mapping.

Action: Specify a different source or target VDisk and issue the command again.

CMMVC5886E The FlashCopy mapping was not created because the source or target virtual disk (VDisk) is associated with the recovery I/O group.

Explanation: The FlashCopy mapping was not created because the source or target virtual disk (VDisk) is associated with the recovery I/O group.

Action: Specify a different source or target VDisk and issue the command again.

CMMVC5888E The action failed because an entity that was specified in the command does not exist.

Explanation: The action failed because an entity that was specified in the command does not exist.

Action: Specify the correct entity and issue the command again.

CMMVC5889E The FlashCopy mapping was not deleted because an entity that was specified in the command does not exist.

Explanation: The FlashCopy mapping was not deleted because an entity that was specified in the command does not exist.

Action: Specify a different entity and issue the command again.

CMMVC5890E The FlashCopy mapping or consistency group was not started because starting consistency group 0 is not a valid operation.

Explanation: The FlashCopy mapping or consistency group was not started because starting consistency group 0 is not a valid operation.

Action: Not applicable.

CMMVC5891E The FlashCopy consistency group was not created because the name is not valid.

Explanation: The FlashCopy consistency group was not created because the name is not valid.

Action: Specify a different name and issue the command again.

CMMVC5892E The FlashCopy consistency group was not created because it already exists.

Explanation: The FlashCopy consistency group was not created because it already exists.

Action: Not applicable.

CMMVC5893E The action failed because an entity that was specified in the command does not exist.

Explanation: The action failed because an entity that was specified in the command does not exist.

Action: Specify the correct entity and issue the command again.

CMMVC5894E The FlashCopy consistency group was not deleted because you are trying to delete consistency group 0 or the name of the consistency group is not valid.

Explanation: The FlashCopy consistency group was not deleted because you are trying to delete consistency group 0 or the name of the consistency group is not valid.

Action: Specify the correct consistency group and issue the command again.

CMMVC5895E The FlashCopy consistency group was not deleted because it contains mappings. To delete this consistency group, a forced deletion is required.

Explanation: The FlashCopy consistency group was not deleted because it contains mappings.

Action: Specify that -force option to delete the consistency group.

CMMVC5896E The FlashCopy mapping was not deleted because the mapping or consistency group is in the preparing state. The mapping or consistency group must be stopped first.

Explanation: The FlashCopy mapping was not deleted because the mapping or consistency group is in the preparing state. The mapping or consistency group must be stopped first.

Action: Stop the consistency group and then issue the command again.

CMMVC5897E The FlashCopy mapping was not deleted because the mapping or consistency group is in the prepared state. The mapping or consistency group must be stopped first.

Explanation: The FlashCopy mapping was not deleted because the mapping or consistency group is in the prepared state. The mapping or consistency group must be stopped first.

Action: Stop the consistency group and then issue the command again.

CMMVC5898E The FlashCopy mapping was not deleted because the mapping or consistency group is in the copying state. The mapping or consistency group must be stopped first.

Explanation: The FlashCopy mapping was not deleted because the mapping or consistency group is in the copying state. The mapping or consistency group must be stopped first.

Action: Stop the consistency group and then issue the command again.

CMMVC5899E The FlashCopy mapping was not deleted because the mapping or consistency group is in the stopped state. To delete the mapping, a forced deletion is required.

Explanation: The FlashCopy mapping was not deleted because the mapping or consistency group is in the stopped state.

Action: Specify the -force option to delete the mapping.

CMMVC5900E The FlashCopy mapping was not deleted because the mapping or consistency group is in the suspended state. The mapping or consistency group must be stopped first.

Explanation: The FlashCopy mapping was not deleted because the mapping or consistency group is in the suspended state. The mapping or consistency group must be stopped first.

Action: Stop the consistency group and then issue the command again.

CMMVC5901E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the preparing state.

Explanation: The FlashCopy mapping was not prepared because the mapping or consistency group is already in the preparing state.

Action: Not applicable.

CMMVC5902E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the prepared state.

Explanation: The FlashCopy mapping was not prepared because the mapping or consistency group is already in the prepared state.

Action: Not applicable.

CMMVC5903E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the copying state.

Explanation: The FlashCopy mapping was not prepared because the mapping or consistency group is already in the copying state.

Action: Not applicable.

CMMVC5904E The FlashCopy mapping was not prepared because the mapping or consistency group is already in the suspended state.

Explanation: The FlashCopy mapping was not prepared because the mapping or consistency group is already in the suspended state.

Action: Not applicable.

CMMVC5905E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the idle state. The mapping or consistency group must be prepared first.

Explanation: The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the idle state.

Action: Prepare the mapping or consistency group and then issue the command again.

CMMVC5906E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the preparing state.

Explanation: The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the preparing state.

Action: Not applicable.

CMMVC5907E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is already in the copying state.

Explanation: The FlashCopy mapping or consistency group was not started because the mapping or consistency group is already in the copying state.

Action: Not applicable.

CMMVC5908E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the stopped state. The mapping or consistency group must be prepared first.

Explanation: The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the stopped state.

Action: Prepare the mapping or consistency group and issue the command again.

CMMVC5909E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the suspended state.

Explanation: The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the suspended state.

Action: Not applicable.

CMMVC5910E The FlashCopy mapping or consistency group was not started because the mapping or consistency group is in the idle state.

Explanation: The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is in the idle state.

Action: Not applicable.

CMMVC5911E The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is in the preparing state.

Explanation: The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is in the preparing state.

Action: Not applicable.

CMMVC5912E The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is already in the stopped state.

Explanation: The FlashCopy mapping or consistency group was not stopped because the mapping or consistency group is already in the stopped state.

Action: Not applicable.

CMMVC5913E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the preparing state.

Explanation: The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the preparing state.

Action: Not applicable.

CMMVC5914E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the prepared state.

Explanation: The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the prepared state.

Action: Not applicable.

CMMVC5915E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the copying state.

Explanation: The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the copying state.

