

IBM TotalStorage Multiple Device Manager



Command-Line Interface User's Guide

Version 1.1

IBM TotalStorage Multiple Device Manager



Command-Line Interface User's Guide

Version 1.1

Note

Before using this information and the product it supports, read the information in "Notices," on page 139.

First Edition (March 2004)

© Copyright International Business Machines Corporation 2004. All rights reserved.

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

Contents

About this guide	v
Command-line interface conventions	v
Accessing the Multiple Device Manager Information Center	x

Chapter 1. Multiple Device Manager overview	1
--	----------

Chapter 2. CLI introduction	3
--	----------

Chapter 3. Performance Manager commands	7
--	----------

percli	8
help	9
chprofile	10
cpprofile	13
cpthresh	14
exit	16
lscollection	17
lsdbpurge	19
lsdev	21
lsfilter	23
lsgauge	25
lshost	27
lslun	29
lsprofile	30
lsthresh	32
mkprofile	35
mkrecom	38
quit	41
rmgauge	42
rmprofile	43
setfilter	44
setessthresh	46
setoutput	49
setsvcthresh	52
showcapacity	55
showdev	56
showdbinfo	57
showgauge	58
showprofile	61
startdbpurge	63
startesscollection	65
startsvccollection	67
stopcollection	69
tracklun	70

Chapter 4. Replication Manager commands	73
--	-----------

repcli	74
help	76
approvecpset	77
chcpset	78
chgrp	80
chtgtpool	83
chsess	85
exit	87
flashsess	88
generatecpset	89
lscpset	90
lsdev	92
lsgrp	94
lspair	96
lspath	98
lsseq	100
lssess	102
lstgtpool	104
lsvol	106
mkgrp	108
mksess	110
mkcpset	112
mktgtpool	114
setoutput	116
rmcpset	119
rmgrp	120
rmsess	121
rmtgtpool	122
showcpset	123
showdev	125
showgrp	126
showsess	128
showtgtpool	130
startsess	131
stopflashsess	132
stopsess	133
suspendsess	134

Chapter 5. How to send your comments	137
---	------------

Appendix. Notices	139
------------------------------------	------------

Legal Information	141
------------------------------------	------------

Trademarks	143
-----------------------------	------------

Index	145
------------------------	------------

About this guide

This section briefly describes the content and the intended audience of this guide.

Content:

This guide tells you how to use the Multiple Device Manager (MDM) command-line interface. It also includes a list and description of the Performance Manager commands and the Replication Manager commands.

Who should use this guide:

This guide is intended for administrators, application programmers, and others who are using the MDM command-line interface (CLI) to perform Performance Manager and Replication Manager functions.

Command-line interface conventions

This topic provides information about command conventions and modes, command format requirements, and other command information for the Multiple Device Manager (MDM) command-line interface (CLI).

Program startup options:

Enter the program name, **perfcli** or **repcli** to start the interactive mode for the corresponding CLI program. The following flags can be used with each CLI program name:

-help or -h or -?

Displays a help screen on how to use the CLI program. If one of the help flags is used with anything other than the command name, the program ignores all other options and simply displays the help screen.

-overview

Provides overview information for using the CLI application.

-script

Initiates script mode so that multiple program commands can be issued consecutively through a saved file.

-ver Displays the version of the CLI program and license information.

Command modes:

You can work with the administrative CLI in one of three modes:

Single-shot mode

If you only want to run a single command, specify the CLI program and the command to execute at the shell prompt, for example:

```
shell> repcli startsess -session_name session1 shell>
```

Interactive mode

If you want to run several commands, start the CLI program and the commands to execute at the shell prompt, for example:

```
shell> repcli chsess session1 -desc "new description" -approve man>
startsess -session_name session1 shell>
```

Script mode

If you want to run a set of commands that you defined in a file, start the CLI program and specify a file that contains the commands, for example:

```
shell> repcli -script ~/bin/containersetup
shell>
```

Note: The MDM CLIs support and interpret comments specified with a pound symbol (#) in the first column.

Command emphasis:

The following typefaces are used to show command emphasis:

boldface	Text in boldface represents command names.
<i>italics</i>	Text in <i>italics</i> 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.

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.

Modifying command input:

You can use these commands to modify command input:

- Using “-f” forces destructive action, such as make a volume even if the LUN already has a label; suppresses confirmation and error messages.
- Using “-quiet” suppresses confirmation prompts and messages (answers *yes* to all confirmation prompts).

Specifying multiple Values and ranges:

When input values can be repeated, delimit the values with a comma with no space in the input string. For example, -vol 3,5,8,9 . When input values include a range of numbers, delimit the range with a hyphen. For example, -vol 3, 5, 7-9.

When input values require multiple value types in one string, delimit value types with a period. For example, if a volume requires a device number and a vol number, use -vol FCA86.3,FCA78.5,FCA96.8. When input values are of different

types but specified in the same flag, use a colon. For example, to specify a minimum and maximum value in the same flag, use `-size min:max`.

Specifying volumes and locations:

Valid volume values include the following:

ESS single volume

ESS:2105.65312:VOL:202F
(ESS:ELEMENTTYPE.DEVICEID:VOL:LSSVOLNUM)

ESS multiple volumes in a range

You can use a hyphen to specify a range of volumes:
ESS:2105.65312:VOL:202A-2F
(ESS:ELEMENTTYPE.DEVICEID:VOL:LSSVOLNUM-VOLNUM)

ESS multiple volumes

ESS:2105.65312:VOL:202A,2C,2F
(ESS:ELEMENTTYPE.DEVICEID:VOL:LSSVOLNUM,VOLNUM,VOLNUM)

Interactive mode Commands:

Each CLI program recognizes these built-in commands in interactive mode:

setoutput

This command specifies various command output format options. The syntax for `setoutput` is as follows: `setoutput [-p on | off] [-r #] [-fmt xml | stanza | default | delim char] [-hdr on | off] [-v on | off]`. See viii for descriptions of each of these `setoutput` options.

All settings specified with `setoutput` remain in effect for the duration of the interactive command session unless reset either with a command option or with `setoutput`. With no options, `setoutput` displays the current settings in the default output format. `setoutput` settings do not apply to help pages. Help pages are shown in text output only. Deleted: Standard

help Displays a list of commands available from the CLI program.

exit Exits the CLI program.

quit Exits the CLI program.

Object naming guidelines:

The requirements for valid user-defined object names are:

- User-defined object names can be 250 characters or fewer, unless otherwise noted.
- Legal characters are A-Z ,a-z, 0-9, dash (-), underscore (_), period (.), and colon (:)
- Object names must contain at least one alphanumeric character. The name must also start with an alphanumeric character.
- Object names cannot begin with a dash (-) and cannot consist of only a dash (that is, objects cannot be named -).
- Object names cannot begin with an underscore (_) and cannot consist of only an underscore (that is, objects cannot be named _).
- Objects cannot begin with a period (.) or colon (:).
- Object names cannot contain any blank space.

- User-defined objects should not be translated or otherwise modified from what the user entered (they should remain case-sensitive).

User-defined descriptions:

The requirements for valid user-defined descriptions are:

- Description strings must be enclosed in (matching) single or double quotation marks, if they contain spaces.
- If a description that is already enclosed in matching quotation marks includes an asterisk, the asterisk must be escaped. For example, `-desc "This is the * pool"`.
- User-defined descriptions can be 256 characters or fewer. They cannot contain any leading blank space.
- User-defined descriptions should not be translated or otherwise modified from what the user entered (they should remain case-sensitive)

Parameter entry:

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. If you are performing an action against a specific object, the object ID or name must be the last argument in the line.

Standard command options:

The standard format parameters set the output format of the listing (`ls`) commands in the administrative CLI. These parameters can be used either in the listing command syntax or in the **setoutput** command. The format settings remain in effect for the duration of the session or until you reset the parameters either by specifying these parameters in a listing command (commands that start with `ls`) or using the **setoutput** command.

- **-p** specifies whether to display one page of text at a time or all text at once.
- **off** displays all text at one time. This is the default value when the **repli** command is run in single-shot mode.
- **on** displays one page of text at time. Pressing any key displays the next page. This is the default value when the command is run in interactive mode.
- **-r number** specifies the number of rows per page to display when the **-p** parameter is on. The default is 24 rows. You can specify a value of 1 - 100.
- **-fmt** specifies the format of the output. You can specify one of the following values:
 - **default** specifies to display output in a tabular format using spaces as the delimiter between the columns. This is the default value.
 - **delim character** specifies to display output in a tabular format using the specified character to separate the columns. If you use a shell metacharacter as the delimiting character, enclose the character in single or double quotation marks. A blank space is not a valid character.
 - **xml** specifies to display output using XML format.
 - **-hdr** specifies whether to display the table header. Use **on** to display the table header. This is the default value. Use **off** to hide the table header.
- **-v** specifies whether to enable verbose mode. Use **off** to disable verbose mode. This is the default value. Use **on** to enables verbose mode.

Standard listing options:

These standard listing options modify command output in any CLI mode:

- ***Isobject-s*** lists only the objects without other columns of information. For example, `lsess -s` lists only the Name header and the session names.
- ***Isobject-l*** lists all the objects with all defined columns, including description.

Confirmation prompts:

When commands might lead to an irrecoverable operation, loss of data, memory drain, long-running task, or impact to concurrent operations, you receive an interactive confirmation prompt that asks if you are sure that you want to continue with the specific action, such as:

Are you sure you want to xxx? Y/N

All confirmation prompts accept the following input: YES, yes, Y, y or NO, no, N, n

Syntax diagram conventions:

A syntax diagram uses symbols to represent the elements of a command and to specify the rules for using these elements. A key word represents the name of a command, flag, parameter, or argument. Required key words indicate the parameters or arguments that must be specified for the command.

To read syntax diagrams, follow the path of the line. Read the diagrams from left-to-right, top-to-bottom, following the main path line.

- Required keywords appear on the main path line. Mutually exclusive required keywords are stacked vertically. Optional key words 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.
- The 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. The -->< symbol indicates the end of the syntax diagram.
- A dash (-) indicates that you want to supply parameters from the stdin file rather than entering parameters.
- An arrow returning to the left means you can repeat the item. A character or space within the arrow means you must separate repeated items with that character or space.
- A stack of items followed by an arrow returning to the left means that you can select more than one item or, in some cases, repeat a single item.
- Commands that contain lengthy groups or a section that is used more than once in a command are shown as separate fragments following the main diagram. The fragment name appears between vertical bars in the diagram. The expanded fragment is also shown between vertical bars after the heading with the same fragment name.
- Italicized, lowercase elements denote variables.

Syntax diagrams use position to indicate required, optional, and default values for keywords, variables, and operands:

- If an element is shown on the line, the element is required. If an element is shown below the line, the element is optional. If an element is shown above the line, the element is the default.
- Required keywords appear on the main path line. Mutually exclusive required keywords are stacked vertically. Optional key words indicate the parameters or arguments that can be selected for the command. Optional keywords appear below the main path line. Mutually exclusive optional keywords are stacked vertically.
- If an operand has a default value, the operand is shown both above and below the main line. A value below the main line indicates that the operand must be specified. Either the default value or another of the values that are shown must also be specified. If an operand is not specified, the default value above the main line is used.
- When one or more items are shown below the main line, all of the items are optional.
- When a keyword, variable, or operand is shown on the main line, it is required.
- When two or more items are in a stack and one of them is on the main line, one item must be specified.

User assistance for commands:

You can get user assistance for the commands that are available in the administrative CLI by using the help command or the command-help parameters. The help command enables you to display a list of available commands.

You can use the command-help parameters supported by each command to display a detailed description of the specified command. The following help parameters are supported:

```
-?
-h
-help
```

Note: When you use a help parameter, all other parameters are ignored.

Accessing the Multiple Device Manager Information Center

This topic explains how to access the Multiple Device Manager (MDM) Information Center.

Prerequisites:

MDM must be installed and configured.

Steps:

Perform the following steps to access the MDM Information Center:

1. In the IBM® Director Task pane, click **Multiple Device Manager**. The MDM tasks are displayed.
2. Invoke any MDM task. An MDM user interface panel opens.
3. Click **Help** on the interface panel. The Information Center opens.

Chapter 1. Multiple Device Manager overview

Multiple Device Manager (MDM) is a storage management software product that provides a single point of control across multiple supported storage devices on a storage area network (SAN). MDM services are provided by components that manage data storage. These components include Device Manager, Replication Manager, and Performance Manager.

MDM provides basic SAN management functions like logical unit (LUN) discovery, allocation, and zoning as well as interoperability with other managers. It also provides basic device configuration functions like LUN creation and allocation for single and multiple devices. MDM provides both an interactive, graphical interface and a full-feature command-line interface (CLI). The CLI provides support for scripting and automation. MDM contains an IBM WebSphere® Application Server, and uses an IBM DB2® database for persistent data storage.

Storage devices that MDM supports:

The following are the storage devices that are supported by MDM:

- IBM TotalStorage® Enterprise Storage Server® (ESS)
- IBM TotalStorage SAN Volume Controller
- IBM FAStT Storage Server

Device Manager:

Device Manager has a dual role as a component of MDM *and* as part of the infrastructure of IBM Director. Device Manager provides the following services:

- Uses the Service Location Protocol (SLP) on the IBM Director to discover storage devices
- Creates *managed objects* to represent these discovered devices
The objects are displayed as individual icons in the Group Contents pane of the IBM Director Console.
- Collects asset, configuration, and availability data from the supported storage devices

You can perform the following tasks with Device Manager:

- View essential information about storage devices
- Examine the relationships between storage devices
- Change storage device configurations

Replication Manager:

Replication Manager provides a graphical interface that you can use to define copy tasks between source and target logical volumes, which are also referred to as managed disks (MDisks). With Replication Manager, you can also create, maintain, and use volume (MDisk) groups. You can use the interface panels to perform the following tasks:

- Define a copy session
- Identify a target
- Create a session wizard

- Select a source group
- Select a copy type
- Select a target pool
- Save a session

The interface panels enable you to perform paths tasks and scheduled replication tasks.

Performance Manager:

Performance Manager provides a graphical interface that you can use to define a task for collecting performance data from a supported storage device.

Performance data-collection tasks are saved as noninteractive tasks and displayed in the IBM Director Task pane under the **Performance Data Collection** task. You can view, change, and delete a performance data-collection task.

Once a performance data-collection task has been saved, you can create a job using the IBM Director **Scheduler** task. You can view, change, and delete jobs containing performance data-collection tasks from the **Scheduler** task. Performance data-collection tasks must be defined and contained in a scheduled job before you can define performance gauges and thresholds.

Chapter 2. CLI introduction

This chapter introduces the Multiple Device Manager (MDM) command-line interface (CLI). It includes an overview of the following topics:

- CLI requirements
- How to use the MDM commands
- CLI limitations
- Performance Manager commands
- Replication Manager commands

CLI requirements:

You must meet the following requirements before you can use the Multiple Device Manager command-line interface:

- Java 1.3.1 must be installed on the machine that has the CLI.
- You must have installed the following components:
 - DB2 Universal Database
 - IBM Director
 - Multiple Device Manager
 - WebSphere Application Server
 - Performance Manager
 - Required only if you are using the Performance Manager perfcli commands.
 - Replication Manager
 - Required only if you are using the Replication Manager repli commands.

See the *IBM TotalStorage Multiple Device Manager: Installation and Configuration Guide* for instructions on installing and configuring MDM. This document is included on your MDM CD.

Note: These components do not have to be installed on the same machine as the CLI.

- MDM must be installed and running, and storage devices must have been discovered.

How to use the MDM commands:

The MDM CLI is comprised of two independent CLIs: perfcli for Performance Manager commands and repli for Replication Manager commands. Each of these CLIs interacts with the MDM Device Manager, the common CLI infrastructure, and the MDM Performance Manager and MDM Replication Manager servers, respectively.

For details about the following and other CLI conventions, see the CLI conventions in the preface:

- Command modes
- Command input
- Standard format parameters

- Object naming guidelines
- Syntax conventions and syntax diagrams

CLI limitations:

The following are some CLI limitations.

- Performance Manager limitations:
 - Data collection tasks such as `startesscollection` and `startsvccollection` that are performed from the CLI are not visible to the MDM graphical user interface (GUI).
 - You can display and remove gauges using the CLI, but first you must create the gauge using the MDM GUI.
 - Thresholds and filters for newly discovered systems are not shown in the CLI until they are viewed in the GUI.
- Replication Manager limitations:
 - The only filter type supported by the `mkssess` command is location, which matches the value set in the `mkgtgtpool` command.
 - When the parameter `-approve` is set (manual or automatic option) in the `mkssess` command, it only affects the `mkcpset` command when the user creates a copy set for the session that was created. The `generatecpset` command automatically generates copy sets in a session. It is not affected, if for instance, manual approval is selected by the `mkcpset` command for a session.

Performance Manager commands:

The Performance Manager component of MDM is responsible for managing and monitoring the performance of the storage devices that MDM supports. The Performance Manager commands are listed in the Performance Manager commands chapter.

All Performance Manager commands must be invoked by first typing `perfcli` followed by the command. The Performance Manager (`perfcli`) commands provide the following functions:

Collect data from devices

Performance Manager collects performance data from storage devices. This retrievable data is stored in a database.

Configure performance thresholds

Performance Manager enables you to set performance thresholds for certain predefined metrics for each device type. Performance Manager detects when these thresholds are exceeded and notifies you of the event.

Extract performance data from the database

You can use the gauge applications to extract performance data from the Performance Manager database. Gauges are comprised of graphical and tabular presentations of performance data.

Purge performance data from the database

Performance Manager enables you to view the size of the performance data in the database and to create a database purge task to free up space in the database or to delete obsolete information. You can choose to purge performance data based on the age of the data, the type of data, and the storage devices associated with the data.

Provide advise about volume creation

The volume performance advisor uses historical performance information to advise you about where volumes should be created. The advisor balances the storage allocation activity across all available resources based on the existing performance of those resources and on the performance needs of the new storage to be allocated.

Provide workload profile recommendations

Performance Manager enables you to create a workload profile that is based on historical performance data or on an existing workload profile. Performance Manager uses these profiles to create a performance recommendation for volume or virtual disk (VDisk) allocation on an IBM storage server.

Replication Manager commands:

Replication Manager is responsible for managing and monitoring the copy-service related functions for storage devices that MDM supports. These commands enable you to replicate data; list detailed information; create, modify, and delete groups and pools; and start and stop replication sessions.

All Replication Manager commands must be invoked by first typing `repcli` followed by the command. The Replication Manager commands are listed in the Replication Manager commands chapter.

Related topics:

- “Command-line interface conventions” on page v
- Chapter 3, “Performance Manager commands,” on page 7
- Chapter 4, “Replication Manager commands,” on page 73

Chapter 3. Performance Manager commands

This section lists the different Performance Manager commands that you can issue. You can use these commands to collect data, list detailed information, start and stop performance tasks, and monitor performance. This section provides an overview of the types of tasks you can perform by issuing the Performance Manager commands.

Performance collection data: Performance collection data encapsulates the information that the Performance Manager needs in order to collect performance data. Each performance task has a unique name. It is associated with a set of devices of like types, has a sampling frequency (a frequency in which performance data can be sampled), and a duration of which the performance data can be sampled on each device associated with the task.

A collection task will in general contain the following information for each device:

- Serial number or unique ID to identify the device
- IP address of the device
- User ID or password needed to access the device (might only be required for some device types, not all.)
- For the Enterprise Storage Server (ESS), the port information

Testing for threshold exceeded conditions: Performance Manager enables you to specify critical and warning threshold values for a subset of performance metrics. This information is saved in the Performance Manager database. Performance Manager provides default threshold values for some of those metrics. For a device and metric combination, you can use Performance Manager panels to enable (default) or disable threshold checking. For each of the metrics pertaining to a device, you can also specify if the Performance Manager should send an alert when a threshold exceeded condition occurs.

Performance-type gauges: The performance-type gauges present sample-level performance data. The frequency at which the performance data is sampled on a device depends on the sampling frequency that you specify during the definition of the performance collection task. The sampling frequency, in general, has a maximum and minimum value depending on the device type. There are two variations of performance gauges:

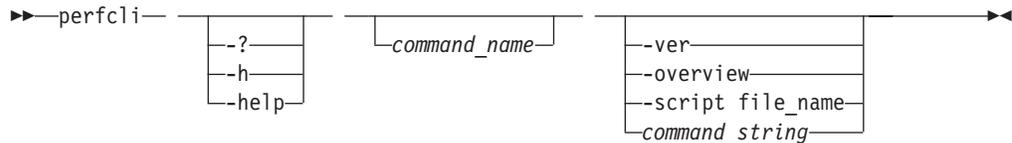
- A static display that presents historical data over time
- A refreshable display that presents near real time data from a device that is currently collecting performance data.

