IBM System Storage SAN Volume Controller

Command-Line Interface User's Guide

Version 4.1.0

SC26-7903-00

IBM System Storage SAN Volume Controller

Command-Line Interface User's Guide

Version 4.1.0

SC26-7903-00

Seventh Edition (June 2006)

This edition replaces SC26-7544-05 and all previous versions of SC26-7544.

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

Contents

Tables

About this guide

The IBM System Storage SAN Volume Controller Command-Line Interface (CLI) User's Guide provides information that helps you configure and use the IBM System Storage SAN Volume Controller.

Who should use this guide

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

Summary of Changes

This document contains terminology, maintenance, and editorial changes.

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

Summary of Changes for SC26-7903-00 SAN Volume Controller Command-Line Interface Guide

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

New information

This topic describes the changes to this guide since the previous edition, SC26-7544-05. The following sections summarize the changes that have since been implemented from the previous version.

This version includes the following new information:

- Adding the following new SAN Volume Controller topics:
  - enablecli
  - lscopystatus
  - lsdiscoverystatus
  - Audit log commands
  - dumpauditlog
  - lsauditlogdumps
  - catauditlog

Changed information

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

Note: In this edition, the brand ″IBM TotalStorage″ has been replaced with ″IBM System Storage″.

The following topics have been updated:

- chhost (added new parameter)

- mkhost added new parameter)
- lshost (added new parameter)
- startstats (updated description to include new statistics)
- addhostiogrp (corrected syntax diagram)
- chfcmap (corrected syntax diagram)
- chrcrelationship (corrected syntax diagram)
- chvdisk (corrected syntax diagram)
- lsfabric (corrected syntax diagram)
- rmhostiogrp (corrected syntax diagram)
- Configuring the PuTTY session for the CLI
- chcluster (updated descriptions to include new feature)
- stopcluster (updated descriptions to include new statistics)
- applysoftware (added new parameter and messages)
- mkvdisk (corrected syntax diagram)
- mkrcrelationship (added new parameter)
- migratetoimage (updated descriptions)

Emphasis

Different typefaces are used in this guide to show emphasis.

The following typefaces are used to show emphasis:

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

SAN Volume Controller library and related publications

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

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

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

SAN Volume Controller library

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

http://www.ibm.com/storage/support/2145

Title	Description	Order number
IBM System Storage SAN Volume Controller: CIM agent Developer's Reference	This reference guide describes the objects and classes in a Common Information Model (CIM) environment.	GA32-0552
IBM System Storage SAN Volume Controller: Command-Line Interface User's Guide	This guide describes the commands that you can use from the SAN Volume Controller command-line interface (CLI).	SC26-7903
IBM System Storage SAN Volume Controller: Configuration Guide	This guide provides guidelines for configuring your SAN Volume Controller.	SC26-7902
IBM System Storage SAN Volume Controller: Host Attachment Guide	This guide provides guidelines for attaching the SAN Volume Controller to your host system.	SC26-7905
IBM System Storage SAN Volume Controller: Installation Guide	This guide includes the instructions the service representative uses to install the SAN Volume Controller.	GC26-7900
IBM System Storage SAN Volume Controller: Planning Guide	This guide introduces the SAN Volume Controller and lists the features you can order. It also provides guidelines for planning the installation and configuration of the SAN Volume Controller.	GA32-0551
IBM System Storage SAN Volume Controller: Service Guide	This guide includes the instructions the service representative uses to service the SAN Volume Controller.	GC26-7901
IBM System Safety Notices	This guide contains the danger and caution notices for the SAN Volume Controller. The notices are shown in English and in numerous other languages.	G229-9054
IBM System Storage Master Console for SAN File System and SAN Volume Controller: Installation and User's Guide	This guide includes the instructions on how to install and use the SAN Volume Controller Console	GC30-4090

Other IBM publications

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

Title	Description	Order number
IBM System Storage Multipath Subsystem Device Driver: User's Guide	This guide describes the IBM System Storage Multipath Subsystem Device Driver Version 1.5 for TotalStorage Products and how to use it with the SAN Volume Controller. This publication is referred to as the *IBM System Storage Multipath Subsystem Device Driver: User's Guide*.	SC30-4131

Related Web sites

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

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

How to order IBM publications

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

The IBM publications center

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

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

Publications notification system

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

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

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

How to send your comments

Your feedback is important to help us provide the highest quality information. If you have any comments about this book or any other documentation, you can submit them in one of the following ways:

- e-mail

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

  starpubs@us.ibm.com

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

- Mail

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

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

Syntax diagrams

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

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

Element	Syntax	Description
Main path line	>>><>() () ()	>>Begins on the left with double arrowheads ()>< and ends on the right with two arrowheads facing each other (). If a diagram is longer than one line, each line to be continued ends with a single> arrowhead () and the next line begins with a single arrowhead. Read the diagrams from left–to–right, top–to–bottom, following the main path line.
Keyword		Represents the name of a command, flag, parameter, or argument. A keyword is not in italics. Spell a keyword exactly as it is shown in the syntax diagram.

Element	Syntax	Description
Required keywords	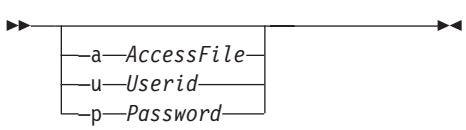	Indicate the parameters or arguments you must specify for the command. Required keywords appear on the main path line. Mutually exclusive required keywords are stacked vertically.
Optional keywords		Indicate the parameters or arguments you can choose to specify for the command. Optional keywords appear below the main path line. Mutually exclusive optional keywords are stacked vertically.
Default value		Appears above the main path line.
Repeatable keyword or value	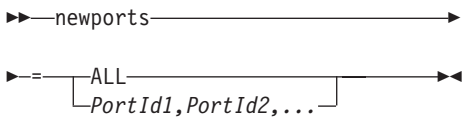	Represents a parameter or argument that you can specify more than once. A repeatable keyword or value is represented by an arrow returning to the left above the keyword or value.
Variable		Represents the value you need to supply for a parameter or argument, such as a file name, user name, or password. Variables are in italics.
Space separator	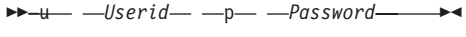	Adds a blank space on the main path line to separate keywords, parameters, arguments, or variables from each other.
Quotation mark delimiters	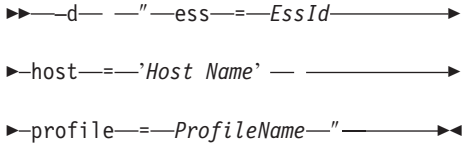	Indicates the start and end of a parameter or argument that contains multiple values. Enclose one or more name–value pairs in a set of double quotation marks for a particular parameter or argument. If the value of a parameter or name–value pair contains a blank or white space, enclose the entire value in a set of single quotation marks.
Equal–sign operator	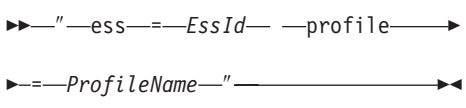	Separates a name from its value in a name–value pair.

Element	Syntax	Description
Syntax fragment	▸▸—Fragment Name———————————▸◂ **Fragment name:** ├—(—*fragment details*—)———————┤	Breaks up syntax diagrams that are too long, too complex, or repetitious. The fragment name is inserted in the main diagram, and the actual fragment is shown below the main diagram.

Terminology

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

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

Name	Object type
Host	host
Virtual disk	vdisk
Managed disk	mdisk
Managed disk group	mdiskgrp
I/O group	iogrp
Node	node
Cluster	cluster
Controller	controller
FlashCopy® mapping	fcmap
FlashCopy consistency group	fcconsistgrp
Metro Mirror relationship	rcrelationship
Metro Mirror consistency group	rcconsistgrp
Unsupported/unknown object	unknown

CLI special characters

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

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

Using wildcards in the SAN Volume Controller CLI

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

The SAN Volume Controller allows the use of the '*' as a wildcard within the arguments of certain parameters. There are some behavioral issues that must be

considered when using wildcards in order to prevent unexpected results. These behavioral issues, and the ways to avoid them, are described below.

1. Running the command while logged onto the node.

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

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

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

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

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

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

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

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

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

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

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

   with double quotes inside single quotes.

Data types and value ranges

The maximum length of any single parameter entered into the command line is 2176 bytes.

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

Data types	Value ranges
filename_arg	This is a (optionally fully qualified) file name. Maximum length is 169 characters. Valid characters consist of the following: • . • / • - • _ • a - z • A - Z • 0 - 9 The field may not contain two consecutive '.', or start with a '.', or end with a '.'.
directory_or_file_filter	Specifies a directory and or filename filter within the specified directory. Valid directory values consist of the following: • /dumps • /dumps/audit • /dumps/configs • /dumps/elogs • /dumps/feature • /dumps/iostats • /dumps/iotrace • /dumps/software The filename filter can be any valid filename with or without the wildcard '*'. The filename filter can be appended to the end of one of the above directory values. Maximum length is 128 characters. Valid characters consist of the following: • * • . • / • - • _ • a - z • A - Z • 0 - 9 The field may not contain two consecutive '.', or start with a '.', or end with a '.'.
filename_prefix	This is a prefix to be used when naming a file. Maximum length is 128 characters. Valid characters consist of the following: • a - z • A - Z • 0 - 9 • - • _

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

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

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

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

The maximum number of values that can be entered into a colon separated list is 128. If more than 128 items are entered into a list an error is returned.

CLI parameters

CLI parameters are found within the syntax diagram.

CLI parameters can be entered in any order except:

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

A valid parameter meets the following requirements:

- Parameters can be entered in any order.
- If a parameter has an associated argument, the argument must *always* follow the parameter.
- A parameter *must* start with a '-'; otherwise, it is assumed to be an argument.
- The maximum length of any single parameter that can be entered into the CLI is 128 bytes.
- An argument can contain multiple data items. The maximum number of data items that you can enter into such a list is 128. For a component list, separate the individual items by a colon.

CLI flags

The flags that are listed below are common to all command-line interface (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.

CLI messages

Ensure that you are familiar with the command-line interface (CLI) messages.

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.

Chapter 1. Preparing the SSH client system

Before you can issue command-line interface (CLI) commands from the host to the cluster, you must prepare the Secure Shell (SSH) client system.

Windows operating systems

If you have purchased the master console hardware and software from IBM, PuTTY for Windows® operating systems has been previously installed for you.

If you are installing the master console on your own hardware with a Windows operating system, 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 OpenSSH 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 but you also need its prerequisite, OpenSSL, from the AIX® toolbox for Linux® applications for Power Systems. For AIX 4.3.3, you can get the software from the AIX toolbox for Linux applications.

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

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

Linux operating systems

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

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

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

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

Preparing the SSH client system to issue CLI commands

In order to issue command-line interface (CLI) commands to the cluster from a host, you must prepare the Secure Shell (SSH) client on the host so that the host is accepted by the SSH server on the cluster.

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

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

1. For the master console and Windows hosts:
   a. Generate an SSH key pair using the PuTTY key generator.
   b. Store the SSH clients public key on the cluster (using a browser pointing to the SAN Volume Controller Console).
   c. Configure the PuTTY session for the CLI.
2. For other types of hosts:
   a. Follow the instructions specific to the SSH client to generate an SSH key pair.
   b. Store the SSH clients public key on the cluster (using a browser pointing to the SAN Volume Controller Console or the CLI from an already established host).
   c. Follow the instructions that are specific to the SSH client to establish an SSH connection to the SAN Volume Controller cluster.

Related tasks

"Configuring the PuTTY session for the CLI" on page 6
You must configure the PuTTY session using the Secure Shell (SSH) key pair that you have generated before you can use the command-line interface (CLI).

"Adding SSH keys for hosts other than the master console" on page 7
You can add Secure Shell (SSH) keys on other hosts.

Related reference

"Configuring the Secure Shell client system" on page 4
Ensure that you configure the Secure Shell (SSH) client system to set up your environment.

Chapter 2. Secure Shell

Secure Shell (SSH) is a client-server network application.

Overview

The SAN Volume Controller cluster acts as the SSH server in this relationship. The SSH client provides a secure environment in which to connect to a remote machine. It uses the principles of public and private keys for authentication.

SSH keys are generated by the SSH software. This includes a public key, which is uploaded and maintained by the cluster and a private key that is kept private to the host that is running the SSH client. These keys authorize specific users to access the administration and service functions on the cluster. Each key is associated with a user-defined ID string that can consist of up to 40 characters. Up to 100 keys can be stored on the cluster. You can also add new IDs and keys or delete unwanted IDs and keys.

Attention: The SAN Volume Controller does not support running multiple SSH sessions concurrently against a single cluster. This can cause the system to lose access to data and cause data to be lost. To avoid running multiple SSH sessions concurrently against a single cluster, do not run scripts that create child processes that run in the background and invoke SAN Volume Controller commands.

Secure Shell (SSH) is a communication vehicle between the host system and the following components:

- The SAN Volume Controller
- The system on which the SAN Volume Controller console is installed

Authenticating SSH logins

When you are using AIX hosts, SSH logins are authenticated on the cluster using the RSA-based authentication that is supported in the OpenSSH client available for AIX. This scheme is based on public-key cryptography, using a scheme known commonly as RSA.

Note: The authentication process for non-AIX hosts systems is similar.

With this scheme (as in similar OpenSSH systems on other host types) the encryption and decryption is done using separate keys. This means it is not possible to derive the decryption key from the encryption key.

Because physical possession of the private key allows access to the cluster, the private key must be kept in a protected place, such as the .ssh directory on the AIX host, with restricted access permissions.

When an SSH client (A) attempts to connect to an SSH server (B), the key pair is needed to authenticate the connection. The key consists of two halves: the public and private keys. The SSH client public key is put onto the SSH Server (B) using some means outside of the SSH session. When the SSH client (A) tries to connect, the private key on the SSH client (A) is able to authenticate with its public half on the SSH server (B).

Running the CLI

In order to use the command-line interface (CLI) or SAN Volume Controller Console system, you must have an SSH client installed on that system and perform the following tasks:

- Generate the SSH key pair on the client system.
- Store the private key from this key pair on the client system.
- Store the SSH public key for the client on the SAN Volume Controller clusters.

The master console has the SSH client software called PuTTY preinstalled. This software provides the Secure Shell (SSH) client function for users logged into the master console who want to invoke the SAN Volume Controller command-line interface (CLI).

If you are installing the master console on your own hardware with a Windows operating system, you can install 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

If you want to run the SAN Volume Controller command-line interface (CLI) from a different system than the master console, you must install an SSH client. You can generate SSH public and private keys using the PuTTY software. You must store the SSH Client public key on all SAN Volume Controller clusters.

Connecting the SAN Volume Controller Console to additional clusters

The master console also has the SAN Volume Controller Console Web server and Common Information Model (CIM) Object manager software preinstalled. This software depends on the PuTTY Secure Shell (SSH) client function for the SAN Volume Controller Console to programmatically access the SAN Volume Controller cluster. The master console comes with PuTTY SSH keys preinstalled. You can generate new PuTTY SSH keys unique to your master console and copy the private SSH key to the SAN Volume Controller Console directory and store the public SSH key on all clusters to which the SAN Volume Controller Console will connect.

You can also install the SAN Volume Controller Console on a Windows 2000 server system which you provide. If you intend to install the SAN Volume Controller Console on a host which you supply, you must install PuTTY first, which is a prerequisite for the SAN Volume Controller Console.

Configuring the Secure Shell client system

Ensure that you configure the Secure Shell (SSH) client system to set up your environment.

The related topics elaborate on each step to configure a PuTTY SSH client system. IBM has preinstalled the PuTTY SSH client software on the master console. You can also install PuTTY on any Windows 2003 server (for master console version 1.4 and later) where you will run the command-line interface (CLI) or where you install the SAN Volume Controller Console. If you have some other SSH client

software to run on another host, follow that software documentation to perform the tasks equivalent to the following steps.

1. Install the SSH client software on the computer that will host the master console. This step is not required for a master console that has PuTTY preinstalled.
2. Generate SSH keys on the SSH client system.
3. Configure the PuTTY session, if required, on the SSH client system.
4. If the client system is the master console, copy the private key into the SAN Volume Controller install directory. If the client system is not the master console, store the private key on the SSH client system.

   Attention: Do not run scripts that create child processes that run in the background and invoke SAN Volume Controller commands. This can cause the system to lose access to data and cause data to be lost.
5. Copy the SSH public key to the master console.
6. Store the SSH client public key on the SAN Volume Controller cluster.

You will perform step 6 to store the SSH client public key on the SAN Volume Controller when you complete the creation of the SAN Volume Controller cluster. Once you have defined a cluster to the SAN Volume Controller Console and have therefore enabled SSH communication to the cluster, you can store additional SSH client public keys on the cluster. You can store additional keys through the SAN Volume Controller Console or the Command-Line Interface.

Related tasks

"Configuring the PuTTY session for the CLI" on page 6
You must configure the PuTTY session using the Secure Shell (SSH) key pair that you have generated before you can use the command-line interface (CLI).

"Adding SSH keys for hosts other than the master console" on page 7
You can add Secure Shell (SSH) keys on other hosts.

Generating a Secure Shell key pair using the Secure Shell client called PuTTY

These are step-by-step instructions for generating Secure Shell (SSH) keys on the PuTTY SSH client system.

1. Start the PuTTY Key Generator to generate public and private keys for SSH client connection to the SSH Server on the SAN Volume Controller cluster. Select **Start ▸ Programs ▸ PuTTY ▸ PuTTYgen** to open the PuTTY Key Generator Graphical User Interface (GUI) window.
2. Use the PuTTY Key Generator GUI window to generate keys:
   a. Select the **SSH2 RSA** radio button.
   b. Leave the number of bits in a generated key value at 1024.
   c. Click **Generate**.

   A message similar to the following is displayed:

   ```
   Please generate some randomness by moving
   the mouse over the blank area.
   ```

   in the section of the GUI labeled Key. The *blank area* indicated by the message is the large blank rectangle on the GUI inside the section of the GUI labeled Key. Continue to move the cursor over the blank area until the progress bar reaches the far right. This generates random characters to create a unique key.

Optionally, if you want to set a passphrase for your key, you can enter a passphrase in the Key passphrase and Confirm passphrase fields. The passphrase will be used to encrypt the key on the disk, therefore, it will not be possible to use the key without first entering the passphrase. More details about the passphrase can be found by clicking on the Help in the puttygen application.

Attention: Do not enter anything in the Key passphrase or Confirm passphrase fields when generating the key pair for the master console.

3. Save the generated SSH keys on your system disk for later use. Two files are generated.
   a. Click **Save public key**. You will be prompted for a name and location for the key. Remember the name and location of the SSH public key you save.

      Note:

      1) For AIX, store the key in the $HOME/.ssh directory.
      2) It is recommended that you use the term `pub` in naming the public key, for example, `pubkey`, to easily differentiate the SSH public key from the SSH private key. You will identify the name and location of the SSH public key to the SAN Volume Controller cluster in a later step.

   b. Click **Save Private key**. You will be prompted with a message similar to the following:

```
Are you sure you want to save this key
without a passphrase to protect it?
 Yes/No
```

      Note: For AIX, store the key in the $HOME/.ssh directory, in the $HOME.ssh/identity file. In the simplest cases, this involves replacing the contents of the identity file with the contents of the key file. However, when using multiple keys, then all of these keys must appear in the identity file.

4. Close the PuTTY Key Generator.

Configuring the PuTTY session for the CLI

You must configure the PuTTY session using the Secure Shell (SSH) key pair that you have generated before you can use the command-line interface (CLI).

Attention: Do not run scripts that create child processes that run in the background and invoke SAN Volume Controller commands. This can cause the system to lose access to data and cause data to be lost.

Perform the following steps to configure a PuTTY session for the CLI:

1. Select **Start** → **Programs** → **PuTTY** → **PuTTY**. The PuTTY Configuration window opens.
2. Click **Session** in the Category navigation tree. The Basic options for your PuTTY session are displayed.
3. Click **SSH** as the Protocol option.
4. Click **Only on clean exit** as the Close window on exit option. This ensures that connection errors are displayed.
5. Click **Connection** → **SSH** in the Category navigation tree. The options controlling SSH connections are displayed.

6. Click **2** as the Preferred SSH protocol version.
7. Click **Connection** → **SSH** → **Auth** in the Category navigation tree. The Options controller SSH authentication are displayed.
8. Click **Browse** or type the location of the SSH private key in the **Private key file for authentication** field.
9. Click **Session** in the Category navigation tree. The Basic options for your PuTTY session are displayed.
10. Click **Default Settings** and then click **Save** as the Load, save or delete a stored session option.
11. Type the name or IP address of the SAN Volume Controller cluster in the **Host Name (or IP Address)** field.
12. Type the port of the SAN Volume Controller cluster in the **Port** field.
13. Type the name that you want to use to associate with this session in the **Saved Sessions** field. For example, you can name the session SVC Cluster 1.
14. Click **Save**.

You have now configured a PuTTY session for the CLI.

Adding subsequent SSH public keys to the SAN Volume Controller

You can add subsequent Secure Shell (SSH) public keys to the SAN Volume Controller from the SSH Public Key Maintenance panel.

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

The SSH key allows the master console (where the SAN Volume Controller Console is running) to access the cluster.

During the cluster creation wizard, you added a SSH key to the cluster. You can add additional SSH keys to grant SSH access to other servers.

Perform the following steps to add additional SSH keys:

1. Click **Clusters** in the portfolio.
2. Click the cluster whose SSH keys you want to maintain.
3. Select **Maintain SSH Keys** from the task list and click **Go**. The SSH Public Key Maintenance panel is displayed.
4. Follow the instructions that are on the SSH Public Key Maintenance panel.
5. Click Add Key when you have completed the SSH Public Key Maintenance panel.

After the initial configuration of the cluster has been performed using the SAN Volume Controller Console and at least one SSH client key has been added, the remainder of the configuration can either be performed using the SAN Volume Controller Console or the command-line interface.

Adding SSH keys for hosts other than the master console

You can add Secure Shell (SSH) keys on other hosts.

Perform the following steps to add SSH keys on hosts other than the master console:

1. Generate the public private key pair on each host that you want to use the SAN Volume Controller command-line interface. See the information that came with your SSH client for specific details about using the key generation program that comes with your SSH client.
2. Copy the public keys from each of these hosts to the master console.
3. Secure copy these public keys from the master console to the cluster.
4. Repeat for each public key copied onto the master console in step 2.

Chapter 3. Copying the SAN Volume Controller software upgrade files using PuTTY scp

PuTTY scp (pscp) provides a file transfer application for secure shell (SSH) to copy files either between two directories on the configuration node or between the configuration node and another host.

To use pscp, you must have the appropriate permissions on the source and destination directories on your respective hosts.

The pscp application is available when you install an SSH client on your host system. You can access the pscp application through a command prompt.

Perform the following steps to use pscp:

1. Start a PuTTY session.
2. Configure your PuTTY session to access your SAN Volume Controller Console cluster.
3. Save your PuTTY configuration session. For example, you can name your saved session SVCPUTTY.
4. Open a command prompt.
5. Issue the following command to set the path environment variable to include the PuTTY directory:

   ```
   set path=C:\Program Files\putty;%path%
   ```

   Where *Program Files* is the directory where PuTTY is installed.
6. Issue the following command to copy the package onto the node where the CLI runs:

   ```
   directory_software_upgrade_files pscp -load saved_putty_configuration
    software_upgrade_file_name admin@cluster_ip_address:home/admin/upgrade
   ```

   Where *directory_software_upgrade_files* is the directory that contains the software upgrade files, *saved_putty_configuration* is the name of the PuTTY configuration session, *software_upgrade_file_name* is the name of the software upgrade file and *cluster_ip_address* is the IP address of your cluster.

   If there is insufficient space to store the software upgrade file on the cluster, the copy process fails. Perform one of the following steps:

   - Issue the **svctask cleardumps** CLI command to free space on the cluster and repeat step 6.
   - Issue the following command from the cluster to transfer the error logs to the master console:

   ```
   pscp -unsafe -load saved_putty_configuration
    admin@cluster_ip_address:/dump/elogs/* your_desired_directory
   ```

   Where *saved_putty_configuration* is the name of the PuTTY configuration session, *cluster_ip_address* is the IP address of your cluster and *your_desired_directory* is the directory where you want to transfer to the error logs.

   After you have transferred the error logs to the master console, repeat step 6.

Chapter 4. Audit log commands

An audit log keeps track of action commands that are issued through an SSH session or the SAN Volume Controller Console.

The audit log entries provide the following information:

- identity of the user who issued the action command
  - if coming from the command-line interface, the user name (administrator or service), and the label associated with the user's public SSH key in the authorized keys file
  - if coming from the native Web pages, the user's identity (admin[web] or service[web]) according to which user name the user authenticated with
  - if coming from the SAN Volume Controller Console, the user's identity (administrator), the label associated with the CIMOM's key in the authorized keys file, and the user name recorded by the CIMOM when the SAN Volume Controller Console user authenticated with the CIMOM
- the name of the actionable command
- the timestamp of when the actionable command was issued configuration node
- the parameters which were issued with the actionable command

The following commands are not documented in the audit log:

- svctask dumpconfig
- svctask cpdumps
- svctask cleardumps
- svctask finderr
- svctask dumperrlog
- svctask dumpinternallog

Other instances of commands which are not documented in the audit log, are the following:

- Commands that fail
- Result code is 0 (sucesss) or 1 (success in progress)
- Result object ID of node type (for the addnode command) is not logged

catauditlog

Use the **catauditlog** command to display the in-memory contents of the audit log.

Syntax

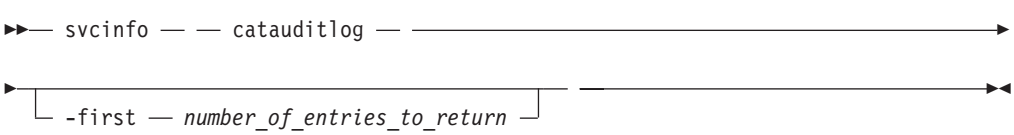

Parameters

-first *number_of_entries_to_return*
Optionally specifies the number of entries in the log. Specify the number of most recent entries that you want to display.

Description

This command lists a specified number of the most recently audited commands.

The in-memory portion of the audit log can hold approximately 1MB of audit information. Depending on the command text size and the number of parameters, 1MB will record approximately 6000 commands.

Once the in-memory audit log has reached its maximum capacity, the log is written to a local file on the configuration node in the /dumps/audit directory. The **svcinfo catauditlog** command only displays the in-memory part of the audit log; the on-disk part of the audit log is in readable text format and does not need any special command to decode it.

The in-memory log entries are then reset and cleared automatically, ready to start accumulating new commands. The on-disk portion of the audit log can then be analyzed at a later date.

The in-memory portion of the audit log can be transferred to an on-disk file using the **svctask dumpauditlog** command. This action clears the in-memory portion of the log.

Possible failures

- There are no error codes.

The user has specified that they want to list the 15 most recent audit log entries.

An invocation example

```
svcinfo catauditlog -delim : -first 15
```

The resulting output

```
audit_seq_no:timestamp:cluster_user:ssh_label:icat_user:result:res_obj_id
:action_cmd
125:060311111800:admin:Joe::0::svctask rmsshkey -key label47 -user admin
126:060311111800:admin:Joe::0::svctask addsshkey -label label48 -file
/home/Joe/id_rsa.pub -user admin
127:060311111800:admin:Joe::0::svctask rmsshkey -key label48 -user admin
128:060311111800:admin:Joe::0::svctask addsshkey -label label49 -file
/home/Joe/id_rsa.pub -user admin
129:060311111800:admin:Joe::0::svctask rmsshkey -key label49 -user admin
130:060311134617:admin:Joe::0::svctask chmdisk -name ca-0 1
131:060311134617:admin:Joe::0::svctask chmdisk -name ca-1 2
132:060311134617:admin:Joe::0::svctask chmdisk -name ca-2 3
133:060311134617:admin:Joe::0::svctask chmdisk -name cb-0 4
134:060311134617:admin:Joe::0::svctask chmdisk -name cb-1 5
135:060311134617:admin:Joe::0::svctask chmdisk -name cb-2 6
136:060311134617:admin:Joe::0::svctask chmdisk -name cc-0 7
137:060311134617:admin:Joe::0::svctask chmdisk -name cc-1 8
138:060311134617:admin:Joe::0::svctask chmdisk -name cc-2 9
139:060311134632:admin:Joe::0::svctask mkmdiskgrp -name custa-mdisks -ext 512
-mdisk ca-0:ca-1:ca-2
```

dumpauditlog

Use the **dumpauditlog** command to reset or clear the contents of the in-memory audit log. The contents of the audit log are sent to a file in the /dumps/audit directory on the current configuration node.

Syntax

Parameters

There are no parameters.

Description

This command dumps the contents of the audit log to a file on the current configuration node. It will also clear the contents of the audit log. This command is logged as the first entry in the new audit log.

Audit log dumps are automatically maintained in the /dumps/audit directory. The local file system space gets utilized by audit log dumps and is limited to 200MB on any node in the cluster. The space limit is maintained automatically by deleting the minimum number of old audit log dump files so that the /dumps/audit directory space is reduced below 200MB. This deletion occurs once per day on every node in the cluster. The oldest audit log dump files are considered to be the ones with the lowest audit log sequence number. Also, audit log dump files with a cluster ID number not matching the current one are considered to be older than any files that do match the cluster ID, regardless of sequence number.

Other than by executing dumps (or copying dump files among nodes) you cannot alter the contents of the audit directory. Each dump file name is generated automatically in the following format:

auditlog_*<firstseq>*_*<lastseq>*_*<timestamp>*_*<clusterid>*

where

- *<firstseq>* is the audit log sequence number of the first entry in the log
- *<lastseq>* is the audit sequence number of the last entry in the log
- *<timestamp>* is the timestamp of the last entry in the audit log being dumped
- *<clusterid>* is the cluster ID at the time the dump was created

The audit log dump files names cannot be changed.

The audit log entries in the dump files contain the same information as displayed by the **svcinfo catauditlog** command, however, the **svctask dumpauditlog** command displays the information with one field per line.

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.

An invocation example

```
svctask dumpauditlog
```

The resulting output

```
No feedback
```

lsauditlogdumps

Use the **lsauditlogdumps** command to list the audit log dumps that are available on the configuration node.

Syntax

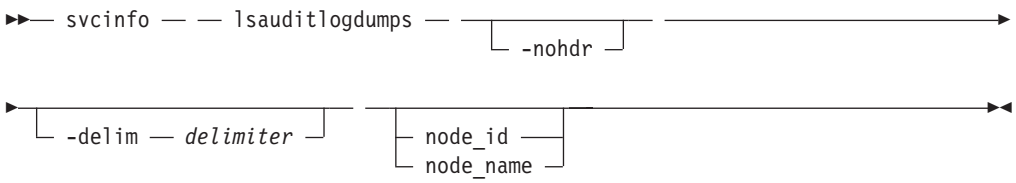

Parameters

-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 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 will list the audit log files in the /dumps/audit directory on the current configuration node or on the designated node.

The cluster creates the audit log automatically. The audit log can also be created manually by issuing the **svctask dumpauditlog** command. The audit log comprises the files listed by the **svcinfo lsauditlogdumps** command. These files are limited to approximately 200MB on each node in the cluster.

When the configuration node moves to a different node in the cluster, any old audit log files will be left on the former configuration node. As with other types of dumps, you can retrieve those files using the **svctask cpdumps** command.

Possible failures

- There are no error messages.

An invocation example

```
svcinfo lsauditlogdumps
```

The resulting output

```
id auditlog_filename
0 auditlog_0_229_060311234532_0000020060013d8a
1 auditlog_230_475_060312234529_0000020060013d8a
2 auditlog_476_491_060313234527_0000020060013d8a
```

Audit log dump file contents

```
...
Auditlog Entry:23
Audit Sequence Number :138
Timestamp :Sat Mar 11 13:46:17 2006
:Epoch + 1142084777
SVC User :admin
SSH Label :Joe
ICAT User :
Result Object ID :
Result Code :0
Action Command :svctask chmdisk -name cc-2 9
Auditlog Entry:24
Audit Sequence Number :139
Timestamp :Sat Mar 11 13:46:32 2006
:Epoch + 1142084792
SVC User :admin
SSH Label :Joe
ICAT User :
Result Object ID :
Result Code :0
Action Command :svctask mkmdiskgrp -name custa-mdisks -ext
512 -mdisk ca-0:ca-1:ca-2
...
```

Chapter 5. Cluster commands

Cluster commands are used to monitor and modify clusters.

There are a number of cluster commands used for various tasks. A cluster is a pair of nodes that provide a single configuration and service interface.

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

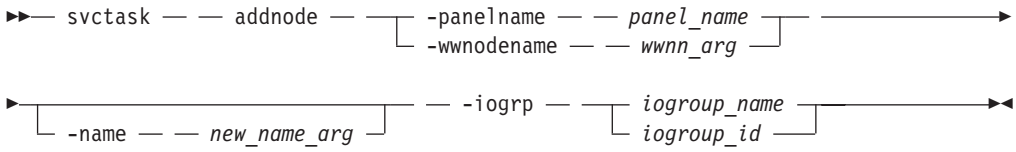

Parameters

-panelname *panel_name*
Identifies the node to be added by the name as it appears on the display panel. 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`.

Node addition completes asynchronously. This means while the node is in adding state, the WWPN is not known and will display as zeroes.

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 node serial number, the WWNN, all WWPN's, and the I/O group to which it has been added. This can prevent possible data corruption 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 the cluster GUI, confirm that the node has previously been a member of the cluster. If it has, follow one of these two procedures:

- 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 **svcinfo lsnode** command. Or,
- if this information is not available, call the support team to add the node back into the cluster without corrupting the data.

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.

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 *Multipath Subsystem Device Driver: User's Guide* for more information.

The SDD driver maintains an association between a vpath and a 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 VDisks.

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 cluster 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 WWNN 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

- CMMVC5707E Required parameters are missing.
- 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].

An invocation example

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

The resulting output

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

chcluster

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

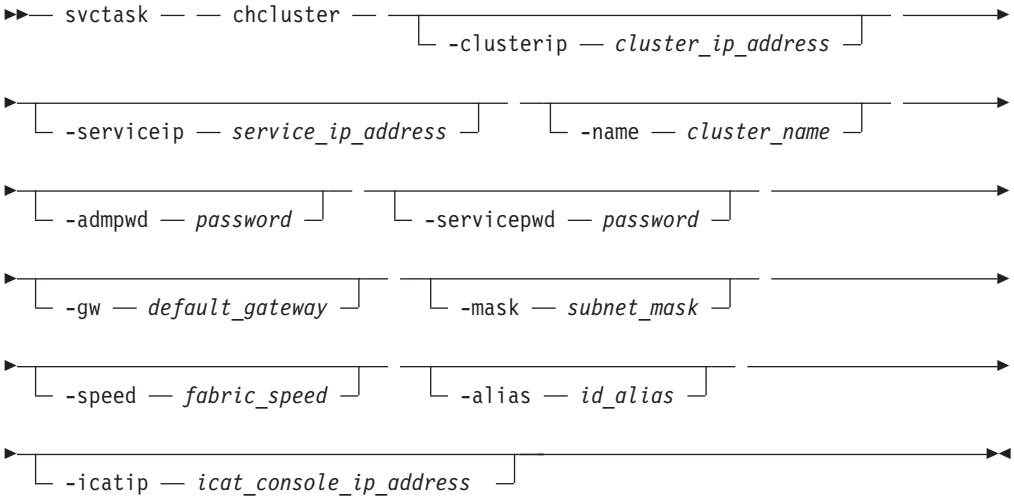

Parameters

-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.

-serviceip *service_ip_address*
Changes the service IP address. This address is the address that must be used if the node has to be started after it has been expelled from the cluster.

-name *cluster_name*
Changes the name of the cluster.

-admpwd *password*
Changes the administrator password. This argument can be specified with or without the password. If the argument is not followed by a password, you will be prompted for the password. When you type the password in response to the prompt, the password will not be displayed.

-servicepwd *password*
Changes the service user password. This argument can be specified with or without the password. If the argument is not followed by a password, you will be prompted for the password. When you type the password in response to the prompt, the password will not be displayed.

-gw *default_gateway*
Changes the default gateway IP address of the cluster.

-mask *subnet_mask*
Changes the subnet mask of the cluster.

-speed *fabric_speed*
Specifies the speed of the fabric to which this cluster is attached. Valid values are 1 or 2 (Gb).

Attention: Changing the speed on a running cluster will break I/O service to attached hosts. Before changing the fabric speed, stop I/O from active hosts and force such hosts to flush any cached data by dismounting volumes (for UNIX host types) or removing drive letters (for Windows host types). Some hosts may need to be rebooted to detect the new fabric speed.