Action: Not applicable.

CMMVC5916E The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the suspended state.

Explanation: The properties of the FlashCopy mapping were not modified because the mapping or consistency group is in the suspended state.

Action: Not applicable.

CMMVC5917E The FlashCopy mapping was not created because there is no memory to create the bitmap.

Explanation: The FlashCopy mapping was not created because there is no memory to create the bitmap.

Action: Not applicable.

CMMVC5918E The FlashCopy mapping was not prepared because the I/O group is offline.

Explanation: The FlashCopy mapping was not prepared because the I/O group is offline.

Action: Not applicable.

CMMVC5919E The FlashCopy mapping or consistency group was not started because the I/O group is offline.

Explanation: The FlashCopy mapping or consistency group was not started because the I/O group is offline.

Action: Not applicable.

CMMVC5920E The FlashCopy mapping was not created because the consistency group is not idle.

Explanation: The FlashCopy mapping was not created because the consistency group is not idle.

Action: Not applicable.

CMMVC5921E The properties of the FlashCopy mapping were not modified because the consistency group is not idle.

Explanation: The properties of the FlashCopy mapping were not modified because the consistency group is not idle.

Action: Not applicable.

CMMVC5922E The FlashCopy mapping was not created because the destination virtual disk (VDisk) is too small.

Explanation: The FlashCopy mapping was not created because the destination virtual disk (VDisk) is too small.

Action: Specify a different VDisk and issue the command again.

CMMVC5923E The FlashCopy mapping was not created because the I/O group is offline.

Explanation: The FlashCopy mapping was not created because the I/O group is offline.

Action: Not applicable.

CMMVC5924E The FlashCopy mapping was not created because the source and target virtual disks (VDisks) are different sizes.

Explanation: The FlashCopy mapping was not created because the source and target virtual disks (VDisks) are different sizes.

Action: Specify a different source and target VDisk that are the same size and issue the command again.

CMMVC5925E The remote cluster partnership was not created because it already exists.

Explanation: The remote cluster partnership was not created because it already exists.

Action: Specify a different remote cluster partnership and issue the command again.

CMMVC5926E The remote cluster partnership was not created because there are too many partnerships.

Explanation: The remote cluster partnership was not created because there are too many partnerships.

Action: Not applicable.

CMMVC5927E The action failed because the cluster ID is not valid.

Explanation: The action failed because the cluster ID is not valid.

Action: Specify the correct cluster ID and issue the command again.

CMMVC5928E The action failed because the cluster name is a duplicate of another cluster.

Explanation: The action failed because the cluster name is a duplicate of another cluster.

Action: Specify a different cluster name and issue the command again.

CMMVC5929E The Remote Copy partnership was not deleted because it has already been deleted.

Explanation: The Remote Copy partnership was not deleted because it has already been deleted.

Action: Not applicable.

CMMVC5930E The Remote Copy relationship was not created because an entity that was specified in the command does not exist.

Explanation: The Remote Copy relationship was not created because an entity that was specified in the command does not exist.

Action: Specify the correct entity and issue the command again.

CMMVC5931E The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is locked.

Explanation: The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is locked.

Action: Unlock the master or auxiliary VDisk and issue the command again.

CMMVC5932E The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is a member of a FlashCopy mapping.

Explanation: The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is a member of a FlashCopy mapping.

Action: Not applicable.

CMMVC5933E The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is in the recovery I/O group.

Explanation: The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is in the recovery I/O group.

Action: Not applicable.

CMMVC5934E The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is in the router mode.

Explanation: The Remote Copy relationship was not created because the master or auxiliary virtual disk (VDisk) is in the router mode.

Action: Not applicable.

CMMVC5935E The action failed because an entity that was specified in the command does not exist.

Explanation: The action failed because an entity that was specified in the command does not exist.

Action: Specify the correct entity and issue the command again.

CMMVC5936E The action failed because an entity that was specified in the command does not exist.

Explanation: The action failed because an entity that was specified in the command does not exist.

Action: Specify the correct entity and issue the command again.

CMMVC5937E The action failed because an entity that was specified in the command does not exist.

Explanation: The action failed because an entity that was specified in the command does not exist.

Action: Specify the correct entity and issue the command again.

CMMVC5938E The Remote Copy consistency group was not deleted because the consistency group contains relationships. To delete the consistency group, the force option is required.

Explanation: The Remote Copy consistency group was not deleted because the consistency group contains relationships.

Action: Specify the `-force` option to delete the consistency group.

CMMVC5939E The action failed because the cluster is not in a stable state.

Explanation: The action failed because the cluster is not in a stable state.

Action: Not applicable.

CMMVC5940E The cluster that contains the auxiliary virtual disk (VDisk) is unknown.

Explanation: The cluster that contains the auxiliary virtual disk (VDisk) is unknown.

Action: Not applicable.

CMMVC5941E The cluster that contains the master virtual disk (VDisk) has too many consistency groups.

Explanation: The cluster that contains the master virtual disk (VDisk) has too many consistency groups.

Action: Not applicable.

CMMVC5942E The cluster that contains the auxiliary virtual disk (VDisk) has too many consistency groups.

Explanation: The cluster that contains the auxiliary virtual disk (VDisk) has too many consistency groups.

Action: Not applicable.

CMMVC5943E The specified relationship is not valid.

Explanation: The specified relationship is not valid.

Action: Specify the correct relationship and issue the command again.

CMMVC5944E The specified consistency group is not valid.

Explanation: The specified consistency group is not valid.

Action: Specify the correct consistency group and issue the command again.

CMMVC5945E The specified master cluster is not valid.

Explanation: The specified master cluster is not valid.

Action: Specify the correct master cluster and issue the command again.

CMMVC5946E The specified auxiliary cluster is not valid.

Explanation: The specified auxiliary cluster is not valid.

Action: Specify the correct auxiliary cluster and issue the command again.

CMMVC5947E The specified master virtual disk (VDisk) is not valid.

Explanation: The specified master virtual disk (VDisk) is not valid.