Displaying gauges: Each gauge that you define with the GUI is identified by a unique gauge name that you can see with the **lsgauge** command. You can then delete a named gauge with **rmgauge**, or see its attributes with **showgauge**.

perfcli

You can use this utility to start the administrative command-line interface (CLI), or perfcli session, to run commands in interactive mode. This utility also runs a single command or runs a set of commands from a script without starting a perfcli session.

Syntax



Parameters

-? | -h | -help *command_name*

Displays help for the specified command (for example, `-h setthresh` displays help for the `setthresh` command). If a command name is not specified, this parameter displays a list of available commands in the Administrative CLI.

Note: Although the `perfcli` utility is listed as a possible command in the help facility, you receive an error if you attempt to run the `perfcli` utility from within `perfcli` (for example, `perfcli> perfcli help`). To get help for the `perfcli` utility, enter `-help perfcli`.

-ver

Displays the current version and licensing information for this utility.

-overview

Displays the overview information about the `perfcli` utility, including command modes, standard command and listing parameters, syntax diagram conventions, and user assistance.

-script *file_name*

Runs the set of command strings in the specified file outside of a `perfcli` session. If you specify this parameter, you must specify a file name.

The format options specified using the `setoutput` command apply to all commands in the script. Output from successful commands routes to `stdout`.

Output from unsuccessful commands route to `stderr`. If an error occurs while one of the commands in the script is running, the script will exit at the point of failure and return to the system prompt.

command_string

Runs the specified command string outside of a `perfcli` session.

Description

You can use the `perfcli` command with no parameters to start the administrative CLI. When you are in the administrative CLI, the `perfcli>` prompt is displayed. If you run this utility with any of the valid parameters, the administrative CLI is not started in interactive mode.

Related topics

help

You can use the **help** command to display a list of the Performance Manager commands or to display a list of commands with syntax.

Syntax

```
▶▶—perfcli— —help—▶▶  
┌──s──┐  
└─l─┘
```

Parameters

- s An optional parameter that displays the list of Performance Manager commands with brief descriptions.
- l An optional parameter that displays a list of the Performance Manager commands with descriptions and syntax.

Description

You can use the **help** command to display a list of the Performance Manager commands or to display a list of commands with syntax.

Examples

An invocation example:

```
perfcli help -s
```

The resulting output:

```
chprofile    Use the chprofile command to modify a workload profile. Use the  
             profile later with the mkrecom command to create a performance  
             recommendation for assigning ESS volumes.  
cprofile    Use the cprofile to copy the settings of a preexisting profile  
            to a new profile.  
...  
...
```

Related topics

chprofile

Use the **chprofile** command to modify a workload profile. Use the profile later with the **mkrecom** command to create a performance recommendation for assigning ESS volumes.

- You must use this command with the **mkrecom** command in order to generate a performance recommendation.
- You cannot save a profile with the same name as a preexisting profile.

Syntax

```
►► perfcli — chprofile — [—desc—"new description"] —————►
► [—io—new_rate] [—trans—new_rsize] [—readcache—new_rpercent] —————►
► [—writedestage—new_rpercent] [—read—new_rpercent] —————►
► [—write—new_rpercent] [—seqread—new_rpercent] —————►
► [—seqwrite—new_rpercent] —starttime—starttime— —endtime—endtime—►
► —startdate—startdate— —enddate—enddate— —original_profile_name————►
```

Parameters

-desc *"new description"*

An optional parameter that changes the user-defined description of the profile to a new description.

-io *new_rate*

An optional parameter that specifies the average I/Os per second per GB over all LUNs.

-trans *new_size*

An optional parameter that specifies the average transfer size in KB.

-readcache *new_percent*

An optional parameter that specifies the percentage of random read cache hits to the total number of random read operations. Valid ranges are 1 - 100.

-writedestage *new_percent*

Specifies the percentage of random write destage operations to the total number of random write destage operations. Valid ranges are 1 - 100.

-read *new_percent*

An optional parameter that specifies the percentage of random reads to the total number of I/O operations. Valid ranges are 1 - 100. Values specified with **-read**, **-write**, **-seqread**, and **-seqwrite** combined cannot exceed 100.

-write *new_percent*

An optional parameter that specifies the percentage of random writes to the total number of I/O operations. Valid ranges are 1 - 100. Values specified with **-read**, **-write**, **-seqread**, and **-seqwrite** combined cannot exceed 100.

- seqread** *new_percent*
An optional parameter that specifies the percentage of sequential reads to the total number of I/O operations. Valid ranges are 1 - 100. Values specified with -read, -write, -seqread, and -seqwrite combined cannot exceed 100.
- seqwrite** *new_percent*
An optional parameter that specifies the percentage of sequential writes to the total number of I/O operations. Valid ranges are 1 - 100. Values specified with -read, -write, -seqread, and -seqwrite combined cannot exceed 100.
- starttime** *starttime*
An optional parameter that specifies the time, in HH:MM format, of the start peak duration..
- endtime** *endtime*
An optional parameter that specifies the time, in HH:MM format, of the end peak duration.
- startdate** *startdate*
An optional parameter that specifies the date, in yyyyymmdd format, when the profile should start.
- enddate** *enddate*
An optional parameter that specifies the date, in yyyyymmdd format, when the profile should end.
- original_profile_name*
An optional parameter that specifies the profile to be changed.

Description

You can use the **chprofile** command to modify a workload profile that can later be used with the **mkrecom** command to create a performance recommendation for assigning ESS volumes.

Return values

Profile *profile_name* successfully modified.

Possible failures

- Profile *profile_name* already exists.
- Profile *profile_name* does not exist.
- The total of the values specified for -read, -write, -seqread, and -seqwrite cannot exceed 100%.
- The predefined profile {0} cannot be modified.

Examples

An invocation example:

```
perfcli> chprofile -read 20 -write 25 -seqread 25 -seqwrite 30 -startdate 20040402
-starttime 08:30 -enddate 20040403 -endtime 13:00 -timemode device myprofile
```

The resulting output:

```
Profile myprofile successfully modified.
```

Related topics

- “mkprofile” on page 35
- “tracklun” on page 70

- “rmprofile” on page 43
- “cprofile” on page 13
- “lsprofile” on page 30
- “showprofile” on page 61
- “mkrecom” on page 38

cpprofile

Use the **cpprofile** to copy the settings of a preexisting profile to a new profile.

- You cannot save a profile with the same name as a preexisting profile.
- You can use the **chprofile** command to change a single attribute in a profile.

Syntax

```
▶▶—perfcli— —cpprofile— —-src—source_profile_name— —new_profile_name—▶▶
```

Parameters

-src *source_profile_name*

Specifies the profile that is being copied.

new_profile_name

Specifies the profile that will adopt the new settings. The name must be unique.

Description

You can use the **cpprofile** command to copy the settings of a preexisting profile to a new profile.

Return values

Profile *source_profile_name* successfully copied.

Possible failures

- Profile *profile_name* already exists.
- Profile *profile_name* does not exist.

Examples

An invocation example:

```
perfcli cpprofile -src profile1 newprofile
```

The resulting output:

```
Profile profile1 successfully copied to newprofile.
```

Related topics

- “mkprofile” on page 35
- “chprofile” on page 10
- “rmprofile” on page 43
- “lsprofile” on page 30
- “showprofile” on page 61
- “mkrecom” on page 38

cpthresh

You can use the **cpthresh** command to copy threshold properties from one device to other devices that are of the same device type.

Syntax

```
▶▶ perfcli —cpthresh— --devtype ess  
svc —srcdev—device_id —————▶▶  
▶▶ —tgtdev—device_id [...]————▶▶
```

Parameters

-devtype *ess* | *svc*

A parameter that displays a choice of device type. The choice can be *ess* or *svc*.

-srcdev *device_id*

Specifies the source device from which the threshold settings are copied. The device ID is the nickname or the model-serial-manufacture of the device.

-tgtdev *device_id* [...]

Specifies the target devices that will adopt the source device's threshold settings. The device ID is the nickname or the model-serial-manufacture of the device. Separate multiple device names by commas. All devices must be of the same device type.

Description

You can use the **cpthresh** command to copy threshold properties from one device to one or more other devices that are of the same device type.

Return values

Threshold settings for device *device_id* successfully copied to device *device_id*.

Possible failures

- Device *device_id* does not exist.
- Source device type does not match target device type.
- Device type and threshold type do not match.

Examples

An invocation example:

```
perfcli cpthresh -srcdev 2105.234345 -tgtdev 2105.123234
```

The resulting output:

```
Threshold settings for device 2105.234345 successfully  
copied to device 2105.123234.
```

Related topics

- “setessthresh” on page 46
- “setsvcthresh” on page 52
- “setfilter” on page 44
- “lsfilter” on page 23

- “lsthresh” on page 32

exit

Use the exit command to exit the perfcli interface.

Syntax

```
▶▶—perfcli— —exit—▶▶
```

Description

Use the exit command to exit a session.

Examples

An invocation example:

```
perfcli exit
```

The resulting output:

```
Are you sure you want to exit? y/n  
Y
```

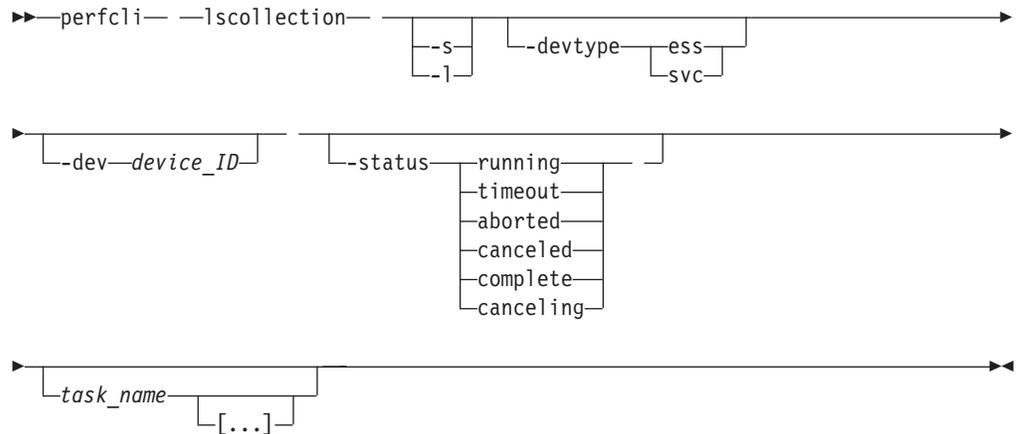
Related topics

- “repli” on page 74

lscollection

You can use the **lscollection** command to display a list of collections and status information for each collection in the list.

Syntax



Parameters

- s An optional parameter that displays the device name only.
- l An optional parameter that displays the default output and description
- devtype **ess** | **svc**
An optional parameter that displays a choice of device type. The choice can be **ess** or **svc**.
- dev *device_ID*
An optional parameter that displays only the nickname or model, serial number, and manufacturer of the specified *device_ID*.
- status **running** | **timeout** | **aborted** | **canceled** | **complete**
An optional parameter that sorts the collection status of a collection. The status can be **running**, **timeout**, **aborted**, **canceled**, and **complete**.

task_name[...]

Specifies the name of the collection task that is to be listed. Separate multiple collection tasks with a comma between each name. If no names are specified, all collection tasks are listed.

Description

You can use the **lscollection** command to display a list of collections and status information for each collection in the list.

Return values

- No collection matched the criteria specified.

- For each collection, the following information is listed:

Column Label	Details	Description
Name	Collection Name	Name that user gave to the performance data collection task at creation
Device Type	Device Type	Device type that the collection task is running on (ESS or SVC)
Status	Status of the collection	Status of collection task, sorted for running, timeout, aborted, canceled, and complete
Length (Hours)	Length of time the collection task ran	Length of time the collection task has been running in hours
Frequency (Minutes)	Frequency	Frequency in minutes at which user set the collection task

Possible failures

- No collection matched the criteria specified.

Examples

An invocation example: The following are examples of the `lscollection` command for a specific device type:

```
perfcli> lscollection -devtype ess
```

The resulting output:

```

Name                                     Device Type Status   Length (hrs) Frequency (min)
-----
ess_26884_5                             ESS           Completed 1             5
ess_28019_10                             ESS           Completed 1             10
ess28019                                 ESS           Completed 1             5
mycollection                             ESS           Aborted   10            10
myveryowncollectionisthisonedontyouthinkwowthisisalongdescript ESS           Aborted   20            20
mycollection2                             ESS           Canceled  10            10
mycollection3                             ESS           Canceled  10            10

```

```
perfcli> lscollection -devtype svc
```

The resulting output:

```

Name           Device Type Status   Length (hrs) Frequency (min)
-----
duc_blackn1   SVC           Aborted 1             15
duc_red1      SVC           Completed 1             15

```

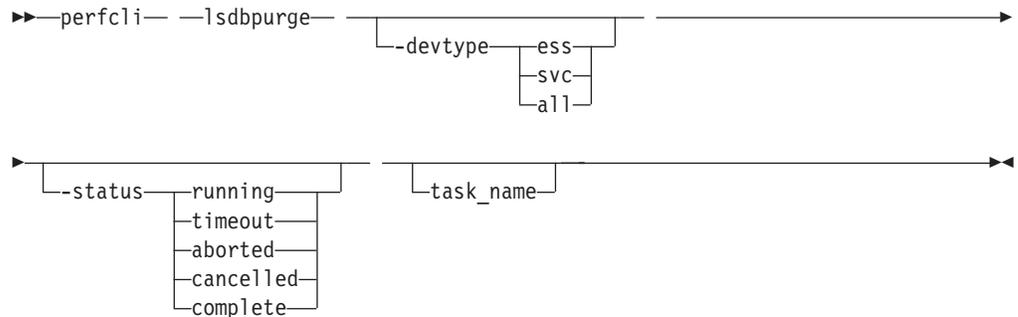
Related topics

- “startesscollection” on page 65
- “startsvccollection” on page 67
- “stopcollection” on page 69

lsdbpurge

Use the **lsdbpurge** command to display the status of database purge tasks running in MDM.

Syntax



Parameters

-devtype *ess* | *svc* | *all*

An optional parameter that displays list of purge tasks by device type. Specify **ess** for Enterprise Storage Server (ESS) or **svc** for SAN Volume Controller. Use **all** to list purge tasks for all device types.

-status *running* | *timeout* | *aborted* | *cancelled* | *complete*

An optional parameter that displays purge tasks for only the purge status type specified.

task_name

An optional parameter that specifies the user-defined task name for the database purge task to be displayed.

Description

Use the **lsdbpurge** command to display the status of database purge tasks running in MDM.

Return values

For each purge task, the following information is listed:

Column Label	Details
Name	The user-defined name of the purge task.
Device type	SVC or ESS, or both for <i>all</i> option.
Status	Purge task status: <i>running</i> , <i>timeout</i> , <i>aborted</i> , <i>canceled</i> , or <i>complete</i> .

Possible failures

- No matches for status type *status type* specified.
- No purge tasks for device type *devtype*.

Examples

An invocation example:

```
perftool lsdbpurge
```

The resulting output:

Name	Device Type	Status
svcTask	SVC	Completed
svcTask2	SVC	Completed
essTask	ESS	Running

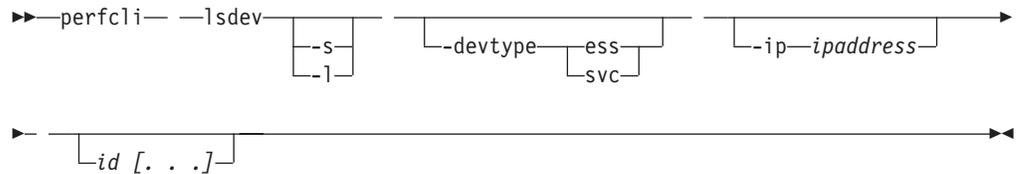
Related topics

- “startdbpurge” on page 63

lsdev

You can use the **lsdev** command to display a summary of storage devices that are supported by MDM.

Syntax



Parameters

- s** An optional parameter that displays only the device ID.
- l** An optional parameter that displays the device ID and the device type.
- devtype *ess* | *svc***
An optional parameter that displays the information for the device type that is specified. Specify **ess** or **svc**.
- ip *ipaddress***
An optional parameter that specifies the device IP address.
- id* [...]**
An optional parameter that displays only the information for the device IDs specified. The Device ID is the nickname or the model-serial-manufacture of the device.

Description

You can use the **lsdev** command to display a summary of storage devices that are supported by MDM.

Return values

The following information is listed for each device:

Column Label	Details
Device ID	The nickname or model-serial-manufacturer of the device.
IP	Device IP address.
Device Type	Device Type: ESS or SVC.

Possible failures

- Device *device_id* not found.
- No device matched the criteria specified.

Examples

An invocation example:
perfcli lsdev -devtype ess

The resulting output:

Device ID	Device IP	Device Type
2105-123123-IBM	9.11.222.10	ESS
2105-123124-IBM	9.11.223.11	ESS
2105-123125-IBM	9.11.224.12	ESS

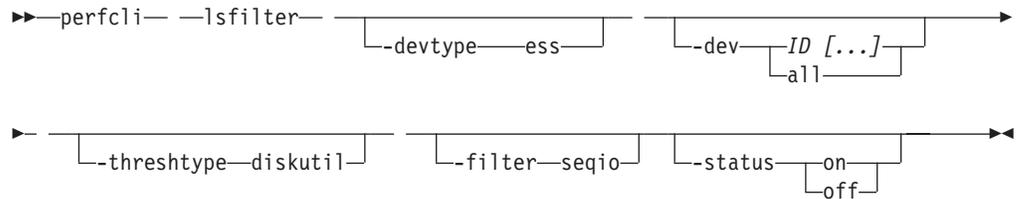
Related topics

- “showdev” on page 56

lsfilter

Use the **lsfilter** command to display the threshold filter settings for all devices specified.

Syntax



Parameters

-devtype *ess* | *svc*

An optional parameter that specifies the device type:

- Enterprise Storage Server (**ess**)
- SAN Volume Controller (**svc**)

-dev *ID [...]* | **all**

An optional parameter that displays only the threshold filter settings for the device IDs that are specified. The Device ID is the nickname or the model-serial-manufacture of the device. Separate multiple device names with a comma in between each device with no white space before or after the comma. All devices must be of the same device type. The value **all** applies the threshold to all devices of the device type specified with the **-devtype** flag. If no device is specified, this parameter displays the threshold filter settings for all devices of the device type specified with the **-devtype** flag.

-threshtype *diskutil*

An optional parameter that filters the list to display only filters for the threshold type specified. The **diskutil** keyword specifies the disk utilization value (the percent of time busy), for each ESS array.

-filter *seqio*

An optional parameter that displays only the filter type that is specified. The **seqio** parameter specifies to make an exception for a sequential I/O percentage filter and its value.

-status *on* | *off*

An optional parameter that specifies whether the specified filter should be enabled or disabled for checking. The sequential I/O filter default is on and enabled for checking.

Description

You can use the **lsfilter** command to list all performance filters and their information for all devices or for a specific device or device type. If the device ID (a unique ID or serial number of a device) and the device type are specified, the information that is returned is a list of all filters for that specific device ID. If the device ID is null but the device type has a value such as ESS or SVC, the information that is returned is a list of filters for all devices of that device type. If

both the device ID and the device type are null, the information that is returned is a list of all filters for all devices and device types.

Return values

Device ID

The device ID is the nickname or manufacturer, model, and serial number of the device.

Device type

The type of the device (ESS or SAN Volume Controller). If null, (device ID should be null too), then it returns all the filters for all the devices of all device types.

Filter name

Displays the name of the filter, for example, *percent of sequential I/Os*.

Threshold name

Displays threshold name, for example, *disk utilization*.

Value Displays the value of the filter.

Status Displays **enabled** or **disabled**, depending on the status of the filter.

The *Percent of sequential I/Os* filter is enabled by default.

Possible failures

- Device *device* does not exist.
- Device type and threshold type do not match.