The fabric speed setting applies only to the 4F2 and 8F2 model nodes in a cluster. The 8F4 nodes automatically negotiate the fabric speed on a per-port basis.

-alias *id_alias*
Does not change the basic ID for the cluster, but does influence the VDisk_UID of every **vdiskhostmap**, both existing and new. These objects appear to have been created for a cluster whose ID matches the alias.

-icatip *icat_console_ip_address*
Changes the IP address used by this cluster. The format of this IP address must be a dotted decimal notation together with the port (for example, 255.255.255.255:8080).

Description

This command modifies specific features of a cluster. The arguments are not mutually exclusive, so multiple features can be changed with a single command invocation.

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.

The service IP address is not used until a node is expelled from the cluster. If this node cannot rejoin the cluster, you can bring the node up in service mode. In this mode, the node can be accessed as a stand-alone node using the service IP address.

If you do not specify any of the options, the command does nothing.

Modifying a password: To change the administrator users password, issue the **svtask chcluster -admpwd <password>** command. To change the service users password, issue the **svtask chcluster -servicepwd <password>** command.

Note: If you do not wish the password to be displayed as you enter the command line then you can omit the new password. The command line tool will then prompt you to enter and confirm the password without the password being displayed.

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.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5789E The cluster was not modified because the IP address, subnet mask, service address, SNMP address, or gateway address is not valid.

An invocation example

```
svctask chcluster -clusterip 217.12.3.11
```

The resulting output

```
No feedback
```

chiogrp

Use the **chiogrp** command to modify the name assigned to an I/O group.

Syntax

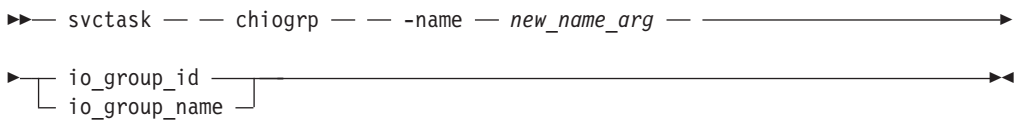

Parameters

-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.

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.

An invocation example

```
svctask chiogrp -name testiogrpone io_grp0
```

The resulting output

```
No feedback
```

chnode

You can use the **chnode** command to change the name, or label, assigned to a node. The name can then be used from then on for subsequent command line tools.

Syntax

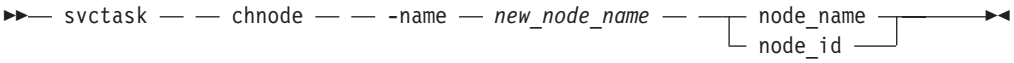

Parameters

-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).

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.

An invocation example

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

The resulting output

```
No feedback
```

cleardumps

The **cleardumps** command cleans the various dump directories on a specified node.

Syntax

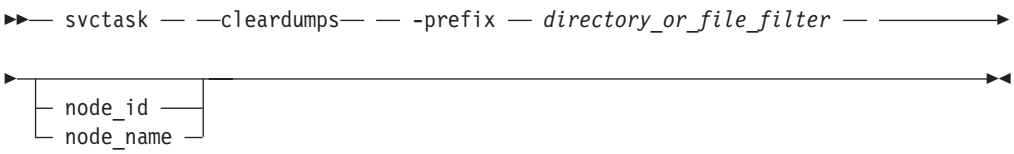

Parameters

-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/audit
- /dumps/configs
- /dumps/elogs
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace
- /home/admin/upgrade

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 list the contents of these directories on the given node by using the **svcinfo lsxxxxdumps** commands.

You can use this command to clean specific files in a given directory by specifying a directory or file name. The wildcard * can be used as part of the file name.

Note: In order to preserve the configuration and trace files, any files matching the following wildcard patterns will not be cleaned:

- *svc.config*
- *.trc
- *.trc.old

Possible failures

- CMMVC5985E 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, or /home/admin/upgrade.

An invocation example

```
svctask cleardumps -prefix /dumps/configs
```

The resulting output

```
No feedback
```

cpdumps

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

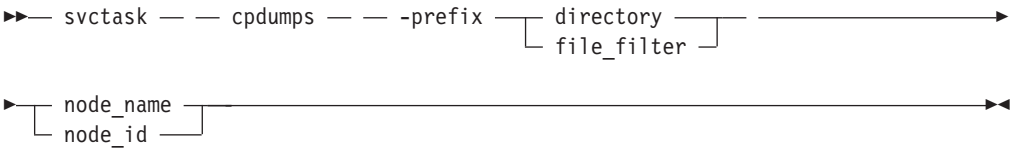

Parameters

-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 (retrieves all files in all subdirectories)
- /dumps/audit
- /dumps/configs
- /dumps/elogs
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace

- /home/admin/upgrade

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 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.

Possible failures

- CMMVC5985E 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, or /home/admin/upgrade.

An invocation example

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

The resulting output

```
No feedback
```

detectmdisk

Use the **detectmdisk** command to manually rescan the fibre-channel network for any new managed disks (MDisks) that might have been added and to rebalance MDisk access across the available controller device ports.

Syntax

Description

This command causes the cluster to rescan the fibre-channel network. The rescan discovers any new MDisks that might have been added to the cluster and

rebalances MDisk access across the available controller device ports. This command also detects if a controller is being decommissioned.

Note: Although it might appear that the **detectmdisk** command has completed, some extra time might be required for it to run. The **detectmdisk** is asynchronous and returns a prompt while the command continues to run in the background.

In general, the cluster automatically detects these disks 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 new storage and the cluster has not detected it, you might need to run this command before the cluster detects 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 cluster, the cluster automatically discovers the back-end controller and integrates the controller to determine what storage is presented to it. The SCSI LUs that are presented by the back-end controller are displayed as unmanaged MDisks. However, if the configuration of the back-end controller is modified after this has occurred, the cluster might be unaware of these configuration changes. This task allows a user to request the cluster to rescan the fibre-channel SAN to update the list of unmanaged MDisks.

Note: The automatic discovery that is performed by the cluster does not write anything to an unmanaged MDisk. It is only when a user instructs the cluster to add an MDisk to an MDisk group or use an MDisk to create an image mode virtual disk that the storage is actually 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.

An invocation example

```
svctask detectmdisk
```

The resulting output

```
No feedback
```

rmnode

The **rmnode** command deletes a node from the cluster. You can enter this command any time after a cluster has been created.

Syntax

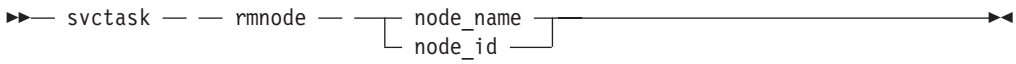

Parameters

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.

Prerequisities:

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 VDisks, where the filter attribute is the I/O group.

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

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

   Note: Any VDisks that are assigned to the I/O group that this node belongs to, will be assigned to the other node in the I/O group; that is, the preferred node will be changed. You cannot change this setting back.
2. Determine the hosts that the VDisks are mapped to by issuing the **svcinfo lsvdiskhostmap** command.
3. Determine if any of the VDisks 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 VDisks, go to step 5.
   - If you *do* want to maintain access to some or all of the VDisks, back up the data or migrate the data to a different (online) I/O group.
4. Determine if you need to turn the power off to the node:
   - If this is the last node in the cluster, you do not need to turn the power off to the node. Go to step 5.
   - If this is *not* the last node in the cluster, turn the power off to the node that you intend to remove. This step ensures that the 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 **Adding a node back into the cluster:** below.
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 *Multipath 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.

Note:

1. If the node you are removing is the configuration node, it may take one or two minutes to complete the command.
2. If the node you are removing is the last node in the cluster, the cluster 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:

Note:

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 be lost if last node in the cluster is deleted. Deleting the configuration node results in the CLI failing over to another node. 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.

- Node serial number
- WWNN

- all WWPNs
- I/O group that contains the node

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

Replacing a faulty node:

You can replace a failed node with a "spare" or replacement node. This would be necessary if the node fails. The cluster will continue to operate with degraded performance until the failed node is repaired. To enhance availability, you might choose to replace the failed node with a "spare", then repair the node offline. However, various procedures must be followed and precautions taken, to replace a failed node without interrupting I/O and without risk to data integrity when the repaired node is reconnected to the SAN fabric. The procedure involves changing the World Wide Node Name (WWNN) of the replacement node. This procedure must be followed with care to avoid duplicate WWNNs which are illegal and can cause data corruption.

Prerequisities:

Before replacing the failed node you must:

- Ensure that the GUI that is used to access the cluster and spare node is running at least version 1.1.1 of the software.
- Know the cluster name that contains the failing node.
- Confirm that a spare node is in the same rack as the cluster containing the failed node.
- Make a record of the last five characters of the original WWNN. This identification is needed if you decide in the future to designate the spare node as a normal node that can be assigned to any cluster. Use the **svcinfo lsnode** command to see the WWNN.

Additional Information

When you replace a node, the following process takes place:

- The node Front Panel ID will change. This is the number printed on the front of the node, used to select the node that is to be added to a cluster.
- The Node Name might change. If you permit the cluster to assign default names when adding nodes to the cluster it will create a new name each time a node is added. If you choose to assign your own names then you must type in the node name that you want to use. If you are using scripts to perform management tasks on the cluster and those scripts use the Node Name then by assigning the original name to a replacement node you will avoid the need to make changes to the scripts following service activity on the cluster.
- The Node ID will change. A new Node ID is assigned each time a node is added to a cluster. The Node ID or the Node Name can be used when performing management tasks on the cluster but if scripts are being used to perform those tasks it is recommended that the Node Name is used in preference to the Node ID since the Node Name will remain unchanged following service activity on the cluster.
- The World Wide Node Name (WWNN) will not change. The WWNN is used to uniquely identify the node and the fiber channel ports. The node replacement

procedure changes the WWNN of the spare node to match that of the failed node. The node replacement procedures must be followed exactly to avoid any duplication of WWNNs.

- The World Wide Port Name (WWPN) of each fiber channel port will not change. The WWPNs are derived from the WWNN that is written to the replacement node as part of this procedure.

Perform the following steps to replace a node:

1. Use the command **svcinfo lsnode** to display the node name. This command, when implemented, prints a detailed list report containing information about all the nodes on a cluster. The failed node will be offline. Note the nodes' name.
2. Again use **svcinfo lsnode** to display the I/O Group name. Note the name of the group.
3. Use the command **svcinfo lsnodevpd** to display the front panel id. Note the id number.
4. Again use the command **svcinfo lsnodevpd** to record the UPS serial number. Note this number.
5. Use the Front Panel ID to locate the failed node. Disconnect all four fiber channel cables from the node.

   Important: The cables must not be reconnected until the node is repaired and the node number has been changed to the default spare node number.
6. Connect the power/signal cable from the spare node to the UPS with the serial number noted in step 1. The signal cable can be plugged into any vacant position on the top row of serial connectors on the UPS. If no spare serial connectors are available on the UPS, disconnect the cables from the failed node. Power-on the spare node. Display the node status on the service panel.

Perform these steps to change the WWNN of the node and add the replacement node:

1. With the node status displayed on the front panel, press and hold the Down button; press and release the Select button; release the Down button. The text "WWNN" is displayed on line-1 of the display. Line-2 of the display contains the last five characters of the WWNN.
2. With the WWNN displayed on the service panel, press and hold the Down button, press and release the Select button, release the Down button. This switches the display into edit mode.
3. Change the displayed number to match the WWNN recorded in step1. To edit the displayed number use the Up and Down buttons to increase or decrease the numbers displayed. Use the left and right buttons to move between fields. When the five characters match the number recorded in step1, press the select button twice to accept the number.
4. Connect the four fiber channel cables that were disconnected from the failed node to the spare node. Delete the offline node.
5. Add the spare node into the cluster. See "Adding a node back into the cluster".
6. Use the Subsystem Device Drive (SSD) management tool on the host systems to verify all paths are now online. See the *Service Guide* menu options for further information.

When the failed node is repaired do not connect the fibre channel cables to it. Connecting the cables might cause data corruption. Perform these steps after the failed node is repaired:

1. Display the node status on the service panel.
2. With the status displayed on the front panel, press and hold the Down button; press and release the Select button; release the Down button. The text "WWNN" is displayed on line-1 of the display; line-2 of the display contains the last five characters of the WWNN.
3. With the WWNN displayed on the service panel, press and hold the Down button, press and release the Select button, release the Down button. This switches the display into edit mode.
4. Change the displayed number to "00000". To edit the displayed number use the Up and Down buttons to increase or decrease the numbers displayed. Use the left and right buttons to move between fields. When the number is set to "00000", press the select button twice to accept the number. Never connect a node with a WWNN of "00000" to the cluster.

This error should not be marked as fixed in the error log until the failed node has been repaired and returned to the customer. If this is not done, the service technician will not be able to easily find the front panel ID of the failed node.

This node can now be used as a spare node. If this node is no longer required as a spare and is to be used for normal attachment to a cluster you must first use the procedure described above to change the WWNN to the number saved when a spare was being created. See "Prerequisites" above. Using any other number might cause data corruption.

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.

An invocation example

```
svctask rmnode 1
```

The resulting output

```
No feedback
```

setclustertime

You can use the **setclustertime** command to set the time for the cluster.

Syntax

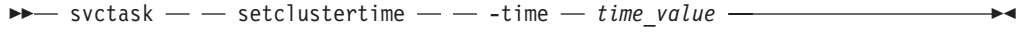

Parameters

-time *time_value*
Specifies the time to which the cluster must be set. This must be in the following format:

MMDDHHmmYYYY

Description

This command sets the time for the cluster.

Possible failures

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

An invocation example

```
svctask setclustertime -time 040509142003
```

The resulting output

```
No feedback
```

setpwdreset

Use the **setpwdreset** command to view and change the status of the password-reset feature for the display panel.

Syntax

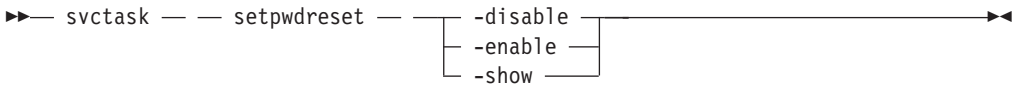

Parameters

-disable
Disables the password-reset feature that is available through the front panel menu system.

-enable
Enables the password-reset feature that is available through the front panel menu system.

-show
Displays the status of the password-reset feature, which is either `enabled` or `disabled`.

Description

The front panel menu system provides an option to reset the administrator password. This option resets the password to a random string that is displayed on the front panel. You can then use this password to access the system. You should change the password at the next login.

Issue the **svctask setpwdreset** command to view and change the status of the password-reset feature for the display panel. Passwords can consist of A - Z, a - z, 0 - 9, and underscore. Make a careful note of the admin password, because without it, you cannot access the cluster.

This command allows you access in case the administrator password is forgotten. If you leave this feature enabled, you should ensure adequate physical security to the cluster hardware.

You can view or change the status of this feature.

Possible failures

- There are no error codes.

An invocation example

```
svctask setpwdreset -show
```

The resulting output

```
Password status: [1]
```

This outcome means that the password or reset feature that is available through the front panel menu system is enabled. If the password status displayed an outcome of [0], this would mean that this feature is disabled.

settimezone

Use the **settimezone** command to set the time zone for the cluster.

Syntax

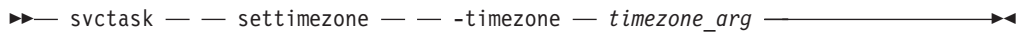

Parameters

-timezone *timezone_arg*
Specifies the time zone to set for the cluster.

Description

This command sets the time zone for the cluster. Use the -timezone parameter to specify the numeric ID of the time zone you want to set. Issue the **svcinfo lstimezones** command to list the time-zones available on the cluster. A list of valid time-zones settings are displayed in a list.

The time zone that this command sets will be used when formatting the error log produced by typing:

```
svctask dumperrlog
```

Note: If you have changed the timezone, you must clear the error log dump directory before viewing the error log through the Web application.

Issue the **svcinfo showtimezone** command to display the current time-zone settings for the cluster. The cluster ID and its associated time-zone are displayed. Issue the **svctask setclustertime** command to set the time for the cluster.

Possible failures

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

An invocation example

```
svctask settimezone -timezone 5
```

The resulting output

```
No feedback
```

startstats

Use the **startstats** command to start the collection of statistics for both virtual disks (VDisks) and managed disks (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 1 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 MDisks, VDisks and node statistics.

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

A maximum of 16 files are stored in the directory at any one time for each statistics file type, for example:

```
Nm_stats_nodepanelname_date_time
Nv_stats_nodepanelname_date_time
Nn_stats_nodepanelname_date_time
m_stats_nodepanelname_date_time
v_stats_nodepanelname_date_time
```

Statistics files beginning with m_stats_* and v_stats_* will not be written if the specified time interval is less than 15 minutes. Statistics files beginning with Nm_stats_*, Nv_stats_* and Nn_stats_* will be written for all time intervals.

Before the 17th file (for each type) is created, the oldest file of that type is deleted.

These files can be listed by using the **svcinfo lsiostatsdumps** command.

The following naming convention is used for these files:

*stats_type*_stats_*nodepanelname_date_time*

Where *stats_type* is m or Nm for MDisks, v or Nv for VDisks, and Nn for node statistics. *nodepanelname* is the current configuration node panel name, *date* is in the format of yymmdd, and *time* is in the format of hhmmss.

The following are examples of MDisk file names:

```
m_stats_000229_031123_072426
Nm_stats_000229_031123_072426
```

The following are examples of VDisk file names:

```
v_stats_000229_031123_072426
Nv_stats_000229_031123_072426
```

The following is an example of a node statistics file name:

```
Nn_stats_000229_031123_072426
```

The statistics that are collected for each MDisk and VDisk are reported in the files named m_stats_*nodepanelname_date_time* and v_stats_*nodepanelname_date_time* include the following statistical information:

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

The statistics that are collected for each MDisk are reported in the file named Nm_stats_*nodepanelname_date_time* include the following statistical information:

- The number of SCSI read and write commands that are processed during the sample period
- The number of blocks of data that are read and written during the sample period
- Per MDisk cumulative read and write external response times in milliseconds
- Per MDisk cumulative read and write queued response times

The statistics that are collected for each VDisk are reported in the file named Nv_stats_*nodepanelname_date_time* include the following statistical information:

- The total number of SCSI read and write commands processed
- The total amount of data read and written
- Cumulative read and write response time in milliseconds
- Statistical information about the read/write cache usage
- Global Mirror statistics including latency

The statistics that are collected for the node from which the statistics file originated and are reported in the file named Nn_stats_*nodepanelname_date_time* include the following statistical information:

- Utilization figure for the node from which the statistic file was obtained
- The amount of data transferred to and received from each port on the node to other devices on the SAN.
- Statistical information about communication to other nodes on the fabric

Note: The v_* and m_* statistics are per-cluster statistics that are only collected on the configuration node. The Nm_*, Nn_* and Nv_* files are per-node statistics that are generated on each node.

Possible failures

- There are no error codes.

An invocation example

```
svctask startstats -interval 25
```

The resulting output

```
No feedback
```

stopcluster

Use the **stopcluster** command to shut down a single node or the entire cluster in a controlled manner. When this command is issued you will be prompted with a confirmation of intent to execute the command.

Syntax

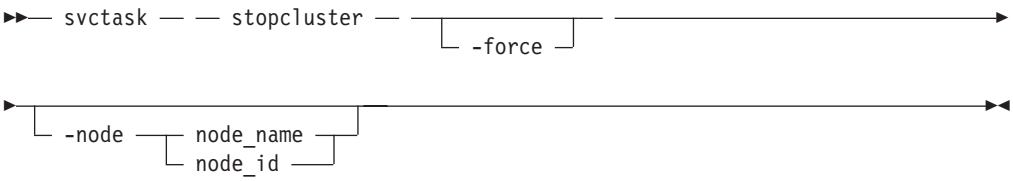

Parameters

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

-node *node_name* | *node_id*
Optionally identifies the node you want to shut down. Specifies the node to be shutdown. 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.

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, Metro Mirror, 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 with 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 cluster is to be removed for more than a few minutes, (for example, if the machine room power is to be shutdown for maintenance), it is

important that the cluster is shutdown before the power is removed. The reason for this is that if the input power is removed from the uninterruptible power supply units without first shutting down the cluster and the uninterruptible power supplies, the uninterruptible power supply units will remain operational and eventually become drained of power.

When input power is restored to the uninterruptible power supplies they will start to recharge but the nodes will not permit any I/O activity to be performed to the virtual disks until the uninterruptible power supply is charged enough to enable all the data on the nodes to be saved in the event of an unexpected power loss. This might take as long as three hours. Shutting down the cluster prior to removing input power to the uninterruptible power supply units will prevent the battery power being drained and will make it possible for I/O activity to be resumed as soon as input power is restored.

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 VDisks 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. 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.

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

If a shutdown command has been sent to the cluster and both cluster and uninterruptible power supply units have powered off, when input power is restored it will be necessary to restart the uninterruptible power supply units by pressing the power button on the uninterruptible power supply front panel.

Ensure that you have stopped all FlashCopy mappings and Metro Mirror relationships. In addition, ensure that all data migration operations and forced deletions have completed 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.

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.

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

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.

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?
```

stopstats

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

Syntax

Description

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

Possible failures

- There are no error codes.

An invocation example

```
svctask stopstats
```

The resulting output

```
No feedback
```

Chapter 6. Backup and restore commands

The following commands are used for backing up and restoring configuration information with the SAN Volume Controller.

backup

Use the **backup** command to backup your configuration. You can enter this command any time after a cluster has been created.

Syntax

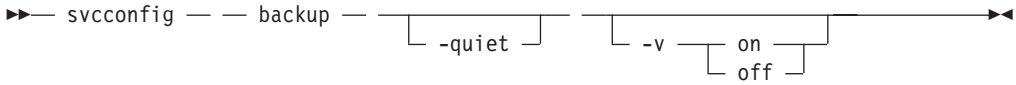

Parameters

-quiet
Suppresses standard output (STDOUT) messages from the console.

-v *on* | *off*
On means verbose messages will be displayed. Off means normal messages (the default) will be displayed.

Description

The **backup** command extracts configuration data from the cluster and saves it to ″svc.config.backup.xml″ in ″/tmp″. A file ″svc.config.backup.sh″ is also produced. This can be studied to see what other commands have been issued to extract information. A log ″svc.config.backup.log″ is also produced. This can be studied for detail as to what has been done and when. This log also includes information on the other commands issued.

Any preexisting file ″svc.config.backup.xml″ is archived as ″svc.config.backup.bak″. (Only one such archive is kept.)

It is recommended that the ″.xml″ file and related ″.key″ files (see limitations below) be immediately moved off-cluster for archiving, and then the files erased from ″/tmp″ using the **clear** command. It is strongly recommended that all objects having default names be changed to have non-default names, as objects with default names cannot be restored faithfully.

The prefix ″_″ (underscore) is reserved for backup and restore command usage, and should not be used in any object names.

The **backup** command has the following limitations:

- ″.key″ SSH public key value files are not produced to go along with the ″.xml″ file in ″/tmp″. Warnings will nevertheless be issued for the missing file or files that should be supplied by the user. These will conform to the template ″svc.config.identifier.user.key″ where *identifier* and *user* are as specified for the **addsshkey** command. You are requested to supply these files if they were used with the **addsshkey** command. If they are not available, then it will be necessary

to install a new set of keys during the cluster restoration process, assuming that cluster restoration is necessary at some time in the future.

Possible failures

- CMMVC6112W *object-type object-name* has a default name
- CMMVC6136W No SSH key file *file-name*
- CMMVC6147E *object-type object-name* has a name beginning with *prefix*

An invocation example

```
svcconfig backup
```

The resulting output

```
No feedback
```

clear

Use the **clear** command to erase files in the ″/tmp″ directory previously produced by other svcconfig commands. You can enter this command any time after a cluster has been created.

Syntax

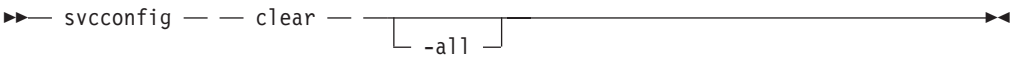

Parameters

-all
Includes ″.key″, ″.bak″ and ″.xml″ files in file clearance; otherwise, clears just ″.log″ and ″.sh″ files. The ″.key″, ″.bak″ and ″.xml″ files contain configuration information, whereas the others do not.

Description

This command clears some or all files in the ″/tmp″ directory produced by svcconfig. Files conform to the template ″svc.config.*″.

Possible failures

- CMMVC6103E Problem file *file-name*: *details*

An invocation example

```
svcconfig clear -all
```

The resulting output

```
No feedback
```

help

Use the **help** command to obtain summary information about the syntax of **svcconfig**. You can enter this command any time after a cluster has been created.

Syntax

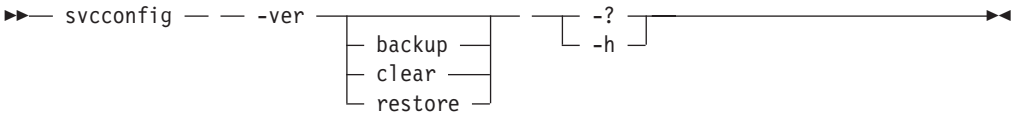

Parameters

-h | -?
Provides general help.

(action) -h | -?
Provides command help: the possible values for (action) are backup, clear, and restore.

-ver
Returns the version number for the **svcconfig** command.

Description

This command provides syntax help for svcconfig.

Possible failures

- CMMVC6100E *-option* not consistent with *action*
- CMMVC6101E *-option* not consistent with *-option*
- CMMVC6102E *-option* and *-optio*n are alternatives
- CMMVC6114E No help for action *action*
- CMMVC6134E No argument for *-option*
- CMMVC6135E Argument *value* for *-option* is not valid
- CMMVC6138E *-option* is required
- CMMVC6141E *-option* does not contain any argument
- CMMVC6149E An action is required
- CMMVC6150E The action *action* is not valid
- CMMVC6151E The option *-option* is not valid
- CMMVC6153E *object* not consistent with *action*

An invocation example

```
svcconfig -ver
svcconfig -?
svcconfig backup -h
```

The resulting output

```
Help text.
```

restore

Use the **restore** command to take information from configuration files in the ″/tmp″ directory and restore the cluster to that configuration. You can only enter this command immediately after a cluster has been created.

Syntax

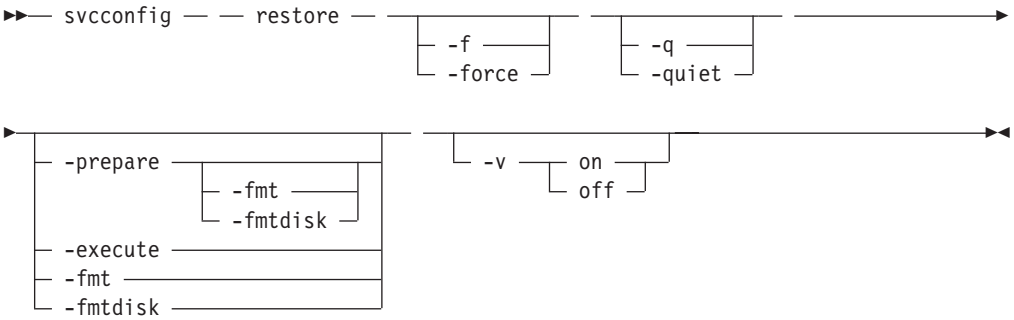

Parameters

-f | force
Forces continued execution where possible.

-q | quiet
Suppresses console output (STDOUT).

-prepare
Checks the current configuration against the information held in ″svc.config.backup.xml″ on the configuration to be restored. Prepares commands for execution in ″svc.config.restore.sh″, and produces a log of events in ″svc.config.restore.prepare.log″.

-fmt | fmtdisk
Includes the -fmtdisk option an all **mkvdisk** commands to be issued.

-execute
Executes the command script ″svc.config.restore.sh″. Produces a log of events in ″svc.config.restore.execute.log″.

-v *on | off*
Produces verbose output (on); the default is regular output (off).

Description

This command restores the target cluster configuration from the ″svc.config.backup.xml″ file, and associated ″.key″ files (if present) in the configuration files directory. If neither the ″prepare″ nor ″-execute″ option is specified, then only a single event log ″svc.config.restore.log″ is produced.

This command will pause for 5 minutes if any nodes are added during this process. You will be informed of this at run-time.

After restoration, a VDisk will comprise a specific list of MDisks. If the relevant MDisk group comprises a larger list, either now or in the future, the restored VDisk will not be able to exploit any MDisks that are not presently in its own list.

The configuration files directory is ″/tmp.″

Possible failures

- CMMVC6105E Different names for source *name* and target *name* clusters
- CMMVC6106E Target cluster has non-default id_alias *value*
- CMMVC6107E *x* io_grp objects in target cluster; *y* are required

- CMMVC6109E Disk controller system with WWNN of *value* not available
- CMMVC6120E Target is not the configuration node
- CMMVC6139E Incorrect XML tag nesting in *file-name*
- CMMVC6142E Existing *object-type object-name* has a non-default name
- CMMVC6143E Required configuration file *file-name* does not exist
- CMMVC6146E Problem parsing *object-type* data: *line*
- CMMVC6147E *object-type object-name* has a name beginning with *prefix*
- CMMVC6148E Target cluster has *actual* object(s) of type *object-type* instead of *required*
- CMMVC6152E vdisk *nam*e instance number *value* is not valid
- CMMVC6155I SVCCONFIG processing completed successfuly
- CMMVC6156W SVCCONFIG processing completed with errors
- CMMVC6165E Target is not the original configuration node with WWNN of *value*

Note: Messages 6155 and 6156 are only shown with ″-v on″.

An invocation example

```
svcconfig restore -prepare
svcconfig restore -execute
```

The resulting output

```
No feedback
```

Chapter 7. Cluster diagnostic and service-aid commands

Cluster diagnostic and service-aid commands are designed to diagnose and find cluster problems.

The SAN Volume Controller 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.

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

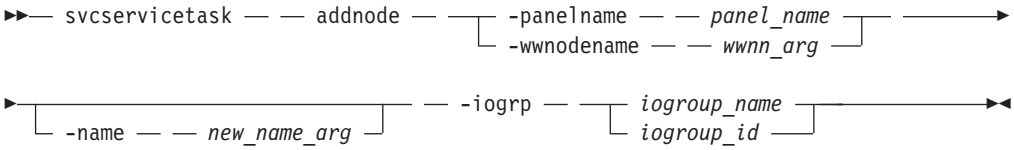

Parameters

-panelname *panel_name*
Identifies the node to be added by the name as it appears on the display panel. 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`.

Node addition completes asynchronously. This means while the node is in adding state, the WWPN is not known and will display as zeroes.

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 node serial number, the WWNN, all WWPN's, and the I/O group to which it has been added. This can prevent possible data corruption 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 the cluster GUI, confirm that the node has previously been a member of the cluster. If it has, follow one of these two procedures:

- 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 **svcinfo lsnode** command. Or,
- if this information is not available, call the support team to add the node back into the cluster without corrupting the data.

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.

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 *Multipath Subsystem Device Driver: User's Guide* for more information.

The SDD driver maintains an association between a vpath and a 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 VDisks.

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 cluster 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 WWNN 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

- CMMVC5707E Required parameters are missing.
- 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].

An invocation example

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

The resulting output

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

applysoftware

The **applysoftware** command upgrades the cluster to a new level of software.

Syntax

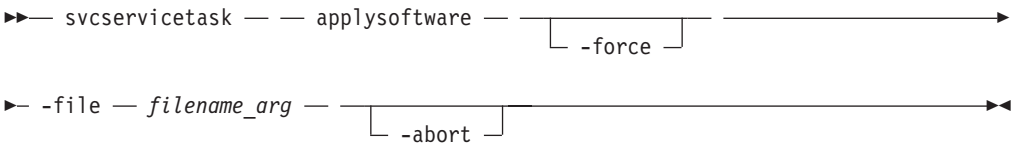

Parameters

-force
Optionally specifies the force flag. The -force flag is needed if any node within an I/O group is not paired. The upgrade process forces the first node in each I/O group to shut down and upgrade. If that node is not paired, then the cluster will become degraded and the data will be lost.

-file *filename_arg*
Specifies the filename of the new software package.

-abort
Optionally specifies the abort flag. The -abort flag is used to force an upgrade to back itself out.

Note: The force flag can be used with the -abort parameter. If one or more nodes are offline, the abort fails unless the -force flag is used.

Description

This command starts the upgrade process of the cluster to a new level of software and can be applied to **svcservicetask** and **svcservicemodetask**. The **applysoftware** command can be used to apply a level of software to the node in both service and non-service modes. In service mode the **applysoftware** command is applied to the specific node in service mode. In non-service mode, the command is applied to the complete cluster.

The software package as specified by the file name must first be copied on to the current configuration node in the /home/admin/upgrade directory. You can use PuTTy secure copy (scp) to copy the file. See "PuTTY scp" for detailed information on this procedure.

The actual upgrade completes asynchronously.

The contents of /home/admin/upgrade can be viewed by using the **svcinfo lssoftwaredumps** command.

Internally, the new package will be moved from the /home/admin/upgrade directory and checksumed. If the package fails the checksum, it will be deleted and the upgrade will fail. Otherwise, the package will be extracted from the directory and the software upgrade will begin.

Possible failures

- 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.
- 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.
- CMMVC6206E The software upgrade failed as a file containing the software for the specified version was not found.
- CMMVC6232E This operation cannot be performed because the cluster is currently aborting the previous software upgrade command.
- CMMVC6233E This operation cannot be performed because the software upgrade is making progress.
- CMMVC6054E The action failed as not all nodes are online.

An invocation example

```
svcservicetask applysoftware -file sanvolumecontroller_update
```

The resulting output

```
No feedback
```

cherrstate

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

Syntax

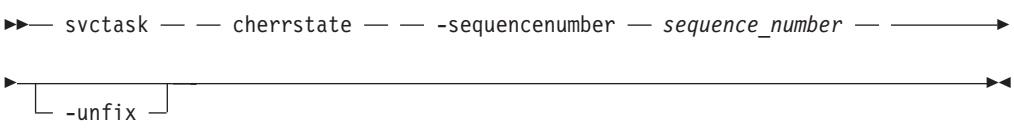

Parameters

-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.

An invocation example

```
svctask cherrstate -sequencenumber 2019
```

The resulting output

```
No feedback
```

clearerrlog

The **clearerrlog** command clears all entries from the error log including status events and any unfixed errors.

Syntax

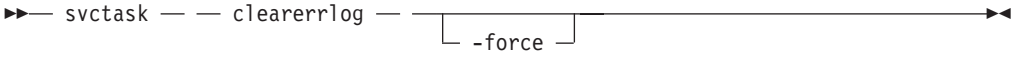

Parameters

-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.

An invocation example

```
svctask clearerrlog -force
```

The resulting output

```
No feedback
```

dumperrlog

The **dumperrlog** command dumps the contents of the error log to a text file.

Syntax

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 within 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.

Files **will not** be deleted from other nodes until you issue the **cleardumps** 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.

An invocation example

```
svcservicetask dumperrlog -prefix testerrorlog
```

The resulting output

```
No feedback
```

enablecli

Use the **enablecli** command to fully enable the CLI for a cluster that has been through a cluster recovery process.

Syntax

Parameters

This command has no parameters.

Description

A cluster that has been through a cluster recovery process has a limited CLI.

The following commands are available:

- svctask cpdumps
- svcinfo ls2145dumps
- svcinfo lserrlogdumps
- svcinfo lsfeaturedumps
- svcinfo lsiostatsdumps
- svcinfo lsiotracedumps
- svcinfo lssoftwaredumps
- svcinfo lssvcdumps

Any attempt to run commands other than those listed above will result in the following error:

```
CMMVC6228E The cluster has been recovered and the CLI functionality is
limited until the cause of the failure has been determined and any
corrective action taken.
```

Contact your service representative.

Possible failures

This command will not fail. It will always enable the CLI.

An invocation example

```
svctask enablecli
```

The resulting output

The user is prompted with the following message:

```
This command will re-enable full functionality of the CLI. This should only
be done once the cause of the cluster failure has been established and
corrected. Do you wish to continue?
```

If you answer yes, the CLI is enabled.

If you answer no, the command does not proceed.

finderr

The **finderr** command analyzes the error log for the highest severity unfixed error.

Syntax

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.

An invocation example