Action: Specify the correct master VDisk and issue the command again.

CMMVC5948E The specified auxiliary virtual disk (VDisk) is not valid.

Explanation: The specified auxiliary virtual disk (VDisk) is not valid.

Action: Specify the auxiliary VDisk and issue the command again.

CMMVC5949E The specified relationship is unknown.

Explanation: The specified relationship is unknown.

Action: Specify a different relationship and issue the command again.

CMMVC5950E The specified consistency group is unknown.

Explanation: The specified consistency group is unknown.

Action: Specify a different consistency group and issue the command again.

CMMVC5951E The operation cannot be performed because the relationship is not a stand-alone one.

Explanation: The operation cannot be performed because the relationship is not a stand-alone one.

Action: Not applicable.

CMMVC5952E The relationship and consistency group have different master clusters.

Explanation: The relationship and consistency group have different master clusters.

Action: Not applicable.

CMMVC5953E The relationship and group have different auxiliary clusters.

Explanation: The relationship and group have different auxiliary clusters.

Action: Not applicable.

CMMVC5954E The master and auxiliary virtual disks (VDisks) are different sizes

Explanation: The master and auxiliary virtual disks (VDisks) are different sizes

Action: Not applicable.

CMMVC5955E The maximum number of relationships has been reached.

Explanation: The maximum number of relationships has been reached.

Action: Not applicable.

CMMVC5956E The maximum number of consistency groups has been reached.

Explanation: The maximum number of consistency groups has been reached.

Action: Not applicable.

CMMVC5957E The master virtual disk (VDisk) is already in a relationship.

Explanation: The master virtual disk (VDisk) is already in a relationship.

Action: Specify a different master VDisk and issue the command again.

CMMVC5958E The auxiliary virtual disk (VDisk) is already in a relationship.

Explanation: The auxiliary virtual disk (VDisk) is already in a relationship.

Action: Specify a different auxiliary VDisk and issue the command again.

CMMVC5959E There is a relationship that already has this name on the master cluster.

Explanation: There is a relationship that already has this name on the master cluster.

Action: Specify a different name and issue the command again.

CMMVC5960E There is a relationship that already has this name on the auxiliary cluster.

Explanation: There is a relationship that already has this name on the auxiliary cluster.

Action: Specify a different name and issue the command again.

CMMVC5961E There is a consistency group that already has this name on the master cluster.

Explanation: There is a consistency group that already has this name on the master cluster.

Action: Specify a different name and issue the command again.

CMMVC5962E There is a consistency group that already has this name on the auxiliary cluster.

Explanation: There is a consistency group that already has this name on the auxiliary cluster.

Action: Specify a different name and issue the command again.

CMMVC5963E No direction has been defined.

Explanation: No direction has been defined.

Action: Not applicable.

CMMVC5964E The copy priority is not valid.

Explanation: The copy priority is not valid.

Action: Not applicable.

CMMVC5965E The virtual disks (VDisks) are in different I/O groups on the local cluster.

Explanation: The virtual disks (VDisks) are in different I/O groups on the local cluster.

Action: Not applicable.

CMMVC5966E The master virtual disk (VDisk) is unknown.

Explanation: The master virtual disk (VDisk) is unknown.

Action: Specify a different master VDisk and issue the command again.

CMMVC5967E The auxiliary virtual disk (VDisk) is unknown.

Explanation: The auxiliary virtual disk (VDisk) is unknown.

Action: Specify a different auxiliary VDisk and issue the command again.

CMMVC5968E The relationship cannot be added because the states of the relationship and the consistency group do not match.

Explanation: The relationship cannot be added because the states of the relationship and the consistency group do not match.

Action: Not applicable.

CMMVC5969E The Remote Copy relationship was not created because the I/O group is offline.

Explanation: The Remote Copy relationship was not created because the I/O group is offline.

Action: Not applicable.

CMMVC5970E The Remote Copy relationship was not created because there is not enough memory.

Explanation: The Remote Copy relationship was not created because there is not enough memory.

Action: Not applicable.

CMMVC5971E The operation was not performed because the consistency group contains no relationships.

Explanation: The operation was not performed because the consistency group contains no relationships.

Action: Not applicable.

CMMVC5972E The operation was not performed because the consistency group contains relationships.

Explanation: The operation was not performed because the consistency group contains relationships.

Action: Not applicable.

CMMVC5973E The operation was not performed because the consistency group is not synchronized.

Explanation: The operation was not performed because the consistency group is not synchronized.

Action: Not applicable.

CMMVC5974E The operation was not performed because the consistency group is offline.

Explanation: The operation was not performed because the consistency group is offline.

Action: Not applicable.

CMMVC5975E The operation was not performed because the cluster partnership is not connected.

Explanation: The operation was not performed because the cluster partnership is not connected.

Action: Not applicable.

CMMVC5976E The operation was not performed because the consistency group is in the freezing state.

Explanation: The operation was not performed because the consistency group is in the freezing state.

Action: Not applicable.

CMMVC5977E The operation was not performed because it is not valid given the current consistency group state.

Explanation: The operation was not performed because it is not valid given the current consistency group state.

Action: Not applicable.

CMMVC5978E The operation was not performed because the relationship is not synchronized.

Explanation: The operation was not performed because the relationship is not synchronized.

Action: Not applicable.

CMMVC5979E The operation was not performed because the relationship is offline.

Explanation: The operation was not performed because the relationship is offline.

Action: Not applicable.

CMMVC5980E The operation was not performed because the master and auxiliary clusters are not connected.

Explanation: The operation was not performed because the master and auxiliary clusters are not connected.

Action: Not applicable.

CMMVC5981E The operation was not performed because the relationship is in the freezing state.

Explanation: The operation was not performed because the relationship is in the freezing state.

Action: Not applicable.

CMMVC5982E The operation was not performed because it is not valid given the current relationship state.

Explanation: The operation was not performed because it is not valid given the current relationship state.

Action: Not applicable.