Examples

An invocation example:

```
perfcli lsfilter -devtype ess -threshtype diskutil -filter seqio
```

The resulting output:

Device ID	Device Type	Filter Name	Threshtype	value	status
ESS2105-123124-IBM	ESS	Seqio	diskutil	80	Enabled

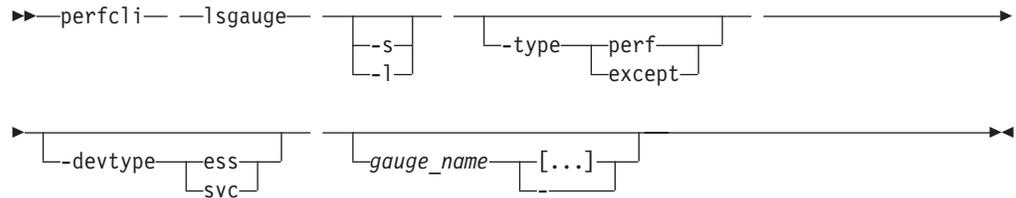
Related topics

- “setessthresh” on page 46
- “setsvcthresh” on page 52
- “setfilter” on page 44
- “cpthresh” on page 14
- “lsthresh” on page 32

lsgauge

You can use the **lsgauge** command to display a list of existing gauge names, gauge type, device type, device ID, date modified, and optionally, description information.

Syntax



Parameters

- s An optional parameter that displays only the gauge name.
- l An optional parameter that displays the name and description of the gauge, and other vital output.
- type **perf** | **except**
An optional parameter that lists only the gauges with the gauge type that is specified. The **perf** keyword displays only performance gauges, and the **except** keyword displays only exception gauges. The truncated list can be further restricted using the **-devtype** flag and gauge name.
- devtype **ess** | **svc**
An optional parameter that lists only the gauges with the device type that is specified. The truncated list can be further restricted using the **-type** flag and gauge name.
- gauge_name* [...] | -
An optional parameter that lists information only for the gauge name that is specified.

Description

You can use the **lsgauge** command to retrieve information pertaining to a temporary gauge definition, resulting from a temporary modification to the properties of a displayed gauge.

Note: You can make modifications to a subset of properties of a gauge that is displayed, and then refresh the display.

Return values

One of the following:

- Gauges matching specified criteria not found.
- The vital output, which displays the following information:
 - Name
 - Type
 - Device ID
 - Device type

- Date modified
- The long output, which displays the following information:
 - Name
 - Type
 - Device type
 - Device ID
 - Date modified
 - Description
- The short output displays only **Name**.

Possible failures

- Gauge *gauge_name* does not exist.

Examples

An invocation example:

```
perfcli lsgauge gauge1
```

The resulting output:

Name	Type	Device Type	Device	Date Modified
gauge1	perf	ESS	ESS2105-123124-IBM	01/02/2005

Related topics

- “rmgauge” on page 42
- “showgauge” on page 58

lshost

You can use the **lshost** command to lists all hosts available on specific storage devices.

Syntax

```
perftool> lshost [-id id] [-wwpn [...]]
```

Parameters

- id *id*

An optional parameter that specifies the device that will have the hosts assigned to them. The device ID is the nickname or model-serial-manufacturer (for example, ESS2105-123123-IBM). Separate multiple device names by commas with no white space before or after each comma. All devices must be of the same device type.

wwpn [. . .] | -

An optional parameter that lists information only for the host names (the world-wide port number) specified. Separate multiple host names with a white space between each name. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin). If no host name is specified, all host names are displayed.

Description

You can use the **lshost** command to lists all hosts available on specific storage devices.

Return values

One of the following types of information is returned:

- No hosts matching the criteria specified were found.
- When the command is successful, the following information displays:
 - Device ID
 - Host name (the world-wide port number)

Possible failures

- Device *id* not found.

Examples

An invocation example:

```
perftool> lshost
```

The resulting output:

```
Device ID          WWPN
=====
2105-2105.26357-IBM 5005076300CDA48D
2105-2105.26357-IBM 5005076300CAA48D
```

2105-2105.26357-IBM 5005076300C1A48D
2105-2105.26357-IBM 5005076300C2A48D
2105-2105.26357-IBM 5005076300C6A48D
2105-2105.26357-IBM 5005076300CEA48D

Related topics

- None.

lslun

You can use the **lslun** command to display information about Performance Manager volumes or logical unit numbers (LUNs).

Syntax

```
▶▶—perftool—lslun—  
└─dev—device_id— [...] ─▶▶
```

Parameters

-dev

An optional parameter that lists LUNs by device ID. The default is *all*.

device_id [...] | -

An optional parameter that specifies the device ID. Separate multiple IDs with a comma between each. Device IDs follow the format used by the `lsdevice` command (for example: 2105.26884).

Description

Use the **lslun** command to display information about Performance Manager volumes or logical unit numbers (LUNs).

Return values

Volume or LUN IDs.

Possible failures

- Volume *device_id* does not exist.
- No LUN for *device_id*.

Examples

An invocation example:

```
perftool lslun -dev 2105.26884
```

The resulting output:

```
Vol  
=====  
ESS:2105.26884:VOL:2221  
ESS:2105.26884:VOL:2208  
ESS:2105.26884:VOL:2212  
ESS:2105.26884:VOL:2204  
ESS:2105.26884:VOL:2216  
ESS:2105.26884:VOL:2206
```

Related topics

lsprofile

You can use the **lsprofile** command to display Performance Volume Advisor profiles.

Syntax

```
▶▶ perfcli —lsprofile— profile_name [...] [-s] ▶▶
```

Parameters

profile_name [...] | -

An optional parameter that lists information for the profile names that are specified. Separate multiple profile names with a white space between each name. Alternatively, when you specify a dash (-), the command accepts user arguments. If no profile name is specified, all profiles are displayed.

-s An optional parameter that displays only the profile name.

-l An optional parameter that displays the profile name, description, and other vital output.

Description

You can use the **lsprofile** command to display Performance Volume Advisor profiles.

Return values

One of the following:

- No profiles matching the criteria specified were found.
- The vital output, which displays the following information:
 - Name
 - Description
- The long output, which displays the following information:
 - Name
 - Description
- The short output displays only **Name**.

Possible failures

- Profile *profile_name* does not exist.

Examples

An invocation example:

```
perfcli lsprofile -s
```

The resulting output:

```
Name
=====
Profile1
Profile2
```

Related topics

- “mkprofile” on page 35
- “chprofile” on page 10
- “cprofile” on page 13
- “rmpfile” on page 43
- “showprofile” on page 61
- “mkrecom” on page 38

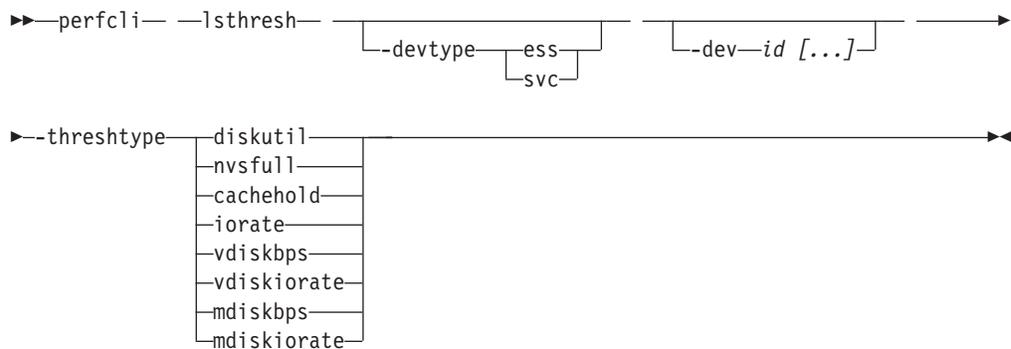
Isthresh

You can use the **Isthresh** command to display a summary of the threshold settings for all devices that are specified. The summary contains a list of the performance thresholds along with pertinent information associated with the threshold including:

- The device and device type with which the threshold is associated
- The name of the performance metric for which a threshold is set
- The critical and warning threshold levels, and whether or not an alert has been set

All devices IDs specified must be of the same device type, and the threshold type must be available for the specified device type.

Syntax



Parameters

-devtype **ess** | **svc**

An optional parameter that specifies the device type.

--dev **id** [...]

An optional parameter that displays only the threshold settings for the device IDs that are specified. The device ID is the nickname or manufacturer, model and serial number of the device. Separate multiple device names with a comma between each device, with no white space before or after the comma.

-threshtype **diskutil** | **nvsfull** | **cachehold** | **iorate** | **vdiskbps** | **vdiskiorate** | **mdiskbps** | **mdiskiorate**

Specifies which ESS threshold is to be set. Optionally, the **diskutil** threshold (ESS) includes the percentage of the sequential I/O filter when you use the **setfilter** command. One of the following ESS thresholds can be specified:

diskutil

The disk utilization value (percent of time busy), for each ESS array.

nvsfull

The nonvolatile storage full value (percent of time full), for each ESS cluster.

cachehold

The average cache hold time (in seconds), for each ESS cluster.

iorate The input-output rate (total number of I/O requests), for each ESS cluster.

vdiskbps

The total number of virtual disk I/s for each I/O group.

vdiskiorate

Virtual disk megabytes per second, for each I/O group.

mdiskbps

Total number of managed disk I/Os for each managed disk group.

mdiskiorate

For each managed disk group, the megabytes per second for each I/O group.

Description

You can use the **lsthresh** command to display a summary of threshold settings for all devices specified. The summary contains a list of the performance thresholds along with pertinent information associated with the threshold.

Return values

No threshold matched the criteria specified.

For each threshold, the following information is listed:

Column Label	Details
Device ID	The nickname or manufacturer, model and serial number of the device.
Device Type	Device Type: ESS .
Threshold name	Threshold value. Specify Diskutil .
Status	On or Off ; the status of the filter.
Warning	Displays the warning value for the threshold, if one exists. You can enter or change the value in this cell. This value is an integer.
Error	Displays the error value for the threshold, if one exists. You can enter or change the value in this cell. This value is an integer.
Alert	This value can be one of the following: None , Warning , or Error . The default is None .

The output includes a row for each device ID and threshold type that you specify.

Possible failures

- Device *device* must be of the same device type.
- Threshold *threshold* type is not available for device *device_id*.

Examples

An invocation example:

```
perfcli lsthresh -devtype ess -threshtype diskutil
```

The resulting output:

Device ID	Device Type	Threshold Name	Status	Warning	Error	Alert
ESS2105-123124-IBM	ESS	Disk Utilization	Enabled	60.0	30.0	None
ESS2105-123125-IBM	ESS	Disk Utilization	Enabled	60.0	30.0	None
ESS2105-123126-IBM	ESS	Disk Utilization	Enabled	50.0	80.0	None

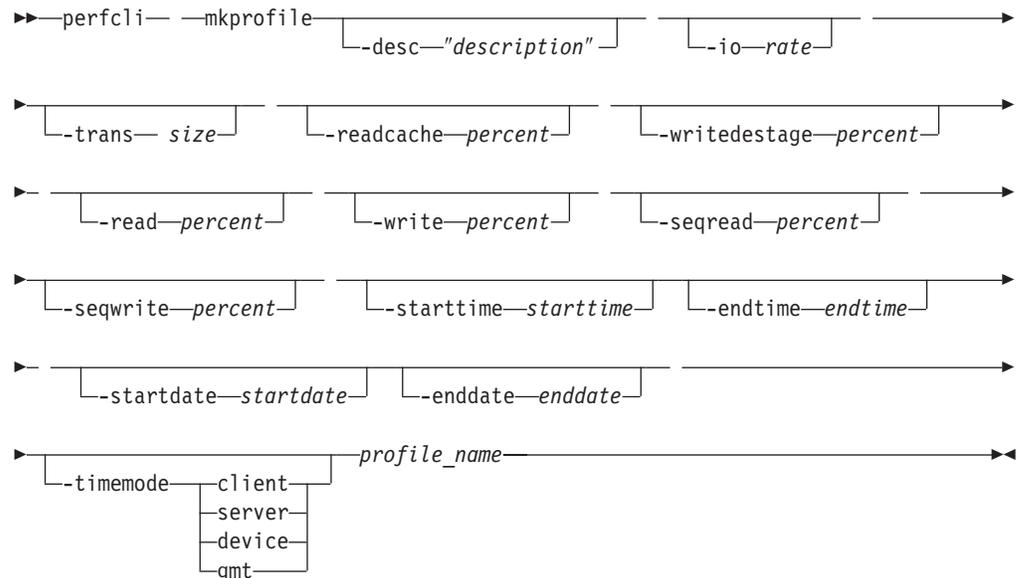
Related topics

- “setessthresh” on page 46
- “setsvcthresh” on page 52
- “setfilter” on page 44
- “lsfilter” on page 23
- “cpthresh” on page 14

mkprofile

Use the **mkprofile** command to create a workload profile that you can use later with the **mkrecom** command to create a performance recommendation for ESS volume allocation.

Syntax



Parameters

-desc "description"

An optional parameter that assigns a user-defined description of the performance data collection task. The maximum character limit is 25.

-trans size

An optional parameter that specifies the average transfer size in KB.

-io rate

An optional parameter that specifies the average I/Os per second per GB over all logical unit numbers.

-readcache percent

An optional parameter that specifies the percentage of random read cache hits to the total number of random read operations. Valid ranges are 0-100. The default value is 50%.

-writedestage percent

An optional parameter that specifies the percentage of random write destages operations to the total number of random write operations. Valid ranges are 0-100. The default value is 50%.

-read percent

An optional parameter that specifies the percentage of random reads to the total number of I/O operations. Valid ranges are 0-100. The default value is 70%. Values specified with -read, -write, -seqread, and -seqwrite combined cannot exceed 100.

-write *percent*

An optional parameter that specifies the percentage of random writes to the total number of I/O operations. Valid ranges are 0-100. The default value is 30%. Values specified with -read, -write, -seqread, and -seqwrite combined cannot exceed 100.

-seqread *percent*

An optional parameter that specifies the percentage of sequential reads to the total number of I/O operations. Valid ranges are 0-100. The default value is 0%. Values specified with -read, -write, -seqread, and -seqwrite combined cannot exceed 100.

-seqwrite *percent*

An optional parameter that specifies the percentage of sequential writes to the total number of I/O operations. Valid ranges are 0-100. The default value is 0%. Values specified with -read, -write, -seqread, and -seqwrite combined cannot exceed 100.

-starttime *starttime*

An optional parameter that specifies the time, in HH:MM format, of the start peak duration..

-endtime *endtime*

An optional parameter that specifies the time, in HH:MM format, of the end peak duration.

-startdate *startdate*

An optional parameter that specifies the date, in yyyyymmdd format, when the profile should start.

-enddate *enddate*

An optional parameter that specifies the date, in yyyyymmdd format, when the profile should end.

-timemode *client | server | device | gmt*

An optional parameter that indicates how the starttime, endtime, startdate, and enddate parameters are interpreted in terms of time zone.

client Dates and times are specified in the time zone of the client.

server Dates and times are specified in the time zone of the server.

device Dates and times are specified in the time zone of the device.

gmt Dates and times are specified in the time zone of the server expressed as Greenwich Mean Time (GMT).

profile_name

A user-assigned name for the profile that is being created.

Description

You can use the **mkprofile** command to recreate a workload profile that can later be used with the **mkrecom** command to create a performance recommendation for ESS volume allocation.

Return values

Profile *profile_name* successfully created.

Possible failures

- Profile *profile_name* already exists.
- Cannot save a profile with the same name as a pre-existing profile.

Examples

An invocation example:

```
percli> mkprofile -desc "my description" -read 20 -write 25 -seqread 25 -seqwrite 30 -startdate 20040214  
-starttime 08:30 -enddate 20040403 -endtime 13:00 -timemode device myprofile
```

The resulting output:

Profile myprofile successfully created.

In this example, the client is in the EST time zone, the server is in the PST time zone, device A is in PST, device B in is CST, the starttime is specified as 13:00, and the endtime is specified as 14:00. The timemode values are interpreted as:

client The server will query the server database for device A and device B data with server timestamps between 10:00 and 11:00, three hours earlier than specified.

device The server will query the database for device A and device B data with device timestamps between 13:00 and 14:00.

server The server will query the database for device A and device B data with server timestamps between 13:00 and 14:00.

gmt Since PST is (GMT - 8), the server will query for device A and device B data with server timestamps between 05:00 and 06:00.

Related topics

- “tracklun” on page 70
- “chprofile” on page 10
- “rmprofile” on page 43
- “cprofile” on page 13
- “lsprofile” on page 30
- “showprofile” on page 61
- “mkrecom” on page 38

and must include device ID. Separate multiple fully-qualified values with commas and no white space between the values. If a fully qualified ID is not available, then the host name and device ID can be indicated by separating each value with a colon.

-exclude LUNs [. . .]

An optional parameter that restricts the volume advisor recommendations from using the LUN or set of LUNs specified for an allocation. The volume advisor considers all available LUNs except those specified with the **-exclude** parameter. All LUNs must be of the same device type. You can specify LUNs at various device levels, such as clusters, adapters, arrays, nodes, MDisk groups and so forth. The **-exclude** parameter cannot be used with the **-include** parameter.

-include LUNs [. . .]

An optional parameter that specifies a LUN or set of LUNs that is to be used during the volume advisor recommendation. The volume advisor only considers the LUNs specified with the **-include** parameter. All LUNs must be of the same device type. You can specify LUNs at various device levels, such as clusters, adapters, arrays, nodes, MDisk groups and so forth. The **-include** parameter cannot be used with the **-exclude** parameter.

-seed value

An optional parameter that specifies a seed value for the random number generator. To allow two consecutive **mkrecom** commands to return the same recommendation when the same **mkrecom** inputs are provided (assuming that the configuration and I/O behavior has not changed in the time between the two **mkrecom** invocations), the same seed value needs to be specified

-quiet

An optional parameter that turns off the confirmation prompt for this command.

Description

Use a profile and the **mkrecom** command to generate a performance LUN advisor allocation recommendation for ESS volumes and, optionally, to apply the recommendation. The **mkrecom** command generates a list of optimal logical unit numbers (LUNs) that can be used for a given storage request.

Notes:

1. Create a workload profile with the **mkprofile** command.
2. Recommendations can apply to either preallocated or nonallocated storage. If the recommendation includes space that is not allocated yet, applying the recommendation automatically creates the LUNs necessary to implement the recommendation.
3. Use the **-seed** parameter to specify a seed to initialize the random number generator. This ensures that not all new volumes get added to the same component when multiple components have the same cost factors associated with them. This means that if you run the advisor twice in a row, with identical input, you will probably get different, but equally good, advice. To ensure that you get the same advice both times in a row, you can specify the same seed value to cause selection of the same random numbers both times you invoke the **mkrecom** command.

Return values

The following message and a table are both returned.

- Volumes recommendation: *xxx* Vdisks or volumes of 50 GB each. (The value for *xxx* is the number of volumes recommended.)
- For each task, the following information is listed in either tabular or multi-row table format:

Column Label	Details
Utilization	Available capacity, such as low, medium, or high.
Device	Device ID (nickname or model-serial-manufacturer)
Location	Array ID
Volume	Volume ID (if volume reused; else a dash)
Size (GB)	Size in GB
RAID Level	RAID-X or JBOD

Possible failures

- Profile `profile_name` does not exist.
- LUNs are not all of the same device type.

Examples

An invocation example:

```
perfcli mkrecom -profile profile4 -space 300 GB
```

The resulting output:

```
Name      Util Device Location Vol Size (GB) RAID Host
=====
profile4  low  ESS6   array4   15   20  RAID-5  HewlittPackard3
```

```
Would you like to apply this recommendation now? y/n
```

```
y
```

```
15 volumes created.
```

Related topics

- “mkprofile” on page 35
- “lsprofile” on page 30
- “showprofile” on page 61

quit

Use the **quit** command to exit the CLI program.

Syntax

▶▶—perfcli— —quit—▶▶

Description

Use the quit command to quit a session.

Examples

An invocation example:

```
perfcli quit
```

The resulting output:

```
Are you sure you want to quit? y/n  
Y
```

Related topics

- “repli” on page 74

rmgauge

Use the **rmgauge** command to delete an existing gauge definition.