```
svctask finderr
```

The resulting output

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

rmnode

The **rmnode** command deletes a node from the cluster. You can enter this command any time after a cluster has been created.

Syntax

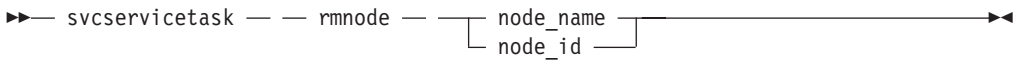

Parameters

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.

Prerequisities:

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 VDisks, where the filter attribute is the I/O group.

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

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

   Note: Any VDisks that are assigned to the I/O group that this node belongs to, will be assigned to the other node in the I/O group; that is, the preferred node will be changed. You cannot change this setting back.
2. Determine the hosts that the VDisks are mapped to by issuing the **svcinfo lsvdiskhostmap** command.
3. Determine if any of the VDisks 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 VDisks, go to step 5.
   - If you *do* want to maintain access to some or all of the VDisks, back up the data or migrate the data to a different (online) I/O group.
4. Determine if you need to turn the power off to the node:
   - If this is the last node in the cluster, you do not need to turn the power off to the node. Go to step 5.
   - If this is *not* the last node in the cluster, turn the power off to the node that you intend to remove. This step ensures that the 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 **Adding a node back into the cluster:** below.
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 *Multipath 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.

Note:

1. If the node you are removing is the configuration node, it may take one or two minutes to complete the command.
2. If the node you are removing is the last node in the cluster, the cluster 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:

Note:

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 be lost if last node in the cluster is deleted. Deleting the configuration node results in the CLI failing over to another node. 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.

- Node serial number
- WWNN

- all WWPNs
- I/O group that contains the node

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

Replacing a faulty node:

You can replace a failed node with a "spare" or replacement node. This would be necessary if the node fails. The cluster will continue to operate with degraded performance until the failed node is repaired. To enhance availability, you might choose to replace the failed node with a "spare", then repair the node offline. However, various procedures must be followed and precautions taken, to replace a failed node without interrupting I/O and without risk to data integrity when the repaired node is reconnected to the SAN fabric. The procedure involves changing the World Wide Node Name (WWNN) of the replacement node. This procedure must be followed with care to avoid duplicate WWNNs which are illegal and can cause data corruption.

Prerequisities:

Before replacing the failed node you must:

- Ensure that the GUI that is used to access the cluster and spare node is running at least version 1.1.1 of the software.
- Know the cluster name that contains the failing node.
- Confirm that a spare node is in the same rack as the cluster containing the failed node.
- Make a record of the last five characters of the original WWNN. This identification is needed if you decide in the future to designate the spare node as a normal node that can be assigned to any cluster. Use the **svcinfo lsnode** command to see the WWNN.

Additional Information

When you replace a node, the following process takes place:

- The node Front Panel ID will change. This is the number printed on the front of the node, used to select the node that is to be added to a cluster.
- The Node Name might change. If you permit the cluster to assign default names when adding nodes to the cluster it will create a new name each time a node is added. If you choose to assign your own names then you must type in the node name that you want to use. If you are using scripts to perform management tasks on the cluster and those scripts use the Node Name then by assigning the original name to a replacement node you will avoid the need to make changes to the scripts following service activity on the cluster.
- The Node ID will change. A new Node ID is assigned each time a node is added to a cluster. The Node ID or the Node Name can be used when performing management tasks on the cluster but if scripts are being used to perform those tasks it is recommended that the Node Name is used in preference to the Node ID since the Node Name will remain unchanged following service activity on the cluster.
- The World Wide Node Name (WWNN) will not change. The WWNN is used to uniquely identify the node and the fiber channel ports. The node replacement

procedure changes the WWNN of the spare node to match that of the failed node. The node replacement procedures must be followed exactly to avoid any duplication of WWNNs.

- The World Wide Port Name (WWPN) of each fiber channel port will not change. The WWPNs are derived from the WWNN that is written to the replacement node as part of this procedure.

Perform the following steps to replace a node:

1. Use the command **svcinfo lsnode** to display the node name. This command, when implemented, prints a detailed list report containing information about all the nodes on a cluster. The failed node will be offline. Note the nodes' name.
2. Again use **svcinfo lsnode** to display the I/O Group name. Note the name of the group.
3. Use the command **svcinfo lsnodevpd** to display the front panel id. Note the id number.
4. Again use the command **svcinfo lsnodevpd** to record the UPS serial number. Note this number.
5. Use the Front Panel ID to locate the failed node. Disconnect all four fiber channel cables from the node.

   Important: The cables must not be reconnected until the node is repaired and the node number has been changed to the default spare node number.
6. Connect the power/signal cable from the spare node to the UPS with the serial number noted in step 1. The signal cable can be plugged into any vacant position on the top row of serial connectors on the UPS. If no spare serial connectors are available on the UPS, disconnect the cables from the failed node. Power-on the spare node. Display the node status on the service panel.

Perform these steps to change the WWNN of the node and add the replacement node:

1. With the node status displayed on the front panel, press and hold the Down button; press and release the Select button; release the Down button. The text "WWNN" is displayed on line-1 of the display. Line-2 of the display contains the last five characters of the WWNN.
2. With the WWNN displayed on the service panel, press and hold the Down button, press and release the Select button, release the Down button. This switches the display into edit mode.
3. Change the displayed number to match the WWNN recorded in step 1. To edit the displayed number use the Up and Down buttons to increase or decrease the numbers displayed. Use the left and right buttons to move between fields. When the five characters match the number recorded in step 1, press the select button twice to accept the number.
4. Connect the four fiber channel cables that were disconnected from the failed node to the spare node. Delete the offline node.
5. Add the spare node into the cluster. See "Adding a node back into the cluster".
6. Use the Subsystem Device Drive (SSD) management tool on the host systems to verify all paths are now online. See the *Service Guide* menu options for further information.

When the failed node is repaired do not connect the fibre channel cables to it. Connecting the cables might cause data corruption. Perform these steps after the failed node is repaired:

1. Display the node status on the service panel.
2. With the status displayed on the front panel, press and hold the Down button; press and release the Select button; release the Down button. The text “WWNN” is displayed on line-1 of the display; line-2 of the display contains the last five characters of the WWNN.
3. With the WWNN displayed on the service panel, press and hold the Down button, press and release the Select button, release the Down button. This switches the display into edit mode.
4. Change the displayed number to “00000”. To edit the displayed number use the Up and Down buttons to increase or decrease the numbers displayed. Use the left and right buttons to move between fields. When the number is set to “00000”, press the select button twice to accept the number. Never connect a node with a WWNN of “00000” to the cluster.

This error should not be marked as fixed in the error log until the failed node has been repaired and returned to the customer. If this is not done, the service technician will not be able to easily find the front panel ID of the failed node.

This node can now be used as a spare node. If this node is no longer required as a spare and is to be used for normal attachment to a cluster you must first use the procedure described above to change the WWNN to the number saved when a spare was being created. See “Prerequisites” above. Using any other number might cause data corruption.

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.

An invocation example

```
svcservicetask rmnode 1
```

The resulting output

```
No feedback
```

setevent

Use the **setevent** command to specify what happens when an error or event is logged to the error log. These settings define what to do when errors and events are logged.

Syntax

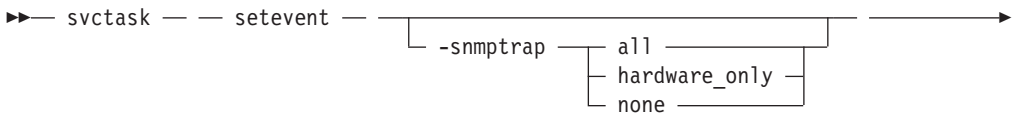

Parameters

-snmptrap *all* | *hardware_only* | *none*
Optionally specifies the SNMP trap setting, that is, when to raise a trap.

-snmpip *ip_address*
Optionally specifies the IP address of the host system that is running the SNMP manager software. This is a colon-separated list of values with up to six items per list.

-community *community*
Optionally specifies the SNMP community string. This is a colon-separated list of values with up to six items per list. The maximum length of the community string that is used in SNMP trap generation cannot be more than 60 characters.

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 argument:

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.
- The IP address of the SNMP manager
- The SNMP community

Possible failures

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

An invocation example

```
svctask setevent -snmptrap all -snmpip 1.2.3.4
-community mysancommunity
```

The resulting output

```
No feedback
```

setlocale

The **setlocale** command changes the locale setting for the cluster. It also changes all interfaces output to the chosen language.

Syntax

```
►►— svcservicetask — — setlocale — — -locale — locale_id ——————————►◄
```

Parameters

-locale *locale_id*
Specifies the locale ID.

Description

This command changes the language in which error messages are displayed as output from the command-line interface. After this, all error messages from the command-line tools will be generated in the chosen language. This command is executed when you request a change of language (locale) and is generally executed from the Web page. Issue the **svcservicetask setlocale** command to change the locale setting for the cluster. It changes all interfaces output to the chosen language. For example, if you wanted to change the English default language to Japanese, type the following:

```
svcservicetask setlocale -locale 3
```

where *3* is the argument that stands for Japanese. The arguments are:

- **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)

Note: This command does not change the front panel display panel settings.

Possible failures

- There are no error codes.

An invocation example

```
svcservicetask setlocale -locale 3
```

The resulting output

```
No feedback
```

svqueryclock

The **svqueryclock** command returns the date, time, and current time-zone of the cluster.

Syntax

```
►►—svqueryclock——————————————————————————————————————————►◄
```

Description

This command returns the date, time and current time-zone of the cluster.

Possible failures

- There are no error codes.

An invocation example

```
svqueryclock
```

The resulting output

```
Mon Nov 25 14:59:28 GMT 2002
```

writesernum

Use the **writesernum** command to write the node serial number into the planar NVRAM.

Syntax

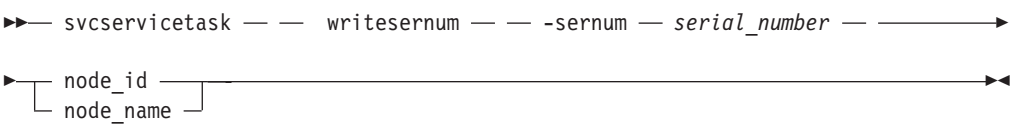

Parameters

-sernum *serial_number*
Specifies the serial number you want to write to the nonvolatile memory of the system planar.

node_id | node_name
Specifies the node where the system planar is located. The serial number will be written to this system planar. This name is not the WWNN.

Description

This command writes the node serial number into the planar NVRAM and then reboots the system. You can find the serial number at the front of the node without having to remove it from the rack. The serial number is located to the left of the right-side thumbscrew that holds the node into the rack. This serial number will usually be seven digits.

Note: Once you have written the serial number to the planar NVRAM, you can issue the **svcinfo lsnodevpd** command to verify that the number is correct. The `system_serial_number` field contains the serial number.

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.

An invocation example

```
svcservicetask writesernum -sernum 1300027 node1
```

The resulting output

```
No feedback
```

Chapter 8. Host commands

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

addhostiogrp

The **addhostiogrp** command enables you to map I/O groups to an existing host object.

Syntax

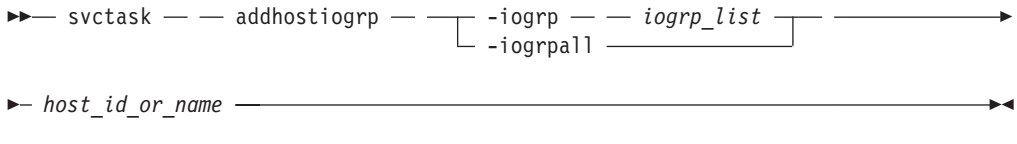

Parameters

-iogrp *iogrp_list*
Specifies a colon separated list of one or more I/O groups that must be mapped to the host. This parameter is mutually exclusive with -iogrpall.

-iogrpall
Specifies that all the I/O groups must be mapped to the specified host. This parameter is mutually exclusive with -iogrp.

host_id_or_name
Identify the host either by ID or name to which the I/O groups must be mapped.

Description

This command will map the list of I/O groups to the specified host object.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5874E The action failed because the host does not exist.
- CMMVC6218E The maximum number of host I/O group pairs for the cluster is already configured.
- CMMVC6220E The maximum number of hosts for one or more I/O groups is already configured.
- CMMVC6224E The host already belongs to one or more of the I/O groups specified.

An invocation example

```
svctask addhostiogrp -iogrpall testhost
```

The resulting output

```
No feedback
```

addhostport

The **addhostport** command adds WWPNs to an existing host object.

Syntax

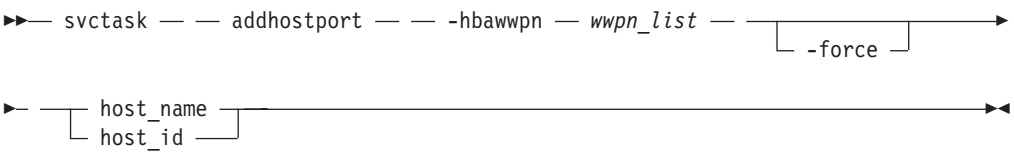

Parameters

-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 -hbawwpn <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 host multipathing device driver documentation 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.

An invocation example

```
svctask addhostport -hbawwpn 210100E08B251DD4 host_one
```

The resulting output

```
No feedback
```

chhost

You can use the **chhost** command to modify the name or type of a host object. This does not affect any current virtual disk-to-host mappings.

Syntax

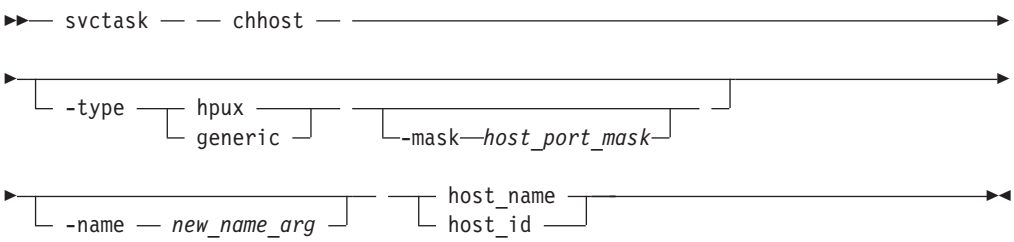

Parameters

-type *hpux | generic*
An optional parameter that specifies the type of host. Valid entries are hpux or generic. The default is generic.

-name *new_name_arg*
Optionally specifies the new name to be assigned to the host object.

-mask *host_port_mask*
An optional mask to indicate which ports on that host are available for use. The mask must be four characters in length and may be made up of a combination of '0' or '1'. '0' indicates that the port can not be used, '1' indicates that it can. The default mask is 1111 (for example, all ports are enabled). The right most bit in the mask corresponds to the lowest numbered cluster WWPN (1 not 4) on a node.

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

Description

This command can change the name of the specified host to a new name, or can change the type of host. 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.

An invocation example

```
svctask chhost -name testhostlode -mask 0011 hostone
```

The resulting output

```
No feedback
```

mkhost

The **mkhost** command creates a logical host object.

Syntax

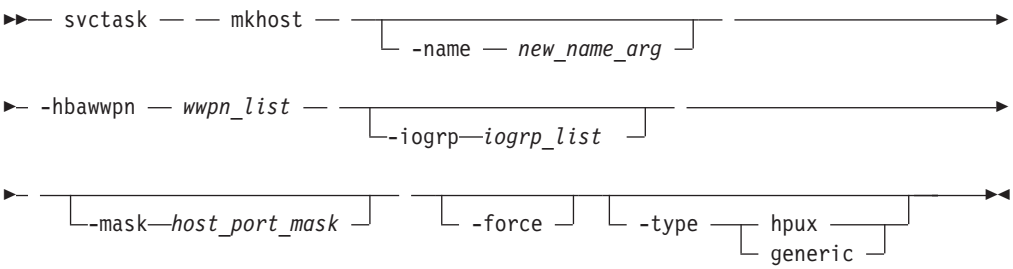

Parameters

-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.

-iogrp *iogrp_list*
Optionally specifies a set of one or more I/O groups that the host accesses VDisks from. I/O groups are specified using their name or ID, separated by a colon. Names and IDs can be mixed in the list. If the parameter is omitted, the host will be associated with all I/O groups.

-mask *host_port_mask*
Optionally specifies which ports the host object can access. The port mask must be four characters in length and may be made up of a combination of '0' or '1'. '0' indicates that the port can not be used, '1' indicates that it can. The default mask is 1111 (for example, all ports are enabled). The right most bit in the mask corresponds to the lowest numbered WWPN (1 not 4) on a node.

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

-type *hpux | generic*
An optional parameter that specifies the type of host. Valid entries are hpux or generic. The default is generic.

Description

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

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 **svcinfo lshbaportcandidate** commands.

This command will fail if you add the host to an I/O group that is associated with more host ports or host objects than is allowed by the limits within the cluster.

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.
- CMMVC6217E The maximum number of hosts for the cluster is already configured.

- CMMVC6218E The maximum number of host I/O group pairs for the cluster is already configured.
- CMMVC6219E The maximum number of WWPNs for the cluster is already configured.
- CMMVC6220E The maximum number of hosts for one or more I/O groups is already configured.
- CMMVC6221E The maximum number of WWPNs for one or more I/O groups is already configured.
- CMMVC6222E The maximum number of WWPNs for the host is already configured.

An invocation example

```
svctask mkhost -name hostone -hbawwpn 210100E08B251DD4 -force -mask 1001
```

The resulting output

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

rmhost

The **rmhost** command deletes a host object.

Syntax

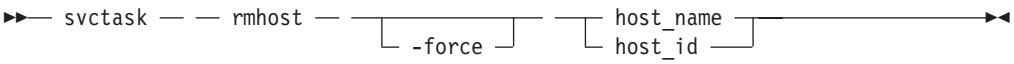

Parameters

-force
Optionally forces the delete. If any mappings still exist between this host and virtual disks (VDisks), the command will fail unless the force flag is specified. When specified, this deletes the mappings before deleting the host object.

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.

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.

- CMMVC5871E The host object was not deleted because there are virtual disk to host mappings for this host. To delete this host, a forced deletion is required.

An invocation example

```
svctask rmhost host_one
```

The resulting output

```
No feedback
```

rmhostiogrp

The **rmhostiogrp** command enables you to delete mappings between one or more I/O groups and a specified host object.

Syntax

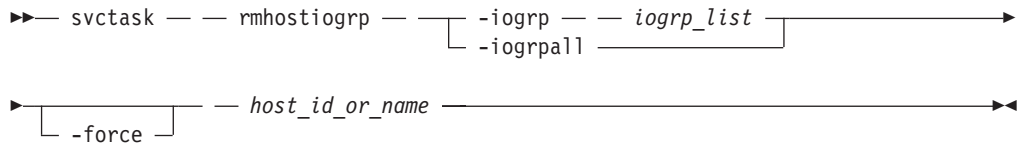

Parameters

-iogrp *iogrp_list*
Specifies a colon separated list of one or more I/O group mappings that must be deleted from the host. This parameter is mutually exclusive with -iogrpall.

-iogrpall
This parameter is mutually exclusive with -iogrp. It specifies that all the I/O group mappings must be deleted from the specified host.

-force
If the removal of a host to I/O group mapping will result in the loss of VDisk to host mappings then the command must fail if the -force flag has *not* been used. The -force flag will, however, override such behavior and force the host to I/O group mapping to be deleted.

host_id_or_name
Identify the host either by ID or name from which the I/O group mappings must be deleted.

Description

This command will delete the mappings between the list of I/O groups and the specified host object.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5874E The action failed because the host does not exist.
- CMMVC6223E The host does not belong to one or more of the I/O groups specified or inferred

An invocation example

```
svctask rmhostiogrp -iogrp 1:2 host0
```

The resulting output

```
No feedback
```

rmhostport

You can use the **rmhostport** command to delete worldwide port names (WWPNs) from an existing host object.

Syntax

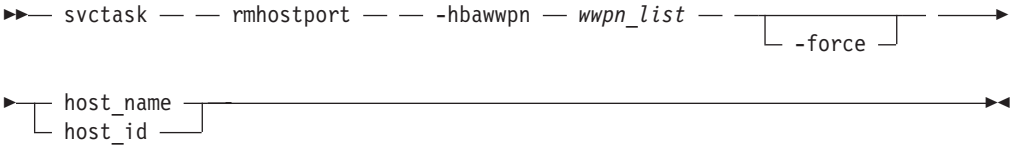

Parameters

-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 -hbawwpn <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 Multipath 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.

An invocation example

```
svctask rmhostport -hbawwpn 210100E08B251DD4 host_one
```

The resulting output

```
No feedback
```

Chapter 9. Virtual disk commands

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

chvdisk

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

Syntax

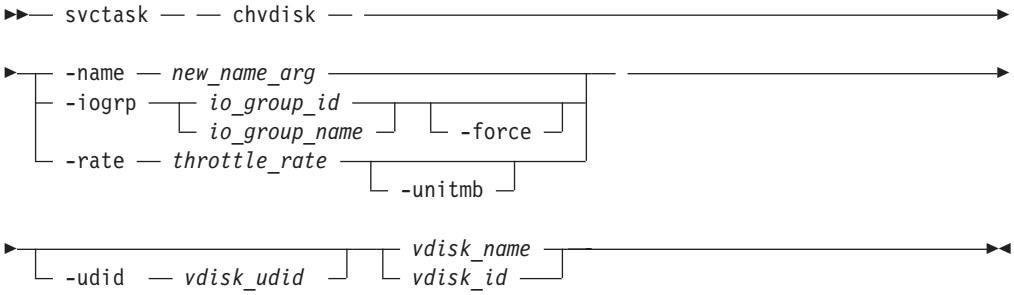

Parameters

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

If the VDisk has a mapping to any hosts, it will not be possible to move the VDisk to an I/O group which does not include any of those hosts.

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

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

-rate *throttle_rate [-unitmb]*
Optionally sets the I/O governing rates for the virtual disk. The default units are I/Os, but they 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.

-udid *vdisk_udid*
Optionally specifys the udid for the disk. Valid options are a decimal number from 0 to 32767, or a hex number from 0 to 0x7FFF. A hex number must be preceded by '0x' (for example, 0x1234). If this parameter is omitted then the default udid is 0.

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

Note: The -iogrp, -rate, -udid and -name parameters are mutually exclusive. 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 at the host level before performing this operation.

Attention: Under no circumstances should you move a VDisk to an offline I/O group. You must ensure the I/O group is online before moving the VDisks to avoid any data loss scenarios.

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

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

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

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

Attention:

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

Make sure that when you migrate a VDisk to a new I/O group, you quiesce all I/O operations for the VDisk. You may need to determine the hosts that are using this VDisk. Any FlashCopy mappings or Metro Mirror relationships that use this VDisk should be stopped or deleted. Issue the following command, to check if the VDisk is part of a relationship or mapping, issue the **svcinfo lsvdisk <vdiskname/id>** command, where <vdiskname/id> is the name or ID of the VDisk.

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

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

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

Possible failures

- CMMVC5756E Cannot perform the request as the object is already mapped.
- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5832E The property of the virtual disk (VDisk) was not modified because an entity that was specified in the command does not exist.
- CMMVC5833E The property of the virtual disk (VDisk) was not modified because there are no nodes in the I/O group.
- CMMVC5834E The I/O group for the virtual disk (VDisk) was not modified because the group is a recovery I/O group. To modify the I/O group, use the force option.
- CMMVC5848E The action failed because the virtual disk (VDisk) does not exist or it is being deleted.
- CMMVC5853E The action failed because there was a problem with the group.
- CMMVC5856E The action failed because the virtual disk (VDisk) does not belong to the specified managed disk (MDisk) group.
- CMMVC5857E The action failed because the managed disk (MDisk) does not exist or it is not a member of the managed disk (MDisk) group.
- CMMVC5858E The action failed because the virtual disk (VDisk) is in the wrong mode, the managed disk (MDisk) is in the wrong mode, or both are in the wrong mode.
- CMMVC5860E The action failed because there were not enough extents in the managed disk (MDisk) group.
- CMMVC5861E The action failed because there were not enough extents on the managed disk (MDisk).
- CMMVC5862E The action failed because the virtual disk (VDisk) is being formatted.
- CMMVC6032E The operation was not performed because one or more of the entered parameters is invalid for this operation.
- CMMVC6076E The VDisk cache is not empty. Wait for the cache to flush or use the force flag to discard contents of the cache.
- CMMVC6223E The host does not belong to one or more of the I/O groups specified or inferred.

An invocation example

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

The resulting output

```
No feedback
```

expandvdisksize

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

Syntax

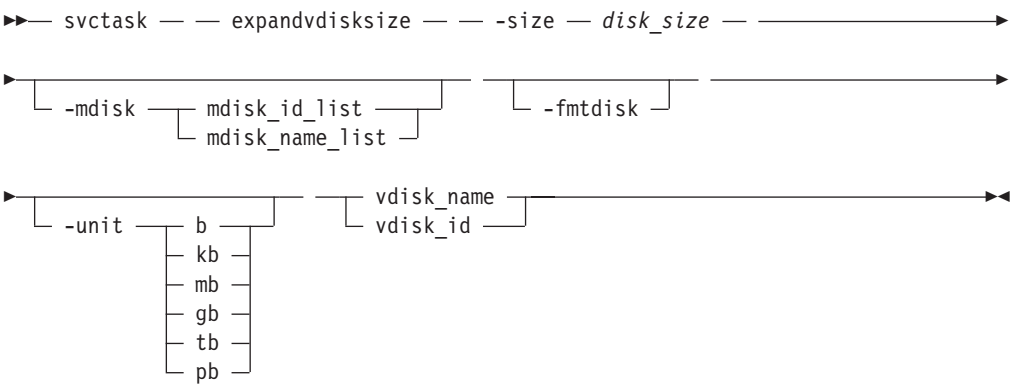

Parameters

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

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

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

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

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

Description

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

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

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5808E The action failed because the managed disk (MDisk) does not exist.
- CMMVC5835E The virtual disk (VDisk) was not expanded because an object that was specified in the command does not exist.
- CMMVC5837E The action failed because the virtual disk (VDisk) is part of a FlashCopy mapping.
- CMMVC5838E The action failed because the virtual disk (VDisk) is part of a Remote Copy mapping.
- CMMVC5848E The action failed because the virtual disk (VDisk) does not exist or it is being deleted.
- CMMVC5853E The action failed because there was a problem with the group.
- CMMVC5856E The action failed because the virtual disk (VDisk) does not belong to the specified managed disk (MDisk) group.
- CMMVC5857E The action failed because the managed disk (MDisk) does not exist or it is not a member of the managed disk (MDisk) group.
- CMMVC5858E The action failed because the virtual disk (VDisk) is in the wrong mode, the managed disk (MDisk) is in the wrong mode, or both are in the wrong mode.
- CMMVC5860E The action failed because there were not enough extents in the managed disk (MDisk) group.
- CMMVC5861E The action failed because there were not enough extents on the managed disk (MDisk).
- CMMVC5862E The action failed because the virtual disk (VDisk) is being formatted.
- CMMVC5998W The virtualized storage capacity exceeds the amount that you are licensed to use.
- CMMVC6036E An invalid action was requested

An invocation example

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

The resulting output

```
No feedback
```

mkvdisk

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

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

Syntax

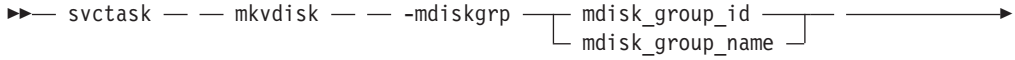

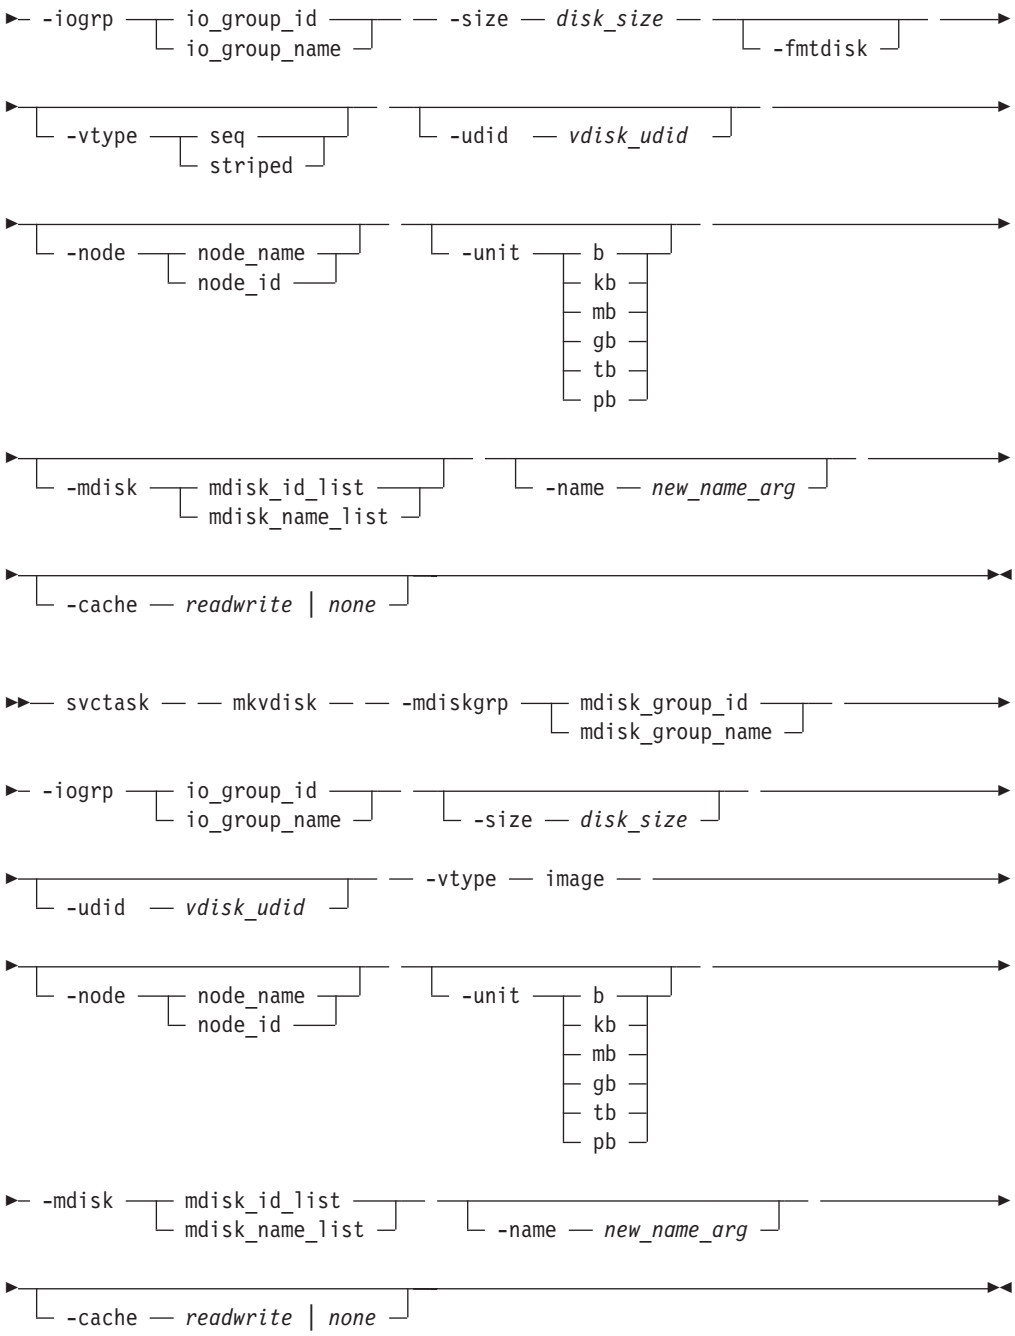

Parameters

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

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

-udid *vdisk_udid*
Optionally specifies the udid for the disk. Valid options are a decimal number from 0 to 32767, or a hex number from 0 to 0x7FFF. A hex number must be preceded by '0x' (for example, 0x1234). If this parameter is omitted then the default udid is 0.

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

-fmtdisk
Optionally specifies that the virtual disk should be formatted before use. The -fmtdisk argument formats (sets to all zeros) the extents that make up this VDisk after it is created. If this parameter is used, the command completes asynchronously and you can query the status with the **svcinfo lsvdiskprogress** command. You cannot use this flag when you are creating an image mode VDisk.

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

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

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

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

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

-cache *readwrite* | *none*
Optionally specifies the caching options for the VDisk. Valid entries are readwrite or none. The default is readwrite. If -cache is not entered, the default is used.

Description

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

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

Decide which I/O group the VDisk should be assigned to. This determines which 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 nodes. Use the **svcinfo lsiogrp** command to show the I/O groups and the number of virtual disks assigned to each I/O group.

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

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

Striped
This is the default policy. If the -vtype is not specified, then this policy is used in its default form. That is, all managed disks in the managed disk group will be used to create the virtual disk. The striping is at an extent level, in a circular fashion, one extent from each managed disk in the group is used. For example, a managed disk group with 10 managed disks uses one extent from each managed disk, then it uses the 11th extent from the first managed disk, and so on.

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

A capacity of 0 is allowed.

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

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

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

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

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

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

The command returns the IDs of the newly created VDisk.

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

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

Possible failures

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

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

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

An invocation example

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

The resulting output

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

An invocation example for creating an image mode VDisk

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

The resulting output

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

mkvdiskhostmap

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

Syntax

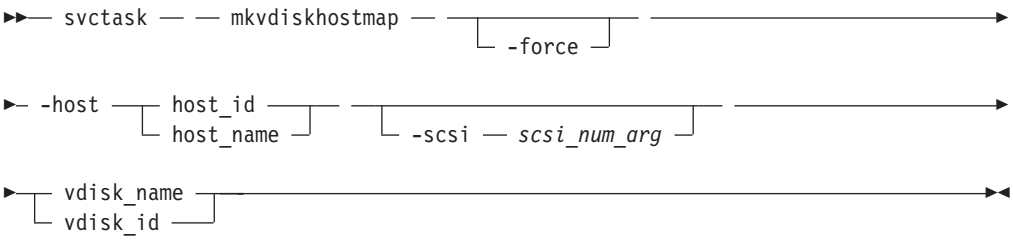

Parameters

-force
The system does not allow multiple VDisk to host assignments. Using the -force flag will override this behavior and allow multiple assignments to be made.

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

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

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

Description

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

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

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

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

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

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

svctask mkvdiskhostmap -host host1 -force 4

svctask mkvdiskhostmap -host host2 -force 4

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

The command will fail if the host object (to which this mapping is being made) is not associated with the I/O group containing the VDisk.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5842E The action failed because an entity that was specified in the command does not exist.
- CMMVC5843E The virtual disk (VDisk)-to-host mapping was not created because the VDisk does not have a capacity greater than zero bytes.
- CMMVC5844E The virtual disk (VDisk)-to-host mapping was not created because the SCSI logical unit number (LUN) ID is not valid.
- CMMVC5862E The action failed because the virtual disk (VDisk) is being formatted.
- CMMVC5874E The action failed because the host does not exist.
- CMMVC5875E The action failed because the virtual disk (VDisk) does not exist.
- CMMVC5876E The virtual disk (VDisk)-to-host mapping was not created because the maximum number of mappings has been reached.
- CMMVC5877E The virtual disk (VDisk)-to-host mapping was not created because the maximum number of SCSI LUNs has been allocated.

- CMMVC5878E The virtual disk (VDisk)-to-host mapping was not created because this VDisk is already mapped to this host.
- CMMVC5879E The virtual disk (VDisk)-to-host mapping was not created because this VDisk is already mapped to this host with this SCSI LUN.
- CMMVC5880E The virtual disk (VDisk)-to-host mapping was not created because the VDisk has a capacity of zero bytes.
- CMMVC6071E This action will result in the creation of multiple mappings. Use the -force flag if you are sure that this is what you wish to do.

An invocation example

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

The resulting output

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

rmvdisk

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

Syntax

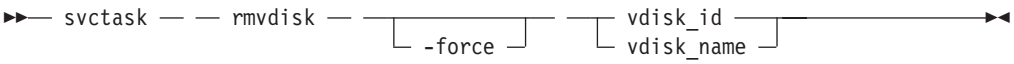

Parameters

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

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

Description

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

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

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

Deleting a managed mode virtual disk

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

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

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

Deleting an image mode virtual disk

When you use this command to delete an image mode virtual disk, the data on the controller logical unit will be consistent with the data that had been available on the image mode virtual disk before it was deleted. In other words, the fast write data will be moved to the controller logical unit. If the force flag is used, then the data is not moved to the controller logical unit.

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

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

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

Possible failures

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

An invocation example