CMMVC5983E The dump file was not created. The file system might be full.

Explanation: The dump file was not created. The file system might be full.

Action: Not applicable.

CMMVC5984E The dump file was not written to disk. The file system might be full.

Explanation: The dump file was not written to disk. The file system might be full.

Action: Not applicable.

CMMVC5986E The tracing of I/O operations was not started because the virtual disk (VDisk) or managed disk (MDisk) failed to return statistics.

Explanation: The tracing of I/O operations was not started because the virtual disk (VDisk) or managed disk (MDisk) failed to return statistics.

Action: Not applicable.

CMMVC5987E Address is not valid.

Explanation: Address is not valid.

Action: Specify a different address and issue the command again.

CMMVC5988E This command should not be issued if you are logged in with a root user ID. Use the admin userid.

Explanation: This command should not be issued if you are logged in with a root user ID. Use the admin userid.

Action: Log off of the root user ID and log in as admin.

CMMVC5989E The FlashCopy consistency group was not started as there are no FlashCopy mappings within the group.

Explanation: The FlashCopy consistency group was not started as there are no FlashCopy mappings within the group.

Action: Include FlashCopy mappings into the consistency group or start a FlashCopy consistency group that includes mappings.

CMMVC5990E The FlashCopy consistency group was not stopped as there are no FlashCopy mappings within the group.

Explanation: The FlashCopy consistency group was not stopped as there are no FlashCopy mappings within the group.

Action: Not applicable.

CMMVC5991E The FlashCopy consistency group was not stopped as there are no FlashCopy mappings within the group.

Explanation: The FlashCopy consistency group was not stopped as there are no FlashCopy mappings within the group.

Action: Not applicable.

CMMVC5992E The Remote Copy consistency group was not stopped as there are no Remote Copy relationships within the group.

Explanation: The Remote Copy consistency group was not stopped as there are no Remote Copy relationships within the group.

Action: Not applicable.

CMMVC5996E The specific upgrade package cannot be installed over the current version.

Explanation: The specific upgrade package cannot be installed over the current version.

Action: Not applicable.

CMMVC5997E The action failed because the capacity of the MDisk is smaller than the extent size of the MDisk group.

Explanation: The action failed because the capacity of the MDisk is smaller than the extent size of the MDisk group.

Action:

- Select an MDisk that has a capacity equal to or larger than the extent size of the MDisk group.
-

- Select a smaller extent size but one that is at least equal in size to the smallest MDisk in the MDisk group.

Note: You can select a smaller extent size only if you are creating an MDisk group. After an MDisk is created, you cannot change the extent size.

CMMVC5998E This command can not be run on a node that is in service mode.

Explanation: This command can not be run on a node that is in service mode.

Action: Not applicable.

CMMVC5998W The virtualized storage capacity exceeds the amount that you are licensed to use.

Explanation: You have attempted to create more virtualized storage capacity than you are licensed to use.

Action: Either reduce the amount of virtualized storage capacity currently in use, or license additional storage capacity, then retry the command.

CMMVC5999W Featurization for this facility has not been enabled.

Explanation: Featurization for this facility has not been enabled.

Action: Not applicable.

CMMVC5999E Undefined error message.

Explanation: Undefined error message.

Action: Not applicable.

CMMVC6000W Featurization for this facility has not been enabled.

Explanation: Featurization for this facility has not been enabled.

Action: Not applicable.

CMMVC6001E The FlashCopy consistency group was not started as there are no FlashCopy mappings within the group.

Explanation: The FlashCopy consistency group was not started as there are no FlashCopy mappings within the group.

Action: Create a FlashCopy within the appropriate group.

CMMVC6002E This command can only be run on a node that is in service mode.

Explanation: This command can only be run on a node that is in service mode.

Action: Not applicable.

CMMVC6003E This command can not be run on a node that is in service mode.

Explanation: This command can only be run on a node that is in service mode.

Action: Not applicable.

CMMVC6004E The delimiter value, %1, is invalid.

Explanation: The delimiter value, %1, is invalid.

Action: Specify a different delimiter.

CMMVC6005E The view request failed as the specified object is not a member of an appropriate group.

Explanation: A view was request on an object that has been incorrectly initialized.

Action: Ensure that the object is correctly initialized before resubmitting the view request.

CMMVC6006E The managed disk (MDisk) was not deleted because the resource was busy.

Explanation: An attempt was made to delete an MDisk from a MDisk group that is being used as a source and destination for migration operations.

Action: Ensure that the MDisk group is not being used for migration operations before issuing the command again.

CMMVC6009E Unable to malloc a block of memory in which to copy the returned data.

Explanation: The command line was unable to allocate a block of memory in which to copy the results of the query.

Action: Free up some memory and issue the command again.

CMMVC6010E Unable to complete the command as there are insufficient free extents.

Explanation: There are not enough free extents to meet the request.

Action: Not applicable.

CMMVC6011E At least one remote cluster partnership has been found. This upgrade package cannot be applied to the current code level until all remote cluster partnerships are deleted.

Explanation: An attempt was made to apply software when a Remote Copy relationship to a remote cluster exists.

Action: Delete the Remote Copy relationship to the remote clusters and issue the command again.

CMMVC6012W The virtualized storage capacity is approaching the amount that you are licensed to use.

Explanation: The action you have requested has been completed. However, you are approaching the limits permitted by the license you have purchased.

Action: Subsequent actions may require that you increase your licensed limits.

CMMVC6013E The command failed because there is a consistency group mismatch on the aux cluster.

Explanation: The action has failed as there was a difference in attributes between the Remote Copy consistency groups involved.

Action: Ensure that the attributes of the two Remote Copy consistency groups match before resubmitting the command.

CMMVC6014E The command failed because the requested object is either unavailable or does not exist.

Explanation: The command failed because the requested object is either unavailable or does not exist.

Action: Ensure that all parameters have been correctly entered. If this is the case the determine why the object is unavailable, then issue the command again.