Syntax

```
▶▶ perfcli — rmgauge — -quiet — gauge_name — [...] ▶▶
```

Parameters

-quiet

An optional parameter that turns off the confirmation prompt for this command.

gauge_name [...] | -

Specifies the gauge names to be deleted. Separate multiple gauge names with a white space between each name. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin).

Description

You can use the **rmgauge** command to delete the definition of a gauge.

Return values

Gauge *gauge_name* successfully deleted.

Possible failures

- Gauge *gauge_name* does not exist.
- Gauge *gauge_name* is currently being displayed in the GUI. You must close the display before deletion is permitted.

Examples

An invocation example:

```
perfcli rmgauge gauge1
```

The resulting output:

```
Are you sure you want to delete gauge gauge_name?
```

```
Gauge gauge_name successfully deleted.
```

Related topics

- “lsgauge” on page 25
- “showgauge” on page 58

rmprofile

Use the **rmprofile** command to delete performance LUN advisor profiles.

Syntax

```
▶▶—perfcli— —rmprofile— _quiet_ —profile_name— [...]▶▶
```

Parameters

-quiet

An optional parameter that turns off the confirmation prompt for this command.

profile_name [. . .] | -

Specifies the user-defined profile name. Multiple names can be specified with a white space between each value. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin).

Description

You can use the **rmprofile** command to delete an existing LUN Advisor workload profile.

Return values

Profile *profile_name* successfully deleted.

Possible failures

.

Examples

An invocation example:

```
perfcli rmprofile myprofilename
```

The resulting output:

```
Are you sure you want to delete profile myprofilename? y/n  
Y  
profile myprofilename successfully deleted.
```

Related topics

- “mkprofile” on page 35
- “chprofile” on page 10
- “cpprofile” on page 13
- “lsprofile” on page 30
- “showprofile” on page 61
- “mkrecom” on page 38

-restore

An optional parameter that restores all default values for the specified filter type. This parameter works only for those filters that have default values.

Description

You can use the **setfilter** command to set or change the threshold filter or exceptions.

Return values

Filter successfully set.

Possible failures

- Device type and threshold type do not match.
- Device *device* does not exist.

Examples

An invocation example:

```
perfcli> setfilter -devtype ess -threshtype diskutil -filter seqio -limit 80 -dev 2105.28019
```

The resulting output:

Filter successfully set.

Related topics

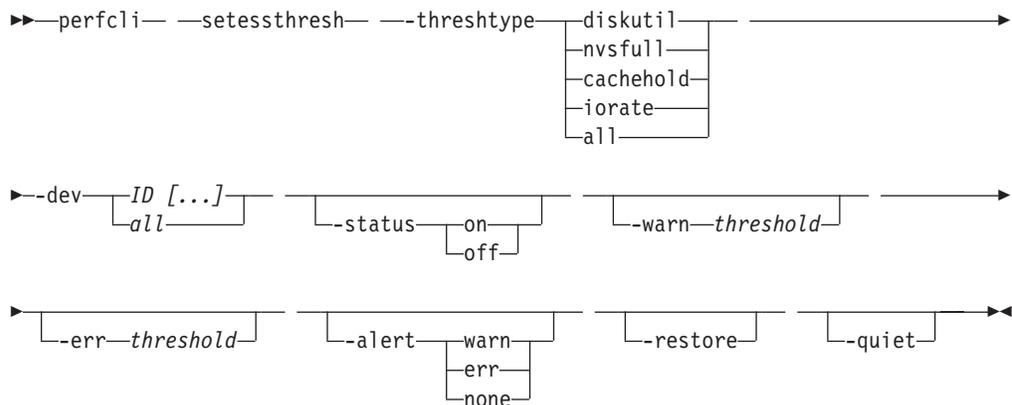
- “setessthresh” on page 46
- “setsvcthresh” on page 52
- “lsfilter” on page 23
- “cpthresh” on page 14
- “lsthresh” on page 32

setessthresh

You can use the **setessthresh** command to set or change the ESS performance thresholds. To modify an existing threshold setting for a specific device, use the **setessthresh** command again. Only new values specified change; unspecified settings are preserved.

- All changes to threshold settings apply immediately, even if a current performance data collection task is already active and running.
- Threshold checking does not occur on a device unless a performance data collection task has been created for that device and the task has been started as part of a performance collection job. Performance data collection jobs, when run, check devices to determine which thresholds are enabled in order to see if any of the performance statistics, for which thresholds are set, exceed the threshold value.
- Enabling a threshold status results in any detected exception data to be recorded in the Performance Manager database. However, if the alert level is set to **none**, a log message is entered into the MDM event log and the exception information is saved in the Performance Manager database. If the alert level is set to anything other than **none**, the Performance Manager sends an internal event to the Director. Users can configure how these alerts should manifest (for example, SNMP, E-mail, and so on) using Director's Event Action Planner.
- All ESS performance thresholds, except for I/O rate, have default critical and warning values, and are enabled for threshold checking by default. None of the SAN Volume Controller thresholds are enabled by default, nor do they have default warning and critical values.

Syntax



Parameters

-threshtype **diskutil** | **nvsfull** | **cachehold** | **iorate** | **all**

Specifies which ESS threshold is to be set. Optionally, the **diskutil** threshold (ESS) includes the percentage of the sequential I/O filter when you use the **setfilter** command. One of the following ESS thresholds can be specified:

diskutil

The disk utilization value (percent of time busy), for each ESS array.

nvsfull

The nonvolatile storage full value (percent of time full), for each ESS cluster.

cachehold

The average cache hold time (in seconds), for each ESS cluster.

iorate The input-output rate (total number of I/O requests), for each ESS cluster.

all Sets all the thresholds.

-dev *id* [...] | all

Specifies the ESS device that will apply the threshold to be set. The device ID is the nickname or model-serial-manufacturer of the device. Multiple devices must be separated by commas. All devices must be of the same device type. When you specify **all** the threshold applies to all ESS devices.

-status on | off

An optional parameter that specifies whether the threshold type should be enabled or disabled for checking. For all threshold types, except the I/O rate, the default is **on** and enabled for checking.

-warn *threshold*

An optional parameter that specifies the warning threshold value. Valid ranges are 0 - 100 for all devices. The default percentage values for each warning threshold type are as follows:

- diskutil = 50
- nvsvfull = 3
- cachehold = 60
- io = no default

-err *threshold*

An optional parameter that specifies the error threshold value. Valid ranges are 0 - 100. The default percentage values for each error threshold type are as follow:

- diskutil = 80
- nvsvfull = 10
- cachehold = 30
- io = no default

-alert warn | err | none

An optional parameter that specifies which thresholds should send an alert if the threshold is exceeded. One of the following alert values can be specified:

1. Warn. An alert is sent when the warning threshold is exceeded.
2. Err. An alert is sent when the warning and error threshold is exceeded.
3. None. No alerts are sent. This is the default.

Note: If the threshold status is enabled for checking, but no alerts are set to be sent, a log message is entered into the Multiple-Device Resource Manager event log, and exception information is recorded in the database.

-restore

An optional parameter that restores all default values for the specified threshold type. This argument works only on those thresholds that have default values.

-quiet

An optional parameter that turns off the confirmation prompt for this command.

Description

You can use the **setessthresh** command to set or change the ESS performance thresholds.

Return values

Threshold successfully set.

Possible failures

- Device *device* does not exist.
- Historical data must exist in the database for threshold checking to be activated. You must use the **mkcollector** command or schedule a data collection task in the GUI.

Examples

An invocation example:

```
perfcli setessthresh -threshtype iorate -dev 2105.111222 -warn 60 -alert err
```

The resulting output:

Threshold successfully set.

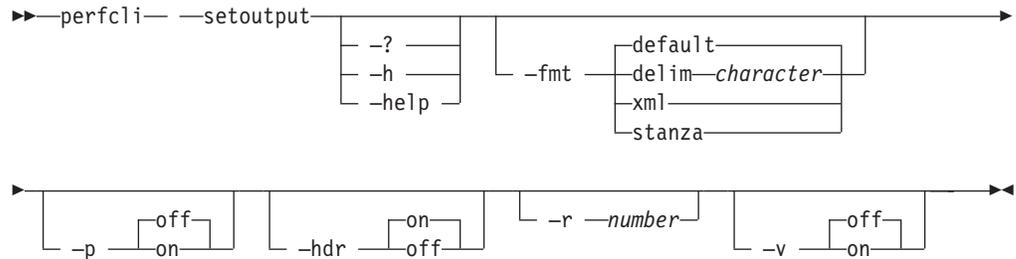
Related topics

- “setsvcthresh” on page 52
- “setfilter” on page 44
- “lsfilter” on page 23
- “cpthresh” on page 14
- “lsthresh” on page 32

setoutput

You can use the **setoutput** command to display current output settings for both perfcli and repcli commands. The output format set by this command remains in effect for the duration of a command session or until the options are reset.

Syntax



Parameters

-? | -h | -help

Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

-fmt

Specifies the format of the output. You can specify one of the following values:

default

Specifies to display output in a tabular format using spaces as the delimiter between the columns. This is the default value. For example:

```
Device ID      Device Type Threshold Name      Status  Warning  Error  Alert
-----
ESS2105.123123.IBM  ESS      Disk Utilization    Enabled  50      80     None
```

delim character

Specifies to display output in a tabular format using the specified character to separate the columns. If you use a shell metacharacter (for example, * or \t) as the delimiting character, enclose the character in single quotation mark (') or double quotation mark ("). A blank space is not a valid character. For example:

```
Device ID,Device Type,Threshold Name,Status,Warning,Error,Alert
=====
ESS2105.123123.IBM,ESS,Disk Utilization,Enabled,50,80,None
```

xml

Specifies to display output using XML format, for example:

```
<IRETURNVALUE>
<INSTANCE CLASSNAME="lsthresh">
<PROPERTY NAME="Device ID" TYPE="string"><VALUE>
  ESS2105.123123.IBM</VALUE>
</PROPERTY>
<PROPERTY NAME="Device Type" TYPE="uint32"><VALUE>ESS</VALUE>
</PROPERTY>
<PROPERTY NAME="Threshold Name" TYPE="uint64"><VALUE>Disk
  Utilization</VALUE>
</PROPERTY>
<PROPERTY NAME="Status" TYPE="uint16"><VALUE>Enabled</VALUE>
</PROPERTY>
<PROPERTY NAME="Warning" TYPE="uint64"><VALUE>d50</VALUE>
</PROPERTY>
```

```

<PROPERTY NAME="Error" TYPE="uint64"><VALUE>80</VALUE>
</PROPERTY>
<PROPERTY NAME="Alert" TYPE="uint16"><VALUE>None</VALUE>
</PROPERTY>
</INSTANCE>
</IRETURNVALUE>

```

stanza Specifies to display output using XML format, for example:

```

Device ID      ESS2105.123123.IBM
Device Type    ESS
Threshold Name Disk Utilization
Status         Enabled
Warning        50
Error          80
Alert          None

```

-p Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value when the **perftool** command is run in single-shot mode.

on Displays one page of text at time. Pressing any key displays the next page. This is the default value when the **perftool** command is run in interactive mode.

-hdr

Specifies whether to display the table header.

on Displays the table header. This is the default value.

off Does not display the table header.

-r number

Specifies the number of rows per page to display when the **-p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

-v Specifies whether to enable verbose mode.

off Disables verbose mode. This is the default value.

on Enables verbose mode.

Description

You can use the **setoutput** command to display current output settings for both **perfcli** and **repcli** commands. The output format set by this command remains in effect for the duration of a command session or until the options are reset either by using this command or by specifying an output-format parameter as part of a command.

Note: The output formats do not apply to help pages.

Return values

Running this command with no parameters displays the current output settings in the default output format, for example:

```
perfcli setoutput
```

```

Paging  Rows  Format  Header  Verbose
-----
off     -      default  on      off

```

Examples

The following perfcli example sets the output format to display in tabular form using a comma as the delimiter without header information.

```
perfcli>setoutput -fmt delim , -hdr off
perfcli>1sthresh -1 -type default
ESS2105.123123.IBM ,ESS,Disk Utilization,Enabled,50,80,None
```

Related topics

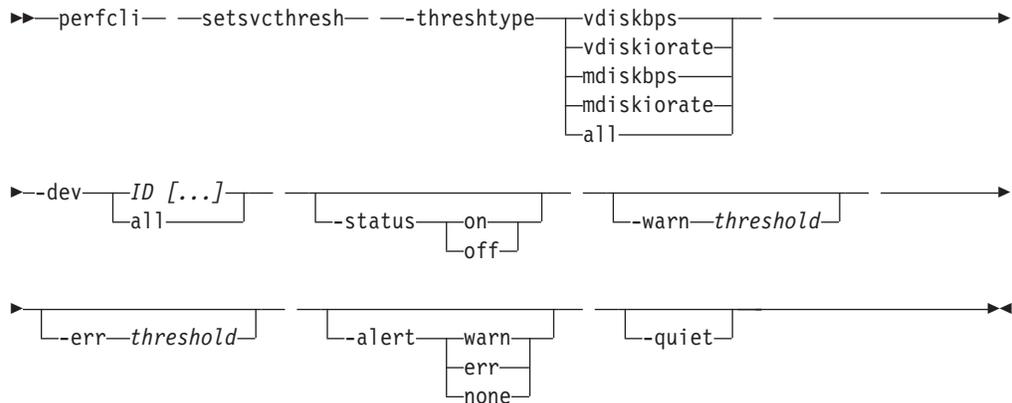
- “perfcli” on page 8
- “repli” on page 74

setsvcthresh

You can use the **setsvcthresh** command to set or change the SAN Volume Controller performance thresholds. To modify an existing threshold setting for a specific device, use the **setsvcthresh** command again. Only new values specified will change; unspecified settings will be preserved.

- All changes to threshold settings apply immediately, even if a current performance data collection task is already active and running.
- Threshold checking does not occur on a device unless a performance data collection task has been created for that device and the task has been started as part of a performance collection job. Performance data collection jobs, when run, check devices for which thresholds are enabled to see if any of the performance statistics, for which thresholds are set, exceed the threshold value.
- Enabling a threshold status results in any detected exception data to be recorded in the Performance Manager database, but if the alert level is set to **none**, a log message is entered into the MDM event log and the exception information is saved in the Performance Manager database. If the alert level is set to anything other than **none**, the Performance Manager sends an internal event to Director. Users can configure how these alerts manifest (for example, SNMP, E-mail, and so on) using Director's Event Action Planner.

Syntax



Parameters

-threshtype vdiskbps | vdiskiorate | mdiskbps | mdiskiorate | all
Specifies which threshold is to be set. Specify one of the following types:

vdiskbps

The total number of virtual disk I/s for each I/O group.

vdiskiorate

Virtual disk megabytes per second, for each I/O group.

mdiskbps

Total number of managed disk I/Os for each managed disk group.

mdiskiorate

For each managed disk group, the megabytes per second for each I/O group.

all Sets all the thresholds.

-dev *id* [...] | all

Specifies the device that will apply the threshold to be set. The device ID is the nickname or model-serial-manufacturer of the device. Separate multiple device IDs by commas. All devices must be of the same device type.

-status on | off

An optional parameter that specifies whether the specified threshold type should be enabled or disabled for checking. The default is **off** and disabled for checking.

-warn *threshold*

An optional parameter that specifies the warning threshold percentage value. Valid ranges are 0 - 100. There are no default values.

-err *threshold*

An optional parameter that specifies the error threshold percentage value. Valid ranges are 0 - 100. There are no default values.

-alert warn | err | none

An optional parameter that specifies which thresholds should send an alert if the threshold is exceeded. One of the following alert values can be specified:

1. Warn. An alert is sent when the warning threshold is exceeded.
2. Err. An alert is sent when the warning and error threshold is exceeded.
3. None. No alerts are sent. This is the default.

Note: If the threshold status is enabled for checking, but no alerts are set to be sent, a log message is still entered into the Multiple-Device Resource Manager event log and exception information is recorded in the database.

-quiet

An optional parameter that turns off the confirmation prompt for this command.

Description

You can use the **setsvcthresh** command to set the threshold properties for the SAN Volume Controller.

Return values

Threshold *thresh_type* successfully set.

return value

An explanation of the return value.

Possible failures

- Device *device* does not exist.

Examples

An example:

```
perfccli setsvcthresh -threshtype mdiskiorate -devid all -status off
```

The resulting output:

```
Threshold successfully set.
```

Related topics

- “setessthresh” on page 46
- “setfilter” on page 44
- “lfilter” on page 23
- “cpthresh” on page 14
- “lsthresh” on page 32

showcapacity

Use the showcapacity command to display the managed capacity for each ESS and SVC device, the sum of managed capacity by device type, and the total of all ESS and SVC device managed storage.

Syntax

```
►►—perfcli— —showcapacity—————▶▶
```

Description

Use the showcapacity command to display the amount of managed storage for each device, the total amount of managed storage by device type, and the total of all managed storage.

Examples

An invocation example:

```
perfcli showcapacity
```

The resulting output:

```
ESS Devices
2105.26884 1921.3964 GB
2105.28019 1600.5969 GB
ESS Capacity      3.52 TB

SVC Devices
00000200622002F0 70236.766 GB
0000020064800318 10.48576 GB
SVC Capacity      70.25 TB

Total MDM Capacity 73.77 TB
```

showdev

You can use the **showdev** command to display device properties. You can also use the *-id* option of this command to display properties for a specific device.

Syntax

```
percli —showdev— --id [...]—————▶
```

Parameters

-id [...]

A required parameter that displays the properties for the device IDs that are specified. The device ID is the nickname or model-serial-manufacturer. Separate multiple device names with a comma in between each name with no white space before or after the comma.

Description

You can use the **showdev** command to display device properties. You can also use the *-id* option of this command to display properties for a specific device.

Return values

One of the following types of information is returned:

Column Label	Details
Device ID	The nickname or model-serial-manufacturer of the device..
IP	Device IP address.
Device Type	Device Type: ESS or SVC.

Possible failures

- No device matched the criteria specified.
- Device *device_id* not found.

Examples

An invocation example:

```
percli showdev 2105-2105.26884-IBM
```

The resulting output:

```
Device ID          IP address Device Type
=====
2105-2105.26884-IBM 9.11.222.10  ESS
```

Related topics

- “lsdev” on page 21

showdbinfo

Use the **showdbinfo** command to display percent full, used space, and free space of the Performance Manager database.

Syntax

```
►—perfcli— —showdbinfo—◄
```

Description

You can use the **showdbinfo** command to display percent full, used space, and free space of the Performance Manager database.

Return values

The output includes the following information in table format:

- Name
- Full (%)
- Used (MB)
- Free (MB)

Possible failures

- Task *task_name* already exists.

Examples

An invocation example:

```
perfcli showdbinfo
```

The resulting output:

```
Name  
Full (%)  
Used (MB)  
Free (MB)
```

Related topics

- “startdbpurge” on page 63

showgauge

You can use the **showgauge** command to display performance output by triggering an existing gauge.

Syntax

```
►►—perfcli— —showgauge— —gauge_name—◀◀
```

Parameters

gauge_name

Specifies the gauge name to be displayed.

Description

You can use the **showgauge** command to display performance output by triggering an existing gauge.

Return values

Performance Historical Gauges: Each gauge displays the following vital output for performance historical gauges.

Gauge Name

The user-defined name from task creation.

Description

The user-defined description.

Device

The device ID of the component that is displayed in the gauge. The device ID is the nickname or model-serial-manufacturer of the device.

Component level

The component level that is identified in task creation.

Component ID

The component ID identified during task creation.

Thresholds

The thresholds that are applied to the selected metrics.

Warning

The warning value.

Time

The time of last data collection for the device or device component (“Last Data Collection”).

After the information above is listed, a table follows with this information:

Date

The date collection was taken.

Time

The time of collection.

metric_name

The name of the metric, for example, Average Cache Holding Time is the heading for this column, and the value of the metric is listed in the table.

This represents the gauge type that you selected during the creation of the definition. [0..greatest value registered by the performance data]

The metric name is the same as the one selected during the creation of the gauge. The number of metric columns is the same as the number of metrics selected during creation of the gauge.

Performance Real Time Gauges: Each gauge displays the following vital output for performance real time gauges.

Gauge Name	The user-defined name from task creation.
Description	The user-defined description.
Device	The device ID of the component displayed in gauge. The device ID is the nickname or model-serial-manufacturer.
Component level	The component level that was identified in task creation.
Component ID	The component ID that was identified during task creation.
<i>metric_name</i>	The value of the metric. The gauge type is selected by the user during the creation of the definition.