```
svctask rmvdisk -force vdisk5
```

The resulting output

```
No feedback
```

rmvdiskhostmap

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

Syntax

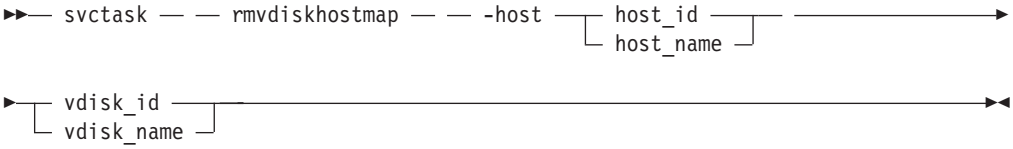

Parameters

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

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

Description

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

This command also deletes a SCSI or persistent reservation that a host has on a VDisk. Once the reservation is removed, a new host is allowed to access the VDisk in the future since the original host no longer has access.

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

Possible failures

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

An invocation example

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

The resulting output

```
No feedback
```

shrinkvdisksize

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

Syntax

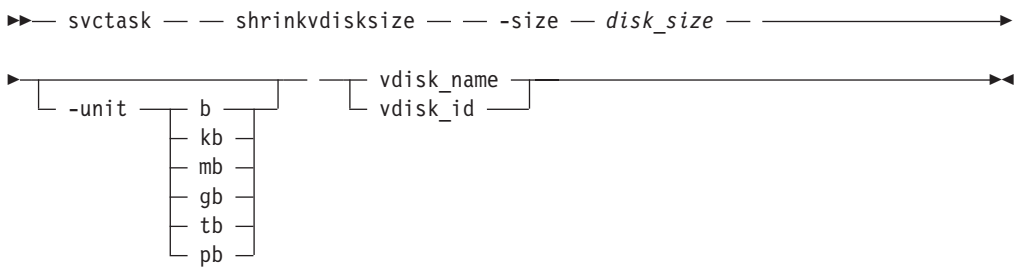

Parameters

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

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

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

Description

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

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

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

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

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

Validate that the VDisk is not mapped to any host objects. If the VDisk is mapped, data is displayed. You can determine the exact capacity of the source or master VDisk by issuing the **svcinfo lsvdisk -bytes <vdiskname>** command. Shrink the

VDisk by the required amount by issuing the **svctask shrinkvdisksize -size <capacitytoshrinkby> -unit <unitsforreduction> <vdiskname/ID>** command.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5836E The virtual disk (VDisk) was not shrunk because it is locked.
- CMMVC5837E The action failed because the virtual disk (VDisk) is part of a FlashCopy mapping.
- CMMVC5838E The action failed because the virtual disk (VDisk) is part of a Remote Copy mapping.
- CMMVC5839E The virtual disk (VDisk) was not shrunk because an entity that was specified in the command does not exist.
- CMMVC5848E The action failed because the virtual disk (VDisk) does not exist or it is being deleted.
- CMMVC5862E The action failed because the virtual disk (VDisk) is being formatted.
- CMMVC6010E Unable to complete the command as there are insufficient free extents.

An invocation example

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

The resulting output

```
No feedback
```

Chapter 10. Managed disk group commands

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

addmdisk

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

Syntax

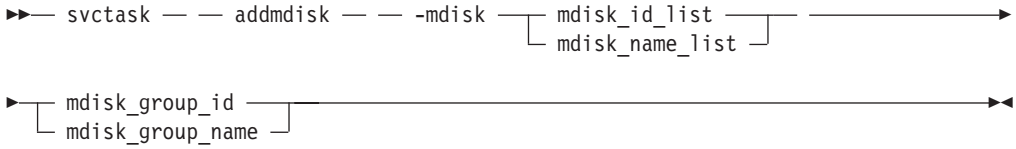

Parameters

-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.

An invocation example

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

The resulting output

```
No feedback
```

chmdiskgrp

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

Syntax

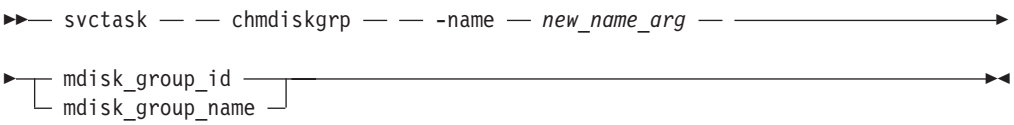

Parameters

-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.

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

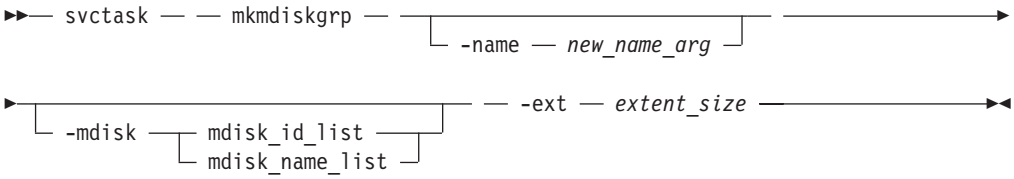

Parameters

-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 **svcinfo lsmdiskcandidate** 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 cluster 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 (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.

An invocation example

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

The resulting output

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

rmmdisk

The **rmmdisk** command deletes a managed disk from a managed disk group. There are some constraints on this command.

Syntax

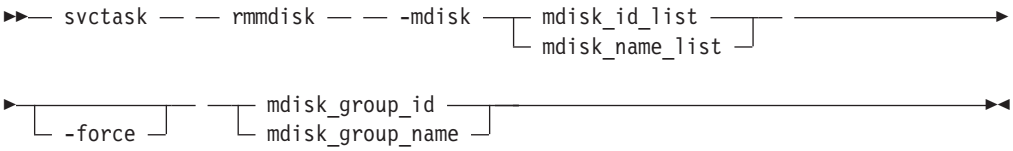

Parameters

-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.

The **rmmdisk** command will fail if there are insufficient free extents on other disks in the mdisk group for the duration of the command. To avoid this problem, do not issue new commands that use extents until **rmmdisk** is completed.

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.

An invocation example

```
svctask rmmdisk -mdisk mdisk12 -force Group3
```

The resulting output

```
No feedback
```

rmmdiskgrp

The **rmmdiskgrp** command deletes a managed disk group. Care should be taken when using this command, because when executed, the command literally 'destroys' the specified managed disk group.

Syntax

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.

An invocation example

```
svctask rmmdiskgrp -force Group3
```

The resulting output

```
No feedback
```

Chapter 11. Managed disk commands

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

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

chmdisk

Use the **chmdisk** command to modify the name of a managed disk.

Syntax

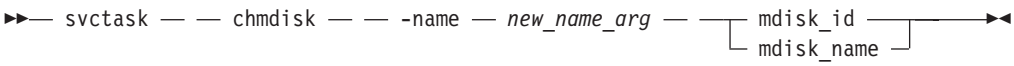

Parameters

-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.

An invocation example

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

The resulting output

```
No feedback
```

includemdisk

Use the **includemdisk** command to include a disk that has been excluded by the cluster.

Syntax

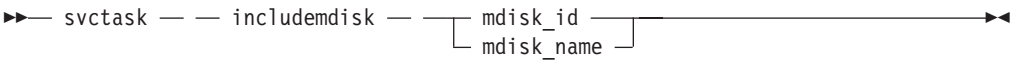

Parameters

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 an MDisk may change its state, regardless of whether the state is reported as excluded.

Note: If an MDisk is in the excluded state, is offline and does not belong to an MDisk group, issuing an include for this MDisk will result in the MDisk record being deleted from the cluster.

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.

An invocation example

```
svctask includemdisk mdisk5
```

The resulting output

```
No feedback
```

setquorum

Use the **setquorum** command to change the managed disks (MDisks) that are assigned as quorum candidate disks.

Syntax

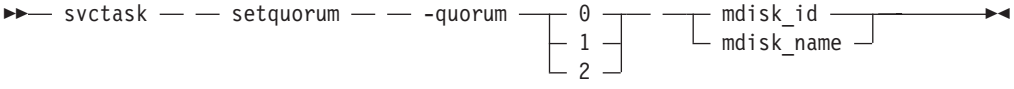

Parameters

-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 to the specified quorum index.

The quorum disk is used by the cluster as a tie breaker when exactly half of the nodes that were previously a member of the cluster are now present. The use of a quorum disk prevents the cluster from being split exactly in half. Both halves either continue to operate or both halves stop. There is only *one* quorum disk; however, the cluster uses three disks as quorum candidate disks. The cluster will select the actual quorum disk from the pool of quorum candidate disks. The quorum candidate disks also hold a copy of important cluster metadata. 256 MB of contiguous space is reserved for this purpose on each quorum candidate disk.

When you issue this command, the MDisk that currently is assigned the quorum index number is set to a non-quorum disk. The cluster will automatically assign quorum indexes.

This command can be used if you want to ensure that a particular set of MDisks continue to be accessible given a cluster split.

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.

An invocation example

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

The resulting output

```
No feedback
```

Chapter 12. FlashCopy commands

The following commands enable you to work with FlashCopy methods and functions with the SAN Volume Controller.

chfcconsistgrp

Use the **chfcconsistgrp** command to modify the name of an existing consistency group.

Syntax

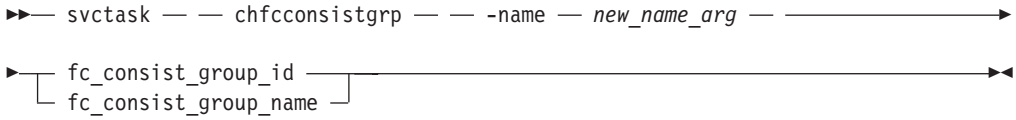

Parameters

-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.

An invocation example

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

The resulting output

```
No feedback
```

chfcmap

Use the **chfcmap** command to modify certain attributes of an existing mapping.

Syntax

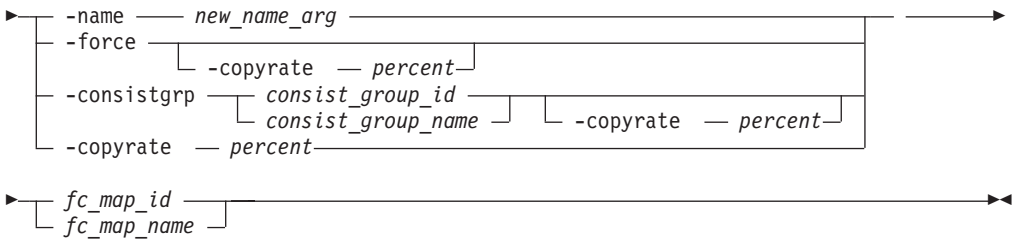

Parameters

-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 percent** 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. Expressed as a percentage. The default is 50.

-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.
- CMMVC6215E The FlashCopy mapping was not created or modified because the consistency group already contains the maximum number of mappings.

An invocation example

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

The resulting output

```
No feedback
```

mkfcconsistgrp

The **mkfcconsistgrp** command creates a new FlashCopy consistency group.

Syntax

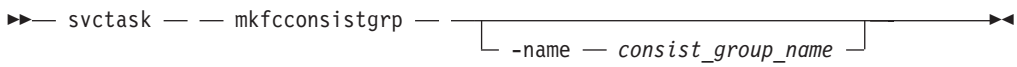

Parameters

-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.

An invocation example

```
svctask mkfcconsistgrp
```

The resulting output

```
FlashCopy Consistency Group, id [1], successfully created
```

mkfcmap

The **mkfcmap** command creates a new FlashCopy mapping, which maps a source virtual disk to a target virtual disk to ready for subsequent copying.

Syntax

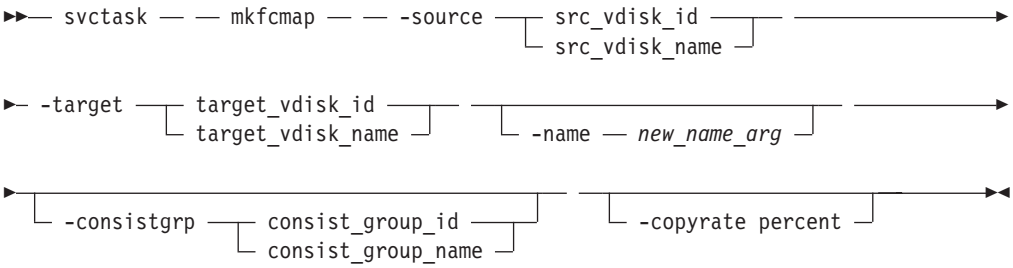

Parameters

-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 -copyrate. Expressed as a percentage. The default is 50.

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 **svcinfo lsvdisk -bytes** command to find the exact size of the source Vdisk that you want to create a target disk of the same size. The source and destination cannot be in an existing mapping. That is, a virtual disk can be either a source or a destination disk in **only one** mapping. A mapping is triggered at the point in time when the copy is required.

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.
- 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.
- CMMVC6215E The FlashCopy mapping was not created or modified because the consistency group already contains the maximum number of mappings.

An invocation example

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

The resulting output

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

prestartfcconsistgrp

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

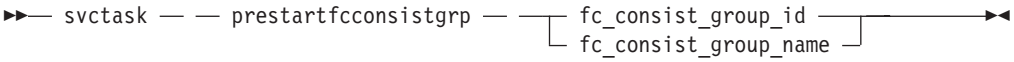

Parameters

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 specify the map_id | name argument.

Description

This command prepares a group of mappings (at a consistency group level) 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.

An invocation example

```
svctask prestartfcconsistgrp 1
```

The resulting output

```
No feedback
```

prestartfcmap

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

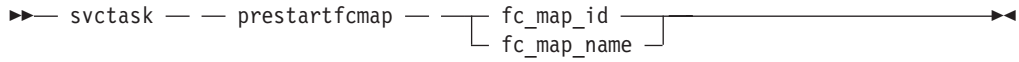

Parameters

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.

Attention: This command may take a considerable amount of time to complete. Please be patient.

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.

An invocation example

```
svctask prestartfcmap 1
```

The resulting output

```
No feedback
```

rmfcconsistgrp

The **rmfcconsistgrp** command deletes a FlashCopy consistency group.

Syntax

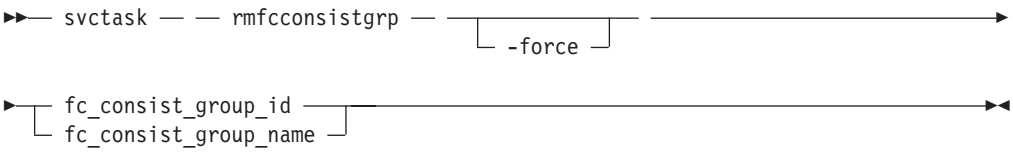

Parameters

-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 specified Flash Copy consistency group. 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.

An invocation example

```
svctask rmfcconsistgrp fcconsistgrp1
```

The resulting output

```
No feedback
```

rmfcmap

The **rmfcmap** command deletes an existing mapping.

Syntax

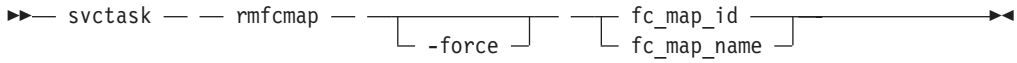

Parameters

-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 affect 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.

An invocation example

```
svctask rmfcmap testmap
```

The resulting output

```
No feedback
```

startfcconsistgrp

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

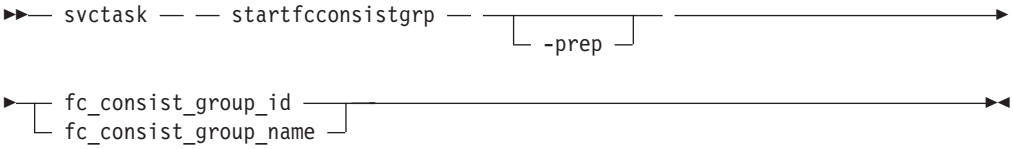

Parameters

-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.

An invocation example

```
svctask startfcconsistgrp -prep 2
```

The resulting output

```
No feedback
```

startfcmap

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

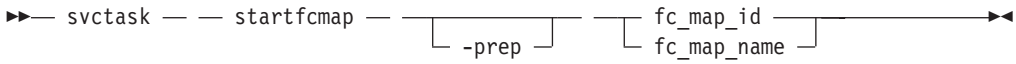

Parameters

-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.

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

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.

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

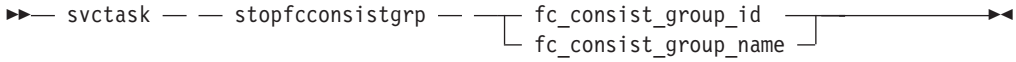

Parameters

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.

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

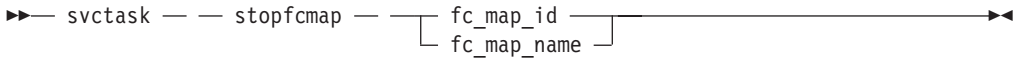

Parameters

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 retriggered.

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.

An invocation example

```
svctask stopfcmap testmapone
```

The resulting output

```
No feedback
```

Chapter 13. Metro Mirror commands

The following commands enable you to work with the Metro Mirror services provided by the SAN Volume Controller.

chpartnership

Use the **chpartnership** command to specify the bandwidth available for background copy in a cluster partnership that has been created for Metro Mirror purposes.

Syntax

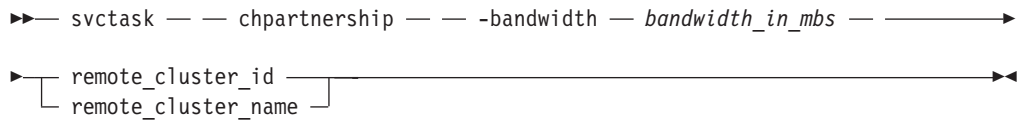

Parameters

-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 Metro Mirror 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.

An invocation example

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

The resulting output

```
No feedback
```

chrcconsistgrp

Use the **chrcconsistgrp** command to modify the name of an existing Metro Mirror consistency group.

Syntax

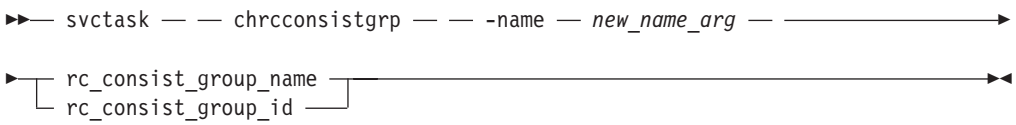

Parameters

-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.

An invocation example

Change the name of the Metro Mirror consistency group called ″rc_testgrp″ to ″rctestone″.

```
svctask chrcconsistgrp -name rctestone rc_testgrp
```

The resulting output

```
No feedback
```

chrcrelationship

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

Syntax

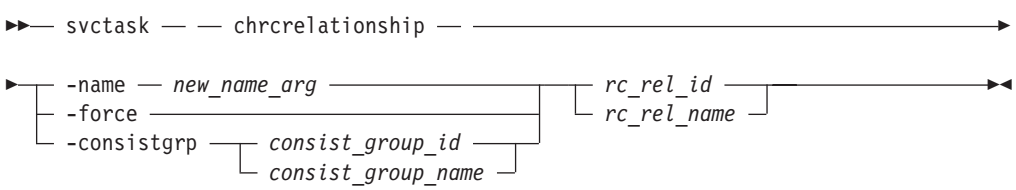

Parameters

-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 three 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 (**chrcrelationship**) 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.

An invocation example

Change the name of the relationship `rccopy1` to `testrel`

```
svctask chrcrelationship -name testrel rccopy1
```

Add relationship `rccopy2` to group called `newgroup`.

```
svctask chrcrelationship -consistgrp newgroup rccopy2
```

Remove relationship `rccopy3` from whichever consistency group it is a member of.

```
svctask chrcrelationship -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 Metro Mirror partnership between the local cluster and a remote cluster.

Syntax

To establish a fully functional Metro Mirror partnership, you must issue this command to both clusters. This step is a prerequisite to creating Metro Mirror relationships between VDisks on the clusters.

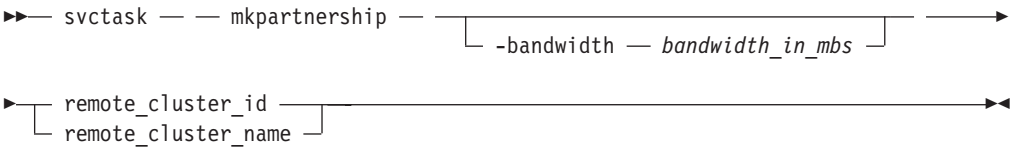

Parameters

-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 Metro Mirror 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 **svcinfo lsclustercandidate** 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 that you specify in the command. To create a two-way partnership, the equivalent **svctask mkpartnership** command must be issued from the other cluster.

Intercluster Metro Mirror relationships can be created between primary VDisks in the local cluster and auxiliary VDisks in the remote cluster. Intracluster relationships can be created between VDisks that reside in a local cluster. The VDisks must belong to the same I/O group within the 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.

An invocation example

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

The resulting output

```
No feedback
```

mkrcconsistgrp

The **mkrcconsistgrp** command creates a new, empty Metro Mirror consistency group.

Syntax

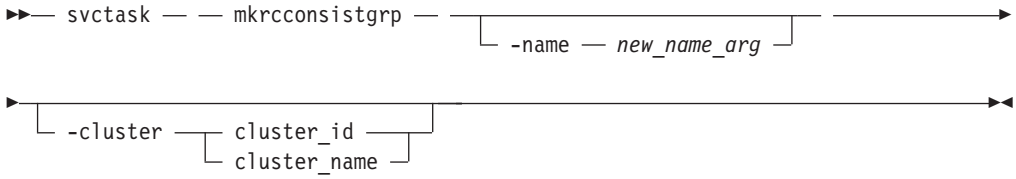

Parameters

-name *new_name_arg*
Optionally specifies a name for the new consistency group.

-cluster *cluster_id | cluster_name*
To create an intercluster consistency group, enter the ID or name of the remote cluster. If -cluster is not specified, then an intracluster 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. Metro Mirror 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.

An invocation example

```
svctask mkrcconsistgrp -name rc_testgrp
```

The resulting output

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

mkrcrelationship

The **mkrcrelationship** command creates a new Metro Mirror relationship with virtual disks (VDisks) in the same cluster (intracluster relationship) or in two different clusters (intercluster relationship).

Syntax

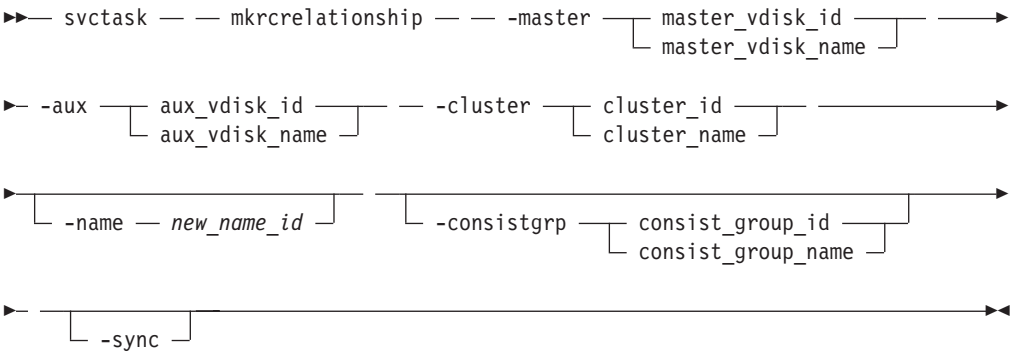

Parameters

-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*
Specifies the ID or name of the remote cluster.

If you are creating an intracluster relationship, enter the ID of the local cluster. The VDisks in the relationship must belong to the same I/O group within the cluster.

If you are creating an intercluster relationship, enter the ID of the remote cluster. To create a relationship in two different clusters, the clusters must be connected at the time that the **svctask mkrcrelationship** 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 Metro Mirror 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 Metro Mirror consistency group. A consistency group is used to ensure that a number of relationships are managed so that, in the event of a disconnnection 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 Metro Mirror relationships stop copying data from the primary to the secondary site, there is no assurance 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 VDdisks into a consistency group. Metro Mirror 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 an assurance 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.
- CMMVC6024E The auxiliary VDisk entered is invalid.
- CMMVC5965E The virtual disks (VDisks) are in different I/O groups on the local cluster.

An invocation example

```
svctask mkrcrelationship -master vdisk1 -aux vdisk2 -name rccopy1
 -cluster 0000020063432AFD
```

The resulting output

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

rmpartnership

The **rmpartnership** command removes a Metro Mirror partnership.

Syntax

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 Metro Mirror partnership will enter a partially configured state and Metro Mirror activity will cease as the relationships become disconnected.

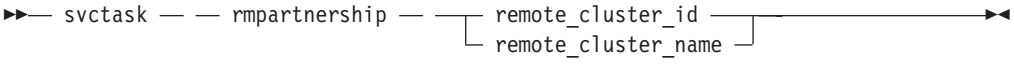

Parameters

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.

An invocation example

```
svctask rmpartnership cluster1
```

The resulting output

```
No feedback
```

rmrcconsistgrp

The **rmrcconsistgrp** command deletes an existing Metro Mirror consistency group.

Syntax

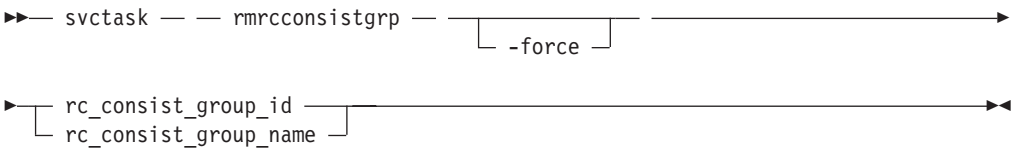

Parameters

-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 want 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.

An invocation example

```
svctask rmrcconsistgrp rctestone
```

The resulting output

```
No feedback
```

rmrcrelationship

The **rmrcrelationship** command deletes an existing Metro Mirror relationship.

Syntax

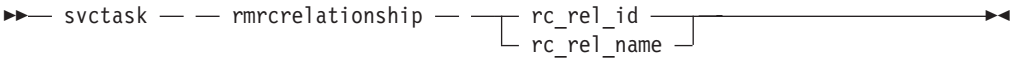

Parameters

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 want to remove the relationship on both clusters, you can issue the **svctask 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 chrcrelationship -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 Metro Mirror 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.

An invocation example

```
svctask rmrcrelationship rccopy1
```

The resulting output

```
No feedback
```

startrcconsistgrp

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

Syntax

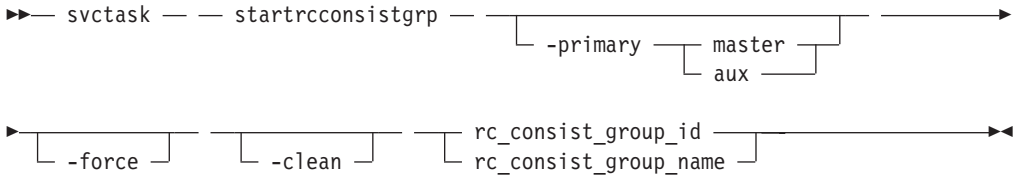

Parameters

-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 Metro Mirror 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 is 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 Metro Mirror 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 synchronizing 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 **stoprcconsistgrp** 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 **startrcconsistgrp** 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.

An invocation example

```
svctask startrcconsistgrp rccopy1
```

The resulting output

```
No feedback
```

startrcrelationship

Use the **startrcrelationship** command to start the Metro Mirror relationship copy process, set the direction of copy if undefined, and optionally mark the secondary VDisk of the relationship as clean.

Syntax

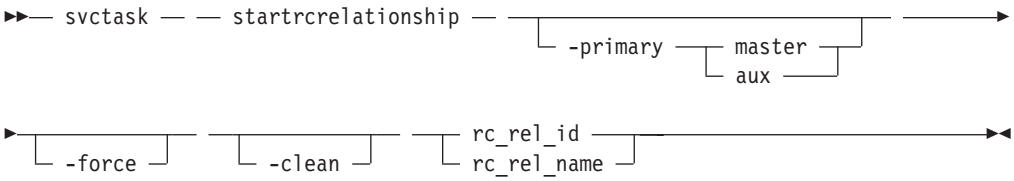

Parameters

-primary *master | aux*
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 Metro Mirror 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 synchronizing 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 **svctask stoprcrelationship** 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 **svctask startrcrelationship** 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.

An invocation example

```
svctask startrcrelationship rccopy1
```

The resulting output

```
No feedback
```

stoprcconsistgrp

You can use the **stoprcconsistgrp** command to stop the copy process in a Metro Mirror 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

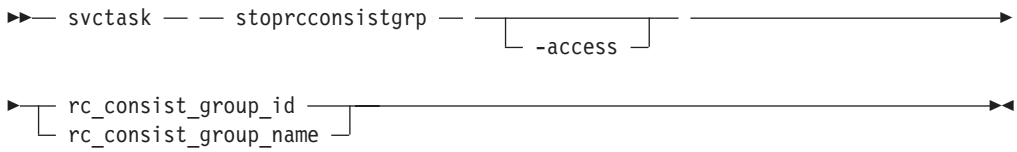

Parameters

-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 startrcconsistgrp** 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 stoprcconsistgrp command to enable write access to the secondary virtual disks within that group.

Initial state	Final state	Notes
InconsistentStopped	InconsistentStopped	
InconsistentCopying	InconsistentStopped	

Initial state	Final state	Notes
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 stoprcconsistgrp** command.
InconsistentDisconnected	unchanged	On the disconnected cluster.
ConsistentDisconnected	ConsistentStopped	On the cluster issuing the **svctask stoprcconsistgrp** 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.

An invocation example

```
svctask stoprcconsistgrp rccopy1
```

The resulting output

```
No feedback
```

stoprcrelationship

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

Syntax

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 startrcrelationship** 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 stoprcrelationship 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 stoprcrelationship** command.
InconsistentDisconnected	unchanged	On the disconnected cluster.
ConsistentDisconnected	ConsistentStopped	On the cluster issuing the **svctask stoprcrelationship** 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.

An invocation example

```
svctask stoprcrelationship 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 Metro Mirror 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

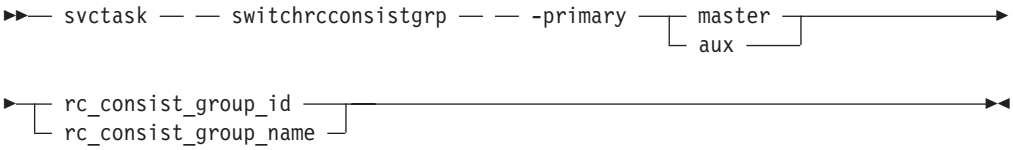

Parameters

-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.

An invocation example

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

The resulting output

```
No feedback
```

switchrcrelationship

Use the **switchrcrelationship** to reverse the roles of primary and secondary virtual disks in a Metro Mirror relationship when that relationship is in a consistent state.

Syntax

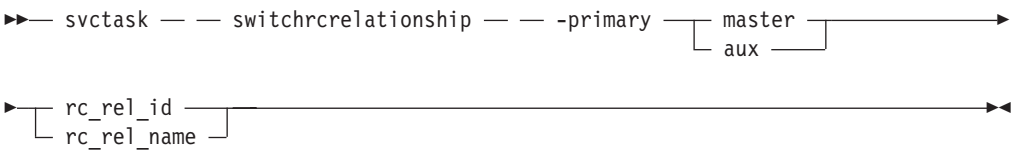

Parameters

-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.

An invocation example

```
svctask switchrcrelationship -primary master rccopy2
```

The resulting output

```
No feedback
```

Chapter 14. Migration commands

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

migrateexts

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

Syntax

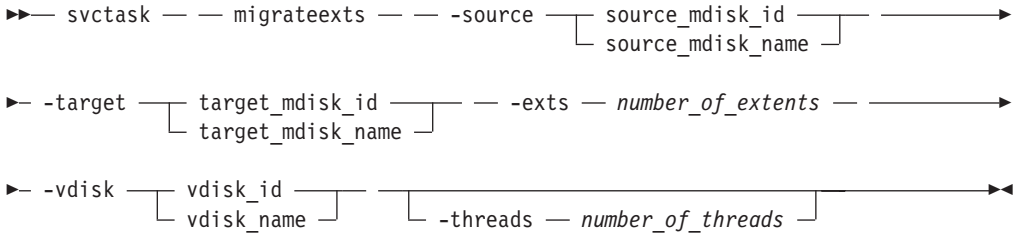

Parameters

-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. You can specify 1 - 4 threads. The default of the number of threads is 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 **svcinfo lsmigrate** command.

The **svctask migrateexts** command will fail if there are insufficient free extents on the target managed disk. To avoid this problem, do not issue new commands that use extents until the extents migration is completed.

Note: Migration activity on a single managed disk is limited to a maximum of 4 concurrent operations. This limit does not take into account whether the managed disk is the source or the destination. If more than four migrations are scheduled for a particular managed disk, further migration operations are queued pending the completion of one of the currently running migrations. If a migration operation is stopped for any reason, a queued migration task could be started. If however a migration is suspended, the current migration continues to use resources and a pending migration is not started. For example, the following setup is a possible initial configuration:

- MDiskGrp 1 has VDisk 1 created in it
- MDiskGrp 2 has VDisk 2 created in it
- MDiskGrp 3 has only one MDisk

With configuration above, the following migrate operations are started:

- Migrate 1 migrates VDisk 1 from MDiskGrp 1 to MDiskGrp 3, running with 4 threads.
- Migrate 2 migrates VDisk 2 from MDiskGrp 2 to MDiskGrp 3, running with 4 threads.

Due to the limitations described above, the two migration operations do not always run at the same speed. MDiskGrp 3 has only one MDisk and the two migration operations have a total of 8 threads trying to access the one MDisk. Four threads are active. The remaining threads are in standby mode waiting for the chance to access the MDisk.

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.

An invocation example

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

The resulting output

```
No feedback
```

migratetoimage

You can use the **migratetoimage** command to migrate the data of the user-specified source virtual disk (either image mode or managed mode) onto the specified unmanaged disk to create an image mode VDisk. The target disk can be in the same or in a different managed disk (MDisk) group to the source disk.

Syntax

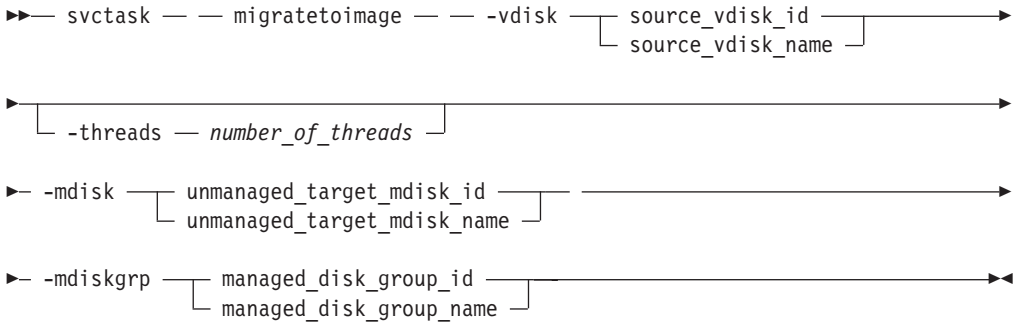

Parameters

-vdisk *source_vdisk_id | name*
Specifies the name or ID of the source VDisk to be migrated.

-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 4.

-mdisk *unmanaged_target_mdisk_id | name*
Specifies the name of the MDisk to which the data must be migrated. This disk must be unmanaged and large enough to contain the data of the disk being migrated.

-mdiskgrp *managed_disk_group_id | name*
Specifies the MDisk group into which the MDisk must be placed, after the migration has completed.

Description

This command migrates the data of the user-specified source virtual disk onto the managed disk specified as the target. At completion of the command, the virtual disk will be classified as an image mode disk.

The managed disk specified as the target must be in an unmanaged state at the time the command is run. Execution of this command will result in the inclusion of the MDisk into the user-specified MDisk group.

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.

- 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.

The example below specifies that the user wants to migrate the data from vdisk1 onto mdisk5 and the MDisk must be put into the MDisk group mdgrp2.

An invocation example

```
svctask migratetoimage -vdisk vdisk1 -mdisk mdisk5 -mdiskgrp mdgrp2
```

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

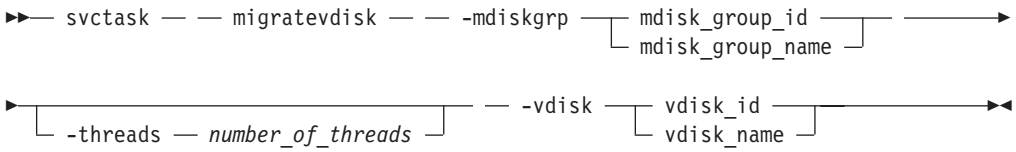

Parameters

-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 4.

-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 **svcinfo lsmigrate** 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.

The **migratevdisk** command will fail if there are insufficient free extents on the targeted managed disk group for the duration of the command. To avoid this problem, do not issue new commands that use extents until the vdisk migration is completed.

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.
- CMMVC5853E The action failed because there was a problem with the MDisk group.
- 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).

An invocation example

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

The resulting output