CMMVC6015E A delete request is already in progress for this object.

Explanation: A delete request is already in progress for this object.

Action: Not applicable.

CMMVC6016E The action failed as there would be, or are, no more disks in the MDisk group.

Explanation: The action failed as there would be, or are, no more disks in the I/O group.

Action: Ensure that all parameters have been correctly entered.

CMMVC6017E %1 contains invalid characters. Ensure that all characters are ASCII.

Explanation: The CLI will only accept ASCII input.

Action: Ensure that all input to the CLI is ASCII, then resubmit the command.

CMMVC6019E The software upgrade failed as a node pended while the upgrade was in progress.

Explanation: The software upgrade failed as a node pended while the upgrade was in progress.

Action: Ensure that all nodes are online and available before restarting the upgrade process.

CMMVC6020E The software upgrade failed as the system was unable to distribute the software package to all nodes.

Explanation: The software upgrade failed as the system was unable to distribute the software package to all nodes.

Action: Ensure that all nodes are correctly zoned and that all nodes are online and can see the other nodes in the cluster. You may also want to check the error log.

CMMVC6021E The system is currently busy performing another request. Please try again later.

Explanation: The requested action failed as the system is processing another request.

Action: Wait a while before resubmitting the request.

CMMVC6022E The system is currently busy performing another request. Please try again later.

Explanation: The requested action failed as the system is processing another request.

Action: Wait a while before resubmitting the request.

CMMVC6023E The system is currently busy performing another request. Please try again later.

Explanation: The requested action failed as the system is processing another request.

Action: Wait a while before resubmitting the request.

CMMVC6024E The auxiliary VDisk entered is invalid.

Explanation: The auxiliary VDisk is entered as a parameter in the CLI is not a valid auxiliary VDisk.

Action: Select a valid auxiliary VDisk and issue the command again.

CMMVC6025E The RC consistency group Master cluster is not the local cluster.

Explanation: The auxiliary VDisk is entered as a parameter in the CLI is not a valid auxiliary VDisk.

Action: Resubmit the command with a consistency group that belongs to the local cluster.

CMMVC6026E The RC consistency group is not in the stopped state.

Explanation: The action failed as the Remote Copy consistency group is not in the stopped state.

Action: Ensure that the Remote Copy consistency group is in the stopped state before resubmitting the command.

CMMVC6027E The RC consistency group is not the primary master.

Explanation: The RC consistency group requested in the command is not the Remote Copy primary master.

Action: Ensure that the parameters have been entered correctly on the command line.

CMMVC6028E This upgrade package cannot be applied to the current software level because it contains changes to the cluster state and there are remote cluster partnership defined.

Explanation: The action failed because there is a connected remote cluster. The upgrade cannot be applied because it would render the remote cluster at a different code level to the remote cluster.

Action: Ensure that the cluster partnership is deconfigured before resubmitting the command. Ensure that you deconfigure the remote cluster and upgrade the code on it before reconfiguring the cluster partnership.

CMMVC6029E All nodes must have identical code level before a concurrent code upgrade can be performed.

Explanation: The concurrent upgrade failed as two or more nodes were at differing code levels. All nodes must be at the same code level before a software upgrade can be performed.

Action: Use the service mode to bring all nodes to the same level before resubmitting the concurrent upgrade.

CMMVC6030E The operation was not preformed because the FlashCopy mapping is part of a consistency group. The action must be performed at the consistency group level.

Explanation: An attempt was made to stop a FlashCopy mapping. This failed as the FlashCopy mapping is part of a consistency group.

Action: Issue the stop command to the FlashCopy consistency group. This will stop all FlashCopies within that group that are in progress.

CMMVC6031E The operation was not preformed because the FlashCopy consistency group is empty.

Explanation: An attempt was made to prestart an empty FlashCopy consistency group.

Action: Not applicable.

CMMVC6032E The operation was not preformed because one or more of the entered parameters is invalid for this operation.

Explanation: An invalid parameter was entered for the command.

Action: If attempting to change the I/O group to which the VDisk belongs, ensure that the VDisk is not already a part of the group.

CMMVC6034E The action failed because the maximum number of objects has been reached.

Explanation: The action failed because the maximum number of objects has been reached.

Action: Not applicable.

CMMVC6071E The virtual disk is already mapped to a host. To create additional virtual disk-to-host mappings, you must use the command-line interface.

Explanation: The virtual disk is already mapped to a host.

Action: Use the command-line interface to create additional mappings.

CMMVC6073E The maximum number of files has been exceeded.

Explanation: The maximum number of files has been exceeded.

Action: Not applicable.

CMMVC6074E The command failed as the extent has already been assigned.

Explanation: The command failed as the extent has already been assigned.

Action: Assign a different extent and issue the command again.

CMMVC6075E The expand failed as the last extent is not a complete extent.

Explanation: The expand failed as the last extent is not a complete extent.

Action: Assign a different extent and issue the command again.

CMMVC6076E The command failed due to an error while flushing the VDisk.

Explanation: The command failed due to an error while flushing the VDisk.

Action: Not applicable.

CMMVC6100E -option not consistent with action

Explanation: The specified option is not supported for this action.

Action: Remove the option and issue the command again.

CMMVC6101E -option not consistent with -option

Explanation: The two specified options cannot be used together.

Action: Remove one of the options and issue the command again.

CMMVC6102E -option and -option are alternatives

Explanation: The two specified options are alternatives and cannot be used together.

Action: Remove one of the options and issue the command again.

CMMVC6103E Problem with file-name: details

Explanation: A problem occurred when opening a file. Determine the cause of the problem and correct it before trying again.

Action: Correc the problem and then issue the command again.

CMMVC6104E Action name not run

Explanation: An unexpected error occurred. Contact your IBM service representative.

Action: Contact your IBM service representative.

CMMVC6105E Different names for source (name) and target (name) clusters

Explanation: The backup configuration cannot be restored to the target cluster because the source and target cluster have different names.