After the information above is listed, a table follows with this information:

Date	The date that the collection was taken.
Time	The time of the collection.
ID	For the maximum thresholds only, number of IDs that exceeded the maximum threshold.
<i>metric_name</i>	The name of the metric, for example, Average Cache Holding Time is the heading for this column, and the value of the metric is listed in the table. This represents the gauge type that you selected during the creation of the definition. [0..greatest value registered by the performance data]. The metric name is the same as the one selected during the creation of the gauge. The number of metric columns is the same as the number of metrics selected during creation of the gauge.

Exception Gauges: Each gauge displays the following vital output for exception gauges.

Gauge Name	The user-defined name from the task creation.
Device	The device ID of the component displayed in gauge. The device ID is the nickname or model-serial-manufacturer.
Thresholds	The thresholds that were applied to the selected metrics.
Warning	The warning value that was given.

Time	The time of last data collection for the device or device component (“Last Data Collection”).
Description	The user-defined description.

After the information above is listed, a table follows with this information:

Date	The date that the collection was taken.
Time	The time of the collection.
<i>metric_name</i>	<p>The name of the metric, for example, Average Cache Holding Time is the heading for this column, and the value of the metric is listed in the table. This represents the gauge type that you selected during the creation of the definition. [0..greatest value registered by the performance data].</p> <p>The metric name is the same as the one selected during the creation of the gauge. The number of metric columns is the same as the number of metrics selected during creation of the gauge.</p>

Possible failures

- Gauge *gauge_name* does not exist.

Examples

An invocation example:

```
perfcli> showgauge essTest5
```

The resulting output:

```
Name essTest5
Description My performance gauge
Device 2105.20870
Component Level Cluster
Component ID 2
Thresholds Cache Holding Time; Warning: 60.0; Error: 30.0; Status: enabled
LastDataCollection Wed Oct 15 14:31:44 PDT 2003
```

```
Date Time Average Cache Holding Time
=====
Sep. 14, 2003 9:51:25 70
Sep. 15, 2003 9:52:25 65
```

Related topics

- “*rmgauge*” on page 42
- “*lsgauge*” on page 25

showprofile

Use the **showprofile** command to display the specific properties of an individual performance workload profile. Use the **lsprofile** command to retrieve a list of available profile names.

Syntax

```
►►—perfcli— —showprofile— —profile_name—◄◄
```

Parameters

profile_name

The name of the profile to be displayed.

Description

You can use the **showprofile** command to display the specific properties of an individual performance workload profile.

Return values

When the command is successful, the following information displays:

- Name
- Description
- Total I/O per sec/GB
- Average Transfer Size (KB)
- Cache Read Hits (%)
- Cache Destage (%)
- Random Reads (%)
- Sequential Reads (%)
- Random Writes (%)
- Sequential Writes (%)

Possible failures

- Profile *profile_name* does not exist.

Examples

An invocation example:

```
perfcli showprofile myprofilename
```

The resulting output:

```
Name  
Description  
Total I/O per sec/GB  
Average Transfer Size (KB)  
Cache Read Hits (%)  
Cache Destage (%)  
Random Reads (%)  
Sequential Reads (%)  
Random Writes (%)  
Sequential Writes (%)
```

Start Peak Activity Date
End Peak Activity Date
Start Peak Activity Time
End Peak Activity Time

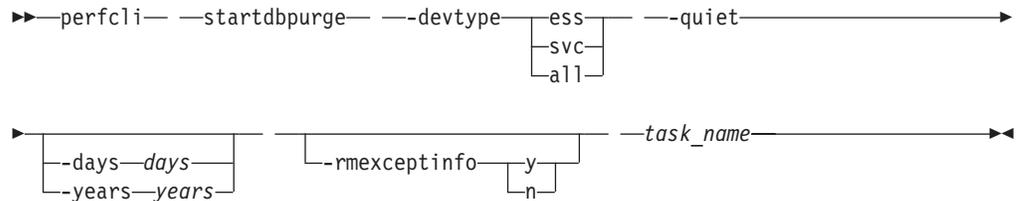
Related topics

- “mkprofile” on page 35
- “chprofile” on page 10
- “cprofile” on page 13
- “rmprofile” on page 43
- “showprofile” on page 61
- “mkrecom” on page 38

startdbpurge

Use the **startdbpurge** command to create a database purge task to be run as a non-interactive task at a later time, using the IBM Director scheduler.

Syntax



Parameters

-devtype *ess* | *svc* | *all*

An optional parameter that clears data for only the device type specified.

-quiet

An optional parameter that turns off the confirmation prompt for this command.

-days *days* | -years *years*

Specifies the number of days or years for which the performance data will be purged. Valid values for *days* are from 1 to 365 days. Valid values for *years* are from 1 to 10 years. If this flag is not specified, performance data for 365 days will be purged. You should specify either the **-days** or the **-years** flag.

-rmexceptinfo *y* | *n*

An optional parameter that specifies whether exception data is to be purged. Specify *y* to purge the data; specify *n* if you do not want the data purged.

task_name

Assigns a user-defined task name for the database purge task. The task name must be unique and cannot be the same as another existing Director task name.

Description

You can use the **startdbpurge** command to create a database purge task to be run as a non-interactive task at a later time, using the IBM Director scheduler.

Return values

- Database purge task *task_name* successfully started.

Possible failures

- Task *task_name* already exists.

Examples

An invocation example:

```
perfcli startdbpurge -devtype ess -perfinfo 5 -unit days mypurgetask
```

The resulting output:

Are you sure you want to purge the database since this is an irrecoverable operation that could be a long-running task and impact concurrent operations? y/n

y

Database purge task mypurgetask successfully started.

Related topics

- “showdbinfo” on page 57

startesscollection

You can use the **startesscollection** command to start a named data collection task on the Performance Manager server for an Enterprise Storage Server (ESS).

Syntax

```
▶—perfcli— —startesscollection— —-devid—device_id [...]— —————▶  
▶--devip—device_ip [...]— — --freq—frequency— --length—hours— —————▶  
▶—collection_name | —————▶▶
```

Parameters

-devid *device_id* [...]

Specifies the nickname or manufacturer, model, and serial number (for example, ESS2105-123123-IBM) of the device to be included in the performance data collection task. Multiple devices must be separated by commas with no white space before or after each comma. All devices must be of the same device type.

-devip *device_ip* [...]

Specifies the device IPs to be included in performance data collection task. Multiple devices must be separated by commas with no white space before or after each comma. All devices must be of the same device type.

-freq *frequency*

Specifies how often the performance data is to be collected in increments of 5 minutes. For ESS devices, options are 5, 10, 15, ..., 60.

-length *hours*

Specifies how long, in increments of 1 hour, the collection task is to run. Valid values are between 1 hour and 720 hours.

collection_name *name*

Specifies the name of the collection task. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin). The name cannot be > 64 characters.

Description

You can use the **startesscollection** command to collect data. After you create and name a data collection task for an ESS, you can start the data collection task in one of these ways:

- Using the **startesscollection** command
- Using the Director Scheduler

Return values

success

Collection *collection_name* was successfully executed.

Possible failures

- Task with same name is running on Performance Manager server.
- All the devices that are part of this task are currently collecting.

Examples

The following example starts a collection against two ESS devices to run every 10 minutes for five hours.

```
perfcli startcollection -devid 1234.12345,6789.67890 -devip 192.168.0.10,192.168.0.11  
-freq 10 -length 5 collection2
```

The following example starts a collection against a single ESS device to run every minute for an hour.

```
perfcli startcollection -devid 1234.12345 -devip 192.168.0.10  
-freq 1 -length 1 collection2
```

The resulting output:

Performance task collection2 successfully started.

Related topics

- “stopcollection” on page 69
- “lscollection” on page 17

Possible failures

- Task with same name is running on Performance Manager server.
- All the devices that are part of this task are currently collecting.
- Performance Manager not customized with the Secure Shell Client (SSH Client) *private* key which enables SAN Volume Controller (SVC) performance data collection.

Examples

The following example starts a collection against two SVC devices to run every 10 minutes for 5 hours.

```
perfcli startsvccollection -dev 1234.12345,6789.67890 -devip  
192.168.0.10,192.168.0.11 -devuser bob,susie -freq 10 -length 5  
collection1
```

The following example starts a collection against a single SVC device to run every hour for 24 hours.

```
perfcli startsvccollection -devid 1237.12348 -devip 195.178.0.20  
-freq 60 -length 24 collectiona
```

Related topics

- “stopcollection” on page 69
- “lscollection” on page 17

stopcollection

You can use the **stopcollection** command to stop a named collection task on the Performance Manager server.

Syntax

```
▶▶—perfcli— —stopcollection— —-devtype— ess  
svc -quiet —name—▶▶
```

Parameters

-devtype *ess* | *svc*

Specifies the device type (ess or svc).

-quiet

Turns off the confirmation prompt for this command.

name

Specifies the name of the collection task.

Description

You can use the **stopcollection** command to stop a named collection task on the Performance Manager server.

Possible failures

- No task with this name is currently running on the Performance Manager server.

Examples

An invocation example:

```
perfcli stopcollection -devtype svc collection3
```

The resulting output:

```
Collection task collection3 successfully stopped.
```

Related topics

- “startesscollection” on page 65
- “startsvccollection” on page 67
- “lscollection” on page 17

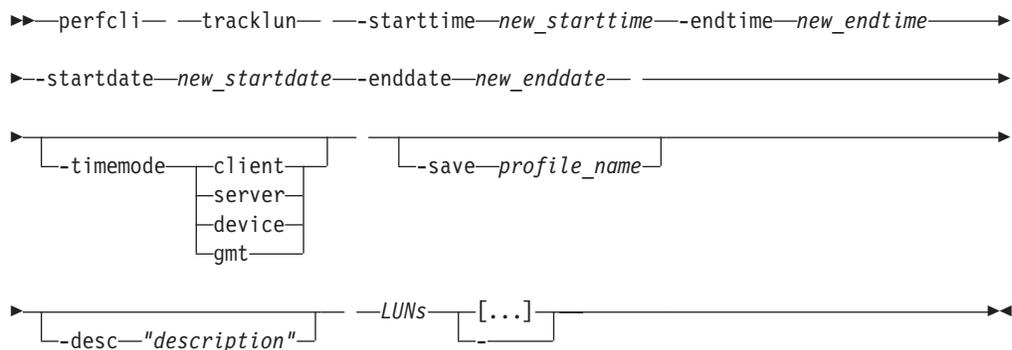
tracklun

Use the **tracklun** command to report historical performance statistics that are necessary to create a workload profile. The profile is created on specified ESS volumes based on a specified activity period.

Optionally, you can choose to save the historical statistics as a performance workload profile to be used with a **mkrecom** command.

- You must use the **mkcollector** command or schedule a data collection task in the graphical user interface (GUI). Historical data must exist in the database for the specified logical unit numbers (LUN) for the specified time.
- You must use the **mkcollector** command with the **mkrecom** command in order to generate a performance recommendation for storing data on a target set of ESS volumes.
- You cannot save a profile with the same name as a preexisting profile.
- You can manually create a performance workload profile using the **mkprofile** command as an alternative to using the **tracklun** command.

Syntax



Parameters

- starttime** *new_starttime*
The time, in HH:MM format, of the start peak duration.
- endtime** *endtime*
The time, in HH:MM format, of the end peak duration.
- startdate** *startdate*
The date, in yyyyymmdd format, when the profile should start.
- enddate** *enddate*
The date, in yyyyymmdd format, when the profile should end.
- timemode** *client* | *server* | *device* | *gmt*
An optional parameter that indicates how the *starttime*, *endtime*, *startdate*, and *enddate* parameters are interpreted in terms of time zone.
 - client** Dates and times are specified in the time zone of the client.
 - server** Dates and times are specified in the time zone of the server.
 - device** Dates and times are specified in the time zone of the device.
 - gmt** Dates and times are specified in the time zone of the server expressed as Greenwich Mean Time (GMT).

-save *profile_name*

An optional parameter that saves the performance data as a performance workload profile that can be used at a later time with the **mkrecom** command.

-desc "*description*"

An optional parameter that specifies a description for the profile when the **-save** parameter is used.

LUNs [...] | -

Specifies the ESS volumes to track. All LUNs must be the same device type. Separate each volume with a white space. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin).

Description

You can use the **tracklun** command to report historical performance statistics necessary to create a workload profile on specified ESS volumes based on a specified activity period. Optionally, you can choose to save the historical statistics as a performance workload profile to be used with a **mkrecom** command.

Return values

Profile *profile_name* is successfully created (if the **-save** parameter was specified).

- Name (if the **-save** parameter was specified)
- Description
- Total I/O per sec/GB
- Average Transfer Size (KB)
- Cache Read Hits (%)
- Cache Destage (%)
- Random Reads (%)
- Sequential Reads (%)
- Random Writes (%)
- Sequential Writes (%)

Possible failures

- Profile *profile_name* does not exist.
- Device type does not match LUN specified.
- Cannot specify **-desc** flag when the **-save** flag is not specified.
- No performance data is available for the LUNs specified. Ensure that the data collection task is scheduled for these LUNs.

Examples

An invocation example:

```
perfli> tracklun -starttime 12:00 -endtime 13:00 -startdate 20040101 -enddate 20040101  
ESS:2105.65302:VOL:1610 -save myprofile2
```

The resulting output:

Profile myprofile2 successfully created.

In this example, the client is in the EST time zone, the server is in the PST time zone, device A is in PST, device B in is CST, the starttime is specified as 13:00, and the endtime is specified as 14:00. The timemode values are interpreted as:

- client** The server will query the server database for device A and device B data with server timestamps between 10:00 and 11:00, three hours earlier than specified.
- device** The server will query the database for device A and device B data with device timestamps between 13:00 and 14:00.
- server** The server will query the database for device A and device B data with server timestamps between 13:00 and 14:00.
- gmt** Since PST is (GMT - 8), the server will query for device A and device B data with server timestamps between 05:00 and 06:00.

Related topics

- “tracklun” on page 70
- “chprofile” on page 10
- “rmprofile” on page 43
- “cpprofile” on page 13
- “lsprofile” on page 30
- “showprofile” on page 61
- “mkrecom” on page 38

Chapter 4. Replication Manager commands

Replication Manager (RM) provides two types of copy services, continuous remote copy and point-in-time copy. These data-copy services maintain consistent copies of data on source volumes that are managed by Multiple Device Manager. This topic lists the different Replication Manager commands that you can issue. These commands enable you to replicate data, to list detailed copy information, to create, modify, and delete sessions, groups, and pools, and to start and stop replication sessions.

Creating a replication session: Replication Manager supports the session concept in which multiple pairs are handled as a consistent unit. You can create and manage copy relationships between source and target volume pairs or source volume groups, and among target pools through a Replication Manager copy session. You use the **mksess** command to create and name a session. To create a session you provide the name of the volume group to be copied and the name to be assigned to the session that manages the copy of this group (the copy instance). You also select the copy type and a set of pool criteria that allows Replication Manager to know from which pool to choose a target. This is the sequence of events for a session:

1. Create a source group for volumes that are to be copied using the **mkgrp** command and create a target pool using the **mktgtpool** command.
2. Use the groups and pools to create a session using the **mksess** command.
3. Use the **startsess** command or the **flashsess** command to start a session.

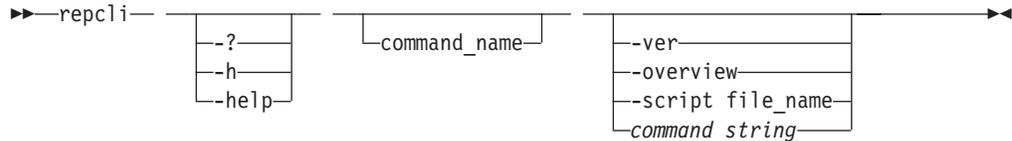
Creating a copy set: A copy set is all volumes that make up a solution for a given copy type. For example, if a copy type calls for four volumes where the copy goes sequentially from A to B and then from B to C, and then from C to D, then the copy set is the sequential, linear, horizontal path and collection of those volumes. A is the original source and B, C, and D are all considered targets. There are two ways that copy sets are created for Replication Manager. One way is with automated target volume selection using the **mktgtpool** command. The other way is where you know what targets you want to go with their source. This is the sequence of events if you want to specify a copy set yourself using the **mkcpset** command:

1. Create a source group for volumes that are to be copied using the **mkgrp** command.
2. Use the groups to create a session using the **mksess** command.
3. Get a list of volumes using the **lsvol** command.
4. Specify the copy sets using the **mkcpset** command.
5. Use the **startsess** command or the **flashsess** command to start a session.

repcli

Use this utility to start the administrative command-line interface (CLI), or the **repcli** session, which is used to run commands in interactive mode. This utility can also run a single command or a set of commands from a script without starting a **repcli** session.

Syntax



Parameters

-? | -h | -help *command_name*

Displays help for the specified command (for example, **-h mktgtpool** displays help for the **mktgtpool** command). If a command name is not specified, this parameter displays a list of available commands in the administrative CLI.

Note: Although the **repcli** utility is listed as a command in the help facility, you receive an error if you attempt to run the utility from within **repcli** (for example, `repcli> repcli help`). For help for the utility, enter `-help repcli .`

-ver

Displays the current version and licensing information for this utility.

-overview

Displays the overview information about the **repcli** utility, including command modes, standard command and listing parameters, syntax diagram conventions, and user assistance.

-script *file_name*

Runs the set of command strings in the specified file outside of a **repcli** session. You must specify a file name.

The format options specified using the **setoutput** command apply to all commands in the script. Output from successful commands routes to stdout.

Output from unsuccessful commands route to stderr. If an error occurs while one of the commands in the script is running, the script exits at the point of failure and return to the system prompt.

command_string

Runs the specified command string outside of a **repcli** session.

Description

You can use the **repcli** utility to start the administrative command-line interface (CLI), or **repcli** session, which is used to run commands in interactive mode. This utility can also run a single command or a set of commands from a script without starting a **repcli** session.

If you run this utility with any of the valid parameters, the administrative CLI is not started in interactive mode. If you run this utility with no parameters, this utility starts the administrative CLI. When you are in the administrative CLI, the `repl i>` prompt is displayed.

Related topics

help

You can use the **help** command to display a list of the Replication Manager commands or to display a list of commands with syntax.

Syntax

```
►► repli — help [ -s ] [ -l ]
```

Parameters

- s An optional parameter that displays the list of Replication Manager commands with brief descriptions.
- l An optional parameter that displays a list of the Replication Manager commands with descriptions and syntax.

Description

You can use the **help** command to display a list of the Replication Manager commands or to display a list of commands with syntax.

Examples

An invocation example:

```
repli help -s
```

The resulting output:

```
approvecpset Use approvecpset to approve a generated copy set as valid.
chcpset      Use chcpset to change volume membership of an existing
              Replication Manager copy set. You can only add volumes
              that are of the same device type.
...
...
```

Related topics

approvecpset

Use **approvecpset** to approve a generated copy set as valid.

Syntax

```
▶▶—repli— —approvecpset— --src—source_volume_id— [ . . . ] —————▶▶  
▶ [ -quiet ] —————▶▶ —session_name————▶▶
```

Parameters

-quiet

An optional parameter that turns off the confirmation prompt for this command.

-src *source_volume_id* [...] | -

Specifies the name of the source volume ID of the copy set to approve. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin).

session_name

Specifies the session name to which the copy set that is being approved belongs.

Description

You can use the **approvecpset** command to approve the generated copysets as valid copysets.

Return values

Copy set *copyset_name* successfully verified.

Possible failures

- Copy set *source volume id* does not exist.

Examples

An invocation example:

```
repli approvecpset  
-src ESS:2105.65302:VOL:1611,ESS:2105.65302:VOL:1612 mysession
```

The resulting output:

```
Copy set mysession successfully verified.
```

Related topics

- “mkcpset” on page 112
- “chcpset” on page 78
- “lscpset” on page 90
- “showcpset” on page 123
- “lsseq” on page 100

chcpset

Use **chcpset** to change volume membership of an existing Replication Manager copy set. You can only add volumes that are of the same device type.

Syntax

```
►—repli— —chcpset— —src—source_volume_id— —————►  
►—tgt—new_target_volume_id—[...]— —session_name—————►
```

Parameters

- src** *source_volume_id*
Specifies the source volume of the copy set that is being changed.
- tgt** *new_target_volume_id* [...]
Specifies the new target or targets for the copy set that is being changed.
- session_name*
Specifies the session name to which the copy set belongs.