```
No feedback
```

Chapter 15. Tracing commands

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

setdisktrace

Use the **setdisktrace** command to set a list of disks of a given type, to include in a disk trace.

Syntax

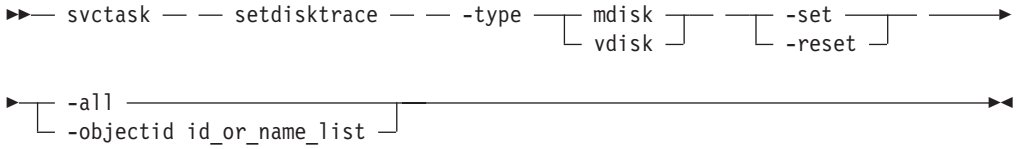

Parameters

-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.

An invocation example

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

The resulting output

```
No feedback
```

settrace

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

Syntax

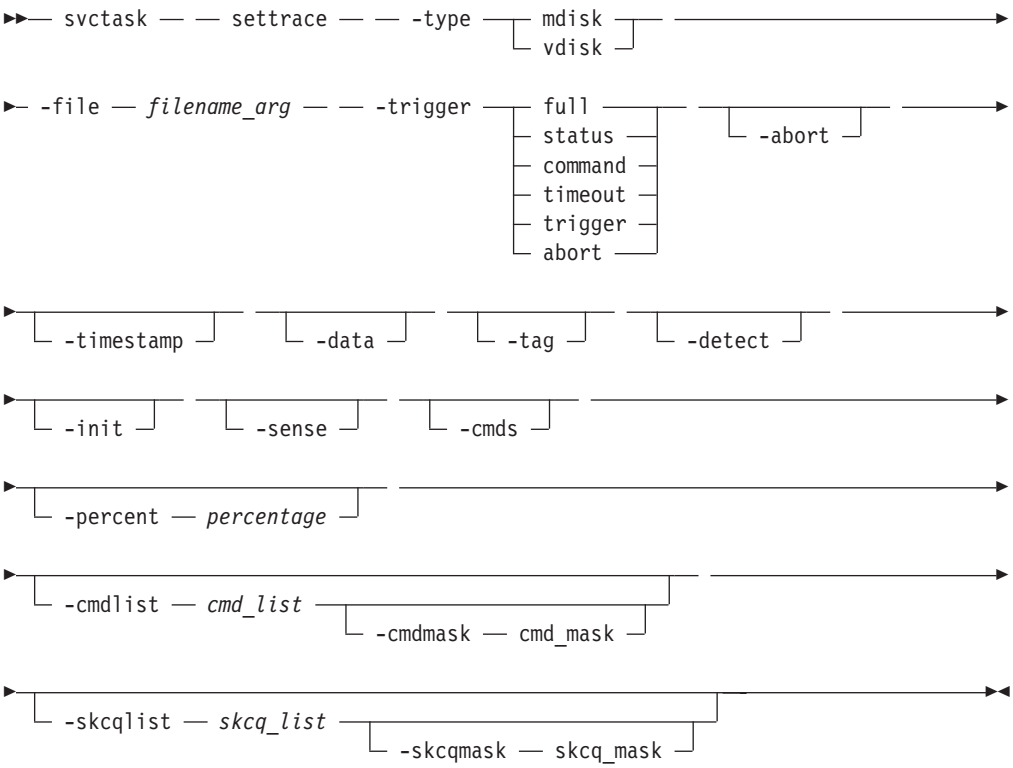

Parameters

-type *mdisk* | *vdisk*
Specifies the type of objects to trace.

-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).

full When the trace buffer is full, stop the trace (valid for both MDisks and VDisks).

status When the given SCSI status (-skcqlist) is reported in sense data (valid for both MDisks and VDisks).

command
When the given SCSI command (-cmdlist) is sent (valid for both MDisks and VDisks).

timeout
When a timeout occurs (only valid for MDisks).

trigger
Keep running until the trigger event, that is, wrap (only valid for MDisks).

abort When an abort occurs (only valid for VDisks).

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

-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>_AAAAAA_YYMMDD_HHMMSS, where *AAAAAA* is the panel name of the node generating the trace file.

-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 VDisks.

-cmds
Optionally specifies the commands flag, which adds the commands data to the trace. This flag is only valid for VDisks 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 to use when generating the trace file. The system appends the node panel name and a timestamp to the file name.

A maximum of 10 trace files will be kept on the cluster. When the eleventh trace is made, the oldest existing trace file is overwritten.

The directory may also hold files retrieved from other nodes. These files are not counted. The cluster will delete the oldest file, when necessary, to maintain the maximum number of files.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC6073E The maximum number of files has been exceeded.

An invocation example

```
svctask settrace -type vdisk -file tracedump -trigger abort
 -percent 100 -abort -timestamp
```

The resulting output

```
No feedback
```

starttrace

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

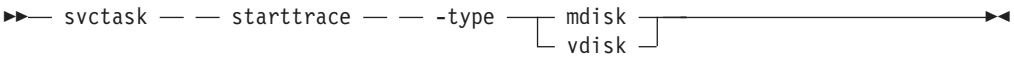

Parameters

-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 **svcinfo lsiotracedumps** 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.

An invocation example

```
svctask starttrace -type vdisk
```

The resulting output

```
No feedback
```

stoptrace

Use the **stoptrace** command to stop the tracing for the given disk type.

Syntax

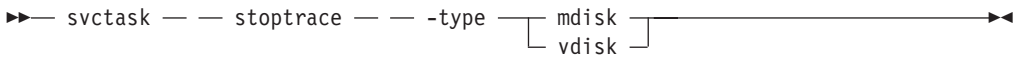

Parameters

-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.

An invocation example

```
svctask stoptrace -type mdisk
```

The resulting output

```
No feedback
```

Chapter 16. Attributes of the -filtervalue argument

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

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

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

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

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

-filtervalue ″vdisk_count<4″

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

-filtervalue capacity=24 -unit mb

The valid input options for the -unit parameter are:

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

Table 1 provides a list of valid filter attributes, as well as descriptions, qualifiers and wildcards for each object type.

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

Table 1. Valid filter attributes

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

Table 1. Valid filter attributes (continued)

Object	Attribute	Valid Qualifiers	Wildcard Valid	Description
node	node_name or name	=	Yes	The node name.
	id	=, <, <=, >, >=	No	The node ID.
	status	=	No	The status of the node. The valid input options for node status are: • adding • deleting • online • offline • pending
	IO_group_name	=	Yes	The I/O group name.
	IO_group_id	=, <, <=, >, >=	No	The I/O group ID.
	hardware	=	No	The hardware type with valid entries are 4F2, 8F2, and 8F4.
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.
	host_count	=, <, <=, >, >=	No	The number of hosts associated with the io_grp.
controller	controller_id or id	=, <, <=, >, >=	No	The controller ID.
mdisk	name	=	Yes	The name of the MDisk.
	id	=, <, <=, >, >=	No	The ID of the MDisk.
	controller_name	=	Yes	The name of the controller the MDisk belongs to.
	status	=	No	The status of the MDisk. The valid input options for MDisk status are: • online • degraded • excluded • offline
	mode	=	No	The mode of the MDisk. The valid input options for MDisk mode are: • unmanaged • managed • image
	mdisk_grp_name	=	Yes	The MDisk group name.
	mdisk_grp_id	=, <, <=, >, >=	No	The MDisk group ID.
	capacity	=, <, <=, >, >=	No	The capacity. Requires the -unit argument.

Table 1. Valid filter attributes (continued)

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

Table 1. Valid filter attributes (continued)

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

Table 1. Valid filter attributes (continued)

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

Table 1. Valid filter attributes (continued)

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

Related reference

"Using wildcards in the SAN Volume Controller CLI" on page xv
You can use wildcards in the SAN Volume Controller Command-Line Interface.

Chapter 17. Overview of the list dump commands

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

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

- /dumps
- /dumps/audit
- /dumps/elogs
- /dumps/feature
- /dumps/iostats
- /dumps/iotrace

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

Audit log dump: An audit log keeps track of action commands that are issued through an SSH session or the SAN Volume Controller Console. To list the audit log files in the /dumps/audit directory on the current configuration node or on the designated node, issue the **svcinfo lsauditlogdumps** command. To list a specified number of the most recently audited commands, you can issue the **svctask catauditlog**. If you want to dump the contents of the audit log to a file on the current configuration node, you can issue the **svctask dumpauditlog** command. This command will also clear the contents of the audit log.

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

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

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

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

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

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

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

Chapter 18. Information commands

The following commands enable you to work with displaying specific types of information with the SAN Volume Controller.

Note: IDs are assigned at run-time by the system and cannot be relied upon to be the same after configuration restoration. Therefore, wherever possible, object names should be used in preference to IDs when working with objects.

caterrlog

The **caterrlog** command returns the contents of the cluster error and event log.

Syntax

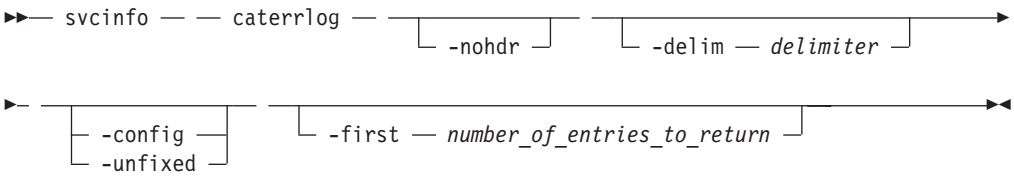

Parameters

-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.

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
```

caterrlogbyseqnum

The **caterrlogbyseqnum** command displays all the errors with a sequence number, or root cause number, as specified by the user.

Syntax

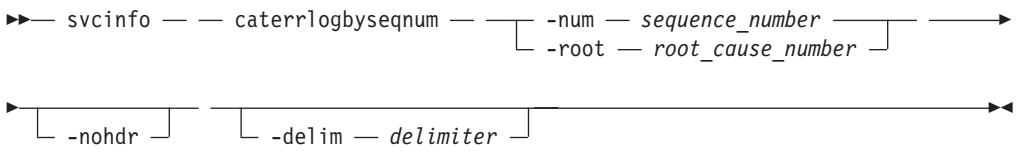

Parameters

-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.

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
```

ls2145dumps

The **ls2145dumps** command obtains a list of node assert dumps and associated output files from the /dumps directory.

Syntax

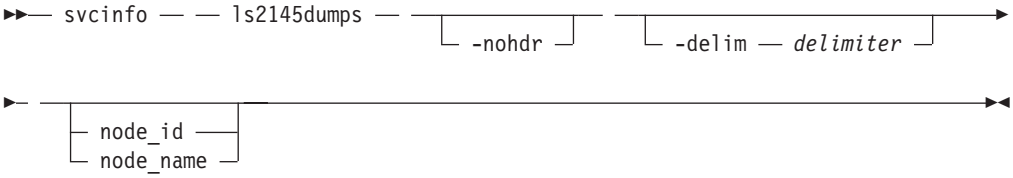

Parameters

-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 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 directory.

Possible failures

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

An invocation example

```
svcinfo ls2145dumps -delim :
```

The resulting output

```
id:2145_filename
0:000108.trc.old
1:dump.000108.030328.144007
2.000108.trc
```

lscopystatus

You can use the **lscopystatus** command to show if any file copies are currently in progress or not.

Syntax

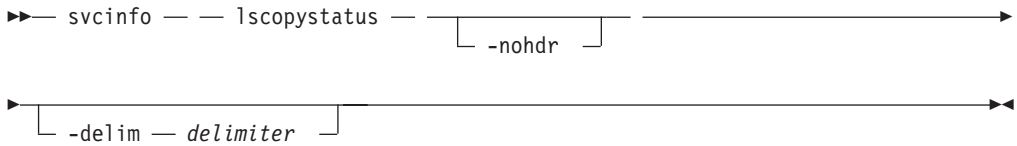

Parameters

-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.

Description

This command returns an indicator to show if a file copy is currently in progress. Only one file can be copied in the cluster at any one time.

Possible failures

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

An invocation example

```
svcinfo lscopystatus
```

The resulting output

```
status
 active
```

lscluster

The **lscluster** command returns a report as a concise list, or a detailed view, of clusters. The list provides possible values that are applicable to the attributes that are displayed as data in the output views.

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

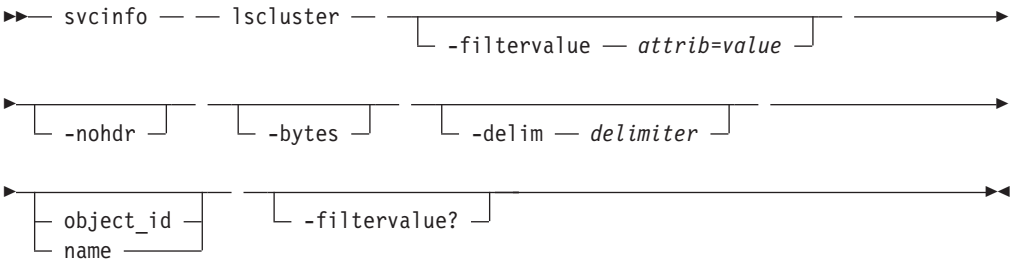

Parameters

-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 lscluster** 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

The location, partnership and bandwidth fields are relevant to Metro Mirror 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. For example, if the partnership has been at least partially established from the local cluster.

You can issue the **svcinfo lscluster** command to display a concise view of the cluster.

```
svcinfo lscluster -delim : 10030a007e5
```

where *10030a007e5* is the name of the cluster. The output from this command will include the following for each cluster on the fabric:

- cluster name
- cluster IP address
- cluster service mode 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.

A concise invocation example

```
svcinfo lscluster -delim :
```

The concise resulting output

```
id:name:location:partnership:bandwidth:cluster_IP_address:
id:name:location:partnership:bandwidth:cluster_IP_address:
cluster_service_ip_address:id_alias
0000020062813ABA:clusterA:local:::9.20.247.210:1.1.1.1:0000020062813ABA
0000020062006746:clusterB:remote:fully_configured:50:9.20.247.211:
1.1.1.1:0000020062006746
```

A detailed invocation example

```
svcinfo lscluster -delim : 10030a007e5
```

The detailed resulting output

```
id:1521071282978998
name:cluster1
location:local
partnership:
bandwidth:
cluster_IP_address:9.20.165.16
cluster_service_IP_address:9.20.165.17total_mdisk_capacity:59.8GB
space_in_mdisk_grps:0
space_allocated to vdisks:0
total_free_space:59.8GB
statistics_status:on
statistics_frequency:300
required_memory:1280
subnet_mask:255.255.255.0
default_gateway:9.20.165.1
cluster_locale:en_US
SNMP_setting:snmp all
SNMP_community:
SNMP_server_IP_address:9.20.165.18
time_zone:522 UTCemail_setting:all
email_id:another@support.com
code_level:1.20abcG
FC_port_speed:1Gb
id_alias:1521071282978998
```

A concise invocation example for a Metro Mirror configuration, where clusterA has issued the mkpartnership to clusterB, and the intercluster bandwidth is set to 50 MB/s.

```
svcinfo lscluster -delim :
```

The concise resulting output

```
id:name:location:partnership:bandwidth:
 cluster IP address:cluster service IP address
0000020062813ABA:clusterA:local:::9.20.247.210:1.1.1.1
0000020062006746:clusterB:remote:
 fully_configured:50:9.20.247.211:1.1.1.1
```

lsclustercandidate

The **lsclustercandidate** command lists the clusters that are available for setting up a two-cluster partnership. This is a prerequisite for creating inter-cluster Metro Mirror relationships.

Syntax

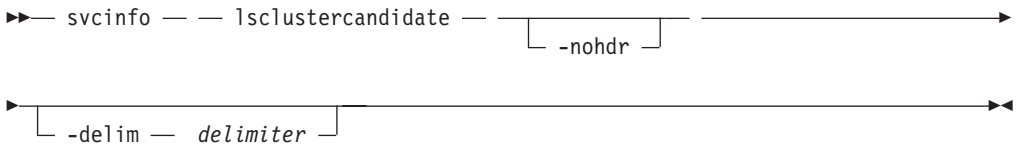

Parameters

-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.

Description

This command returns a list of clusters that are available as candidate partner clusters to form a Metro Mirror 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.

An invocation example

```
svcinfo lsclustercandidate
```

The resulting output

```
id               configured     cluster_name
0000010034E0F430 no             ldcluster26
```

lscontroller

The **lscontroller** command returns a concise list, or a detailed view, of controllers visible to the cluster.

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

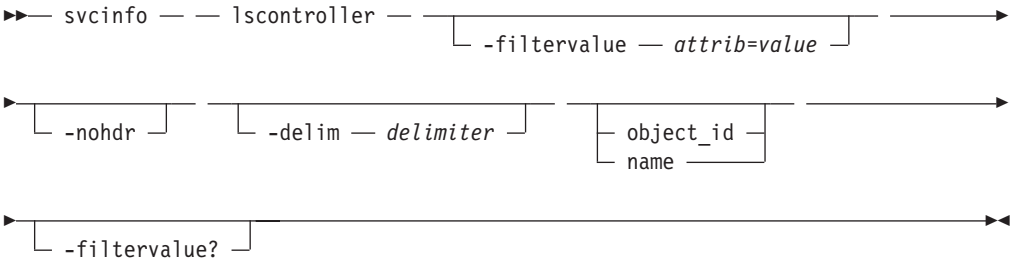

Parameters

-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.

-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 the name shown on the cluster: 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.

path_count is the number of MDisks multiplied by the number of controller ports, through which those MDisks are visible.

max_path_count is the highest possible value that path_count could reach for that controller type.

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.

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
```

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
```

lsdiscoverystatus

You can use the **lsdiscoverystatus** command to determine if a discovery operation is in progress or not.

Syntax

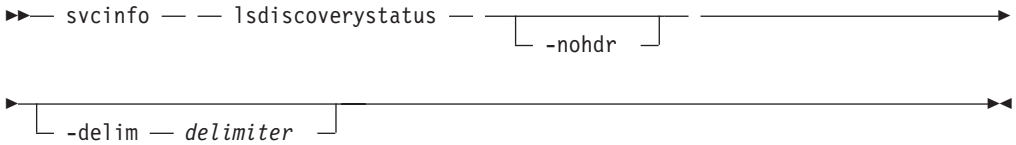

Parameters

-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.

Description

This command displays the following:

- active: if there is a discovery operation in progress at the time that the command is issued
- inactive: if there are no discovery operations in progress at the time that the command is issued

Possible failures

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

An invocation example

```
svcinfo lsdiscoverystatus
```

The resulting output

```
status
inactive
```

lserrlogbyfcconsistgrp

The **lserrlogbyfcconsistgrp** command displays errors and events in the log related to FlashCopy consistency groups.

Syntax

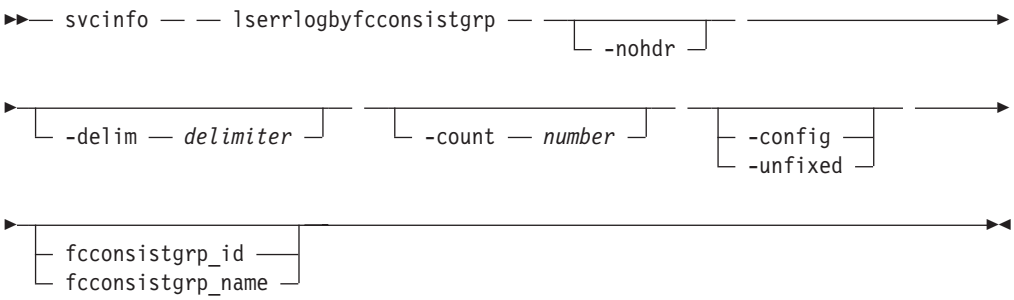

Parameters

-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.

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
```

lserrlogbyfcmap

The **lserrlogbyfcmap** command displays a list of the errors and events in the log related to FlashCopy mappings.

Syntax

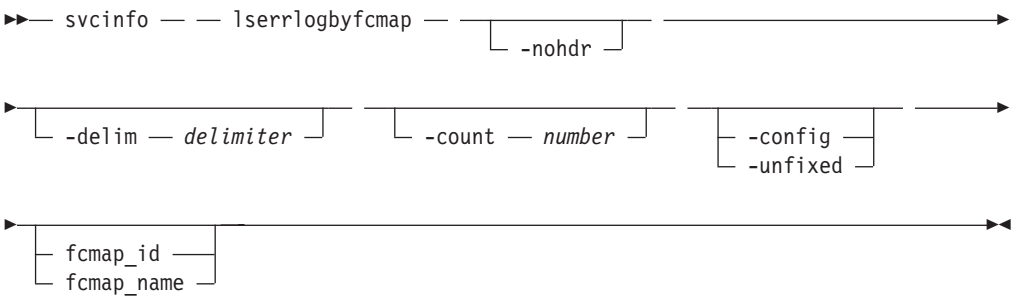

Parameters

-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.

fcmap_id | fcmap_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.

An invocation example

```
svcinfo lserrlogbyfcmap -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
```

lserrlogbyhost

The **lserrlogbyhost** command displays a list of the errors and events in the log related to hosts.

Syntax

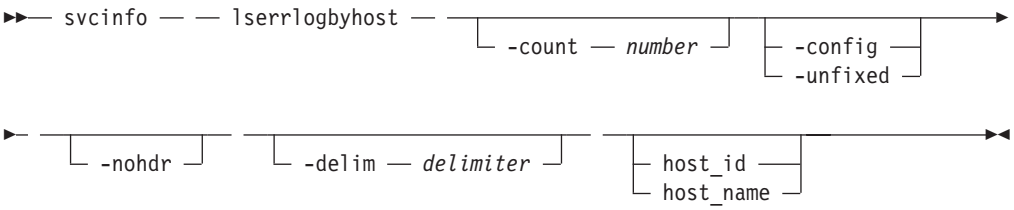

Parameters

-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.

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
```

lserrlogbyiogrp

The **lserrlogbyiogrp** command displays a list of the errors and events in the log related to I/O groups.

Syntax

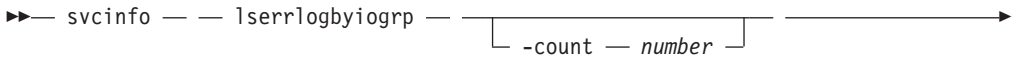

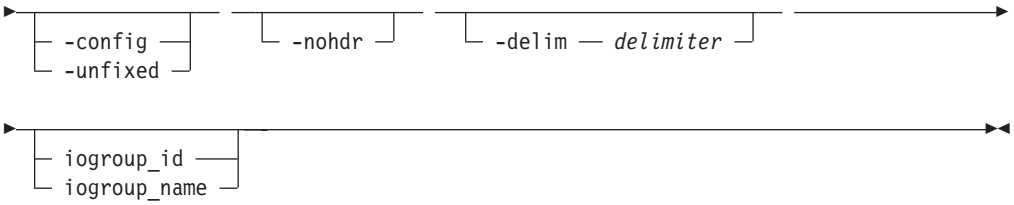

Parameters

-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.

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

The **lserrlogbymdisk** command displays a list of the errors and events in the log related to a specific MDisk.

Syntax

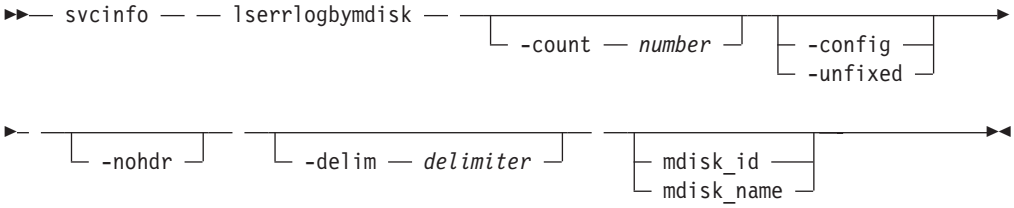

Parameters

-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 a specific 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.

An invocation example

```
svcinfo lserrlogbymdisk -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
```

lserrlogbymdiskgroup

The **lserrlogbymdiskgroup** commands display a list of the errors and events in the log related to MDisk groups.

Syntax

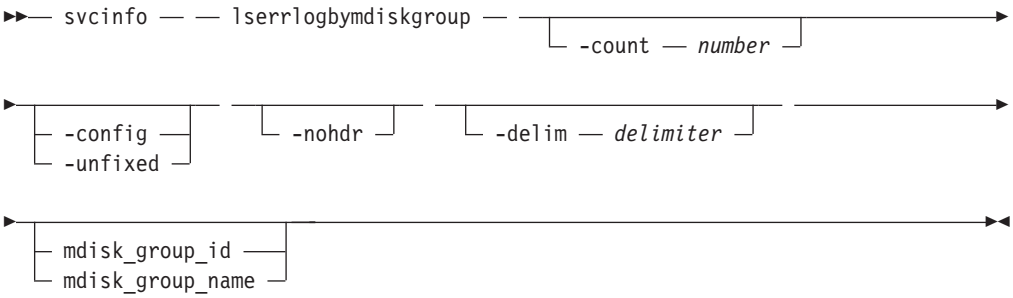

Parameters

-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.

An invocation example

```
svcinfo lserrlogbymdiskgrp -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
```

lserrlogbynode

The **lserrlogbynode** command displays a list of the errors and events in the log related to nodes.

Syntax

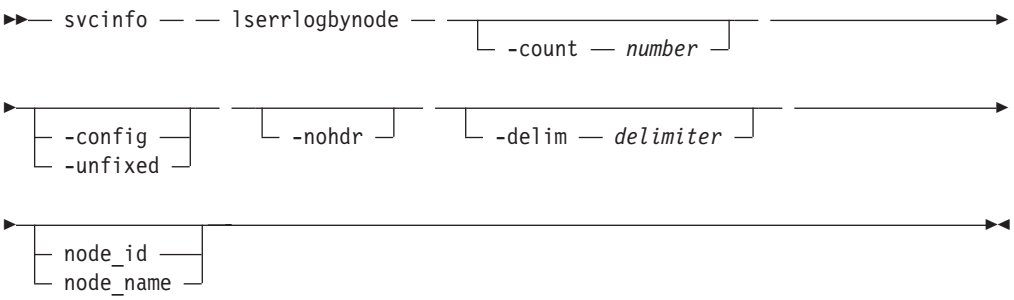

Parameters

-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.

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
```

lserrlogbyrcconsistgrp

You can use the **lserrlogbyrcconsistgrp** command to display the error log by Metro Mirror consistency group.

Syntax

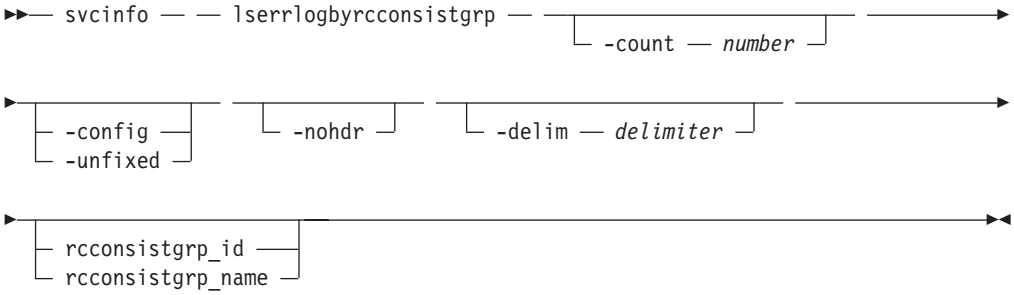

Parameters

-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.

rcconsistgrp_id | rcconsistgrp_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 Metro Mirror 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.

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

The **lserrlogbyrcrelationship** command displays a list of the errors and events in the log related to Metro Mirror relationships.

Syntax

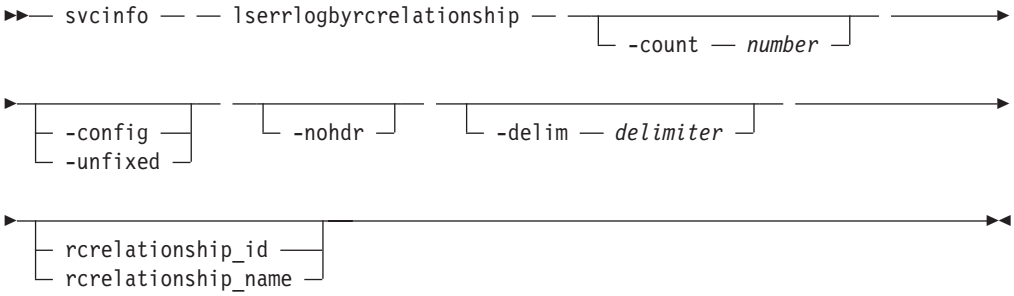

Parameters

-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.

rcrelationship_id | rcrelationship_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 Metro Mirror 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.

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
```

lserrlogbyvdisk

The **lserrlogbyvdisk** command displays a list of the errors and events in the log related to VDisks.

Syntax

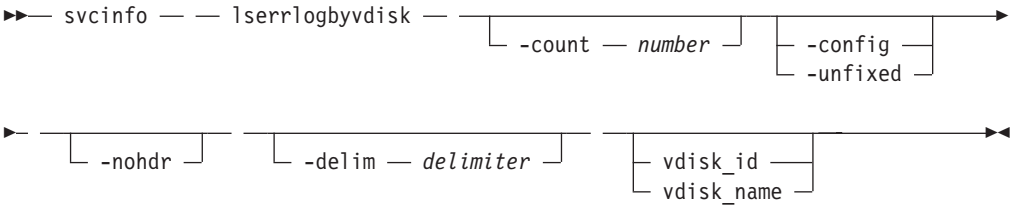

Parameters

-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 VDisks. 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.

An invocation example

```
svcinfo lserrlogbyvdisk -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

The **lserrlogdumps** commands return a list of error log dumps in the /dumps/elogs directory. These dumps are created as a result of issuing the **svctask dumperrlog** command.

Syntax

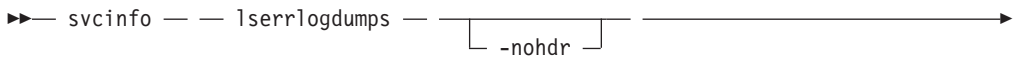

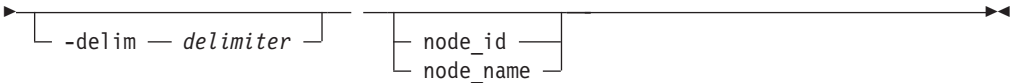

Parameters

-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.

An invocation example

```
svcinfo 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
```

lsfabric

The **lsfabric** command can be used to display connectivity between nodes and other controllers and hosts.

Syntax

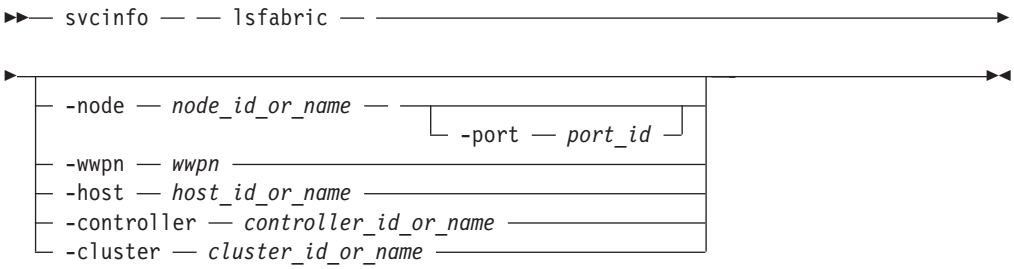

Parameters

-node *node_id_or_name*
Optionally specifies a node ID or name. Mutually exclusive to all other parameters except -port, which is an optional parameter that can be used in conjunction with -node. Using this parameter by itself displays output for all ports of a specified node. Output is given for each port in turn.

-port *port_id*
Optionally specifies a port ID. An optional parameter that can only be used together with the -node parameter. The parameter displays a concise view of all the WWPNs currently logged into a specified node and port. Valid data is a number in the range 1 - 4 that matches the port with the same number in the VPD or the actual hex WWPN of the local port.

-wwpn *wwpn*
Optionally specifies a WWPN. The parameter is mutually exclusive to all other parameters. This parameter displays a list of all ports which have a login to the specified WWPN.

-host *host_id_or_name*
Optionally specifies a host name or ID. This parameter is mutually exclusive to all other parameters. This command and parameter is equivalent to issuing the **svcinfo lsfabric -wwpn wwpn** command for every configured WWPN of the specified host. Output is sorted by remote WWPNs then cluster WWPNs. For example, a host with 2 ports zoned to 1 port of every node in a 8 node cluster should produce 16 lines of output.

-controller *controller_id_or_name*
Optionally specifies a controller ID or name. The parameter is mutually exclusive to all other parameters. This command and parameter is equivalent to issuing the **svcinfo lsfabric -wwpn wwpn** command for every configured WWPN of the specified controller. Output is sorted by remote WWPNs then cluster WWPNs. For example, a controller with 4 ports connected to a 8 node cluster with 2 counter part SANs should produce 64 lines of output.

-cluster *cluster_id_or_name*
Optionally specifies a cluster ID or name. The parameter is mutually exclusive to all other parameters. This command is equivalent to issuing the **svcinfo lsfabric -wwpn wwpn** command for every known WWPN in the specified cluster. Output is sorted by remote WWPNs then cluster WWPNs. This command can be used to check the state of connections within the local cluster or between the local and remote cluster. When the local cluster ID or name is specified, each node-to-node connection will be listed twice: once from each end. For example, a 8 node cluster with 2 counter part SANs should produce 8 nodes * 7 other nodes * 2 SANs * 4 point-to-point logins = 448 lines of output.

Description

This command can be issued with any of the parameters as described above in order display a limited subset of information.

If the command is issued without any parameters, then it will provide output for every node. Output is given for each node and then each port in turn.

Values for the type and state fields are:

- state active

  Note: For a host it means that SCSI commands were issued within the last 5 minutes. For a node it means that the node ports can see other ports. For controllers it means that commands were issues within the last 10 seconds.
- state inactive
- type host/node/controller/node/unknown

You can issue this command if you want to view all connectivity information available to your cluster.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an object that was specified in the command does not exist.

An invocation example

```
svcinfo lsfabric -delim :
```

The resulting output Each row of output will contain the following columns:

```
remote_wwpn: remote_nportid:
  local_wwpn: local_port:
 local_nportid: state: name: type
```

lsfcconsistgrp

The **lsfcconsistgrp** commands return a concise list or a detailed view, of FlashCopy consistency groups visible to the cluster. This information is useful for tracking FlashCopy consistency groups.

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

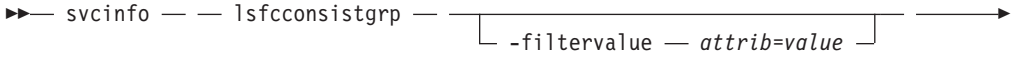

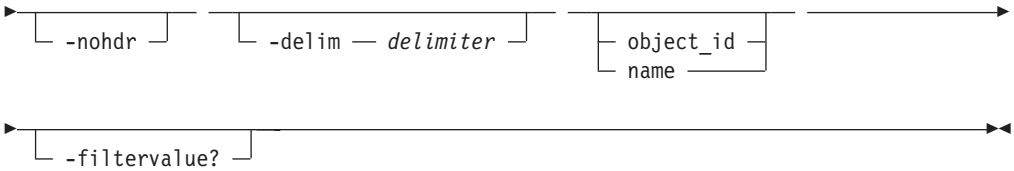

Parameters

-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.

-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 lsfcconsistgrp** 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.