Action: Perform one of the following actions: (1) Use a different backup configuration. (2) Delete the cluster and recreate it with the same name as that stored in the backup configuration file.

CMMVC6106W Target cluster has non-default id_alias alias

Explanation: The id_alias of the target cluster has a non-default target. Clusters should have a default value. The non-default value suggests the cluster is customized and is not suitable for restoration. Restoration changes the id_alias.

Action: Change the id_alias to a default value and issue the command again.

CMMVC6107E x io_grp objects in target cluster; y are required

Explanation: The number of I/O groups in the target cluster is not sufficient to accommodate the I/O groups defined in the backup configuration file. Determine why there are not enough I/O groups.

Action: Correct the problem and issue the command again.

CMMVC6108I Disk controller system with a WWNN of *wwnn* found.

Explanation: A disk controller system has been found with the required WWNN.

Action: Not applicable.

CMMVC6109E Disk controller system with a WWNN of *wwnn* not available.

Explanation: A disk controller system has been found with the required WWNN. Ensure that the required disk controller system is available to the cluster.

Action: Ensure that the required disk controller system is available to the cluster and issue the command again.

CMMVC6110E Bad code level *level*

Explanation: An unexpected error occurred.

Action: Report the details to your IBM service representative.

CMMVC6111E Unable to determine code_level for the cluster from *level*

Explanation: The code level of the cluster could not be determined. The code level should be of the form, x.y.z, where x, y, and z are integers.

Action: If you cannot determine the cause of the problem, contact your IBM service representative.

CMMVC6112W *object-type object-name* has a default name

Explanation: An object in the cluster has a default name. This can cause problems when restoring a cluster as default names are changed during restoration. Object IDs are also changed during restoration.

Action: Choose an appropriate name for each object in the cluster and issue the command again.

CMMVC6113E Remote command failed with return code: *details*

Explanation: An attempt to run a command remotely failed using secure communications.

Action: Determine the cause of the problem and issue the command again.

CMMVC6114E No help for action *action*

Explanation: There is no help for the requested topic.

Action: Not applicable.

CMMVC6115W Feature *property* mismatch: *value1* expected; *value2* found

Explanation: The features in the backup configuration file and the target cluster do not match. There should be an exact match between the two. Nevertheless, the restore of the configuration can continue.

Action: Not applicable.

CMMVC6116I Feature *property* checks out

Explanation: The features in the backup configuration file and the target cluster are an exact match.

Action: Not applicable.

CMMVC6117E *fix-or-feature* is not available

Explanation: An unexpected error occurred.

Action: Contact your IBM service representative.

CMMVC6118I *type* with *property value* [and *property value*] found

Explanation: An object in the cluster has been found with the correct properties.

Action: Not applicable.

CMMVC6119E *type* with *property value* [and *property value*] not found

Explanation: An object in the cluster with the correct properties has not been found. Restoration cannot proceed without the object.

Action: Determine why the object cannot be found. Ensure that the object is available and issue the command again.

CMMVC6120E Target is not the configuration node

Explanation: The target is not the configuration node.

Action: Redirect the action against the configuration node and issue the command again.

CMMVC6122E No *type* with *property value* present in table

Explanation: An unexpected error occurred.

Action: Contact your IBM service representative.

CMMVC6123E No *property* for *type name*

Explanation: An unexpected error occurred.

Action: Contact your IBM service representative.

CMMVC6124E No *type* with *property value*

Explanation: An unexpected error occurred.

Action: Contact your IBM service representative.

CMMVC6125E No *unique ID* for *type name*

Explanation: An unexpected error occurred.

Action: Contact your IBM service representative.

CMMVC6126E No *type* with *unique ID value*

Explanation: An unexpected error occurred.

Action: Contact your IBM service representative.

CMMVC6128W *details*

Explanation: The files in the directory could not be listed.

Action: Determine why they could not be listed and correct the problem and issue the command again.

CMMVC6129E *vdisk-to-host map objects have vdisk_UID values that are not consistent*

Explanation: All of the VDisk-to-host map objects do not have the same number for the VDisk LUN instance. Therefore, there is a possibility the backup configuration file is corrupt. The LUN instance number should be the same for all VDisk-to-host map objects that are associated with a specific VDisk. The LUN instance number is incorporated into the VDisk_UID property.

Action: Determine why the LUN instance number is not the same and correct the problem and issue the command again.

CMMVC6130W *Inter-cluster property will not be restored*

Explanation: The restoration of inter-cluster objects is not supported.

Action: Not applicable.

CMMVC6131E *No location cluster information*

Explanation: An unexpected error occurred.

Action: Contact your IBM service representative.

CMMVC6132E *The property of type name is value1; should be value2*

Explanation: An object has a property with an incorrect value. The property most likely reflects the state of the object.

Action: Change the state to the required value and issue the command again.

CMMVC6133E *Required type property property not found*

Explanation: An unexpected error occurred.

Action: Contact your IBM service representative.

CMMVC6134E *No argument for -option*

Explanation: No argument has been supplied for an option that requires an argument.

Action: Supply an argument and issue the command again.

CMMVC6135E *Invalid argument value for -option*

Explanation: An argument has been supplied for an option that is not valid.

Action: Supply an valid argument and retry.

CMMVC6138E *-option is required*

Explanation: An option is missing. The option might be listed as optional, but circumstances make the option mandatory.

Action: Supply the option and issue the command again.

CMMVC6139E *Object (description) is required*

Explanation: The command requires that an object or target be specified. An example might be the IP address for a configuration node or a name.

Action: Specify the missing object or target and issue the command again.

CMMVC6140E *No default name for type type*

Explanation: An unexpected error occurred.

Action: Contact you IBM service representative.

CMMVC6141E *-option does not contain any argument*

Explanation: An argument has been supplied for an option that does not contain any.

Action: Remove the argument and issue the command again.

CMMVC6142W *Existing object-type object-name has a non-default name*

Explanation: An object in the target cluster has a non-default name. The non-default value suggests the cluster is customized and is not suitable for restoration.