Description

You can use **chcpset** to change volume membership of an existing Replication Manager copy set. You can only add volumes of the same device type.

Return values

Column Label	Details
Source Volume	Name of the source volume
Target Volume	Specifies the new target or targets for the copy set that is being changed
Session Name	Name of the session to which the copy set belongs

Possible failures

- Nothing to modify.
- Volume *volume* already belongs to session *session_name*.
- Source *volume* does not currently belong to session *session_name*.
- Volumes are not all of the same device type.
- Volume *volume* does not exist.
- You cannot delete the last volume in the copy set. You must delete the copy set itself.

Examples

An invocation example:

The resulting output:

```
repli chcpset -src volumeA -tgt new_volumeB sessionX  
copyset successfully changed.
```

Related topics

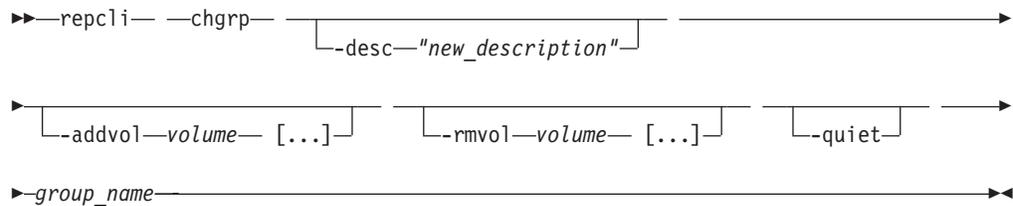
- “mkcset” on page 112
- “rmcset” on page 119
- “approvecset” on page 77
- “lscset” on page 90
- “showcset” on page 123

chgrp

Use the **chgrp** command to change the description or the volume membership of an existing Replication Manager group.

- You must specify a group name that is not the same name as another Replication Manager group.
- You can only add volumes that are of the same device type.
- If you delete volumes in a group, it can affect current sessions.

Syntax



Parameters

-desc "new_description"

An optional parameter that specifies the new user-defined description of the group. The maximum alphanumeric character length for this value is 250.

-addvol volume [...]

An optional parameter that specifies new volumes that you want to add to the group specified. All volumes must be of the same device type, but can belong to multiple groups. Valid volume values include those in the following list:

- Device type
- Component type
- Device ID
- Subsystem ID
- Logical subsystem ID
- Volume ID

What parameter values to specify:	How to code the parameter:
To specify more than one volume value:	ESS:ELEMENTTYPE.SSID:VOL:LSSVOLNUM
To specify a single volume:	-vol ESS:2105.65312:VOL:202A
To specify the logical subsystem so that all of the volumes in the specified logical subsystem will be included: You can choose volumes from any number of LSSs in an ESS. Note: The same volume can reside in multiple groups.	-vol ESS:2105.65312:VOL:20
To specify multiple volumes on different ESSs:	-vol ESS:2105.65312:VOL:202A,2C,2F
To specify a range of volumes for group membership, use a hyphen instead of a comma between volumes:	-vol ESS:2105.65312:VOL:202A-2F

What parameter values to specify:	How to code the parameter:
To specify single volumes and a range of volumes in the same logical subsystem, you can alternate between commas and hyphens, for example:	-vol ESS:2105.65312:VOL:202A,2C,2E-30

-rmvol *volume* [...]

An optional parameter that specifies volumes that you want to remove from the group specified. All volumes must be of the same device type, but can belong to multiple groups. Valid volume values include those in the following list:

- Device type
- Component type
- Device ID
- Subsystem ID
- Logical subsystem ID
- Volume ID

What parameter values to specify:	How to code the parameter:
To specify more than one volume value:	ESS:ELEMENTTYPE.SSID:VOL:LSSVOLNUM
To specify a single volume:	-vol ESS:2105.65312:VOL:202A
To specify the logical subsystem so that all of the volumes in the specified logical subsystem will be included: You can choose volumes from any number of LSSs in an ESS. Note: The same volume can reside in multiple groups.	-vol ESS:2105.65312:VOL:20
To specify multiple volumes on different ESSs:	-vol ESS:2105.65312:VOL:202A,2C,2F
To specify a range of volumes for group membership, use a hyphen instead of a comma between volumes:	-vol ESS:2105.65312:VOL:202A-2F
To specify single volumes and a range of volumes in the same logical subsystem, you can alternate between commas and hyphens, for example:	-vol ESS:2105.65312:VOL:202A,2C,2E-30

-quiet

An optional parameter that turns off the confirmation prompt for this command.

group_name

Specifies the existing name of the group to be changed.

Description

You can use the **chgrp** command to change the description or the volume membership of an existing Replication Manager group.

Return values

Possible failures

- Nothing to modify.
- Volume *volume* already belongs to group *group_name*.

- Volume *volume* does not currently belong to group *group_name*.
- Volumes are not all of the same device type.
- You cannot delete the last volume in the group. You must delete the group itself.

Examples

An invocation example:

```
repli chgrp -addvol ESS:2105.65312:VOL:202A-2F group1
```

The resulting output:

```
Group group1 successfully modified.
```

Related topics

- “mkgrp” on page 108
- “rmgrp” on page 120
- “lsgrp” on page 94
- “showgrp” on page 126
- “mksess” on page 110
- “lssess” on page 102
- “mktgtpool” on page 114
- “lstgtpool” on page 104

chtgtpool

Use the **chtgtpool** command to change description, name, or volume membership of an existing Replication Manager pool.

- The pool name must be unique; it cannot be the same name as another Replication Manager pool name.
- Volumes must be of the same device type.
- A volume cannot belong to multiple groups.
- Deleting volumes in a group can affect current sessions.

Syntax

```
▶▶ replcli — chtgtpool — [ -desc "new_description" ] —————▶▶
▶ [ -addvol volume [...] ] [ -rmvol volume [...] ] —————▶▶
▶ [ -loc new_pool_location ] [ -quiet ] — pool_name —————▶▶
```

Parameters

-desc *"new_description"*

An optional parameter that specifies the new user-defined description of the pool. The maximum alphanumeric character length for this value is 250.

-addvol *volume [...]*

An optional parameter that specifies new volumes to add to the pool specified. All volumes must be of the same device type and cannot belong to multiple pools.

-rmvol *volume [...]*

An optional parameter that specifies volumes that are to be removed from the pool specified.

-loc *new_pool_location*

An optional parameter that specifies the location of the pool.

-quiet

An optional parameter that turns off the confirmation prompt for this command.

pool_name

Specifies the existing name of the pool to be changed.

Description

You can use the **chtgtpool** command to change description, name, or volume membership of an existing Replication Manager pool.

Return values

Pool *pool_name* successfully modified.

Possible failures

- Pool *pool_name* does not exist.
- Pool *pool_name* already exists.
- Volume *volume* does not currently belong to pool *pool_name*.
- Volume *volume* already belongs to pool *pool_name*. You can only add volumes that do not already belong to the pool specified.
- Volume *volume* already belongs to pool *pool_name*. A volume cannot belong to more than one pool.
- Volume *volume* already belongs to pool *pool_name*. A volume cannot belong to more than one pool.
- Volumes are not all of the same device type.

Examples

An invocation example:

```
repcli chtgtpool -loc newlocation mypool
```

The resulting output:

```
Pool mypool successfully modified.
```

Related topics

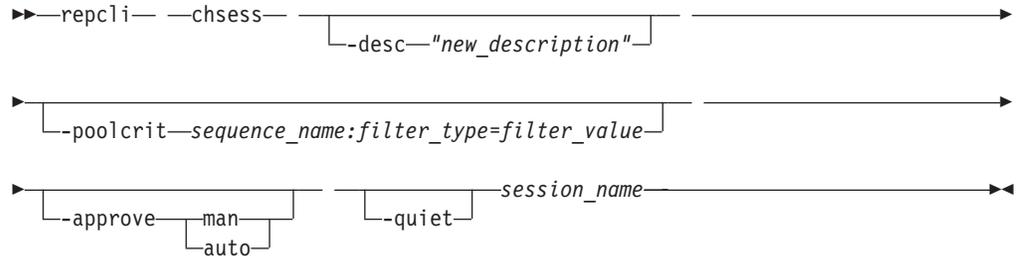
- “mktgtpool” on page 114
- “rmtgtpool” on page 122
- “lstgtpool” on page 104
- “showtgtpool” on page 130
- “mksess” on page 110
- “lssess” on page 102
- “mkgrp” on page 108
- “lsgrp” on page 94
- “lsvol” on page 106

chsess

Use the **chsess** to change the description or pool information of an existing Replication Manager session.

- The session name must be unique; it cannot have the same name as another Replication Manager session.
- To change the copy type or group information, you must delete the session and create a new one.

Syntax



Parameters

-desc *"new_description"*

An optional parameter that specifies the new user-defined description of the session. The maximum alphanumeric character length for this value is 250.

-poolcrit *sequence_name:filter_type=filter_value*

Optionally specifies the target Replication Manager pool(s) based on sequence name and filter value. Sequence name and filter type are separated from each other using a colon.

-approve **man** | **auto**

An optional parameter that specifies whether approval is manual or automatic. Specify **auto** to set the session so that all copy sets that are created from that point, using the automated target volume selection, are marked as approved. Specify **man** to go through an additional approval step.

-quiet

An optional parameter that turns off the confirmation prompt for this command.

session_name

Specifies the existing session name to be changed.

Description

You can use the **chsess** to change the description or pool information of an existing Replication Manager session.

Return values

Session session_name

Successfully modified. Pool changes will only apply to future target selections. (Returned only if **-tgtpool** is used.)

Possible failures

- Nothing to modify.
- Session *session_name* does not exist.
- Session *session_name* already exists.
- Pool *pool_name* already belongs to session *session_name*.
- Pool *pool_name* does not currently belong to session *session_name*.
- Pools are not all of the same device type.

Examples

An invocation example:

```
repli chsess -approve man mysession
```

The resulting output:

```
Session mysession successfully modified.
```

Related topics

- “mksess” on page 110
- “lssess” on page 102
- “showsess” on page 128
- “rmsess” on page 121
- “chgrp” on page 80
- “chtgtpool” on page 83

exit

Use the exit command to exit the repcli interface.

Syntax

►►—repcli— —exit—◄◄

Description

Use the exit command to exit a session.

Examples

An invocation example:

```
repcli exit
```

The resulting output:

```
Are you sure you want to exit? y/n  
Y
```

Related topics

- “repcli” on page 74

flashsess

You can use the **flashsess** command at any point during the life of a session, once that session has been defined, to start FlashCopy® sessions.

Use the **startsess** command to start peer-to-peer remote copy (PPRC) sessions.

Syntax

```
►►—repli— —flashsess— —session_name— [ . . . ] — [ -quiet ] —◀◀
```

Parameters

session_name [. . .] | -

Specifies the session name to be activated. Separate multiple session names with a white space between each name. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin).

-quiet

An optional parameter that turns off the confirmation prompt for this command.

Description

You can use the **flashsess** command to start a flash copy session.

Return values

Session *session_name* successfully started or activated.

Possible failures

- Session *session_name* does not exist.

Examples

An invocation example:

```
repli flashsess mysession
```

The resulting output:

```
Session mysession successfully started.
```

Related topics

- “chsess” on page 85
- “lssess” on page 102
- “showsess” on page 128
- “rmsess” on page 121
- “startsess” on page 131
- “stopsess” on page 133
- “suspendsess” on page 134

generatecpset

Use the **generatecpset** command to automatically generate copy sets in a session. After generating a copy set, use the **lscpset** command to list the copy sets that were generated.

Syntax

```
►—repli— —generatecpset—session_name—◄
```

Parameters

session_name

Specifies the session name that the copy sets are being generated from.

Description

You can use the **generatecpset** command to automatically generate copy sets in a session. After generating a copy set, use the **lscpset** command to list the copy sets that were generated.

Return values

Possible failures

- Nothing to modify.
- Session *session_name* does not exist.

Examples

An invocation example:

```
repli generatecpset session1
```

The resulting output:

```
Copyset successfully added to session1.
```

Related topics

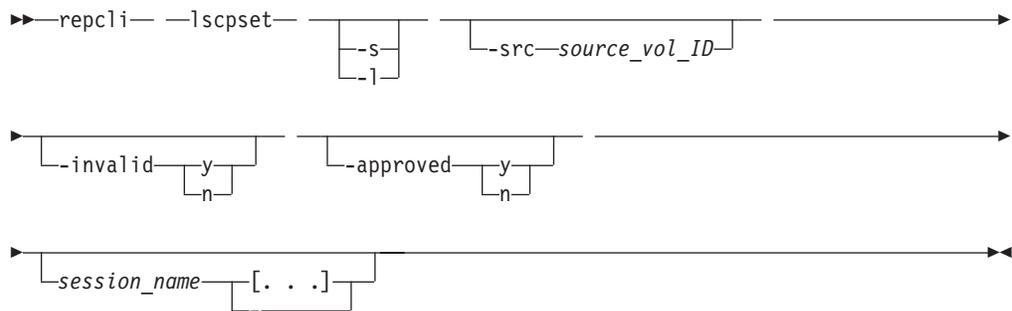
- “mkcpset” on page 112
- “rmcpset” on page 119
- “lscpset” on page 90
- “showcpset” on page 123

Iscpset

Use the **Iscpset** command to list the source and targets of the copy service pairs (copy sets) and their status.

- To see volumes in a copy set, use the **showcopyset** command.
- To see status of volumes in a copy set, use the **lsvol** command.

Syntax



Parameters

- s** An optional parameter that displays copy set name only.
- l** An optional parameter that displays the copy set name and other valid output.
- src** *source_vol_ID*
An optional parameter that displays only the copy sets of the source volumes supplied.
- invalid** **y | n**
An optional parameter that displays only the copy sets that are in a valid state.
- approved** **y | n**
An optional parameter that displays only the copy sets that are in an approved state.
- session_name* **[. . .]** | -
Specifies the session name to which a list of copy sets belongs. Separate multiple session names with a comma between each name. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin).

Description

You can use the **Iscpset** command to list the source and targets of the copy set pairs and their status.

Return values

One of the following types of information is returned.

- No copysets match criteria specified.
-

Column Label	Details
Source Volume	Name of the source volume.

Column Label	Details
Session	Name of the session.
Vols (#)	Number of volumes associated with the copy set.
Invalid	Yes or no to specify whether or not the copy set is valid.
Approved	Yes or no to specify whether or not the copy set is approved.
Last message issued	Last message issued.

Possible failures

- Source volume ID *source_vol_ID* does not exist.
- No copy sets match the criteria specified.
- Session *session_name* does not exist.

Examples

An invocation example:

```
repcli lscpset -s session_name sessionx
```

The resulting output:

```
Sessions matching the criteria specified were not found.
Source Volume
Session
```

Related topics

- “mkcpset” on page 112
- “chcpset” on page 78
- “rmcpset” on page 119
- “approvecpset” on page 77

lsdev

Use the **lsdev** command to display a summary of supported storage devices.

Syntax

```
→ repcli lsdev [-s] [-l] [-devtype ess] [-ip ipaddress]
→ [id [. . .]]
```

Parameters

- s** An optional parameter that displays only the device ID.
- l** An optional parameter that displays the device ID and the device type.
- devtype ess**
An optional parameter that displays information for the device type that is specified. .
- ip ipaddress**
An optional parameter that specifies the device IP address.
- id [...]**
An optional parameter that displays information only for the device IDs that are specified. The device ID is the nickname or model-serial-manufacturer.

Description

You can use the **lsdev** command to display a summary of supported storage devices.

Return values

The following information is listed for each device:

Column Label	Details
Device ID	The nickname or model-serial-manufacturer of the device.
IP	Device IP address.
Device Type	Device Type: ESS.

Possible failures

- Device *device_id* not found.
- No device matched the criteria specified.

Examples

An invocation example:

```
repcli lsdev -devtype ess
```

The resulting output:

Device ID	Device IP	Device Type
2105-123123-IBM	9.11.222.10	ESS
2105-123124-IBM	9.11.223.11	ESS
2105-123125-IBM	9.11.224.12	ESS

lsgrp

Use the **lsgrp** command to display Replication Manager groups and their associations. To see volumes in a group, use the **showgrp** command.

Syntax

```
▶▶ repcli -lsgrp [-s] [-l] group_name [. . .]
```

Parameters

- s An optional parameter that displays only the group name.
 - l An optional parameter that displays group name and other valid output.
- group_name* [. . .] | -
An optional parameter that displays only the groups with the name specified. Separate multiple group names with a comma between each name. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin).

Description

You can use the **lsgrp** command to display Replication Manager groups and their associations.

Return values

No groups match criteria specified.

For each group, the following information is listed:

Column Label	Details
Name	User-defined name of the group.
Volumes	Number of volumes included in the group.
Description	User-defined description of the group.

Possible failures

- Group *group_name* does not exist.

Examples

An invocation example:

```
repcli lsgrp -l -cptype flash groupA
```

The resulting output:

```
Name      Volumes      Description
=====
groupA    24           flashcopy-volumes-for-ESS
```

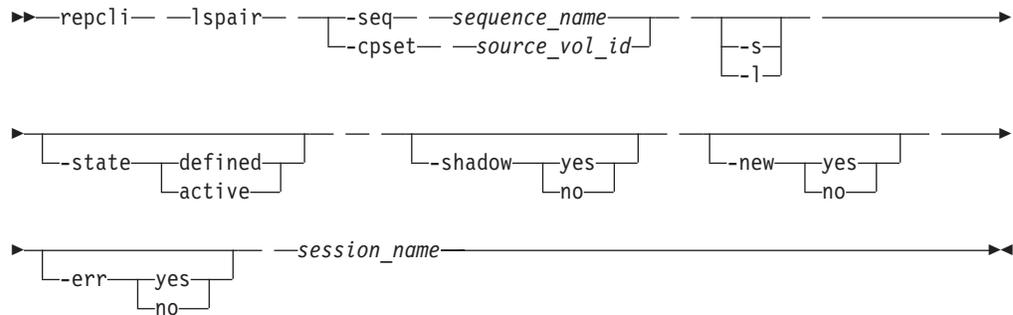
Related topics

- “mkgrp” on page 108
- “chgrp” on page 80
- “rmgrp” on page 120
- “showgrp” on page 126
- “lsvol” on page 106

Ispair

Use the **Ispair** command to list the source and target of the copy service pairs and their status.

Syntax



Parameters

- seq** *sequence_name*
Displays only pairs of the sequence name specified. Mutually exclusive with **-cpset**.
 - cpset** *source_vol_id*
Specifies the source volume ID of the copysset on which you want a list of pairs. Mutually exclusive with **-seq**.
 - s** An optional parameter that displays information about only pairs.
 - l** An optional parameter that displays the default output, including pairs.
 - state** **defined** | **active**
An optional parameter that shows the state as **defined** or **active**.
 - shadow** **yes** | **no**
An optional parameter that displays only pairs that are in the new state.
 - new** **yes** | **no**
An optional parameter that displays only pairs that are in the new state specified.
 - err** **yes** | **no**
An optional parameter that displays only pairs that are in the error state.
- session_name*
The session name by which the pairs are identified.

Description

You can use the **Ispair** command to list the source and target of the copy service pairs and their status.

Return values

For each pair, the following information is listed:

Column Label	Details
SourceID	Volume ID of source volume. Displayed in short output along with Target Volume ID
TargetID	Volume ID of target volume. Displayed in short output along with Source Volume ID.
State	Defined or Active.
Recov	Yes or No to indicate if this pair can be considered recoverable.
Shadow	Yes or No to indicate if this pair can be considered shadowing.
New	Yes or No to indicate if this pair can be considered a new pair.
Copyset	Source volume id of the copyset that the pair is associated with.
Timestamp	Date and time when the volume was suspended.
Last result	The last message issued for this pair. If message ends in <i>E</i> or <i>W</i> then the pair is considered an exception pair

Possible failures

- Sequence name specified unknown.
- Session name specified unknown.

Examples

An invocation example:

```
repli lspair -s
```

The resulting output:

```
SourceID                TargetID
=====
ESS:2105.65312:VOL:2002  ESS:2105.65312:VOL:2103
```

Related topics

- “mkcpset” on page 112
- “chcpset” on page 78
- “rmcpset” on page 119
- “approvecpset” on page 77
- “lscpset” on page 90

lspath

Use the **lspath** command to display paths between Enterprise Storage Servers (ESSs). You can then use this data for a remote copy.