A concise invocation example

```
svcinfo lsfcconsistgrp -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

```
svcinfo lsfcconsistgrp -delim : 1
```

The detailed resulting output

```
id:1
name:ffccg0
status:idle_or_copied
```

lsfcmap

The **lsfcmap** command generates a list containing concise information about all FlashCopy mappings visible to the cluster, or detailed information for a single FlashCopy mapping.

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 FlashCopy Mapping.

Syntax

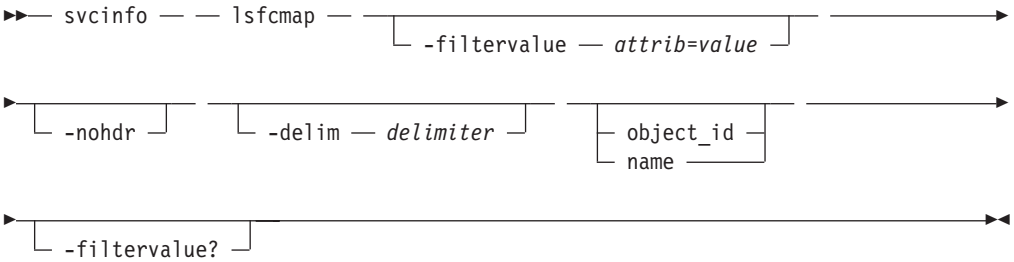

Parameters

-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.

-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 lsfcmap** 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.

A concise invocation example

```
svcinfo lsfcmap -delim :
```

The concise 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:ffcmap1:0:vdisk0:1:vvdisktwo:::idle_or_copied::75
```

A detailed invocation example

```
svcinfo lsfcmap -delim : 0
```

The detailed resulting output

```
id:0
name:ffcmap1
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
```

lsfcmapcandidate

The **lsfcmapcandidate** command lists all the VDisks that can be source or destinations for FlashCopy. For example, those that are not already in a mapping.

Syntax

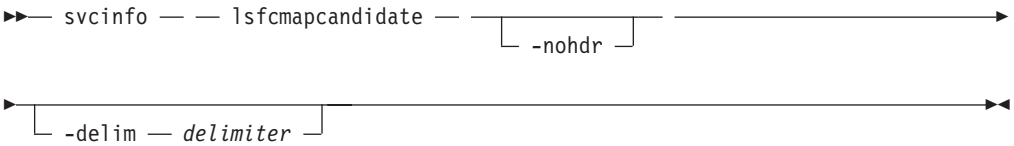

Parameters

-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.

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.

An invocation example

```
svcinfo lsfcmapcandidate
```

The resulting output

```
id
2
3
4
```

lsfcmapprogress

The **lsfcmapprogress** command returns the progress of the background copy of a FlashCopy mapping. This is displayed as a percentage completed value.

Syntax

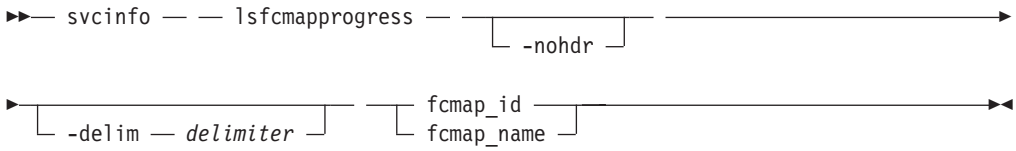

Parameters

-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.

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.

An invocation example

```
svcinfo lsfcmapprogress 0
```

The resulting output

```
id              progress
0               0
```

lsfeaturedumps

The **lsfeaturedumps** command returns a list of dumps in /dumps/feature. These dumps are created as a result of issuing the **svctask dumpinternallog** command.

Syntax

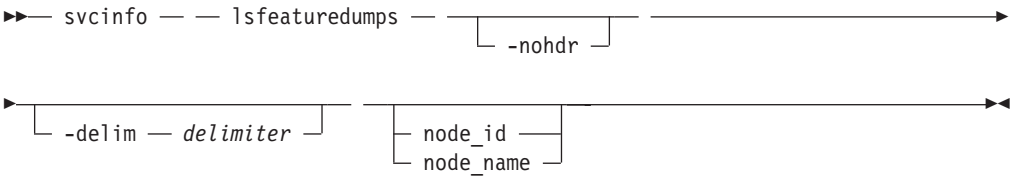

Parameters

-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 featurization dumps. These dumps are created as a result of issuing the **svctask dumpinternallog** 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.

An invocation example

```
svcinfo lsfeaturedumps
```

The resulting output

```
id              feature_filename
0               feature.txt
```

lsfreeextents

The **lsfreeextents** command lists the number of free extents available on a specified MDisk.

Syntax

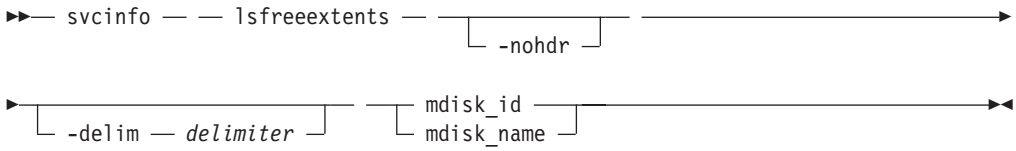

Parameters

-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.

An invocation example

```
svcinfo lsfreeextents 2
```

The resulting output

```
id 2
number_of_extents 4372
```

lshbaportcandidate

The **lshbaportcandidate** command lists all of the unconfigured, logged in, host bus adapter (HBA) ports. This information is used to find open HBA ports.

Syntax

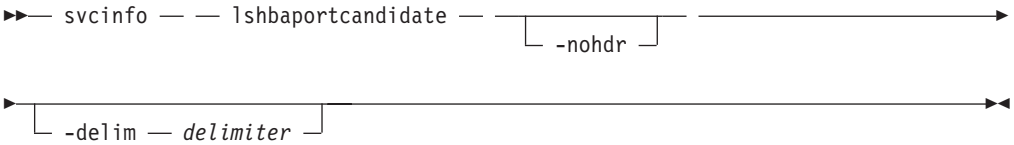

Parameters

-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.

Description

This command returns a list of unconfigured, logged in, HBA ports.

Note: The **svcinfo lshbaportcandidate** command presents a list of host HBA ports that are logged in to nodes. However, there are situations when the information presented might include host HBA ports that are no longer logged in or even part of the SAN fabric. For example, a host HBA port is unplugged from a switch but **scvinfo lshbaportcandidate** still shows the WWPN that is logged in to all nodes. If this occurs, the incorrect entry is removed when another device is plugged in to the same switch port that previously contained the removed host HBA port.

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.

An invocation example

```
svcinfo lshbaportcandidate
```

The resulting output

```
id
210100E08B2520D4
```

lshost

The **lshost** command generates a list with concise information about all the hosts visible to the cluster and detailed information about a single host.

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

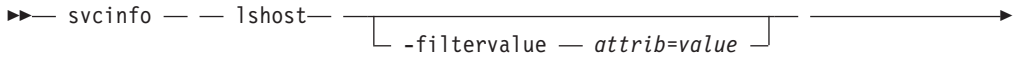

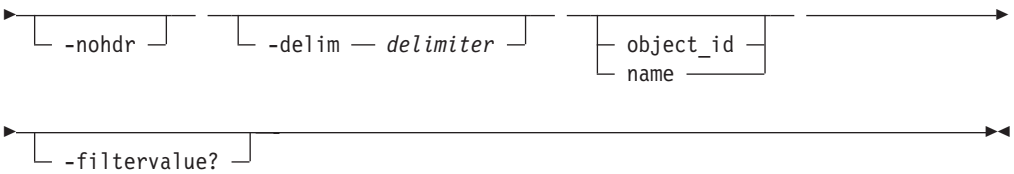

Parameters

-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.

-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 lshost** command are:

- host_name
- host_id
- port_count
- name
- id
- iogrp_count

Description

This command will return a concise list or a detailed view, of hosts visible to the cluster.

The following list provides the different states for a host:

- offline The host is offline if one or more I/O groups with VDisk mappings does not have a login for the specified WWPN.
- degraded The host is degraded if one or more nodes with VDisk mappings does not have a login for the specified WWPN.
- inactive The host is inactive if all the nodes with VDisk mappings have a login for the specified WWPN, however, no nodes have seen any SCSI commands from the WWPN in the last 5 minutes.
- active The host is active if all the nodes with VDisk mappings have a login for the specified WWPN, however, at least one node has seen SCSI commands from the WWPN in the last 5 minutes.

If a host does not have any VDisk mappings, then it is either reported as offline or inactive.

Note: The **svcinfo lshost** command presents a list of host HBA ports that are logged in to nodes. However, there are situations when the information presented may include host HBA ports which are no longer logged in or even part of the SAN fabric. For example: A host HBA port is unplugged from a switch but **scvinfo lshost** still shows the WWPN logged in to all nodes. If this occurs, the incorrect entry will be removed when another device is plugged in to the same switch port that previously contained the removed host HBA port.

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.

A concise invocation example

```
svcinfo lshost -delim :
```

The concise resulting output

```
id:name:port_count:iogrp_count
0:host0:1:0
1:host1:1:0
```

A detailed invocation example

```
svcinfo lshost -delim : 1
```

The detailed resulting output

```
id:1
name:host1
port_count:1
type:generic
iogrp_count:0
mask:1111
WWPN:000000000001AABB
node_logged_in_count:0
state:inactive
```

lshostiogrp

The **lshostiogrp** command returns a list of all the I/O groups associated to a specified host object.

Syntax

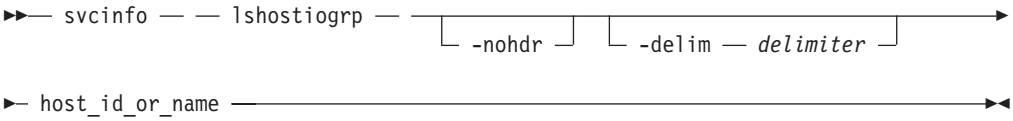

Parameters

-nohdr
By default, headings are displayed for each column of data in a concise and detailed view. Using the -nohdr parameter will suppress the display of these headings.

Note: If no data displays, for example, an empty view has been returned, then headings are also not displayed whether the -nohdr parameter was issued or not.

-delim *delimiter*
In a concise view, all columns of data are separated by a space. 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. In a detailed view, the data is separated from its header by a colon.

host_id_or_name
The name or ID of the host for which the list of I/O groups is required.

Description

This command returns a list of all the I/O groups mapped to the specified host.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an object that was specified in the command does not exist.

An invocation example

```
svcinfo lshostiogrp -delim : hostone
```

The resulting output

```
id:name
0:io_grp0
1:io_grp1
```

lshostvdiskmap

You can use the **lshostvdiskmap** command to obtain a list of the virtual disks that are mapped (visible) to a given host. These are the virtual disks that have been mapped to the specified host, meaning they are visible to the specified host.

Syntax

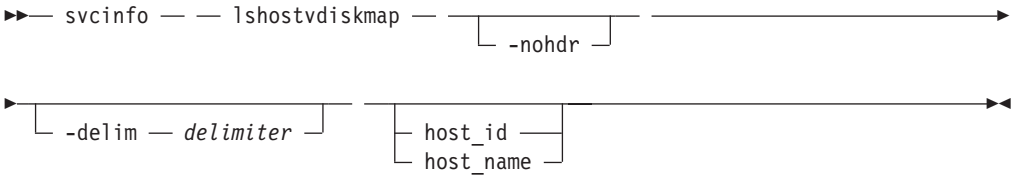

Parameters

-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 command 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 cluster 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 cluster 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 cluster that these ports belong to. The ports are stored as part of the detailed view, so you must list each host in turn by issuing the following:

   ```
   svcinfo lshost <name/id>
   ```

   where <*name/id*> is the name or ID of the host. Check for matching WWPNs.

   Note: You should name your hosts accordingly, for example, if the actual host is called *orange* you should also name the host object defined to the cluster as *orange*.

Now that you have the <host name> as defined to the cluster and the <vpath serial number>, issue the following command:

```
svcinfo lshostvdiskmap <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.

An invocation example

```
svcinfo lshostvdiskmap -delim : 2
```

The resulting output

```
id:name:SCSI_id:vdisk_id:vdisk_name:wwpn:vdisk_UID
2:host2:0:10:vdisk10:0000000000000ACA:600507680195800150000000000000000A
2:host2:1:11:vdisk11:0000000000000ACA:600507680195800150000000000000000B
2:host2:2:12:vdisk12:0000000000000ACA:600507680195800150000000000000000C
2:host2:3:13:vdisk13:0000000000000ACA:600507680195800150000000000000000D
2:host2:4:14:vdisk14:0000000000000ACA:600507680195800150000000000000000E
```

lsiogrp

The **lsiogrp** command will return a concise list or a detailed view, of I/O groups visible to the cluster.

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

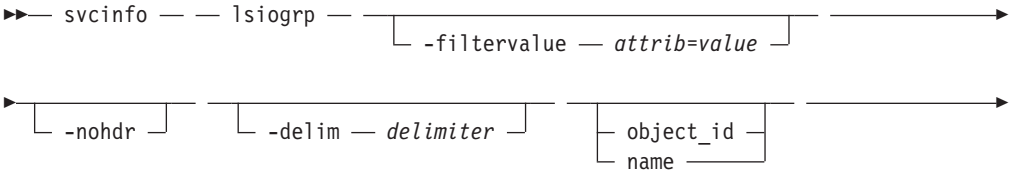

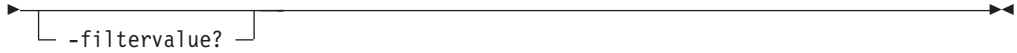

Parameters

-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.

-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 lsiogrp**command are:

- HWS_name
- HWS_unique_id
- node_count
- name
- id
- host_count

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.

A concise and detailed invocation example

```
svcinfo lsiogrp -delim :
```

The concise and detailed resulting output

```
id:name:node_count:vdisk_count:host_count
0:io_grp0:1:0:0
1:io_grp1:0:0:0
2:io_grp2:0:0:0
3:io_grp3:0:0:0
4:recovery_io_grp:0:0:0
```

A detailed and detailed invocation example

```
svcinfo lsiogrp -delim : 2
```

The detailed and detailed resulting output

```
id:2
name:io_grp2
node_count:0
vdisk_count:0
host_count:0
```

lsiogrphost

The **lsiogrphost** returns a list of the hosts that are mapped to a specified I/O group.

Syntax

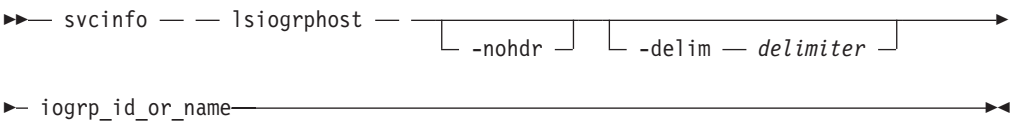

Parameters

-nohdr
By default, headings are displayed for each column of data in a concise and detailed view. Using the -nohdr parameter will suppress the display of these headings.

Note: If no data displays, for example, an empty view has been returned, then headings are also not displayed whether the -nohdr parameter was issued or not.

-delim *delimiter*
In a concise view, all columns of data are separated by a space. 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. In a detailed view, the data is separated from its header by a colon.

iogrp_id_or_name
The ID or name of the I/O group for which a list of all mapped hosts is required.

Description

The **lsiogrphost** command returns a list of hosts that are mapped to a specified I/O group.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC5804E The action failed because an object that was specified in the command does not exist.

An invocation example

```
svcinfo lsiogrphost -delim : 0
```

The resulting output

```
id:name
0:hostzero
1:hostone
```

lsiogrpcandidate

You can use the **lsiogrpcandidate** command to list the I/O groups that can have nodes added to them.

Syntax

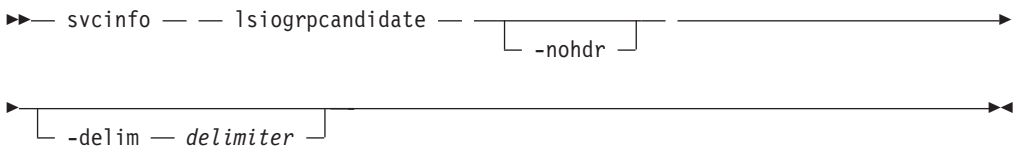

Parameters

-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.

Note: If there is no data to be separated, then the delimiter is not used whether the -nohdr option was used or not.

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.

An invocation example

```
svcinfo lsiogrpcandidate
```

The resulting output

```
id
0
1
2
3
4
```

lsiostatsdumps

The **lsiostatsdumps** command returns a list of dumps in the /dumps/iostats directory. These dumps are created as a result of issuing the **svctask startstats** command.

Syntax

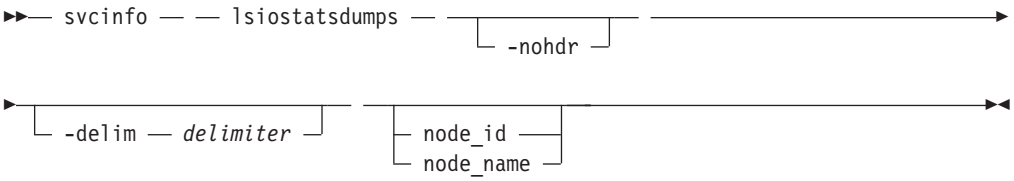

Parameters

-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.

An invocation example

```
svcinfo lsiostatsdumps
```

The resulting output

```
id              iostat_filename
0               v_stats_mala75_031123_072426
1               m_stats_mala75_031123_072425
```

lsiotracedumps

You can use the **lsiotracedumps** command to return a list of files in the /dumps/iotrace directory.

Syntax

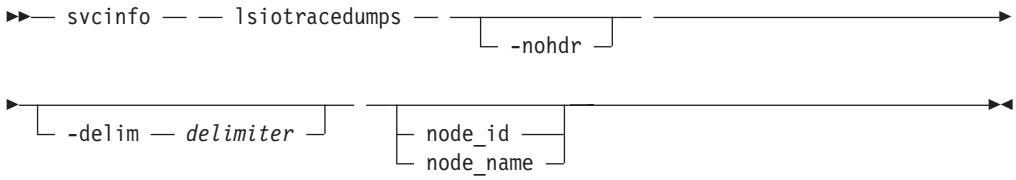

Parameters

-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 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.

An invocation example

```
svcinfo lsiotracedumps
```

The resulting output

```
id               iotrace_filename
0                c1_mala75_030405_092155
1                c2_mala75_030405_092156
2                c3_mala75_030405_092158
3                c4_mala75_030405_092159
4                c5_mala75_030405_092201
```

lslicense

The **lslicense** command returns the current license (featurization) settings for the cluster. The settings are defined as copy services status and the capacity of virtual storage licensed for use by this cluster.

Syntax

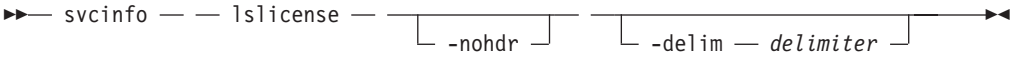

Parameters

-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.

-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 lslicense** 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
- Metro Mirror: 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.

An invocation example

```
svcinfo lslicense
```

The resulting output

```
feature_flash on
feature_remote on
feature_num_gb 32
```

lsmdisk

The **lsmdisk** command can return a concise list or a detailed view of MDisks visible to the cluster, or detailed information about a single managed disk.

Syntax

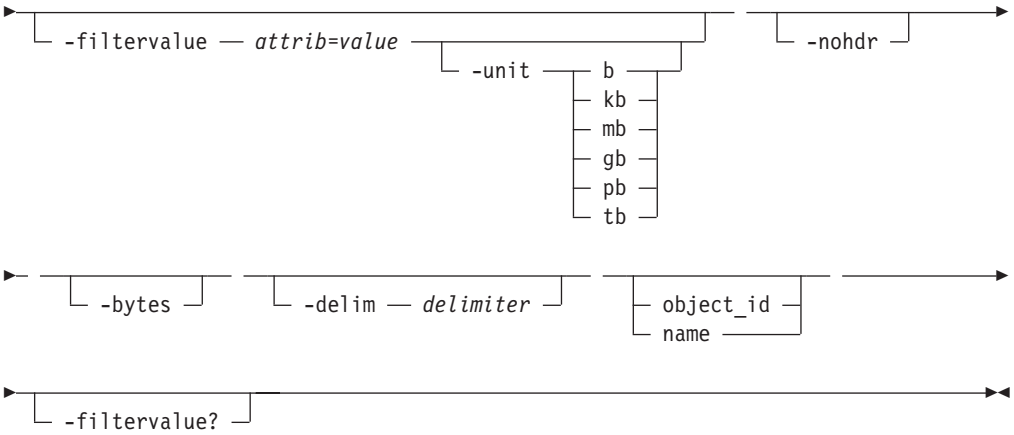

Parameters

-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 cluster, the cluster will automatically discover the back-end controller and will integrate the controller to determine what storage it is presented to the node. 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 cluster may be unaware of these configuration changes. This task allows a user to request the cluster to re-scan the fibre-channel SAN to update the list of unmanaged MDisks.

Note: The automatic discovery performed by the cluster does not write anything to a unmanaged MDisk. It is only when a the user instructs the cluster 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.

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.

Determining the relationship between MDisks and RAID arrays or LUNs: Show the detailed view of the given MDisk <mdiskname>, by issuing the following command:

```
svcinfo lsmdisk <mdiskname>
```

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 following command:

```
svcinfo lscontroller <controllername>
```

where *<controllername>* is the name of the controller.

Note: Remember the vendor ID, product ID, and WWNN. Use these to determine what is being presented to the MDisk.

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.

A concise invocation example

```
svcinfo lsmdisk -delim :
```

The concise resulting output

```
id:name:status:mode:mdisk_grp_id:mdisk_grp_name:capacity:
ctrl_LUN_#:controller_name:UID
0:mdisk0:online:managed:0:mdiskgrp0:68.4GB:0000000000000000:controller0:
20000004cf2422aa0000000000000000000000000000000000000000000000000000
1:mdisk1:online:managed:0:mdiskgrp0:68.4GB:0000000000000000:controller1:
20000004cf1fd19d0000000000000000000000000000000000000000000000000000
2:mdisk2:online:managed:0:mdiskgrp0:68.4GB:0000000000000000:controller2:
20000004cf2425310000000000000000000000000000000000000000000000000000
3:mdisk3:online:managed:0:mdiskgrp0:68.4GB:0000000000000000:controller3:
20000004cf1fd7a00000000000000000000000000000000000000000000000000000
4:mdisk4:online:unmanaged:::68.4GB:0000000000000000:controller4:
20000004cf2413080000000000000000000000000000000000000000000000000000
5:mdisk5:online:unmanaged:::68.4GB:0000000000000000:controller5:
20000004cf2412ca0000000000000000000000000000000000000000000000000000
6:mdisk6:online:unmanaged:::68.4GB:0000000000000000:controller6:
20000004cf2421230000000000000000000000000000000000000000000000000000
7:mdisk7:online:unmanaged:::68.4GB:0000000000000000:controller7:
20000004cf2412ac0000000000000000000000000000000000000000000000000000
8:mdisk8:online:unmanaged:::68.4GB:0000000000000000:controller8:
20000004cf2421de0000000000000000000000000000000000000000000000000000
9:mdisk9:online:unmanaged:::68.4GB:0000000000000000:controller9:
20000004cf1fd7cc00000000000000000000000000000000000000000000000000000
10:mdisk10:online:unmanaged:::68.4GB:0000000000000000:controller10:
20000004cf1ffe880000000000000000000000000000000000000000000000000000
11:mdisk11:online:unmanaged:::68.4GB:0000000000000000:controller11:
20000004cf2412aa0000000000000000000000000000000000000000000000000000
12:mdisk12:online:unmanaged:::68.4GB:0000000000000000:controller12:
20000004cf2423740000000000000000000000000000000000000000000000000000
13:mdisk13:online:unmanaged:::68.4GB:0000000000000000:controller13:
20000004cf2419fe0000000000000000000000000000000000000000000000000000
14:mdisk14:online:unmanaged:::68.4GB:0000000000000000:controller14:
20000004cf2419320000000000000000000000000000000000000000000000000000
```

A detailed invocation example

```
svcinfo lsmdisk -delim : 2
```

The detailed resulting output

```
id:2
name:mdisk2
status:online
mode:unmanaged
mdisk_grp_id:
mdisk_grp_name:
capacity:68.4GB
quorum_index:
block_size:512
controller_name:controller2
ctrl_type:4
ctrl_WWNN:20000004CF242531
controller_id:2
path_count:1
max_path_count:1
ctrl_LUN_#:0000000000000000
UID:20000004cf2425310000000000000000000000000000000000000000000000000000
preferred_WWPN:22000004CF242531
active_WWPN:22000004CF242531
```

The following define the ctrl_type fields:

0 Device that has not yet been processed.

1 Device is a node in which that cluster has not yet been identified.

2 Device is a node within this cluster.

3 Device is a node within a remote cluster.

4 Device is a back-end storage controller.

5 Device is a storage or other fabric device that is incompatible with the cluster.

lsmdiskcandidate

The **lsmdiskcandidate** command lists all of the unmanaged MDisks, by MDisk ID.

Syntax

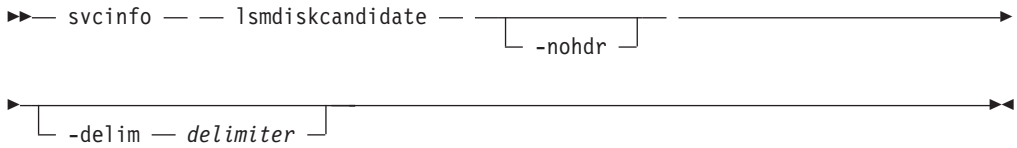

Parameters

-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.

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 cluster the cluster will automatically discover the back-end controller and will integrate the controller to determine what storage it is presented to the node. 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 cluster may be unaware of these configuration changes. This task allows a user to request the cluster to re-scan the fibre-channel SAN to update the list of unmanaged MDisks.

Note: The automatic discovery performed by the cluster does not write anything to a unmanaged MDisk. It is only when a the user instructs the cluster 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.
- CMMVC5804E The action failed because an entity that was specified in the command does not exist.

An invocation example

```
svcinfo lsmdiskcandidate
```

The resulting output

```
id
5
6
7
8
9
10
11
12
13
14
```

lsmdiskextent

The **lsmdiskextent** command return the extent allocation between managed disks and virtual disks. The command returns a list, in which each entry contains a VDisk ID and the number of extents.

Syntax

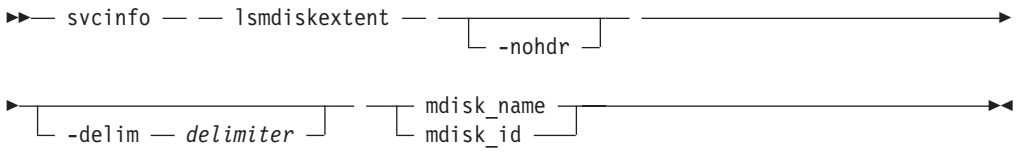

Parameters

-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 VDisks and MDisks and the number of extents provided by each MDisk: You can determine the number of extents that 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: You can 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.
- 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 (MDisk) is not used by any virtual disk (VDisk).
- 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 (MDisk) or virtual disk (VDisk).
- CMMVC6005E The view request failed as the specified object is not a member of an appropriate group.

An invocation example

```
svcinfo lsmdiskextent 2
```

The resulting output

```
id              number_of_extents
1               1
2               1
```

lsmdiskgrp

The **lsmdiskgrp** command returns a concise list or a detailed view, of MDisk groups visible to the cluster.

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

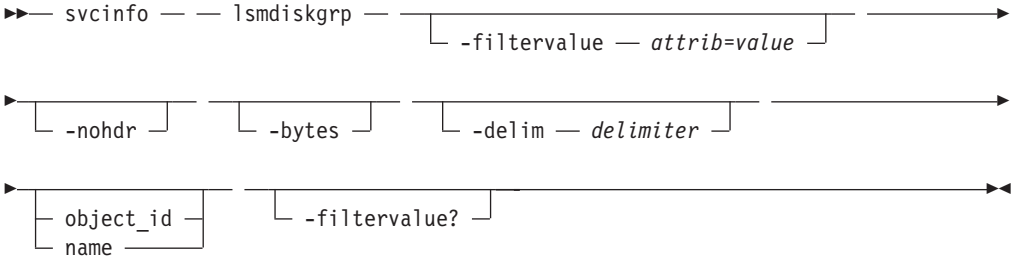

Parameters

-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.

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
```

lsmdiskmember

The **lsmdiskmember** command returns a list of VDisks that are using extents on the specified MDisk. This is defined as the virtual disks that are using extents on the managed disk specified by the ID.

Syntax

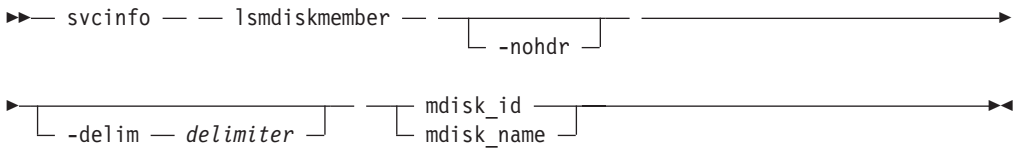

Parameters

-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.

An invocation example

```
svcinfo lsmdiskmember 1
```

The resulting output

```
id
0
```

lsmigrate

The **lsmigrate** command shows the progress of all the data migration operations currently in progress.

Syntax

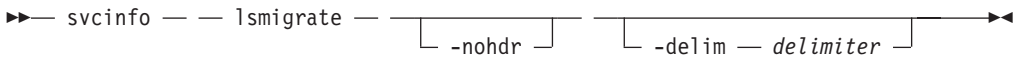

Parameters

-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.

If you use multiple threads to migrate data, the progress will increment when all threads have completed the migration of an extent. For large extent sizes with many threads, this can result in quite large increments in the percentage progress.

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.

An invocation example

```
svcinfo lsmigrate -delim :
```

The resulting output

```
migrate_type:MDisk_Group_Migration
progress:96
migrate_source_vdisk_index:33
migrate_target_mdisk_grp:4
max_thread_count:4
```

lsnode

The **lsnode** command will return a concise list, or a detailed view, of nodes visible to the cluster.

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

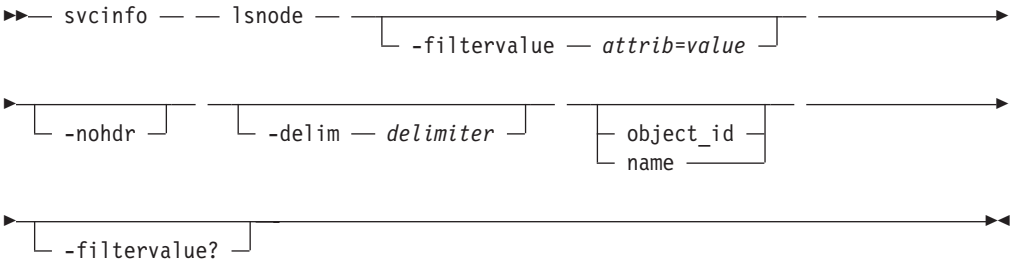

Parameters

-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.

-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 lsnode** command are:

- node_name
- id
- status
- IO_group_name

- IO_group_id
- name
- hardware

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, flushing, pending, online, adding, deleting
config_ node	no, yes
port_status	active, inactive, not installed
hardware	4F2, 8F2, 8F4, other

Determining a node's WWPNs: List the nodes in the cluster by issuing the following command:

```
svcinfo lsnode
```

Note: Remember the node name or ID as you will need it in the next step. For the node or nodes in question, issue the following command:

```
svcinfo lsnode <nodename/id>
```

where <*nodename/id*> is the node name or ID.

Note: Remember the four port ID's (WWPNs).

Attention: 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.

A concise invocation example

```
svcinfo lsnode -delim :
```

The concise resulting output

```
id:name: UPS_serial_number: WWNN:
 status: IO_group_id: IO_group_name:
 config_node: UPS_unique_id: hardware
1:node1: UPS_SN: 5005076801000024:
         online: 0: io_grp0: yes:
 1000000000000024: other
```

A detailed invocation example

```
svcinfo lsnode -delim = 1
```

The detailed resulting output

```
id=1
name=node1
UPS_serial_number=UPS_Fake_SN
WWNN=50050768010007E5
status=online
IO_group_id=0
IO_group_name=io_grp0
partner_node_id=
partner_node_name=
config_node=yes
UPS_unique_id=10000000000007E5
port_id=50050768011007E5
port_status=active
port_speed=2
port_id=50050768012007E5
port_status=inactive
port_speed=2
port_id=50050768013007E5
port_status=not_installed
port_speed=2
port_id=50050768014007E5
port_status=not_installed
port_speed=2
hardware=8F2
```

lsnodecandidate

The **lsnodecandidate** command lists all of the nodes that are not assigned to a cluster.

Syntax

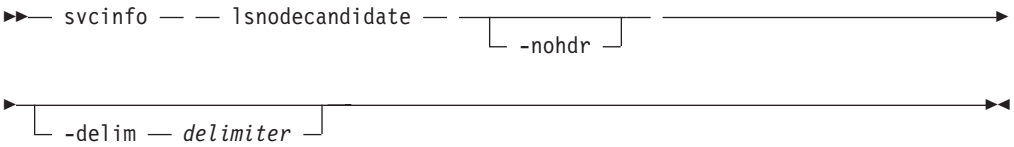

Parameters

-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.

-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.

An invocation example

```
svcinfo lsnodecandidate -delim :
```

The resulting output

```
id: panel_name: UPS_serial_number:
 UPS_unique_id: hardware
50050768010000EF: oneg55:
 UPS_Fake_SN:100000000000000EF: other
```

lsnodevpd

The **lsnodevpd** command returns the vital product data (VPD) for the given node.

Syntax

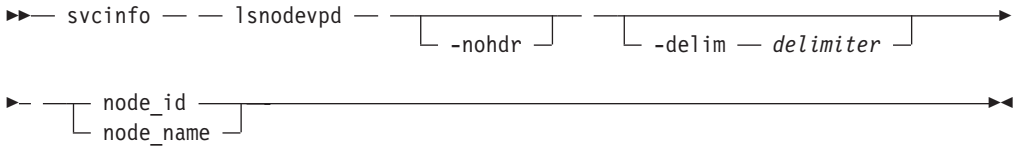

Parameters

-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.

An invocation example

```
svcinfo lsnodevpd 1
```

The resulting output

```
id 1

system board: 17 fields

part_number Unknown
system_serial_number 550117N
number_of_processors 2
number_of_memory_slots 4
number_of_fans 0
number_of_FC_cards 1

number_of_scsi/ide_devices 2
BIOS_manufacturer IBM
BIOS_version -[QAE115AUS-1.01]-
BIOS_release_date 08/16/2001
system_manufacturer IBM
system_product eserver xSeries 342 -[86691RX]-
planar_manufacturer IBM
power_supply_part_number Unknown
CMOS_battery_part_number Unknown
power_cable_assembly_part_number Unknown
service_processor_firmware N/A
```

```
processor: 6 fields
processor_location Processor 1
number_of_caches 2

manufacturer GenuineIntel
version Pentium III
speed 1000
status Enabled
processor cache: 4 fields
type_of_cache Internal L1 Cache
size_of_cache (KB) 32

type_of_cache Internal L2 Cache
size_of_cache (KB) 256

processor: 6 fields
processor_location Processor 2
number_of_caches 2
manufacturer GenuineIntel
version Pentium III
speed 1000
status Enabled

processor cache: 4 fields
type_of_cache Internal L1 Cache
size_of_cache (KB) 32

type_of_cache Internal L2 Cache
size_of_cache (KB) 256
memory module: 16 fields
part_number 33L5039
device_location J1
bank_location Slot1 in bank 1
size (MB) 1024
part_number 33L5039
device_location J4
bank_location Slot2 in bank 1
size (MB) 1024

part_number N/A
device_location J2
bank_location Slot1 in bank 2
size (MB) 0

part_number N/A
device_location J3
bank_location Slot2 in bank 2
size (MB) 0

FC card: 5 fields
part_number 64P7783
port_numbers 1 2
device_serial_number VSI 0000AD3F4
manufacturer Agilent
device DX2
device: 15 fields
part_number Unknown
bus ide0
device 0
model LG CD-ROM CRN-8245B
revision 1.13
serial_number
approx_capacity 0
part_number Unknown
bus scsi
device 0
```

```
device_vendor IBM-ESXS
model ST318305LC    !#
revision 6C48
serial_number 3JKQ93B903196C48
approx_capacity 8
software: 5 fields
code_level 00000000
node_name node1
ethernet_status 1
WWNN 0x50050768010007e5

id 1

front panel assembly: 3 fields
part_number Unknown
front_panel_id lynn02
front_panel_locale en_US

UPS: 10 fields
electronics_assembly_part_number FakElec
battery_part_number FakBatt
frame_assembly_part_number FakFram
input_power_cable_part_number FakCabl
UPS_serial_number UPS_Fake_SN
UPS_type Fake UPS
UPS_internal_part_number UPS_Fake_PN
UPS_unique_id 0x10000000000007e5
UPS_main_firmware 1.4
UPS_comms_firmware 0.0
```

lsrcconsistgrp

The **lsrcconsistgrp** command will return a concise list or a detailed view, of Metro Mirror consistency groups visible to the cluster.

The list report style can be used to obtain two styles of report.

1. The attributes of the Metro Mirror consistency group, plus the ID and name of every relationship which is in the group.
2. The detailed information about a single Metro Mirror consistency group.