Action: Specify a different non-default name and issue the command again.

CMMVC6143E *Required configuration file file-name does not exist*

Explanation: A file is missing that is critical for successful operation.

Action: Not applicable.

CMMVC6144W *Object with default name name restored as substitute-name*

Explanation: An object with a default name has been restored with a different name. Ensure that you account for this name change when using the restored cluster in the future. To avoid this problem in the future, choose

an appropriate name for each object in the cluster.

Action: Choose an appropriate name for each object in the cluster.

CMMVC6145I Use the restore -prepare command first

Explanation: This advisory is given prior to CMMVC6103E when an intermediate file is missing, presumed not created.

Action: Not applicable.

CMMVC6146E Problem parsing *object-type* data: *line*

Explanation: An unexpected error occurred.

Action: Contact your IBM service representative.

CMMVC6147E *type name* has a name beginning with *prefix*

Explanation: An object has been encountered that has a name beginning with a reserved prefix. The only valid reason for an object with this kind of name is: a restoration command did not complete successfully.

Action: Ensure that no objects use the reserved prefix in their name, and issue the command again.

CMMVC6148E Target cluster has *n-actual* objects of *type type* instead of *n-required*

Explanation: The target cluster does not have the required number of certain types of objects.

Action: Correct the problem and issue the command again.

CMMVC6149E An action is required

Explanation: An action is required to run the command.

Action: Supply an action, and issue the command again.

CMMVC6150E The *action*, *action*, is not valid

Explanation: An action has been supplied that is not valid.

Action: Supply a valid action, and issue the command again.

CMMVC6151E The *-option* option is not valid

Explanation: An option has been supplied that is not valid.

Action: Supply a valid action, and issue the command again.

CMMVC6152E *vdisk name instance number instance* is not valid

Explanation: The VDisk could not be restored because the instance number (in hex) is not valid.

Action: Contact your IBM service representative.

CMMVC6153E *object* not consistent with *action*

Explanation: The specified object is not supported for the action.

Action: Remove the object and issue the command again.

CMMVC6154E Required *object-type* property *property-name* has a null value

Explanation: An unexpected error occurred.

Action: Contact your IBM service representative.

CMMVC6200E The action failed because of incompatible software.

Explanation: The software package is incompatible with the software on one or more switches.

Action: Determine from the additional information where the incompatibility lies. Upgrade those components to versions compatible with the new SAN Volume Controller for Cisco MDS 9000 package. Contact your IBM Service Representative for assistance.

CMMVC6201E The node could not be added, because incompatible software: status code [%1].

Explanation: The node that was being added need to have its software package replaced by the software package in use by the cluster, however the cluster's software package is not compatible with the switch or Cisco MDS 9000 Caching Services Module where the node resides.

Action: Either change the code levels of the components identified to be compatible with the cluster software level, upgrade the cluster to be compatible with those components, or choose another node to add to the cluster. Contact your IBM Service Representative to help determine the levels required.

CMMVC6202E The cluster was not modified because the IP address is not valid.

Explanation: The IP address specified is not valid, and the cluster was not modified.

Action: Not applicable.

CMMVC6203E The action failed because the directory that was specified was not one of the following directories: /dumps, /dumps/iostats, /dumps/iotrace, /dumps/feature, /dumps/configs, /dumps/elogs, /dumps/ec or /dumps/pl.

Explanation: The specified action failed because the specified directory not /dumps, /dumps/iostats, /dumps/iotrace, /dumps/feature, /dumps/configs, /dumps/elogs, /dumps/ec, or /dumps/pl.

Action: Specify /dumps, /dumps/iostats, /dumps/iotrace, /dumps/feature, /dumps/configs, /dumps/elogs, /dumps/ec, or /dumps/pl, and issue the command again.

Appendix A. Accessibility

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

Features:

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

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

Navigating by keyboard:

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

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

Accessing the publications:

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

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

Related topics:

- “Related publications” on page v

Appendix B. SAN Volume Controller CLI compatibility

This guide introduces the IBM® TotalStorage® SAN Volume Controller Storage Software for Cisco MDS 9000 Command-Line Interface (CLI).

SAN Volume Controller Storage Software for Cisco MDS 9000 retains compatibility with SAN Volume Controller 1.1.0 CLI with the exceptions noted below. The commands and parameters that you can use with SAN Volume Controller 1.1.0 CLI that do not appear in this guide are accepted but have no effect. Commands that you can use with SAN Volume Controller 1.1.0 that are not accepted by SAN Volume Controller Storage Software for Cisco MDS 9000 are:

- svcservicetask applysoftware
- svcservicemodetask applysoftware

You cannot specify /home/admin/upgrade as a directory filter with the following commands:

- svctask cleardumps
- svctask cpdumps
- svcservicemodetask clean

Related topics:

- “Related publications” on page v

Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Related topics:

- "Trademarks"

Trademarks

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

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

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

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

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

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

Glossary

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

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

The following cross-references are used in this glossary:

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

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

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

A

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

C

cluster. In SAN Volume Controller Storage Software for Cisco MDS 9000, 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 Remote Copy relationship, a copy of a secondary virtual disk (VDisk) that is identical to the primary VDisk from the viewpoint of a host system, even if a power failure occurred while I/O activity was in progress.

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

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.

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

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

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

F

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

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

FC. See *fibre channel*.

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

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

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

host ID. In SAN Volume Controller Storage Software for Cisco MDS 9000, a numeric identifier assigned to a group of host fibre-channel ports for the purpose of logical unit number (LUN) mapping. For each host ID,

there is a separate mapping of Small Computer System Interface (SCSI) IDs to virtual disks (VDisks).

host zone. A zone defined in the storage area network (SAN) fabric in which the hosts can and address the SAN Volume Controller Storage Software for Cisco MDS 9000s.

I

IBM 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). A storage server product.