To create, remove, or restore ESS paths, use the ESS Specialist interface.

Syntax

```
repcli lspath [-l] [-s] [-devtype ess]
```

Parameters

- s An optional parameter that displays the target logical subsystem (LSS) or cluster, and source LSS or cluster only.
- l An optional parameter that displays the target LSS or cluster, and source LSS or cluster, and other valid output.
- devtype ess
An optional parameter that displays information about the device type.

Description

You can use the **lspath** command to display paths between Enterprise Storage Servers (ESSs) clusters. You can then use this data for a remote copy. To create, remove, or restore ESS paths, use the ESS Specialist interface.

Return values

No paths match criteria specified.

The following information is listed:

Column Label	Details
Dev	The nickname or model-serial-manufacturer (short output)
Source	Origin of the path. For ESS this is an LSS.
Target	Target of the path. For ESS this is an LSS.
Type	ESCON® or Fibre Channel
Status	Established or Suspended

Possible failures

.

Examples

An invocation example:

```
repcli lspath -devtype ess
```

The resulting output:

No paths match criteria specified.
Storage subsystem
Source
Target
Type
Status

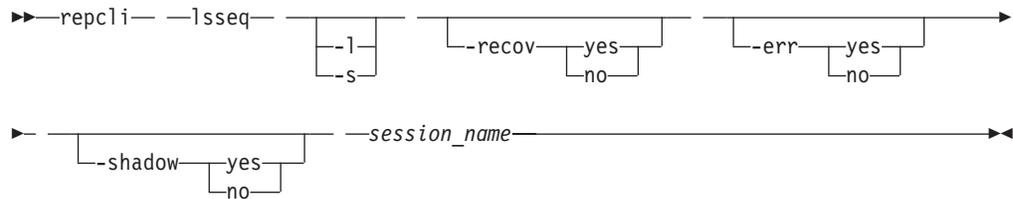
Related topics

- “mksess” on page 110
- “mktgtpool” on page 114
- “lsvol” on page 106

Isseq

You can use the **Isseq** command to display sequences in a Replication Manager session.

Syntax



Parameters

- s** An optional parameter that displays volumes only.
 - l** An optional parameter that displays all valid output. This is the default.
 - recov yes | no**
An optional parameter that indicates whether any sequences in the session can be considered recoverable.
 - shadow yes | no**
An optional parameter that indicates whether or not the sequence is shadowing (copying) the data.
 - err yes | no**
An optional parameter that shows only sessions that have errors or no errors.
- session_name*
Specifies the session name to be activated.

Description

You can use the **Isseq** command to display sequences in a Replication Manager session.

Return values

No sequence match criteria specified.

For each sequence, the following information is listed in either tabular or in a multi-row table format:

Column Label	Details
Name	System generated text string used to identify a sequence.
Recov	Yes or No.
Err	Yes or No.
Shadow	Yes or No.
Err Vols	Integer value. Total number of volumes in an exception state.
Recov Pairs	Integer value.

Column Label	Details
Shadow pairs	Integer value.
Total Pairs	Integer value.
Recov Timestamp	Indicates to what time session is recoverable. Includes both date and time. For a point-in-time copy, this is the time when the copy was taken. For a continuous, synchronous, remote copy, this is the time at which the freeze-and-run was issued. This field is blank if recov=no .

Possible failures

Examples

An invocation example:

```
repli lsseq -s session2
```

The resulting output:

```
Volumes                timestamp
=====
ESS:VOL:2105.25551:11:12  2003:12:13 09:11:24
```

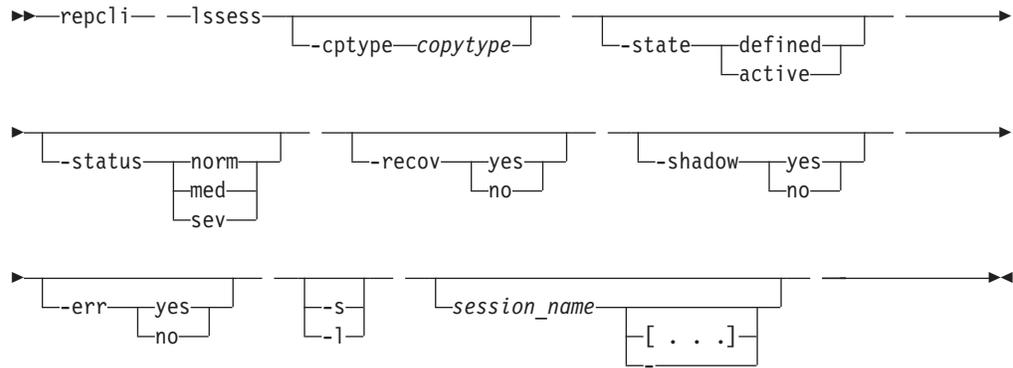
Related topics

- “lscpset” on page 90
- “approvecpset” on page 77

Issess

Use the **Issess** command to display Replication Manager sessions and their associated status. To see properties such as group, pool, and copy type for a specific session, use the **showsess** command.

Syntax



Parameters

-s An optional parameter that displays only the session name.

-l An optional parameter that displays all valid output.

cptype *copytype*

An optional parameter that displays only the sessions with the copy type specified.

-recov **yes** | **no**

An optional parameter that is set to **yes** or **no** to indicate whether the session can be considered recoverable, based on whether any sequences in the session can be considered recoverable.

-shadow **yes** | **no**

An optional parameter that indicates whether any part of the session is shadowing data.

-err **yes** | **no**

An optional parameter that shows only sessions that have errors or no errors.

session_name [. . .] | -

An optional parameter that displays only the sessions with the session name specified. Separate multiple session names with a comma between each name. If no session name is specified, all sessions are displayed unless another filter is used.

Description

You can use the **Issess** command to display Replication Manager sessions and their associated status.

Return values

For each session, the following information is listed:

Column Label	Cell value	Details
Name	Sessions Name	User-defined name of the session.
Status	Status Level	Unknown, Normal, Low, Medium, Severe, Fatal
State	State the session is in	Defined and Active
Group	User defined Group name	Group Name
Type	[pprc, flash]	Copytype
Recover	[yes, no]	Is Recoverable
Shadow	[yes, no]	Is shadowed
Err	[yes, no]	Has errors

Possible failures

- Session *session_name* does not exist.

Examples

An invocation example:

```
replcli lssess -l
```

The resulting output:

```
Name      Type  Status  State  Group  Recov  Err  Shadow  Timestamp
=====
Session1  flash normal  active GroupA yes   no   no   2003:12:13 09:11:24
Session2  pprc  normal  active GroupB yes   no   no   2003:12:13 09:11:25
```

Related topics

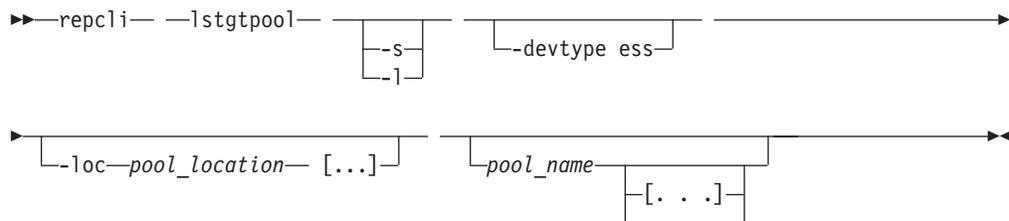
- “mksess” on page 110
- “chsess” on page 85
- “rmsess” on page 121
- “showsess” on page 128
- “lsvol” on page 106

Istgtpool

Use the **lstgtpool** command to display Replication Manager pools and pool memberships.

- To see volumes in a pool, use the **showtgtpool** command.
- To see status of volumes in a pool, use the **lsvol** command.

Syntax



Parameters

- s An optional parameter that displays only the pools.
- l An optional parameter that displays pools, location, description, and number of volumes.
- devtype *ess*
An optional parameter that Filters by device type.
- loc *pool_location* [...]
An optional parameter that displays only the pools with the location type specified. The maximum value is 250 alphanumeric characters. Separate multiple location names with a comma between each location
- pool_name* [...] | -
An optional parameter that displays only the pools with name specified. Separate multiple pool names with a comma between each name. If no pool name is specified, all pools are displayed unless another filter is used. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin).

Description

You can use the **lstgtpool** command to display Replication Manager pools and pool memberships.

Return values

One of the following types of output is returned:

- No pools match criteria specified.
- For each pool or pool membership, the following default information is listed as either tabular or in a multi-row table:

Column Label	Details
Name	Pool name assigned by user when pool was created
Dev Type	Device type

Column Label	Details
Location	Site location of the pool

- For each pool or pool membership, the following long output is listed in either tabular or in a multi-row table format:

Column Label	Details
Name	Pool name assigned by user when pool was created
Dev Type	Device type
Location	Site location of the pool
Description	Description entered by the user when pool was created

- For each pool or pool membership, the short output includes only **Name**.

Possible failures

- Pool *pool_name* does not exist.
- Invalid location.

Examples

An invocation example:

```
repli lstgtpool -l
```

The resulting output:

```
Name   Dev type  Location  Desc
=====
Pool1   svc       loc3     svc_pool
Pool2   svc       loc5     svc_pool
```

Related topics

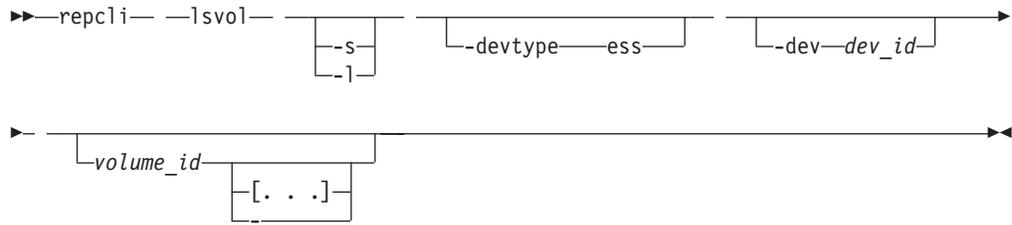
- “mktgtpool” on page 114
- “chtgtpool” on page 83
- “rmtgtpool” on page 122
- “showtgtpool” on page 130
- “lssess” on page 102
- “lsvol” on page 106

lsvol

Use the **lsvol** command to display volume information and status. This information can assist you in choosing available volumes for groups and pools, and determine which volumes are in a suspend or error state. The same volume can reside in multiple groups, but not multiple pools.

To create ESS logical unit numbers (LUN) and create Virtual logical unit numbers, you must use ICAT interfaces.

Syntax



Parameters

- s** An optional parameter that displays only volumes.
- l** An optional parameter that displays all valid output for **lsvol** command.
- devtype ess**
An optional parameter that allows a filter by device type. Specify **ess** for Enterprise Storage Server (ESS).
- dev dev_id**
An optional parameter that displays the nickname or model, serial number, and manufacturer.
- volume_id [...] | -**
Specifies the volume ID for a volume. Volume data is listed for this volume. The same volume can reside in multiple groups but not multiple pools. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin).

Description

You can use the **lsvol** command to display volume information and status to assist in choosing available volumes for groups and pools, and determining which volumes are in a suspend or error state

Return values

For each volume, the following information is listed:

Column Label	Details
ID	Volume ID
Dev	The nickname or model-serial-manufacturer
Dev Type	ESS
LSS/IO Group	Displays LSS ID for ESS device type

Column Label	Details
Vol Type	CKD or FB
Size	Volume size
Unit	Measure that the capacity is given in
Status	Volume status

Possible failures

- No volumes match criteria specified.
- Used invalid flag.

Examples

An invocation example:

```
repcli lsvol -l -devtype ess
```

The resulting output:

```

ID                Dev      Dev Type  LSS/Node Pair  Vol Type  Size  Unit
=====
ESS:2105.65312:VOL:202A 2105.65312 ESS      202A          FB      20000 GB

```

```

Volume Name
Device type
LSS
Location
Format-FB or CKD
Size
Timestamp

```

Related topics

- “mkssess” on page 110
- “lssess” on page 102
- “mkgrp” on page 108
- “lsgrp” on page 94

group_name

Assigns a user-defined group name for the volumes. The group name must be unique; it cannot be the same as another existing Replication Manager group name. The maximum alphanumeric character length for this value is 250.

Description

You can use the **mkgrp** command to create a Replication Manager group of volumes that can be used as target copy relationships.

Return values

return value

Group *group_name* successfully created.

Possible failures

- Group *group_name* already exists.
- Not all volumes specified are of the same device type.

Examples

An invocation example

```
repcli mkgrp -vol ESS:2105.65312:VOL:202A,2C,2E-30 group1
```

The resulting output

```
Group group1 successfully created.
```

Related topics

- “chgrp” on page 80
- “rmgrp” on page 120
- “chsess” on page 85
- “showsess” on page 128

mk sess

Use the **mk sess** to create a Replication Manager session specifying copy relationships among groups.

- For the source group, you cannot specify a group that is already being managed by another session. A group cannot be assigned to more than one session at any given time.
- If no groups exist, use the **mk group** command to create a group.
- If no target pools exist, use the **mk tgt pool** command to create a pool.
- The session is in a stopped state after initial creation. Use the **start sess** or the **flash sess** command to start the copy process.
- Any errors that result from the creation of the session, for example, not enough targets for sources or no path established, are reflected in the **ls copy set** command output.

Syntax

```
►► replcli — mk sess — -cptype copytype — -srcgrp group_name —————►
► -poolcrit sequence_name:filter_type=filter_value —————►
└─approve ┌man
           └auto
► ┌-desc "description" ─── session_name ───────────────────────────►
```

Parameters

-cptype *pprc | flash*

Specifies the copy type for the session. Specify either **flash** for FlashCopy® or **pprc** for continuous, synchronous, remote copy.

-srcgrp *group_name*

Specifies the source group for the session. Choose only one group per session. Do not specify a group that is being managed by another session.

-poolcrit *seq_name:filter_type=filter_value*

Specifies the criteria for target Replication Manager pools based on sequence name and filter value. The sequence name is *Remote Target* for PPRC or *Flashed Target* for FlashCopy. The filter type is the location. Use a colon to separate the sequence name, filter type, and filter. Specify **loc** as the filter value.

-approve **man** | **auto**

An optional parameter that validates the target. Specify **man** for manual; specify **auto** for automatic. The default is automatic.

-desc *"description"*

An optional parameter that assigns a user-defined description of the Replication Manager group. The maximum alphanumeric character length is 250.

session_name

Assigns a user-defined session name. The session name must be unique; it cannot be the same as another existing Replication Manager session name. The maximum alphanumeric character length is 250.

Description

You can use the `mkssess` command to create a Replication Manager session specifying copy relationships among groups.

Return values

Session *session_name* successfully created.

Possible failures

- Session *session_name* already exists.
- Group *group_name* does not exist.
- Pool *pool_name* does not exist.
- Group *group_name* is already managed in another session.
- Incorrect user name or password.
- Incorrect port.
- Target location is not valid or incorrect.
- Target explicit pool not valid or incorrect.

Examples

An invocation example:

```
repcli> mkssess -cptype pprc -srcgrp mygroup -poolcrit "Remote Target:location=rtp" -approve auto mysession
```

The resulting output:

```
session mysession successfully created
```

Related topics

- “`chsess`” on page 85
- “`lssess`” on page 102
- “`showsess`” on page 128
- “`rmsess`” on page 121
- “`mkgrp`” on page 108
- “`mktgtpool`” on page 114

mkcpset

You can use **mkcpset** to create a Replication Manager copy set when you want to define your own source-to-target mapping. Use the **lsvol** command to get a list of volumes for the copy sets. (If you use the **mktgtpool** command and then the **mksess** command, the copy sets are created automatically.) You can only add volumes of the same device type.

Syntax

```
▶▶ repcli — mkcpset — —src—source_volume_id— —————▶▶
▶▶ —tgt—target_volume_id—[...]— —session_name————▶▶
```

Parameters

-src *source_volume_id*
Specifies the source volume of the copy set that is being created.

-tgt *target_volume_id* [...]
Specifies the target volumes for the copy set that is being created.

session_name
Specifies the session name to which the copy set will belong.

Description

You can use **mkcpset** to create a Replication Manager copy set when you want to define your own source-to-target mapping. Use the **lsvol** command to get a list of volumes for the copy sets. (If you use the **mktgtpool** command and then the **mksess** command, the copy sets are created automatically.) You can only add volumes of the same device type.

Return values

Possible failures

- Nothing to modify.
- Volume *volume* already belongs to session *session_name*.
- Volume *volume* does not exist .
- Volumes are not all of the same device type.

Examples

An invocation example:

```
repcli mkcpset -src ESS:2105.65302:VOL:1610
-tgt ESS:2105.65303:VOL:1611 session1
```

The resulting output:

```
Copyset successfully added to session1.
```

Related topics

- “lsvol” on page 106
- “chcpset” on page 78
- “rmcpset” on page 119

- “approvecpset” on page 77
- “lscpset” on page 90
- “showcpset” on page 123
- “lsseq” on page 100

mktgtpool

Use the **mktgtpool** command to create a Replication Manager pool of volumes that can be used in source copy relationships.

- Pool names must be unique across Replication Manager groups; two pools cannot have the same name.
- A volume can reside in only one Replication Manager pool at a time.

Syntax

```
►► repcli — mktgtpool — —vol—vol_id— [...] — [—loc—location—] —►►  
[—desc—"description"—] —pool_name—►►
```

Parameters

-vol *vol_id* [...]

Specifies the volumes to be included in the pool membership for the session. All volumes must be of the same device type

-desc "*description*"

An optional parameter that assigns a user-defined description of the Replication Manager pool. The maximum alphanumeric character length is 250.

-loc *location*

An optional parameter that specifies the location of the pool.

pool_name

Assigns a user-defined pool name for the volumes. The pool name must be unique; it cannot be the same as another existing Replication Manager pool name. The maximum alphanumeric character length is 250.

Description

You can use the **mktgtpool** command to create a Replication Manager pool of volumes that can be used in source copy relationships.

Return values

Pool *pool_name* successfully created.

Possible failures

- Pool *pool_name* already exists.
- Not all volumes specified are of the same device type.

Examples

An invocation example

```
repcli mktgtpool -vol ESS:2105.65312:VOL:202A,2C,2E-30 pool1
```

The resulting output

```
Pool pool1 successfully created.
```

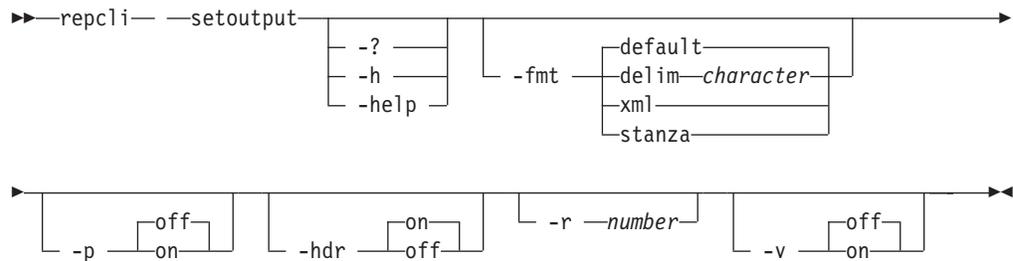
Related topics

- “chtgtpool” on page 83
- “chsess” on page 85
- “showsess” on page 128
- “rmtgtpool” on page 122

setoutput

You can use the **setoutput** command to display current output settings for both perfcli and repli commands. The output format set by this command remains in effect for the duration of a command session or until the options are reset.