Syntax

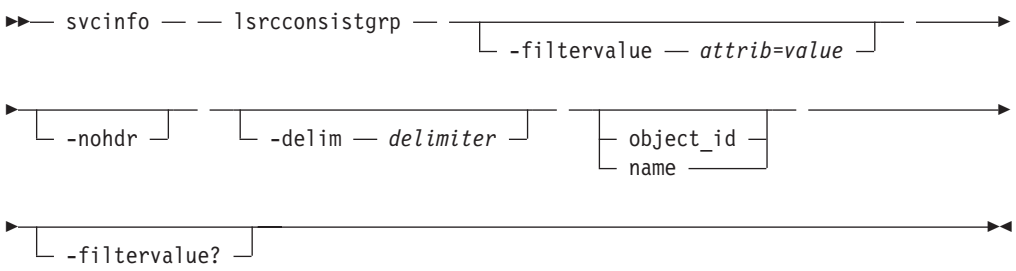

Parameters

-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.

-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 lsrcconsistgrp** 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 Metro Mirror 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 Metro Mirror 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.

A concise invocation example

```
svcinfo lsrcconsistgrp -delim :
```

The concise resulting output

```
id:name:master_cluster_id:master_cluster_name:aux_cluster_id:aux_cluster_name:pr
 imary:state:relationship_count
 248:jdemo_BA_cons1:0000020060406746:clusterB:0000020061413ABA:clusterA:master:co
 nsistent_stopped:2
 249:rccstgrp0:0000020061413ABA:clusterA:0000020061413ABA:clusterA::empty:0
 250:jdemo_BA_cons2:0000020060406746:clusterB:0000020061413ABA:clusterA:master:in
 consistent_stopped:1
 251:BA_cons1:0000020060406746:clusterB:0000020061413ABA:clusterA:master:consiste
 nt_stopped:4
 252:AB_cons2:0000020061413ABA:clusterA:0000020060406746:clusterB::empty:0
 253:AB_cons1:0000020061413ABA:clusterA:0000020060406746:clusterB:aux:consistent_
 stopped:3
 254:AA_cons2:0000020061413ABA:clusterA:0000020061413ABA:clusterA::empty:0
 255:AA_cons1:0000020061413ABA:clusterA:0000020061413ABA:clusterA:master:consiste
 nt_synchronized:2
```

A detailed invocation example

```
svcinfo lsrcconsistgrp -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
```

lsrcrelationship

The **lsrcrelationship** command will return a concise list or a detailed view, of Metro Mirror relationships visible to the cluster.

The list report style can be used to obtain two styles of report.

1. A list containing concise information about all the Metro Mirror relationships visible to the cluster. (Each entry in the list corresponds to a single Metro Mirror relationship.)

2. The detailed information about a single Metro Mirror relationship.

Syntax

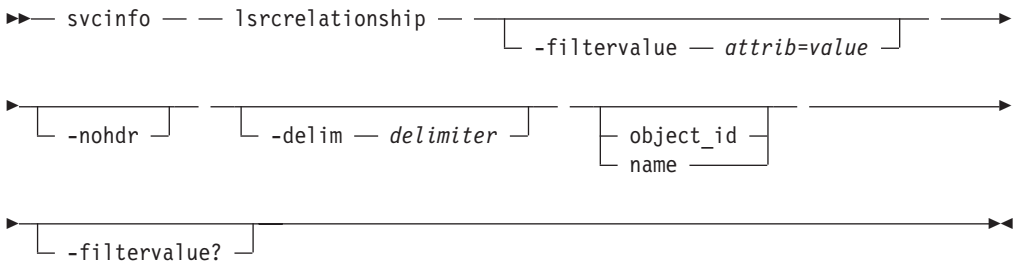

Parameters

-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.

-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 lsrcrelationship** 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 Metro Mirror 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, n/a
freeze time	The time in YY/MM/DD/HH/MM format.
status	online, primary_offline, secondary_offline
sync	n/a, in_sync, out_of_sync

Note: The names of the Metro Mirror 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.

A concise and detailed invocation example

```
svcinfo lsrcrelationship -delim : -filtervalue name=j*
```

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
 45:jrel_AB1:0000020061413ABA:clusterA:45:jdisk_B8:0000020060406746:clusterB:38:j
 disk_B1:master:::consistent_stopped:50:
 46:jrel_AB2:0000020061413ABA:clusterA:46:jdisk_A2:0000020060406746:clusterB:39:j
 disk_B2:master:::consistent_stopped:50:
 47:jrel_AB3:0000020061413ABA:clusterA:47:jdisk_A3:0000020060406746:clusterB:40:j
 disk_B3:master:::consistent_stopped:50:
 48:jrel_AB4:0000020061413ABA:clusterA:48:jdisk_A4:0000020060406746:clusterB:41:j
 disk_B4:master:::consistent_synchronized:50:
 49:jrel_BA_1:0000020060406746:clusterB:42:jdisk_B5:0000020061413ABA:clusterA:49:
```

```
jdisk_A5:master:248:jdemo_BA_cons1:consistent
_stopped:50:
50:jrel_BA_2:0000020060406746:clusterB:43:jdisk_B6:0000020061413ABA:clusterA:50:
jdisk_A6:master:248:jdemo_BA_cons1:consistent
_stopped:50:
51:jrel_BA_3:0000020060406746:clusterB:44:jdisk_B7:0000020061413ABA:clusterA:51:
jdisk_A7:master:250:jdemo_BA_cons2:inconsiste
nt_stopped:50:0
52:jrel_BA_4:0000020060406746:clusterB:45:jdisk_B8:0000020061413ABA:clusterA:52:
jdisk_A8:master:::inconsistent_stopped:50:0
```

A detailed invocation example

```
svcinfo lsrcrelationship -delim : AB_2
```

The detailed resulting output

```
id:9
name:AB_2
master_cluster_id:0000020061413ABA
master_cluster_name:clusterA
master_vdisk_id:9
master_vdisk_name:stripe9
aux_cluster_id:0000020060406746
aux_cluster_name:clusterB
aux_vdisk_id:9
aux_vdisk_name:stripe9_b
primary:master
consistency_group_id:
consistency_group_name:
state:consistent_stopped
bg_copy_priority:50
progress:
freeze_time:2003/07/05/08/26/46
status:secondary_offline
sync:in_sync
```

lsrcrelationshipcandidate

The **lsrcrelationshipcandidate** command lists VDisks that are eligible to form Metro Mirror relationships. You can list suitable VDisks on the local or remote cluster.

Syntax

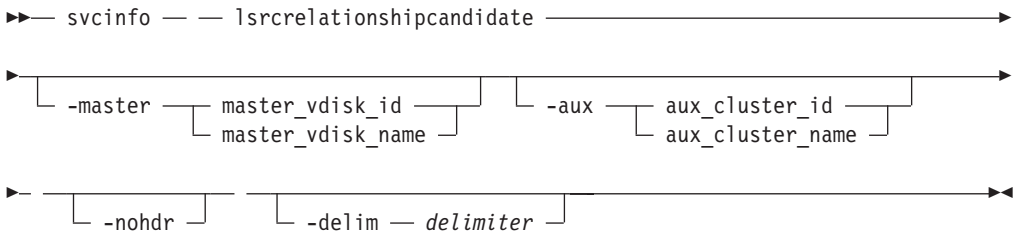

Parameters

-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 Metro Mirror 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.

An invocation example

```
svcinfo lsrcrelationshipcandidate -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 Metro Mirror relationship. When the initial background copy process for a relationship has completed, null will be displayed for the progress of that relationship.

Syntax

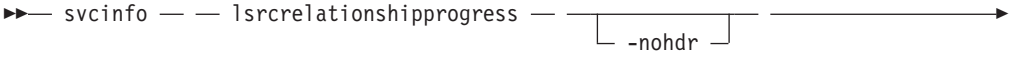

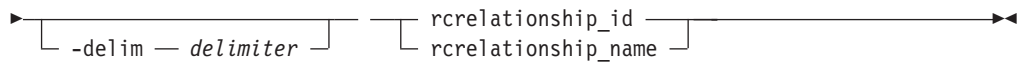

Parameters

-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.

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 Metro Mirror 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.

An invocation example

```
svcinfo lsrcrelationshipprogress -delim : 0
```

The resulting output

```
id:progress
0:58
```

lssoftwaredumps

You can use the **lssoftwaredumps** command to return a list of software packages from the /home/admin/upgrade directory.

Syntax

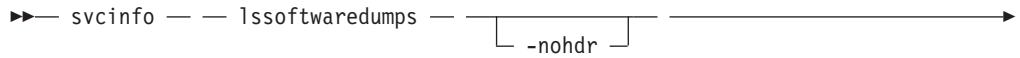

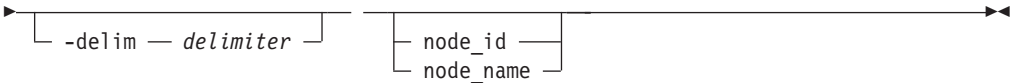

Parameters

-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 software upgrade packages. These packages are copied across as a result of software upgrades. If you do not specify a node, the packages available on the configuration node will be listed. The command will display files from the /home/admin/upgrade directory.

Possible failures

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

An invocation example

```
svcinfo lssoftwaredumps
```

The resulting output

```
id               software_filename
0                s1_mala75_030405_092143
1                s2_mala75_030405_092145
2                s3_mala75_030405_092146
```

lssshkeys

You can use the **lssshkeys** command to return a list of SSH keys that are available on the cluster.

Syntax

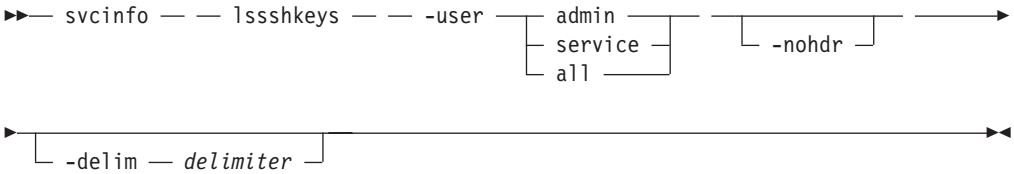

Parameters

-user *admin | service | all*
Specifies if the user wants to see a list of keys that can be used by a service user only, by an admin user only or the list for both user types.

-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 all the keys available on the cluster for the given userid.

Possible failures

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

An invocation example

```
svcinfo lssshkeys -user all -delim :
```

The resulting output

```
id:userid_key:identifier
1:admin:admin
```

lstimezones

The **lstimezones** command lists the timezones available on the cluster. Each timezone is assigned and ID that can be used in the **svctask settimezone** command.

Syntax

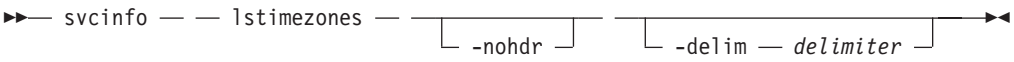

Parameters

-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.

Description

This command returns a list of all the timezones available on the cluster. Each timezone is assigned an ID. This ID can be used in the **svctask settimezone** command.

Possible failures

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

An invocation example

```
svcinfo lstimezones
```

The resulting output

```
id timezone
0 Africa/Abidjan
1 Africa/Accra
2 Africa/Addis_Ababa
3 Africa/Algiers
4 Africa/Asmera
5 Africa/Bamako
6 Africa/Bangui
```

lsvdisk

The **lsvdisk** command will return a concise list or a detailed view, of VDisks visible to the cluster.

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

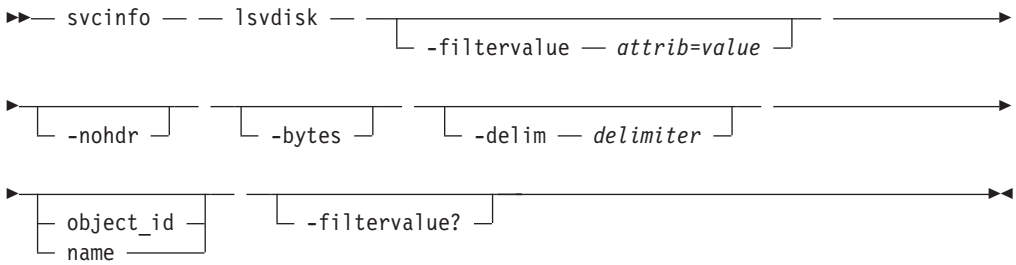

Parameters

-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 **svcinfo lsvdisk** command are:

- vdisk_name
- vdisk_id

- vdisk_UID
- 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, degraded
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
cache	read-write | none

The VDisk is offline and unavailable if one of the following take place:

- If both nodes in the I/O group are missing
- None of the nodes in the I/O group that are present can access the VDisk.

A VDisk is reported as degraded if the following happened:

- One of the nodes in the I/O group is missing.
- One of the nodes in the I/O group cannot access all the MDisks in the MDisk group that the VDisk spans. In this case MDisks will also be shown as degraded and the DMPs for MDisks should be followed to resolve the problem.
- The fast write cache pins data for one or more VDisks in the I/O group and is unable to perform a failback until the situation is resolved. An error log indicating that the cache has pinned data is displayed. Follow the DMP for this error log to resolve the problem. The most common causes of pinned data are the following:
  - One or more VDisks in an I/O group is offline due to an asymmetric failure and has pinned data in the cache. Asymmetric failures can occur because of SAN fabric faults or misconfiguration, back-end controller faults or

misconfiguration or because repeated errors has lead to the cluster excluding access to a MDisk through one or more nodes.
- One or more VDisks in an I/O group is offline due to a problem with a flash 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.

A concise invocation example

```
svcinfo lsvdisk -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:vdisk_UID
0:vdisk0:0:io_grp0:degraded:0:mdiskgrp0:16.0MB:striped:::::60050768017F06BF7800000000000000
1:vdisk1:0:io_grp0:degraded:0:mdiskgrp0:16.0MB:striped:::::60050768017F06BF7800000000000001
2:vdisk2:0:io_grp0:degraded:0:mdiskgrp0:16.0MB:striped:::::60050768017F06BF7800000000000002
3:vdisk3:0:io_grp0:degraded:0:mdiskgrp0:16.0MB:striped:::::60050768017F06BF7800000000000003
```

A detailed invocation example

```
svcinfo lsvdisk -delim : 1
```

The detailed resulting output

```
svcinfo lsvdisk -delim : 1
id:1
name:vdisk1
IO_group_id:0
IO_group_name:io_grp0
status:degraded
mdisk_grp_id:0
mdisk_grp_name:mdiskgrp0
capacity:16.0MB
type:striped
formatted:no
mdisk_id:
mdisk_name:
FC_id:
FC_name:
RC_id:
RC_name:
vdisk_UID:60050768017F06BF7800000000000001
throttling:0
preferred_node_id:1
fast_write_state:empty
cache:readwrite
udid:4326
```

lsvdiskextent

You can use the **lsvdiskextent** command to list the number of extents per MDisk making up the VDisk and to determine the number of extents used by each VDisk. The number of extents being provided by each MDisk is displayed.

Syntax

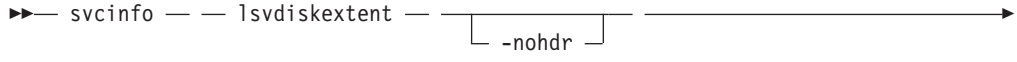

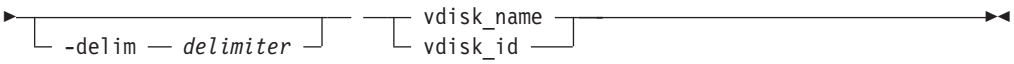

Parameters

-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_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 command returns 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: You can determine the number of extents that 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 command returns a table of MDisk IDs and the corresponding number of extents that 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 command returns 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: You can 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 command 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.
- 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 (MDisk) is not used by any virtual disk (VDisk).
- 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 (MDisk) or virtual disk (VDisk).

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. These hosts have the specified virtual disk mapped to them, that is, the virtual disk is visible to these hosts.

Syntax

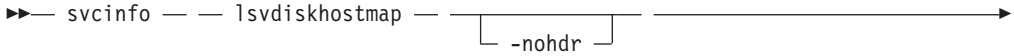

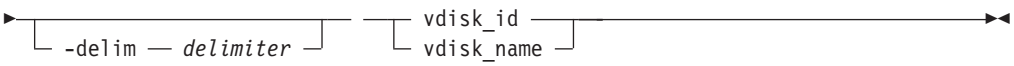

Parameters

-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 cluster 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.

The concise invocation example

```
svcinfo lsvdiskhostmap -delim : 3
```

The concise resulting output

```
id:name:SCSI_id:host_id:host_name:wwpn:vdisk_UID
3:vdisk3:0:2:host2:0000000000100ABC:60050768018A0001500000000000003
```

lsvdiskmember

The **lsvdiskmember** command returns a list of MDisks that are members of the specified VDisk.

Syntax

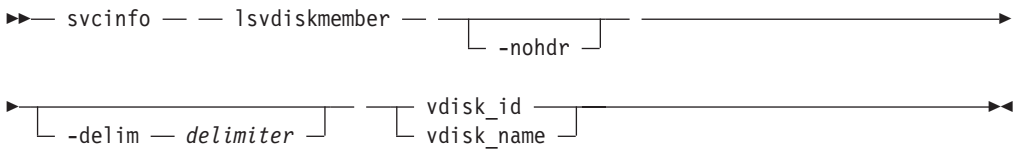

Parameters

-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 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.

An invocation example

```
svcinfo lsvdiskmember 1
```

The resulting output

```
id
2
```

lsvdiskprogress

The **lsvdiskprogress** command tracks the progress during new virtual disk formatting.

Syntax

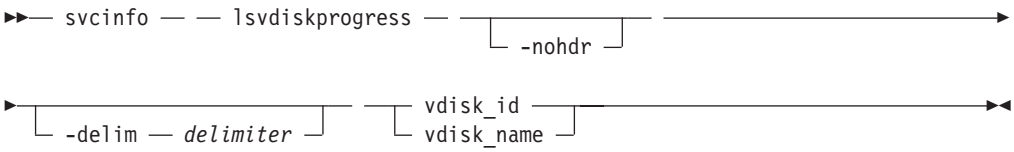

Parameters

-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 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.

An invocation example

```
svcinfo lsvdiskprogress -delim : 0
```

The resulting output

```
id:progress
0:58
```

showtimezone

Use the **showtimezone** command to display the current timezone settings for the cluster.

Syntax

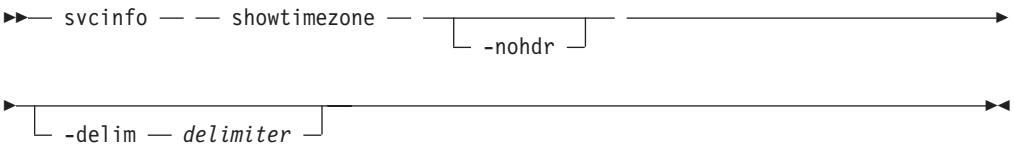

Parameters

-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.

-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 single time-zone and its associated ID. This is the current time-zone setting for the cluster. A list of available time-zones can be viewed by running the **svcinfo lstimezones** command. The time-zone can be changed by running the **svctask settimezone** command.

Possible failures

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

An invocation example

```
svcinfo showtimezone -delim :
```

The resulting output

```
id:timezone
522:UTC
```

Chapter 19. Error log commands

The following commands enable you to work with error logs for the SAN Volume Controller.

finderr

The **finderr** command analyzes the error log for the highest severity unfixed error.

Syntax

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.

An invocation example

```
svcservicetask finderr
```

The resulting output

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

dumperrlog

The **dumperrlog** command dumps the contents of the error log to a text file.

Syntax

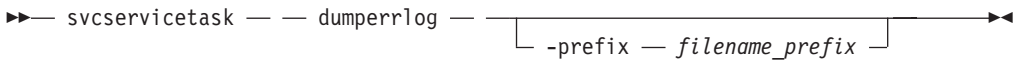

Parameters

-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 within 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.

Files **will not** be deleted from other nodes until you issue the **cleardumps** 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.

An invocation example

```
svcservicetask dumperrlog -prefix testerrorlog
```

The resulting output

```
No feedback
```

clearerrlog

The **clearerrlog** command clears all entries from the error log including status events and any unfixed errors.

Syntax

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.

An invocation example

```
svcservicetask clearerrlog -force
```

The resulting output

```
No feedback
```

cherrstate

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

Syntax

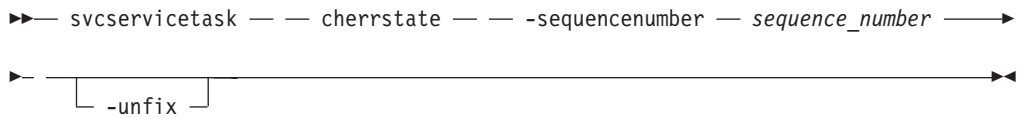

Parameters

-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.

An invocation example

```
svcservicetask cherrstate -sequencenumber 2019
```

The resulting output

```
No feedback
```

setevent

Use the **setevent** command to specify what happens when an error or event is logged to the error log. These settings define what to do when errors and events are logged.

Syntax

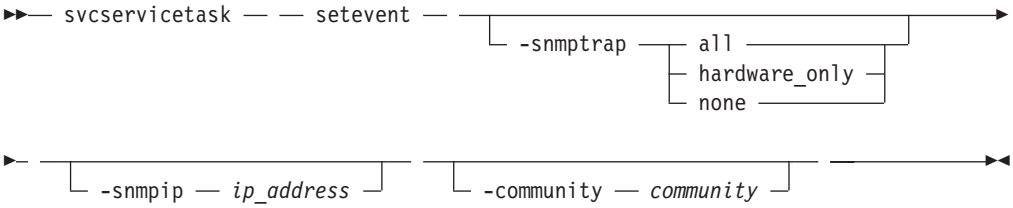

Parameters

-snmptrap *all* | *hardware_only* | *none*
Optionally specifies the SNMP trap setting, that is, when to raise a trap.

-snmpip *ip_address*
Optionally specifies the IP address of the host system that is running the SNMP manager software. This is a colon-separated list of values with up to six items per list.

-community *community*
Optionally specifies the SNMP community string. This is a colon-separated list of values with up to six items per list. The maximum length of the community string that is used in SNMP trap generation cannot be more than 60 characters.

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 argument:

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.
- The IP address of the SNMP manager
- The SNMP community

Possible failures

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

An invocation example

```
svcservicetask setevent -snmptrap all
```

The resulting output

```
No feedback
```

Chapter 20. Featurization commands

The following commands enable you to work with featurization functions with the SAN Volume Controller.

chlicense

The **chlicense** command changes the licensed settings of the cluster.

Syntax

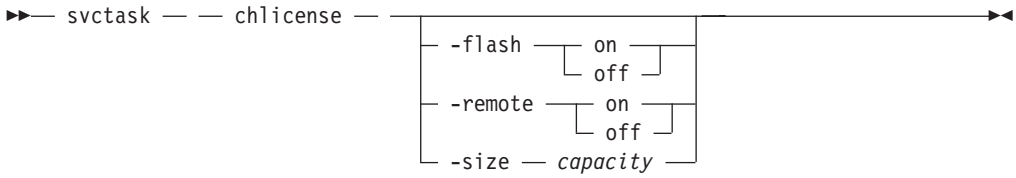

Parameters

-flash *on | off*
Optionally specifies if flash copy has been licensed for this cluster. Default is off.

-remote *on | off*
Optionally specifies if Metro Mirror 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 Metro Mirror. 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 message from the command line tool. This will not stop you from creating and expanding virtual disks. When usage reaches (and exceeds) 100% capacity, errors will be placed in the featurization log. Again, this will not prevent you from creating and expanding virtual disks.

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 90% or more of, 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.
- CMMVC5998W The virtualized storage capacity exceeds the amount that your are licensed to use.

An invocation example

```
svctask chlicense -flash on
```

The resulting output

```
No feedback
```

dumpinternallog

The **dumpinternallog** command dumps the contents of the featurization error and event log to a file on the current configuration node.

Syntax

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.

Before making any entries, the featurization log contains zeroes. A dump of this log by using the **svctask dumpinternallog** command would result in an empty file.

This file may be requested by 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.

An invocation example

```
svctask dumpinternallog
```

The resulting output

```
No feedback
```

Chapter 21. Secure Shell key commands

The following commands enable you to work with Secure Shell (SSH) with the SAN Volume Controller.

addsshkey

The **addsshkey** command installs a new SSH key on the cluster. The key file must first be copied onto the cluster.

Syntax

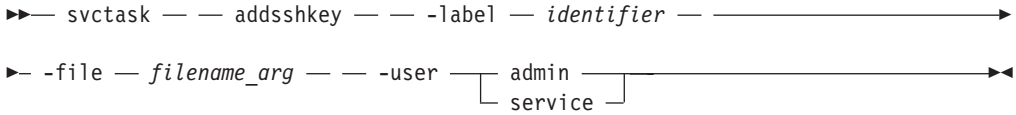

Parameters

-label *identifier*
Specifies a new identifier to associate with this key. Max length = 30 characters.

-file *filename_arg*
Specifies the name of the file containing the SSH key.

-user *admin | service*
Specifies which userid the SSH key will be applied to.

Description

You must first copy the key file to the cluster in the /tmp directory using secure copy (scp).

The **svctask addsshkey** command moves the key file from the /tmp directory, to the required location and activates it for the given user. Once the key has been activated, you will be able to invoke commands on the cluster over SSH using the userid that was specified from the host for which the key was generated. Alternatively you can run an interactive SSH session from the specified host using the userid specified.

The identifier can be used to subsequently identify the key when all keys are listed using the **svcinfo lssshkeys** command, or if the key is to be deleted using the **svctask rmsshkey** command.

The identifier is also specified as the key label which is used in the Audit Log. The Audit Log uses the key label to identify the commands that were issued through an SSH session which was established with an associated key. If you want to use the Audit Log, you must ensure that there is a 1:1 ratio mapping from the key label identifier to the key for each authorized user.

Attention: After you add a cluster, close the Maintaining SSH Keys panel.

You can issue the **svcinfo lssshkeys** command to list the SSH keys that are available on the cluster. You can issue the **svctask addsshkey** command to install a new SSH key on the cluster. The key file must first be copied onto the cluster. Each key is associated with an ID string that you define that can consist of up to 30 characters. Up to 100 keys can be stored on a cluster. You can add keys to provide either administrator access or service access. For example, type the following:

```
svctask addsshkey -user service -file /tmp/id_rsa.pub -label testkey
```

where */tmp/id_rsa.pub* is the name of the file that the SSH key will be saved in and *testkey* is the label to associate with this key.

You can issue the **svctask rmsshkey** command to remove an SSH key from the cluster. You can issue the **svctask rmallsshkeys** command to remove all of the SSH keys from the cluster.

Possible failures

- CMMVC5786E The action failed because the cluster is not in a stable state.
- CMMVC6008E The key already exists.

An invocation example

```
svctask addsshkey -user service -file /tmp/id_rsa.pub -label testkey
```

The resulting output

```
No feedback
```

rmallsshkeys

Use the **rmallsshkeys** command to remove an SSH key. The command will remove all the SSH keys that are applicable to the specified -user argument.

Syntax

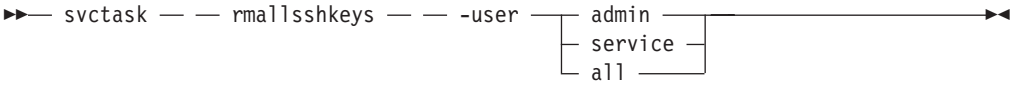

Parameters

-user *admin* | *service* | *all*
The -user argument is used to specify which userid the SSH key will be applied to; if you specify all, all the SSH keys in the cluster will be removed.

Description

This command will remove all the SSH keys that are applicable to the specified -user argument.

Attention: After you add a cluster, close the Maintaining SSH Keys panel.

You can issue the **svcinfo lssshkeys** command to list the SSH keys that are available on the cluster. You can issue the **svctask addsshkey** command to install a new SSH key on the cluster. The key file must first be copied onto the cluster. Each key is associated with an ID string that you define that can consist of up to 30

characters. Up to 100 keys can be stored on a cluster. You can add keys to provide either administrator access or service access. For example, type the following:

```
svctask addsshkey -user service -file /tmp/id_rsa.pub -label testkey
```

where */tmp/id_rsa.pub* is the name of the file that the SSH key will be saved in and *testkey* is the label to associate with this key.

You can issue the **svctask rmsshkey** command to remove an SSH key from the cluster. You can issue the **svctask rmallsshkeys** command to remove all of the SSH keys from the cluster.

Possible failures

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

An invocation example

```
svctask rmallsshkeys -user service
```

The resulting output

```
No feedback
```

rmsshkey

Use the **rmsshkey** command to remove an SSH key specified by the -key argument.

Syntax

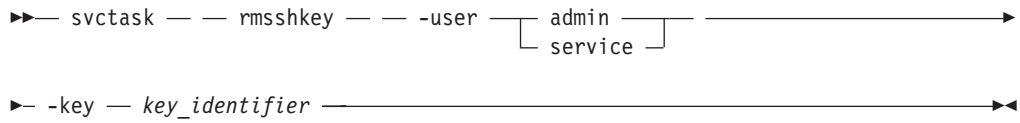

Parameters

-user *admin | service*
The -user argument is used to specify from which userid the SSH key will be deleted.

-key *key_identifier*
The label of the key to be deleted.

Description

When executed this command will remove the SSH key specified by the -key argument.

CAUTION:
If you want to change your SSH key, you must add your new key *before* **removing the old key. If you have an SSH session currently in use, once the SSH key that was used to establish that session is removed, the session will be unable to issue any further commands. This will prevent inaccurate audit log entries.**

Attention: After you add a cluster, close the Maintaining SSH Keys panel.

You can issue the **svcinfo lssshkeys** command to list the SSH keys that are available on the cluster. You can issue the **svctask addsshkey** command to install a new SSH key on the cluster. The key file must first be copied onto the cluster. Each key is associated with an ID string that you define that can consist of up to 30 characters. Up to 100 keys can be stored on a cluster. You can add keys to provide either administrator access or service access. For example, type the following:

```
svctask addsshkey -user service -file /tmp/id_rsa.pub -label testkey
```

where */tmp/id_rsa.pub* is the name of the file that the SSH key will be saved in and *testkey* is the label to associate with this key.

You can issue the **svctask rmsshkey** command to remove an SSH key from the cluster. You can issue the **svctask rmallsshkeys** command to remove all of the SSH keys from the cluster.

Possible failures

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

An invocation example

```
svctask rmsshkey  -key testkey -user service
```

The resulting output

```
No feedback
```

Chapter 22. Service mode commands

Service mode commands perform tasks when the node is in service mode. Some of these tasks are specifying node software, cleaning dump directories and dumping the contents of an error log to a file.

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.

applysoftware

The **applysoftware** command upgrades the cluster to a new level of software.

Syntax

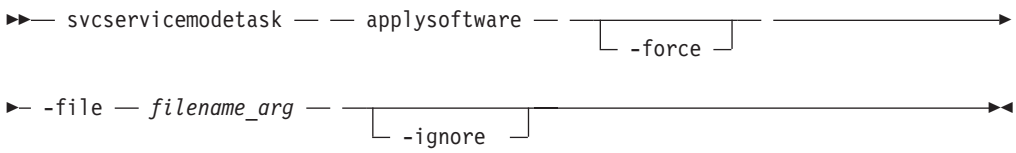

Parameters

-force
Optionally specifies the force flag. The -force flag is needed if any node within an I/O group is not paired. The upgrade process forces the first node in each I/O group to shut down and upgrade. If that node is not paired, then the cluster will become degraded and the data will be lost.

-file *filename_arg*
Specifies the filename of the new software package.

-ignore
Optional parameter that, if used, will result in all prerequisite checking being bypassed and all hardened data to be deleted. This option should be used with care.

Description

This command starts the upgrade process of the cluster to a new level of software and can be applied to **svcservicetask** and **svcservicemodetask**. The **applysoftware** command can be used to apply a level of software to the node in both service and

non-service modes. In service mode the **applysoftware** command is applied to the specific node in service mode. In non-service mode, the command is applied to the complete cluster.

The software package as specified by the file name must first be copied on to the current configuration node in the /home/admin/upgrade directory. You can use PuTTy secure copy (scp) to copy the file. See "PuTTY scp" for detailed information on this procedure.

The actual upgrade completes asynchronously.

The contents of /home/admin/upgrade can be viewed by using the **svcinfo lssoftwaredumps** command.

Internally, the new package will be moved from the /home/admin/upgrade directory and checksumed. If the package fails the checksum, it will be deleted and the upgrade will fail. Otherwise, the package will be extracted from the directory and the software upgrade will begin.

Possible failures

- There are no error codes.

An invocation example

```
svcservicemodetask applysoftware -file newsoftware
```

The resulting output

```
No feedback
```

cleardumps

The **cleardumps** command cleans the various dump directories on a specified node.

Syntax

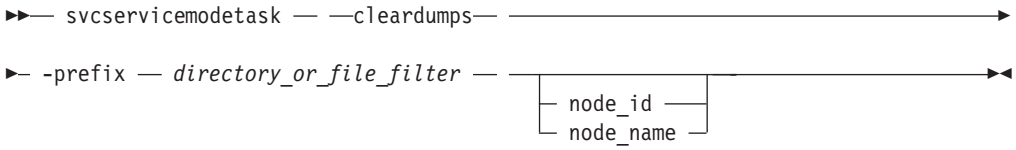

Parameters

-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/audit
- /dumps/configs
- /dumps/elogs
- /dumps/feature
- /dumps/iostats

- /dumps/iotrace
- /home/admin/upgrade

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 list the contents of these directories on the given node by using the **svcinfo lsxxxxdumps** commands.

You can use this command to clean specific files in a given directory by specifying a directory or file name. The wildcard * can be used as part of the file name.

Note: In order to preserve the configuration and trace files, any files matching the following wildcard patterns will not be cleaned:

- *svc.config*
- *.trc
- *.trc.old

Possible failures

- There are no error codes associated with the **svcservicemodetask cleardumps** command.

An invocation example

```
svcservicemodetask cleardumps -prefix /dumps/configs
```

The resulting output

```
No feedback
```

dumperrlog

The **dumperrlog** command dumps the contents of the error log to a text file.

Syntax

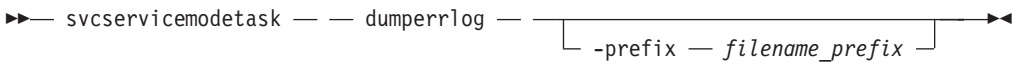

Parameters

-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 within 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.

Files **will not** be deleted from other nodes until you issue the **cleardumps** 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.

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

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.

An invocation example

```
svcservicemodetask exit
```

The resulting output

```
[SSH / webserver connections terminate so an error message to the effect of
'connection lost' may be displayed, or 'CLIENT RECEIVED SERVER DOWN
NOTIFICATION']
```

Chapter 23. Service mode information commands

Service mode information commands perform information gathering tasks when the node is in service mode.

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.

ls2145dumps

The **ls2145dumps** command obtains a list of node assert dumps and associated output files from the /dumps directory.

Syntax

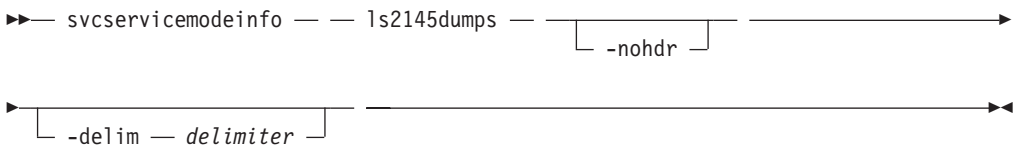

Parameters

-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.

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 directory.

Possible failures

- There are no error codes.

An invocation example

```
svcservicemodeinfo ls2145dumps -delim :
```

The resulting output

```
id:2145_filename
0:000108.trc.old
1:dump.000108.030328.144007
2.000108.trc
```

lsclustervpd

You can use the **lsclustervpd** command to return the vital product data (VPD) for the cluster to which the node belonged.

Syntax

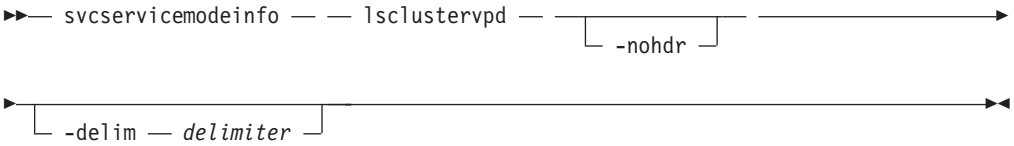

Parameters

-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.