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

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

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

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

node rescue. In SAN Volume Controller Storage Software for Cisco MDS 9000, the process by which a SAN Volume Controller Storage Software for Cisco MDS 9000 node that has no valid software installed on

its hard disk drive can copy the software from another SAN Volume Controller Storage Software for Cisco MDS 9000 node connected to the same fibre-channel fabric.

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 Remote Copy, 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 Storage Software for Cisco MDS 9000, 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.

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

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. A form of storage array in which two or more identical copies of data are maintained on separate media. (S)

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. 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)

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

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

S

SAN. See *storage area network*.

SAN Volume Controller Storage Software for Cisco MDS 9000 fibre-channel port fan in. The number of hosts that can see any one SAN Volume Controller Storage Software for Cisco MDS 9000 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 Remote Copy, the virtual disk (VDisk) in a relationship that contains a copy of data written by the host application to the primary VDisk.

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

SNMP. See *Simple Network Management Protocol*.

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

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

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

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

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

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 which uniquely defines each hardware element and firmware element in the subsystem.

W

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

WWNN. See *worldwide node name*.

WWPN. See *worldwide port name*.

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

Index

A

- accessibility 327
 - keyboard 327
 - shortcut keys 327
- adding
 - nodes 30
- analyzing error logs 36, 262

C

- CLI (command-line interface)
 - preparing SSH client systems 2
- clusters
 - commands 5
 - diagnostic and service aid
 - commands 29
 - dumping configurations 12, 18, 270
 - modifying 9
 - shutting down 24
- command-line interface (CLI)
 - preparing SSH clients 2
- commands
 - nodes
 - Remote Copy 105
- configuring
 - error logs 34, 35, 263, 264
 - secure shell (SSH) 3
 - SSH (secure shell) 3
- consistency groups, FlashCopy 92
 - deleting 96
 - starting 99
 - stopping 103
- controllers
 - command 166, 291
- creating
 - FlashCopy
 - consistency groups 89
 - mappings 90, 94
 - managed disk groups 73
 - Remote Copy
 - consistency groups 112

D

- deleting
 - FlashCopy
 - mappings 97
 - managed disk groups 77
 - nodes 19, 37
 - Remote Copy
 - consistency groups 117
- disability 327
- disks
 - migrating 136
- dump files
 - listing 158, 164, 189, 211, 213, 278, 282, 283, 284, 285, 286
- dumping cluster configurations 164

F

- filtering
 - FlashCopy
 - consistency groups 191
 - mappings 193, 196, 197
- FlashCopy
 - command 85

G

- glossary 333

H

- HBAs (host bus adapters)
 - configuration 203
- hosts
 - adding ports to 42
 - commands 41
 - creating 45
 - deleting 47
 - deleting ports from 48
 - viewing 204

I

- I/O groups
 - renaming 10
- information
 - center v
 - commands 153

K

- keyboard 327
- shortcut keys 327

M

- managed disk (MDisk) groups
 - viewing 226
- managed disks (MDisks)
 - removing from a managed disk
 - group 75
 - viewing 217
- mappings, FlashCopy
 - starting 101
 - stopping 104
- marking errors as fixed 33, 265
- MDisk (managed disk) groups
 - adding
 - managed disks 70
 - command 69
 - renaming 72
- MDisks (managed disks)
 - command 79
 - including 81
 - renaming 80

- MDisks (managed disks) (*continued*)
 - setting quorum 82
- messages
 - CLI (command-line interface) 295
- migrating 231
- migration 129
- modifying
 - FlashCopy
 - consistency groups 86
 - mappings 87
 - Remote Copy
 - consistency groups 107, 127

N

- nodes
 - adding 6
 - renaming 11
 - shutting down 24
 - viewing 232
- notices 331

O

- ordering publications vii

P

- partnerships, Remote Copy
 - creating 110
 - deleting 116
 - modifying 106
- preparing
 - SSH client system
 - overview 1
 - to issue CLI commands 2
- publications
 - ordering vii

R

- related information v
- relationships, Remote Copy
 - creating 113
 - deleting 118
 - modifying 108, 128
 - starting 121
 - stopping 125

S

- secure shell (SSH) 3
 - client system
 - overview 1
 - preparing to issue CLI
 - commands 2
 - configuring 3
- security
 - overview 3

- service
 - mode
 - command 271
 - information command 277
- shortcut keys 327
- SSH (secure shell) 3
 - client system
 - overview 1
 - preparing to issue CLI
 - commands 2
 - configuring 3
- starting
 - Remote Copy
 - consistency groups 119
 - statistics collections 23
- statistics
 - stopping 27
- stopping
 - Remote Copy
 - consistency groups 123
- support
 - Web sites viii

T

- throttling 52
- tracing
 - commands 135, 261
- trademarks 332
- troubleshooting
 - using error logs 39, 266

V

- VDisks (virtual disks)
 - commands 51
 - creating 57
 - creating virtual disk-to-host
 - mappings 61
 - deleting 63
 - deleting virtual disk-to-host
 - mappings 65
 - migrating 132
 - shrinking 66
 - viewing 249, 252, 254, 256, 258
- viewing
 - clusters 160
 - featurization settings 267
 - vital product data (VPD) 280
 - hosts
 - port details 206
 - I/O groups 208
 - nodes 287
 - vital product data (VPD) 288
 - VPD (vital product data) 236
 - Remote Copy
 - consistency groups 239
 - relationships 242, 245, 247

W

- Web sites viii

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

IBM TotalStorage SAN Volume Controller
Storage Software for Cisco MDS 9000
Command-Line Interface User's Guide
Version 1.1

Publication No. SC26-7555-00

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

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

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

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

Please tell us how we can improve this book:

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

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

Name

Address

Company or Organization

Phone No.



Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

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

POSTAGE WILL BE PAID BY ADDRESSEE

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



Fold and Tape

Please do not staple

Fold and Tape



Printed in USA

SC26-7555-00



Spine information:

IBM TotalStorage SAN Volume
Controller
Storage Software for Cisco MDS
9000



Command-Line Interface User's Guide

Version 1.1