Syntax



Parameters

-? | -h | -help

Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

-fmt

Specifies the format of the output. You can specify one of the following values:

default

Specifies to display output in a tabular format using spaces as the delimiter between the columns. This is the default value. For example:

```
Device ID      Device Type Threshold Name      Status  Warning  Error  Alert
=====
ESS2105.123123.IBM  ESS      Disk Utilization  Enabled  50      80     None
```

delim *character*

Specifies to display output in a tabular format using the specified character to separate the columns. If you use a shell metacharacter (for example, * or \t) as the delimiting character, enclose the character in single quotation mark (') or double quotation mark ("). A blank space is not a valid character. For example:

```
Device ID,Device Type,Threshold Name,Status,Warning,Error,Alert
=====
ESS2105.123123.IBM,ESS,Disk Utilization,Enabled,50,80,None
```

xml

Specifies to display output using XML format, for example:

```
<IRETURNVALUE>
<INSTANCE CLASSNAME="lsthresh">
<PROPERTY NAME="Device ID" TYPE="string"><VALUE>
  ESS2105.123123.IBM</VALUE>
</PROPERTY>
<PROPERTY NAME="Device Type" TYPE="uint32"><VALUE>ESS</VALUE>
</PROPERTY>
<PROPERTY NAME="Threshold Name" TYPE="uint64"><VALUE>Disk
  Utilization</VALUE>
</PROPERTY>
<PROPERTY NAME="Status" TYPE="uint16"><VALUE>Enabled</VALUE>
</PROPERTY>
<PROPERTY NAME="Warning" TYPE="uint64"><VALUE>d50</VALUE>
</PROPERTY>
```

```

<PROPERTY NAME="Error" TYPE="uint64"><VALUE>80</VALUE>
</PROPERTY>
<PROPERTY NAME="Alert" TYPE="uint16"><VALUE>None</VALUE>
</PROPERTY>
</INSTANCE>
</IRETURNVALUE>

```

stanza Specifies to display output using XML format, for example:

```

Device ID      ESS2105.123123.IBM
Device Type    ESS
Threshold Name Disk Utilization
Status         Enabled
Warning        50
Error          80
Alert          None

```

-p Specifies whether to display one page of text at a time or all text at once.

off Displays all text at one time. This is the default value when the **perftool** command is run in single-shot mode.

on Displays one page of text at time. Pressing any key displays the next page. This is the default value when the **perftool** command is run in interactive mode.

-hdr

Specifies whether to display the table header.

on Displays the table header. This is the default value.

off Does not display the table header.

-r number

Specifies the number of rows per page to display when the **-p** parameter is on. The default is 24 rows. You can specify a value from 1 to 100.

-v Specifies whether to enable verbose mode.

off Disables verbose mode. This is the default value.

on Enables verbose mode.

Description

You can use the **setoutput** command to display current output settings for both **perfcli** and **repli** commands. The output format set by this command remains in effect for the duration of a command session or until the options are reset either by using this command or by specifying an output-format parameter as part of a command.

Note: The output formats do not apply to help pages.

Return values

repli setoutput

Paging	Rows	Format	Header	Verbose
off	-	default	on	off

Examples

The following **repli** example sets the output format to display in tabular form using a comma as the delimiter without header information.

```
repli>setoutput -fmt delim , -hdr off
```

Related topics

rmcpset

Use the **rmcpset** command to delete an existing Replication Manager copy set.

Syntax

```
► repli — rmcpset -quiet -src source_volume_id [. . .]
-session_name ►
```

Parameters

-quiet

An optional parameter that turns off the confirmation prompt for this command.

-src *source_volume_id* [. . .] | -

Specifies the name of the source volume ID of the copy set to delete. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin).

session_name

Specifies the session name from which the copy set is being deleted.

Description

You can use the **rmcpset** command to delete an existing Replication Manager copy set.

Return values

Copy set *copyset_name* successfully deleted.

Possible failures

- Copy set *copyset_name* does not exist.

Examples

An invocation example:

```
repli rmcpset copyset1 session1
```

```
Are you sure you want to delete copyset1? y/n
```

```
y
```

```
Copyset copyset1 successfully deleted.
```

Related topics

- “mkcpset” on page 112
- “chcpset” on page 78
- “approvecpset” on page 77
- “lscpset” on page 90
- “showcpset” on page 123
- “lsseq” on page 100

rmgrp

Use the **rmgrp** command to delete an existing Replication Manager group.

Syntax

```
►► repcli — rmgrp — -quiet — group_name — [...] ►►
```

Parameters

-quiet

An optional parameter that turns off the confirmation prompt for this command.

group_name [...] | -

Specifies the name of the group to delete. Separate multiple group names with a white space between each name. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin).

Description

You can use the **rmgrp** command to delete an existing Replication Manager group.

Return values

Group *group_name* successfully deleted.

Possible failures

- Group *group_name* does not exist.

Examples

An invocation example:

```
repcli rmgrp group1
```

The resulting output:

```
Are you sure you want to delete group group1? Y/N
```

```
Y  
Group group1 successfully deleted.
```

Related topics

- “mkgrp” on page 108
- “chgrp” on page 80
- “lsgrp” on page 94
- “showgrp” on page 126
- “rmsess” on page 121
- “stopsess” on page 133

rmssess

Use the **rmssess** command to delete an existing Replication Manager session.

Syntax

```
►► repcli — rmssess — -quiet — session_name — [. . .] ►►
```

Parameters

-quiet

An optional parameter that turns off the confirmation prompt for this command.

session_name [. . .] | -

Specifies the session name to delete. Separate multiple session names with a white space between each name. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin).

Description

You can use the **rmssess** command to delete an existing Replication Manager session.

Return values

Session *session_name* successfully deleted.

Possible failures

- Session *session_name* does not exist.
- Cannot delete a session currently in progress. Stop the session before attempting to delete it.

Examples

An invocation example:

```
repcli rmssess session2
```

The resulting output:

```
Are you sure you want to delete session session2? Y/N  
Y  
Session session2 successfully deleted.
```

Related topics

- “chsess” on page 85
- “lssess” on page 102
- “showsess” on page 128
- “stopsess” on page 133

rmtgtpool

Use the **rmtgtpool** command to delete an existing Replication Manager pool.

Syntax

```
► repcli — rmtgtpool -quiet pool_name [...] ◀
```

Parameters

-quiet

An optional parameter that turns off the confirmation prompt for this command.

pool_name [...] | -

Specifies the pool name to delete. Separate multiple pool names with a white space between each name. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin).

Description

You can use the **rmtgtpool** command to delete an existing Replication Manager pool.

Return values

Pool *pool_name* successfully deleted.

Possible failures

- Pool *pool_name* does not exist.

Examples

An invocation example:

```
repcli rmtgtpool pool2
```

The resulting output:

```
Are you sure you want to delete pool pool2? Y/N  
Y  
Pool pool2 successfully deleted.
```

Related topics

- “mktgtpool” on page 114
- “chtgtpool” on page 83
- “lstgtpool” on page 104
- “showtgtpool” on page 130
- “rmsess” on page 121
- “stopsess” on page 133

showcpset

Use the **showcpset** command to display properties for a copy set.

Syntax

```
►—repli— —showcpset— —-src—source_volume_id— —session_name—►
```

Parameters

-src *source_volume_id*

Specifies the name of the source volume ID for a display of its properties.

session_name

Specifies the session name to which the copy set belongs.

Description

You can use the **showcpset** command to display properties for a copy set.

Return values

For each copy set, the following information is listed:

Column Label	Details
Src Vol	Source volume name
Session	Session name
Volumes	Volumes associated with the copy set
Invalid	Whether or not it is invalid
Approved	Whether the copy set is approved
Last Message Issued	Last message issued

Possible failures

- Source volume ID *source_vol_id* does not exist.

Examples

An invocation example:

```
repli showcpset -src ESS:2105.22220:VOL:3010 session6
```

The resulting output:

```
Name           Session      Volumes           Status  Description
=====
ESS:2105.22220:VOL:3010 session6  ESS:2105.22220:VOL:3010  Approved
```

Related topics

- “mkcpset” on page 112
- “chcpset” on page 78
- “rmcpset” on page 119
- “approvecpset” on page 77
- “lscpset” on page 90

- “lsseq” on page 100

showdev

Use the **showdev** command to display device properties.

Syntax

```
▶▶—repli— —showdev— —-id [...]—▶▶
```

Parameters

-id [...]

A required parameter that displays the threshold settings for the device IDs that are specified. The device ID is the nickname or model-serial-manufacturer. Separate multiple device names with a comma in between each name with no white space before or after the comma.

Description

You can use the **showdev** command to display device properties.

Return values

One of the following types of information is returned:

- No device matched the criteria specified.
- For each device, the following information is listed:

Column Label	Details
Device ID	The nickname or model-serial-manufacturer of the device.
IP	Device IP address.
Device Type	Device Type: ESS.

Possible failures

- Device *device-id* not found.

Examples

An invocation example:

```
repli showdev 2105-2105.26884-IBM
```

The resulting output:

```
Device ID           Device IP   Device Type  
=====
```

2105-2105.26884-IBM 9.11.222.10 ESS

showgrp

Use the **showgrp** command to display properties of Replication Manager groups.

Syntax

```
▶▶reptool —showgrp —group_name [ . . . ]▶▶
```

Parameters

group_name [. . .] | -

Displays only the groups with the name that is specified. Separate multiple group names with a white space between each name. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin). If no group name is specified, all groups are displayed unless another filter is used.

Description

You can use the **showgrp** command to display properties of Replication Manager groups.

Return values

Either the following message or a table is returned.

- No groups match criteria specified.
- For each group, the following information is listed:

Column Label	Details
Name	User-defined name of the group.
Vols	List of Volumes in the group.
Desc	User-defined description of the group.

Possible failures

- Group *group_name* does not exist.

Examples

An invocation example:

```
repli> showgrp test5
```

The resulting output:

```
Name          Volumes          Description
=====
test5          ESS:2105.65312:VOL:202A  Test Group 5
```

Related topics

- “mkgrp” on page 108
- “chgrp” on page 80
- “rmgrp” on page 120
- “lsgrp” on page 94

- “lsvol” on page 106
- “lssess” on page 102

showsess

Use the **showsess** command to display properties for a selected session, showing the session name, description, group managed, pools used, and copy type.

To see a list of volumes in a group, use the **showgrp** command. To see a list of volumes in a pool, use the **showtgtpool** command.

Syntax

```
►►—repli— —showsess— —session_name—◄◄
```

Parameters

session_name

Specifies the desired session for display of its properties.

Description

You can use the **showsess** command to display properties for a selected session, showing the session name, description, group managed, pools used, and copy type.

Return values

For each session, the following information is listed:

Column Label	Details
Name	Session Name.
Copy type	Point in Time Copy, Continuous Synchronous Remote Copy.
State	Defined or Active.
Status	Unknown, Normal, Low, Medium, Severe, or Fatal.
Group	Name of group managed
Approval status	Contains [Automatic, Manual]. Displays whichever is selected during session creation.
Pool Criteria	Location.
Shadow	Yes or no.
Recov	Yes or no.
Description	Use-defined session description.

Possible failures

- Standard Syntax Error.
- Session *session_name* does not exist.

Examples

An invocation example:

```
repli> showsess SS
```

The resulting output:

Name	SS
Description	-
Copy Type	flash
Source Group	oscarGroup
Target Pool Criteria	Flashed Target:location=%
Approval	Manual
State	Defined
Status	Normal
Recovery	No
Shadow	No
Error	No
Suspend Type	PointInTime

AWN007080I Command completed successfully.

Related topics

- “mkssess” on page 110
- “chssess” on page 85
- “lssess” on page 102
- “rmssess” on page 121

showtgtpool

Use the **showtgtpool** command to display properties for a selected pool, showing pool name, site location, volumes, and description.

To see the status of volumes in a pool, use the **lsvol** command.

Syntax

```
►►—repli— —showtgtpool— —pool_name—◄◄
```

Parameters

pool_name

Specifies the name of the pool for a display of its properties.

Description

You can use the **showtgtpool** command to display properties for a selected pool, showing pool name, site location, volumes, and description.

Return values

For each pool, the following information is listed:

Column Label	Details
Name	User-defined name of the pool.
Location	User-defined name of the session that manages the pool.
Volumes	List of volumes in the pool.
Description	User-defined description of the pool.

Possible failures

- Pool *pool_name* does not exist.

Examples

An invocation example:

```
repli> showtgtpool testpool7
```

The resulting output:

```
Name      Location  Volumes      Description
=====
test5     session4  ESS:2105.65312:VOL:202A  Test Group 5
```

Related topics

- “mktgtpool” on page 114
- “chtgtpool” on page 83
- “rmtgtpool” on page 122
- “lsvol” on page 106

startsess

Use the **startsess** command at any point during the life of a session once that session has been defined, to start a session. This command is intended for peer-to-peer remote copy (PPRC) sessions.

Use the **flashsess** command to start FlashCopy® sessions.

Syntax

```
►►—repli— —startsess— —session_name— [ . . . ] — [ -quiet ] —►►
```

Parameters

-quiet

An optional parameter that turns off the confirmation prompt for this command.

session_name [. . .] | -

Specifies the session name to be activated. Separate multiple session names with a white space between each name. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin).

Description

You can use the **startsess** command to start a session.

Return values

Session *session_name* successfully started or activated.

Possible failures

- Session *session_name* does not exist.

Examples

An invocation example:

```
repli startsess mysession
```

The resulting output:

```
Session mysession successfully started.
```

Related topics

- “chsess” on page 85
- “lssess” on page 102
- “showsess” on page 128
- “rmsess” on page 121
- “stopsess” on page 133
- “flashsess” on page 88
- “suspendsess” on page 134

stopflashsess

Use the **stopflashsess** command at any point during the life of a flash copy session once that session is in an ACTIVE state. This command does a withdraw of the relationship on the hardware for a flash copy session.

Syntax

```
►► repli — stopflashsess —                    — session_name — [. . .] ◄◄  
                                          -quiet-
```

Parameters

-quiet

An optional parameter that turns off the confirmation prompt for this command.

session_name [. . .] | -

Specifies the session name to be stopped. Separate multiple session names with a white space between each name. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin).

Description

You can use the **stopflashsess** command to stop a flash copy session.

Return values

Session *session_name* successfully stopped.

Possible failures

- Session *session_name* does not exist.
- Session *session_name* is already stopped or inactive.

Examples

An invocation example:

```
repli stopsess mysession
```

The resulting output:

```
Are you sure you want to stop session session_name? Y/N  
y  
Session mysession successfully stopped.
```

Related topics

- “chsess” on page 85
- “lssess” on page 102
- “showsess” on page 128
- “rmsess” on page 121
- “startsess” on page 131
- “suspendsess” on page 134

stopsess

Use the **stopsess** command at any point during the life of a session once that session is in an ACTIVE state. This command does a withdraw of the relationship on the hardware.

Syntax

```
►► repli — stopsess — [-quiet] — session_name — [. . .] ►►
```

Parameters

-quiet

An optional parameter that turns off the confirmation prompt for this command.

session_name [. . .] | -

Specifies the session name to be stopped. Separate multiple session names with a white space between each name. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin).

Description

You can use the **stopsess** command to stop a PPRC session.

Return values

Session *session_name* successfully stopped.

Possible failures

- Session *session_name* does not exist.
- Session *session_name* is already stopped or inactive.

Examples

An invocation example:

```
repli stopsess mysession
```

The resulting output:

```
Are you sure you want to stop session session_name? Y/N  
y  
Session mysession successfully stopped.
```

Related topics

- “chsess” on page 85
- “lssess” on page 102
- “showsess” on page 128
- “rmsess” on page 121
- “startsess” on page 131
- “suspendsess” on page 134

suspendsess

Use the **suspendsess** command to suspend a started session.

To restart a session, use the **startsess** command or the **flashsess** command..

Syntax

```
▶—repli— —suspendsess— —quiet— —type—suspend_type— —————▶  
▶—session_name— [ . . . ] —————▶
```

Parameters

-quiet

An optional parameter that turns off the confirmation prompt for this command.

-type *suspend_type*

Specifies the type of session to suspend. You can choose to freeze a PPRC session by specifying **consist** (for consistent), or stop a session by specifying **immed** (for immediately).

session_name [. . .] | -

Specifies the session name to be suspended. Separate multiple session names with a white space between each name. Alternatively, use the dash (-) to specify that input for this parameter comes from an input stream (stdin).

Description

You can use the **suspendsess** command to suspend a started session.

Return values

Session *session_name* successfully suspended.

Possible failures

- Session *session_name* does not exist.
- Session *session_name* is already stopped or inactive.

Examples

An invocation example:

```
repli suspendsess -type immed mysession
```

The resulting output:

```
Are you sure you want to suspend mysession? Y/N  
y  
Session mysession successfully suspended.
```

Related topics

- “chsess” on page 85
- “lssess” on page 102
- “showsess” on page 128

- “rmssess” on page 121
- “stopssess” on page 133

Chapter 5. 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 or fax

Fill out the Readers' Comments form (RCF) at the back of this book. Return it by mail or fax (1-919-254-0343), or give it to an IBM representative. If the RCF has been removed, you can address your comments to:

International Business Machines Corporation
RCF Processing - Department CGFA
P.O. Box 12195
Research Triangle Park, NC 27709-9990
U.S.A.

Appendix. Notices

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

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

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

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

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

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

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

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

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

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

Legal Information

The topics in this category provide legal information about Multiple Device Manager, including notices and trademarks.

Trademarks

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

- DB2
- DB2 Universal Database
- Enterprise Storage Server
- ESCON
- FlashCopy
- IBM
- TotalStorage
- WebSphere

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

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

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

Index

A

approvecpset command 77

C

chcpset command 78

chgrp command 80

chprofile command 10

chsess command 85

chtgtpool command 83

commands

Performance Manager 7

chprofile 10

cpprofile 13

cpthresh 14

lscollection 17

lsfilter 23

lsgauge 25

lsprofile 9, 30, 76

lsthresh 32

mkprofile 35

mkrecom 38

perfcli 8

rmgauge 42

rmprofile 43

setessthresh 46

setfilter 44

setoutput 49

setsvcthrsh 52

showdbinfo 57

showgauge 58

showprofile 61

startdbpurge 63

stopcollection 69

tracklun 70

Replication Manager

chsess 85

chtgtpool 83

flashsess 88

generatecpset 89

lscpset 90

lsgrp 94

lspair 96

lsseq 100

lssess 102

lsvol 106

mkcpset 112

mkgrp 108

mksess 110

mktgtpool 114

rmcpset 119

rmgrp 120

rmsess 121

rmtgtpool 122

showcpset 123

showdevice 125

showgrp 126

showsess 128

showtgtpool 130

startsess 131

commands (*continued*)

Replication Manager (*continued*)

stopsess 132, 133

suspendsess 134

cpprofile command 13

cpthresh command 14

F

flashsess command 88

G

generatecpset command 89

L

lscollection command 17

lscpset 90

lsdevice command 21, 92

lsfilter command 23

lsgauge command 25

lsgrp command 94

lshost command 27

lspair command 96

lspath command 98

lsprofile command 9, 30, 76

lsseq command 100

lssess command 102

lstgtpool command 104

lsthresh command 32

lsvol command 106

M

MDM (Multiple Device Manager)

overview 1

mkcpset command 112

mkgrp command 108

mkprofile command 35

mkrecom command 38

mksess command 110

mktgtpool command 114

N

notices

legal 139

P

perfcli command

overview 8

R

repcli command 74

rmcpset command 119

rmgauge command 42

rmgrp command 120

rmprofile command 43

rmsess command 121

rmtgtpool command 122

S

setessthresh command 46

setfilter command 44

setoutput command 49, 116

setsvcthrsh command 52

showcpset command 123

showdbinfo command 57

showdevice command 56, 125

showgauge command 58

showgrp command 126

showprofile command 61

showsess command 128

showtgtpool command 130

startdbpurge command 63

startesscollection command 65

startsess command 131

startsvccollection command 67

stopcollection command 69

stopsess command 132, 133

suspendsess command 134

T

tracklun command 70

trademarks 143

Readers' Comments — We'd Like to Hear from You

IBM TotalStorage Multiple Device Manager
Command-Line Interface User's Guide
Version 1.1

Publication No. SC26-7585-00

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>				

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>				
Complete	<input type="checkbox"/>				
Easy to find	<input type="checkbox"/>				
Easy to understand	<input type="checkbox"/>				
Well organized	<input type="checkbox"/>				
Applicable to your tasks	<input type="checkbox"/>				

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? Yes No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name

Address

Company or Organization

Phone No.



Fold and Tape

Please do not staple

Fold and Tape



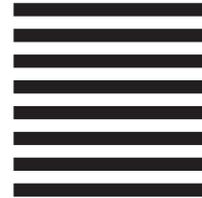
NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Design & Information Development
Department CGF
P.O. Box 12195
Research Triangle Park, NC
USA 27719-9990



Fold and Tape

Please do not staple

Fold and Tape



Printed in USA

SC26-7585-00