Description

This command displays the VPD of the cluster to which the node belonged.

Possible failures

- There are no error codes.

An invocation example

```
svcservicemodeinfo lsclustervpd
```

The resulting output

```
id 000001002FF007E5
name kkk
location local
partnership
bandwidth 0
cluster_IP_address 0.0.0.0
cluster_service_IP_address 1.1.1.1
total_mdisk_capacity 0
space_in_mdisk_grps 0
space_allocated_to_vdisks 0
total_free_space 0
statistics_status off
statistics_frequency 15
required_memory 2048
cluster_locale en_US
SNMP_setting all
SNMP_community
SNMP_server_IP_address 0.0.0.0
subnet_mask 0.0.0.0
default_gateway 0.0.0.0
time_zone 522 UTC
email_setting all
email_id
code_level 00000000
FC_port_speed 1Gb
```

lserrlogdumps

The **lserrlogdumps** commands return a list of error log dumps in the /dumps/elogs directory. These dumps are created as a result of issuing the **svctask dumperrlog** command.

Syntax

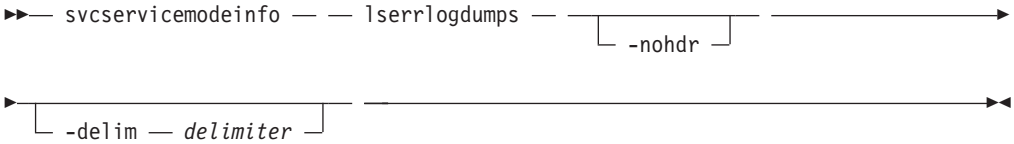

Parameters

-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.

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

- There are no error codes.

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

The **lsfeaturedumps** command returns a list of dumps in /dumps/feature. These dumps are created as a result of issuing the **svctask dumpinternallog** command.

Syntax

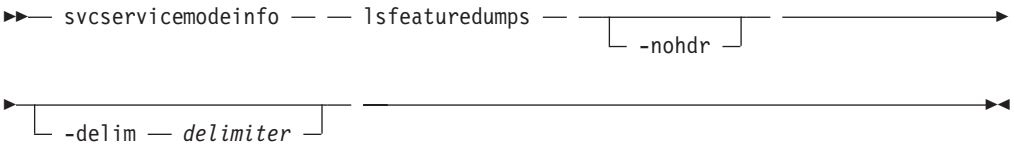

Parameters

-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.

Description

This command returns a list of featurization dumps. These dumps are created as a result of issuing the **svctask dumpinternallog** 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

- There are no error codes.

An invocation example

```
svcservicemodeinfo lsfeaturedumps
```

The resulting output

```
id              feature_filename
0               feature.txt
```

lsiostatsdumps

The **lsiostatsdumps** command returns a list of dumps in the /dumps/iostats directory. These dumps are created as a result of issuing the **svctask startstats** command.

Syntax

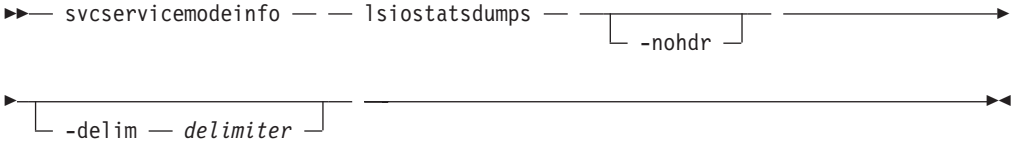

Parameters

-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.

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

- There are no error codes.

An invocation example

```
svcservicemodeinfo lsiostatsdumps
```

The resulting output

```
id              iostat_filename
0               v_stats_mala75_031123_072426
1               m_stats_mala75_031123_072425
```

lsiotracedumps

You can use the **lsiotracedumps** command to return a list of files in the /dumps/iotrace directory.

Syntax

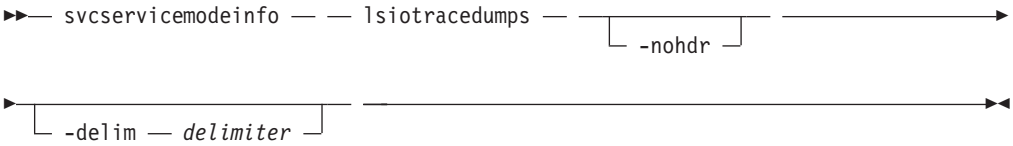

Parameters

-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.

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

- There are no error codes.

An invocation example

```
svcservicemodeinfo lsiotracedumps
```

The resulting output

```
id               iotrace_filename
0                c1_mala75_030405_092155
1                c2_mala75_030405_092156
2                c3_mala75_030405_092158
3                c4_mala75_030405_092159
4                c5_mala75_030405_092201
```

lsnodes

The **lsnodes** command will return an annotated list of the nodes in the cluster in service mode.

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

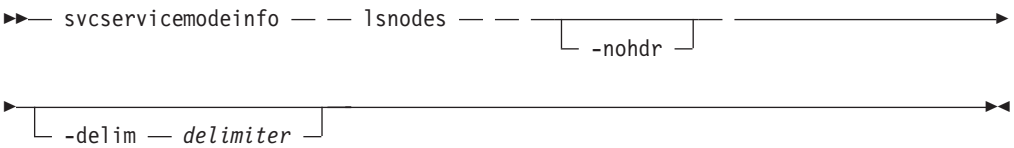

Parameters

-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.

A concise invocation example

```
svcservicemodeinfo lsnodes -delim :
```

The concise resulting output

```
id:WWNN:front_panel_id:node_name:cluster:fabric
1:50050768010007E5:lynn02:node1:yes:yes
```

lsnodevpd

The **lsnodevpd** command returns the vital product data (VPD) for the given node.

Syntax

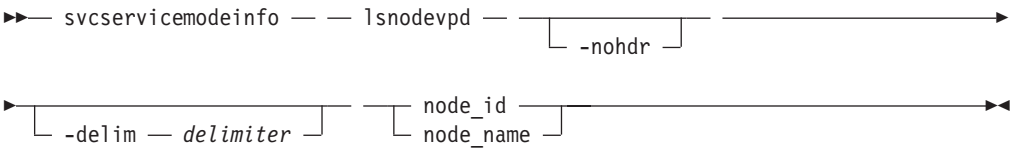

Parameters

-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

- There are no error codes.

An invocation example

```
svcservicemodeinfo lsnodevpd id 1
```

The resulting output

```
id 1

system board: 17 fields

part_number Unknown
system_serial_number 550117N
number_of_processors 2
number_of_memory_slots 4
number_of_fans 0
number_of_FC_cards 1

number_of_scsi/ide_devices 2
BIOS_manufacturer IBM
BIOS_version -[QAE115AUS-1.01]-
BIOS_release_date 08/16/2001
system_manufacturer IBM
system_product eserver xSeries 342 -[86691RX]-
planar_manufacturer IBM
power_supply_part_number Unknown
CMOS_battery_part_number Unknown
power_cable_assembly_part_number Unknown
service_processor_firmware N/A

processor: 6 fields
processor_location Processor 1
number_of_caches 2

manufacturer GenuineIntel
version Pentium III
speed 1000
status Enabled
processor cache: 4 fields
type_of_cache Internal L1 Cache
size_of_cache (KB) 32

type_of_cache Internal L2 Cache
size_of_cache (KB) 256

processor: 6 fields
processor_location Processor 2
number_of_caches 2
manufacturer GenuineIntel
version Pentium III
speed 1000
status Enabled

processor cache: 4 fields
type_of_cache Internal L1 Cache
size_of_cache (KB) 32

type_of_cache Internal L2 Cache
size_of_cache (KB) 256
memory module: 16 fields
part_number 33L5039
device_location J1
bank_location Slot1 in bank 1
size (MB) 1024
part_number 33L5039
device_location J4
bank_location Slot2 in bank 1
size (MB) 1024

part_number N/A
```

```
device_location J2
bank_location Slot1 in bank 2
size (MB) 0

part_number N/A
device_location J3
bank_location Slot2 in bank 2
size (MB) 0

FC card: 5 fields
part_number 64P7783
port_numbers 1 2
device_serial_number VSI 0000AD3F4
manufacturer Agilent
device DX2
device: 15 fields
part_number Unknown
bus ide0
device 0
model LG CD-ROM CRN-8245B
revision 1.13
serial_number
approx_capacity 0
part_number Unknown
bus scsi
device 0
device_vendor IBM-ESXS
model ST318305LC    !#
revision 6C48
serial_number 3JKQ93B903196C48
approx_capacity 8
software: 5 fields
code_level 00000000
node_name node1
ethernet_status 1
WWNN 0x50050768010007e5

id 1

front panel assembly: 3 fields
part_number Unknown
front_panel_id lynn02
front_panel_locale en_US

UPS: 10 fields
electronics_assembly_part_number FakElec
battery_part_number FakBatt
frame_assembly_part_number FakFram
input_power_cable_part_number FakCabl
UPS_serial_number UPS_Fake_SN
UPS_type Fake UPS
UPS_internal_part_number UPS_Fake_PN
UPS_unique_id 0x10000000000007e5
UPS_main_firmware 1.4
UPS_comms_firmware 0.0
```

lssoftwaredumps

You can use the **lssoftwaredumps** command to return a list of software packages from the /home/admin/upgrade directory.

Syntax

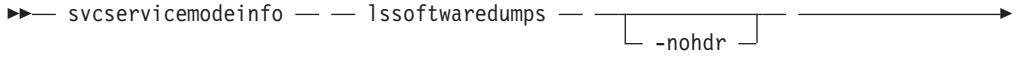

Parameters

-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.

Description

This command returns a list of software upgrade packages. These packages are copied across as a result of software upgrades. If you do not specify a node, the packages available on the configuration node will be listed. The command will display files from the /home/admin/upgrade directory.

Possible failures

- There are no error codes.

An invocation example

```
svcservicemodeinfo lssoftwaredumps
```

The resulting output

```
id               software_filename
0                s1_mala75_030405_092143
1                s2_mala75_030405_092145
2                s3_mala75_030405_092146
```

Chapter 24. Controller command

The following command enables you to work with controllers in the SAN Volume Controller.

chcontroller

You can use the **chcontroller** command to modify the name of a controller.

Syntax

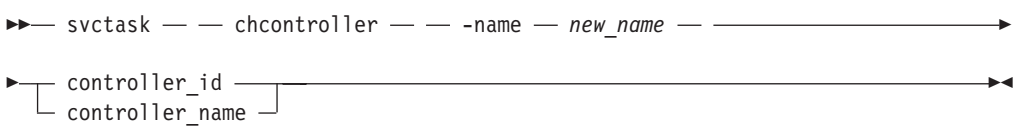

Parameters

-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 cluster.

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 cluster 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 **svcinfo lscontroller** 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:

```
svcinfo lsmdisk -filtervalue mode=unmanaged:controller_name=<new_name>
```

These MDisks should correspond with the RAID arrays or partitions you have created. Remember the field controller LUN number. This corresponds with the LUN number you assigned to each of the arrays or partitions.

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.

An invocation example

```
svctask chcontroller -name newtwo 2
```

The resulting output

```
No feedback
```

Chapter 25. Command-line interface 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 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

Their 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

'The value used in the -filtervalue parameter was not a recognized filter value. Refer to the help for a list of valid filters.

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.

CMMVC5770E The file for the ssh key is not valid.

Explanation

The file for the ssh key is not valid.

Action

Specify a different file 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.

CMMVC5774E The userid is not valid.

Explanation

The userid is not valid.

Action

Specify a different userid 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 in the same power domain.

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.

CMMVC5780E The action could not be completed using the Remote Cluster name. Use the Remote Cluster Unique ID instead.

Explanation

The unique ID of the remote cluster is required for this command.

Action

Specify the unique ID of the remote cluster, and issue the command again.

CMMVC5781E The cluster ID specified is invalid.

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.

CMMVC5788E The service IP address is not valid.

Explanation

The service IP address is not valid.

Action

Specify the correct service IP address and issue the command again.

CMMVC5789E The cluster was not modified because the IP address, subnet mask, service address, SNMP address, or gateway address is not valid.

Explanation

The cluster was not modified because the IP address, subnet mask, service address, SNMP address, or gateway address is not valid.

Action

Specify all correct attributes and issue the command again.

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

A previous configuration command might not yet have completed. Wait for the previous command to complete and then re-try this command.

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 or sequential-mode virtual disk (VDisk) was not created because more than one managed disk (MDisk) is specified.

Explanation

The image-mode or sequential-mode VDisk was not created because more than one MDisk is specified.

Action

Specify a different MDisk and reissue the command.

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

Commands may still be running in the background.

Action

Wait for all commands to complete. Use the svcinfo lsmigrate command to view any migrates running in the background.

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 Metro Mirror mapping.

Explanation

The action failed because the virtual disk (VDisk) is part of a Metro Mirror mapping.

Action

Specify a different VDisk that is not part of a Metro Mirror 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 Metro Mirror 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 Metro Mirror 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

The VDisk or MDisk does not belong to the MDisk group specified.

Action

Select the correct MDisk group and retry.

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 an done 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 svcinfo

lsfreeextents <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 Metro Mirror mapping.

Explanation

The FlashCopy mapping was not created because the source or target virtual disk (VDisk) cannot be a member of a Metro Mirror 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.

CMMVC5887E The FlashCopy mapping was not created because the source or target virtual disk (VDisk) must not be in router mode.

Explanation

The FlashCopy mapping was not created because the source or target virtual disk (VDisk) must not be in router mode.

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 the 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 stopped 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 Metro Mirror partnership was not deleted because it has already been deleted.

Explanation

The Metro Mirror partnership was not deleted because it has already been deleted.

Action

Not applicable.

CMMVC5930E The Metro Mirror relationship was not created because an entity that was specified in the command does not exist.

Explanation

The Metro Mirror 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 Metro Mirror relationship was not created because the master or auxiliary virtual disk (VDisk) is locked.

Explanation

The Metro Mirror 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 Metro Mirror relationship was not created because the master or auxiliary virtual disk (VDisk) is a member of a FlashCopy mapping.

Explanation

The Metro Mirror relationship was not created because the master or auxiliary virtual disk (VDisk) is a member of a FlashCopy mapping.

Action

Not applicable.

CMMVC5933E The Metro Mirror relationship was not created because the master or auxiliary virtual disk (VDisk) is in the recovery I/O group.

Explanation

The Metro Mirror relationship was not created because the master or auxiliary virtual disk (VDisk) is in the recovery I/O group.

Action

Not applicable.

CMMVC5934E The Metro Mirror relationship was not created because the master or auxiliary virtual disk (VDisk) is in the router mode.

Explanation

The Metro Mirror 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 Metro Mirror consistency group was not deleted because the consistency group contains relationships. To delete the consistency group, the force option is required.

Explanation

The Metro Mirror 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 Metro Mirror relationship was not created because the I/O group is offline.

Explanation

The Metro Mirror relationship was not created because the I/O group is offline.

Action

Not applicable.

CMMVC5970E The Metro Mirror relationship was not created because there is not enough memory.

Explanation

The Metro Mirror 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

Specify the Force option when starting the consistency group.

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.

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.

CMMVC5985E 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, or /home/admin/upgrade.

Explanation

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
- /home/admin/upgrade

Action

Specify one of the above directories and issue the command again.

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 operation was not performed because the relationship is offline.

Explanation

The operation was not performed because the relationship is offline.

Action

Not applicable.

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 Metro Mirror consistency group was not stopped as there are no Metro Mirror relationships within the group.

Explanation

The Metro Mirror consistency group was not stopped as there are no Metro Mirror relationships within the group.

Action

Not applicable.

CMMVC5993E The specific upgrade package does not exist.

Explanation

The specific upgrade package does not exist.

Action

Not applicable.

CMMVC5994E Error in verifying the signature of the upgrade package.

Explanation

The system could not verify the signature of the upgrade package due to the following reasons:

- There is not enough space on the system to copy the file.
- The package is incomplete or contains errors.

Action

If the copy failed with an error indicating that there was insufficient space on the system, free up additional space on your system. Otherwise, ensure that the cluster time and date stamp on the signature is correct. (For example, the time and date cannot be in the future.)

CMMVC5995E Error in unpacking the upgrade package.

Explanation

An error occurred when the system was unpacking the upgrade package. The most likely cause of this error is lack of system space.

Action

Reboot the node and unpack the upgrade package again.

CMMVC5996E The specific upgrade package cannot be installed over the current version.

Explanation

The upgrade package is not compatible with the current version or with your system.

Action

Check the available upgrade packages and find the correct upgrade package for your current version and for your system. If the upgrade package is correct for your system, check the version requirements for the package. You may have to upgrade the current version to an intermediate version before you upgrade to the latest version. (For example, if your current version is 1 and you are trying to upgrade to version 3, you may need to upgrade to version 2 before applying the version 3 upgrade.)

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 that 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. (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. Nevertheless, the action you have requested has been completed.

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.

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 not 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.

CMMVC6007E The two passwords that were entered do not match.

Explanation

The two passwords entered for verification of your password change were not the same.

Action

Re-enter the passwords.

CMMVC6008E The key already exists.

Explanation

An attempt was made to load a duplicate SSH key.

Action

Not applicable.

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 Metro Mirror relationship to a remote cluster exists.

Action

Delete the Metro Mirror 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 Metro Mirror consistency groups involved.

Action

Ensure that the attributes of the two Metro Mirror 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.

CMMVC6018E The software upgrade pre-install process failed.

Explanation

The software upgrade failed as there was an error during the preprocessing. The package is either invalid or corrupted.

Action

Ensure the package is a valid upgrade package. Download the package from the source location again as it may have been corrupted during a network transfer.

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 Metro Mirror consistency group is not in the stopped state.

Action

Ensure that the Metro Mirror 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 in not the Metro Mirror 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.

CMMVC6033E The action failed due to an internal error.

Explanation

An internal error caused the action to fail.

Action

Not applicable.

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.

CMMVC6035E The action failed as the object already exists.

Explanation

An operation was requested to create an object that already exists.

Action

Ensure that the name you are attempting to apply to a new object does not exist, or change the name before re-issuing the command.

CMMVC6036E An invalid action was requested.

Explanation

The action failed because it is not a valid action with the command that was issued.

Action

Issue an action that is valid with the command.

CMMVC6037E The action failed as the object is not empty.

Explanation

The action failed because an object was specified.

Action

Issue the command again, and do not specify an object.

CMMVC6038E The action failed as the object is empty.

Explanation

The action failed because an object was not specified.

Action

Specify an object, and issue the command again.

CMMVC6039E The action failed as the object is not a member of a group.

Explanation

The action failed because the object is not a member of a group.

Action

Specify an object that is part of a group, and issue the command again.

CMMVC6040E The action failed as the object is not a parent.

Explanation

The action failed because the object is not a parent object.

Action

Specify an object that is a parent, and issue the command again.

CMMVC6041E The action failed as the cluster is full.

Explanation

The action failed because the cluster is full.

Action

Remove data from the cluster, and issue the command again.

CMMVC6042E The action failed as the object is not a cluster member.

Explanation

The action failed because the object is not a member of the cluster.

Action

Specify an object that is a member of the cluster, and issue the command again.

CMMVC6043E The action failed as the object is a member of a group.

Explanation

The action failed because the object is a member of a group.

Action

Specify an object that is not a member of a group, and issue the command again.

CMMVC6044E The action failed as the object is a parent.

Explanation

The action failed because the object is a parent object.

Action

Specify an object that is not a parent object, and issue the command again.

CMMVC6045E The action failed as the force flag was not entered.

Explanation

The action failed because the -force option was not entered.

Action

Specify the -force option in the command.

CMMVC6046E The action failed as too many candidates were selected.

Explanation

The action failed because too many candidates were specified.

Action

Specify fewer candidates in the command.

CMMVC6047E The action failed as too few candidates were selected.

Explanation

An action was requested with too few candidate objects.

Action

Determine the correct number of candidates required for the specific command and re-issue the command.

CMMVC6048E The action failed as the object is busy.

Explanation

The action failed because the object is busy.

Action

Not applicable.

CMMVC6049E The action failed as the object is not ready.

Explanation

The action failed because the object is not ready.

Action

Not applicable.

CMMVC6050E The action failed as the command was busy.

Explanation

The action failed because the command is busy.

Action

Not applicable.

CMMVC6051E An unsupported action was selected.

Explanation

The action failed because it is not valid with the command.

Action

Specify an action that is valid with the command.

CMMVC6052E The action failed as the object is a member of a Flash Copy mapping.

Explanation

The object is a member of a FlashCopy mapping, thus it cannot be deleted.

Action

Specify an object that is not a member of a FlashCopy mapping, or remove the object from the FlashCopy mapping.

CMMVC6053E An invalid WWPN was entered.

Explanation

An invalid World Wide Port Name (WWPN) was specified.

Action

Specify a valid WWPN.

CMMVC6054E The action failed as not all nodes are online.

Explanation

The action 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.

CMMVC6055E The action failed as an upgrade is in progress.

Explanation

The action failed because a software upgrade is in progress.

Action

Wait for the software upgrade to complete, and then issue the command again.

CMMVC6056E The action failed as the object is too small.

Explanation

The action failed because the object is too small.

Action

Specify a different object, and issue the command again.

CMMVC6057E The action failed as the object is the target of a Flash Copy mapping.

Explanation

The object is the target of a FlashCopy mapping, thus it cannot be deleted.

Action

Specify an object that is not the target of a FlashCopy mapping, or remove the object from the FlashCopy mapping.

CMMVC6058E The action failed as the object is in the recovery HWS.

Explanation

An attempt was made to perform an operation on a node that is in the recovery IO group.

Action

Get the node into one of the other IO Groups and re-issue the command.

CMMVC6059E The action failed as the object is in an invalid mode.

Explanation

The action failed because the object is in the wrong mode.

Action

Check that the object is in the correct mode, and issue the command again.

CMMVC6060E The action failed as the object is being deleted.

Explanation

The action failed because the object is being deleted.

Action

Not applicable.

CMMVC6061E The action failed as the object is being resized.

Explanation

The action failed because the object is being resized.

Action

Check that the object is in the correct mode, and issue the command again.

CMMVC6062E The action failed as the object is being moved between HWS.

Explanation

An attempt was made to perform an action against an object that is currently being moved between IO groups.

Action

Re-issue the command when the move operation has completed.

CMMVC6063E The action failed as there are no more disks in the group.

Explanation

An attempt was made to perform an action against a group that contained no disks.

Action

Either add disks to the group and re-issue the command, or select another group against which to execute the action.

CMMVC6064E The action failed as the object has an invalid name.

Explanation

An attempt was made to create or rename an object using an invalid name.

Action

Use a name that meets the naming standards and re-issue the command.

CMMVC6065E The action failed as the object is not in a group.

Explanation

An attempt was made to perform an action on an object that was not in an appropriate group.

Action

Ensure that the object is a member of an appropriate group and re-issue the command.

CMMVC6066E The action failed as the system is running low on memory.

Explanation

The system is running low on memory.

Action

Not applicable.

CMMVC6067E The action failed as the SSH key was not found.

Explanation

An attempt was made to perform an action using an SSH key that does not exist.

Action

Re-issue the command using a key that does exist.

CMMVC6068E The action failed as there are no free SSH keys.

Explanation

An attempt was made to use an SSH key when there are no free SSH keys.

Action

Upload additional keys and re-issue the command.

CMMVC6069E The action failed as the SSH key is already registered.

Explanation

An attempt was made to register an SSH key that was already registered.

Action

Not applicable.

CMMVC6070E An invalid or duplicated parameter, unaccompanied argument, or incorrect argument sequence has been detected. Ensure that the input is as per the help.

Explanation

The parameters entered for a command were invalid.

Action

Correct the parameters and re-issue the command.

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.

CMMVC6072E Incompatible software.

Explanation

The software version on one or more nodes is incompatible with the new version.

Action

Refer to the compatibility requirements for the software version you are adding. Update the cluster to meet the compatibility requirements, and then perform the upgrade.

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.

CMMVC6077E WARNING - Unfixed errors should be fixed before applying software upgrade. Depending on the nature of the errors, they may cause the upgrade process to fail. It is highly recommended to fix these errors before proceeding. If you cannot fix a particular error, please contact your support center.

Explanation

Unfixed errors should be fixed before applying software upgrade. Depending on the nature of the errors, they may cause the upgrade process to fail. It is highly recommended to fix these errors before proceeding.

Action

If you cannot fix the error, contact your support center.

CMMVC6078E The action failed because the object is in an invalid mode.

Explanation

An attempt was made to perform an action against an object in a mode that did not allow for that action to be performed.

Action

Get the object into a suitable mode and re-issue the command.

CMMVC6098E The copy failed as the specified node is the configuration node.

Explanation

The copy failed because the specified node is the configuration node.

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

Correct the problem and then issue the command again.

CMMVC6104E Action name not run

Explanation

An unexpected error occurred. Contact your support center.

Action

Contact your support center.

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

Explanation

An unexpected error occurred.

Action

Report the details to your support center.

CMMVC6111E Cluster code_level could not be determined 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 support center.

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 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 match forproperty

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 support center.

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.

CMMVC6121E No cluster id or id_alias in backup configuration

Explanation

Neither the cluster id_alias or ID can be extracted from the backup configuration file.

Action

If you cannot determine the cause of the problem, contact your support center.

CMMVC6122E No type with property value present in table

Explanation

An unexpected error occurred.

Action

Contact your support center.

CMMVC6123E No property for type name

Explanation

An unexpected error occurred.

Action

Contact your support center.

CMMVC6124E No type with property value

Explanation

An unexpected error occurred.

Action

Contact your support center.

CMMVC6125E No unique ID for type name

Explanation

An unexpected error occurred.

Action

Contact your support center.

CMMVC6126E No type with unique ID value

Explanation

An unexpected error occurred.

Action

Contact your support center.

CMMVC6127I SSH key identifier for user already defined; will not be restored

Explanation

An identical SSH key for this user is already defined on the cluster. Therefore, the key in the backup file will not be restored.

Action

Specify a different SSH key and issue the command again.

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 support center.

CMMVC6132E An object of a given type has a property with an incorrect value. The operation cannot proceed until the property has the correct value. Take administrative action to change the value and try again.

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 support center.

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 Argument value for -option is not valid

Explanation

An argument has been supplied for an option that is not valid.

Action

Supply an valid argument and retry.

CMMVC6136W No SSH key file file-name

Explanation

A file that contains an SSH key is not present and will not be restored. The backup operation will continue.

Action

No action is required. You may have to manually restore the key.

CMMVC6137W No SSH key file file-name; key not restored

Explanation

An SSH key cannot be restored because the file containing it is not present. The restore operation will continue.

Action

After the restore is complete, locate the file containing the key, and perform one of the following actions: (1) Rename the file so that it has the correct name and issue the command again. (2) Restore the key manually using the svctask addsshkey command.

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 Incorrect XML tag nesting in filename

Explanation

There is a problem with the content of a configuration file. There is a problem parsing the XML in the file, because the XML records are not consistent. The file may be corrupt or has been truncated.

Action

Replace this copy with a good copy and try again. If the problem persists, contact your support center.

CMMVC6140E No default name for type type

Explanation

An unexpected error occurred.

Action

Contact your support center.

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.

CMMVC6142E Existing object-type object-name has a non-default name

Explanation

An object in the target default cluster has a non-default name. This suggests that the cluster was customized. The cluster is therefore not suitable for restoration.

Action

Reset the cluster as per the instructions for restoring the cluster configuration, and try 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 bas 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 support center.

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 support center.

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 support center.

CMMVC6155I SVCCONFIG processing completed successfully

Explanation

Only information and warning messages are issued.

Action

Not applicable.

CMMVC6156W SVCCONFIG processing completed with errors

Explanation

Processing was not successful.

Action

Not applicable.

CMMVC6164E The SVCCONFIG CRON job, which runs overnight on a daily overnight, has failed.

Explanation

The SVCCONFIG CRON job, which runs overnight on a daily overnight, has failed.

Action

Resolve any hardware and configuration problems that you are experiencing on the 2145 cluster. If the problem reoccurs contact your software support center for assistance.

CMMVC6165E Target is not the original configuration node with WWNN of value.

Explanation

A backup configuration can only be restored to the original configuration node.

Action

Recreate the default cluster with the correct configuration node, and issue the command again.

CMMVC6166E A property of an object has changed during svcconfig restore -execute.

Explanation

A property of an object has changed during svcconfig restore -execute. The integrity of the restoration cannot be guaranteed.

Action

Retry the command from svcconfig restore -prepare.

CMMVC6181E The target cluster contains an object that has a counterpart in the configuration to be restored, and has the correct ID.

Explanation

The target cluster contains an object that has a counterpart in the configuration to be restored, and has the correct ID, however the indicated property has an unexpected value.

Action

Check that the correct (matching) backup configuration file (svc.config.backup.xml) is being provided and if it is, use the -force option to ignore the discrepancy. Otherwise, provide the correct file and try again.

CMMVC6182W An object that does not contribute to the fabric of the configuration cannot be restored because its configuration does not permit it to be created.

Explanation

An object that does not contribute to the fabric of the configuration cannot be restored because its configuration does not permit it to be created. For example, a host can only be created if it has at least one port.

Action

N/A

CMMVC6186W An IO group was restored with a different ID value.

Explanation

An IO group was restored with a different ID value. This can occur when the configuration node is different to the one used to create the original cluster. This affects the SCSI Inquiry value for the IO group.

Action

N/A

CMMVC6202E The cluster was not modified because the IP address is not valid.

Explanation

An attempt was made to change the IP address of a cluster to an invalid address.

Action

Correct the address and re-issue the command.

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/config, /dumps/elogs, /dumps/ec or /dumps/pl.

Explanation

An attempt was made to clear a file from, or copy a file to, an invalid directory.

Action

Ensure that the command accesses a valid directory.

CMMVC6204E The action failed as the resulting disk size would be less than, or equal to, zero.

Explanation

An attempt was made to shrink a disk, however the resulting size would have been less than or equal to zero.

Action

Not applicable

CMMVC6206E The software upgrade failed as a file containing the software for the specified MCP version was not found.

Explanation

There are two files required to successfully complete a software upgrade. One file contains the files that make up the base operating system, while the other file contains the 2145 software. This message appears if the OS version is incompatible with the 2145 software.

Action

Upload two compatible files and reissue the command.

CMMVC6207E The action failed because the virtual disk (VDisk) is part of a Metro Mirror mapping.

Explanation

An action was performed against a VDisk that is part of a Metro Mirror mapping.

Action

Remove the VDisk from the Metro Mirror mapping before re-issuing the command.

CMMVC6208E The action failed because the virtual disk (VDisk) is part of a FlashCopy mapping.

Explanation

An action was performed against a VDisk that is part of a FlashCopy mapping.

Action

Remove the VDisk from the FlashCopy mapping before re-issuing the command.

CMMVC6211E The command failed as a migrate to image was in progress.

Explanation

An attempt was made to execute a command against a VDisk that was involved in a migrate to image operation.

Action

Wait for the migration to complete and re-issue the command.

CMMVC6215E The FlashCopy mapping was not created or modified because the consistency group already contains the maximum number of mappings.

Explanation

An attempt was made to create a FlashCopy mapping in, or move a FlashCopy mapping to, a consistency group that has the maximum number of FlashCopy mappings that it can contain.

Action

Create or move the FlashCopy mapping in another consistency group or remove an existing FlashCopy mapping from the desired group and then re-issue the command.

CMMVC6216E The Metro Mirror relationship was not created because the master or auxiliary virtual disk (VDisk) is a member of a Metro Mirror mapping.

Explanation

The Metro Mirror relationship was not created because the master or auxiliary virtual disk (VDisk) is a member of a Metro Mirror mapping.

Action

Select a different VDisk to make up the mapping.

CMMVC6226E The action was not completed because the cluster has reached the maximum number of extents in MDisk Groups.

Explanation

The cluster has reached the maximum number of extents in the MDisk Group, therefore, the action did not complete.

Action

Select a different cluster and issue the command again.

CMMVC6228E The cluster was recovered and the CLI functionality is limited until the cause of the failure is determined and any corrective action taken.

Explanation

The cluster was recovered and the CLI functionality is limited.

Action

Contact your support center.

Accessibility

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

Features

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

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

Navigating by keyboard

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

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

Accessing the publications

You can view the publications for the SAN Volume Controller in Adobe Portable Document Format (PDF) using the Adobe Acrobat Reader. The PDFs are provided at the following Web site:

http://www.ibm.com/storage/support/2145

Related reference

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

Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Trademarks

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

- AIX
- BladeCenter
- Enterprise Storage Server
- FlashCopy
- IBM
- IBM eServer
- IBM TotalStorage
- IBM System Storage
- System p5
- System z9
- System Storage
- TotalStorage
- xSeries

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

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

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

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

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

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

Glossary

This is the glossary for the SAN Volume Controller.

A

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

asymmetric virtualization
A virtualization technique in which the virtualization engine is outside the data path and performs a metadata-style service. The metadata server contains all the mapping and locking tables while the storage devices contain only data. See also *symmetric virtualization*.

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

C

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

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

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

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

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

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

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

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 run from within the SAN Volume Controller application and are documented in the service guide.

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

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

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

E

error code
A value that identifies an error condition.

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

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

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

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

F

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

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

FC See *fibre channel*.

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

FlashCopy mapping
A relationship between two virtual disks.

FlashCopy relationship
See *FlashCopy mapping*.

G

Global Mirror
An asynchronous copy service that enables host data on a particular source virtual disk (VDisk) to be copied to the target VDisk that is designated in the relationship.

H

HBA See *host bus adapter*.

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

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

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

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

I

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

idling

- The status of a pair of virtual disks (VDisks) that have a defined copy relationship for which no copy activity has yet been started.
- In a Global Mirror relationship, that state that indicates that the master virtual disks (VDisks) and auxiliary VDisks are operating in the primary role. Consequently, both VDisks are accessible for write I/O operations.

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

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

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

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

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

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

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

I/O See *input/output.*

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

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

IP See *Internet Protocol*.

L

LBA See *logical block address*.

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

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

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

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

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

LU See *logical unit*.

LUN See *logical unit number*.

M

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

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

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

mapping
See *FlashCopy mapping*.

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

MDisk
See *managed disk*.

migration
See *data migration*.

mirrorset

- IBM definition: See *RAID-1*.
- HP definition: A RAID storageset of two or more physical disks that maintain a complete and independent copy of the data from the virtual disk. This type of storageset has the advantage of being highly reliable and extremely tolerant of device failure. Raid level 1 storagesets are referred to as mirrorsets.

N

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

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

O

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

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

P

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

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

pend To cause to wait for an event.

point-in-time copy
The instantaneous copy that the FlashCopy service makes of the source virtual disk (VDisk). In some contexts, this copy is known as a T_0 *copy*.

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

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

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

Q

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

R

RAID See *redundant array of independent disks*.

RAID 1

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

RAID 5

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

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

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

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

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

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

S

SAN See *storage area network*.

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

SCSI See *Small Computer Systems Interface*.

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

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

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

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).

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

SNMP
See *Simple Network Management Protocol*.

SSH See *Secure Shell*.

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

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

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

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

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

superuser authority
The level of access required to add users.

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

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

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

T

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

U

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

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

V

valid configuration
A configuration that is supported.

VDisk See *virtual disk*.

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

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

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

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

W

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

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.

WWNN
See *worldwide node name*.

WWPN
See *worldwide port name*.

Index

Special characters

A

B

C

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

IBM System Storage SAN Volume Controller
Command-Line Interface User's Guide
Version 4.1.0

Publication No. SC26-7903-00

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

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	☐	☐	☐	☐	☐

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

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	☐	☐	☐	☐	☐
Complete	☐	☐	☐	☐	☐
Easy to find	☐	☐	☐	☐	☐
Easy to understand	☐	☐	☐	☐	☐
Well organized	☐	☐	☐	☐	☐
Applicable to your tasks	☐	☐	☐	☐	☐

Please tell us how we can improve this book:

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

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you. IBM or any other organizations will only use the personal information that you supply to contact you about the issues that you state on this form.

Name

Address

Company or Organization

Phone No.

Readers' Comments — We'd Like to Hear from You
SC26-7903-00

Cut or Fold
Along Line

Fold and Tape | **Please do not staple** | Fold and Tape

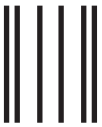

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

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

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Information Development
Department 61C
9032 South Rita Road
Tucson, Arizona 85755-4401

Fold and Tape | **Please do not staple** | Fold and Tape

Cut or Fold
Along Line

Printed in USA

SC26-7903-00

Spine information:

IBM System Storage SAN Volume Controller

Command-Line Interface User's Guide

Version 4.1.0