Dell EqualLogic Host Integration Tools for Linux

Version 1.7

Installation and User's Guide

Notes, cautions, and warnings

NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

April 2020

Rev. 110-6312-EN-R1

Contents

Preface

About This Guide

Dell EqualLogic Host Integration Tools for Linux (HIT/Linux) enable you to configure and manage Dell EqualLogic PS Series groups from a Linux host.

Revision History

Document Number 110-6312-EN

Revision	Date	Description
R1	April 2020	Version 1.7 initial release

Audience

The information in this document is intended for administrators responsible for installing and using the HIT/Linux platform. Administrators are not required to have extensive network or storage system experience. However, it is helpful to understand:

- Linux operating system administration, especially disk and network management
- Basic networking concepts
- Current network environment

Related Documentation

For detailed information about PS Series arrays, groups, volumes, array software, and host software, log into the Documentation page at the customer support site:

https://eqlsupport.dell.com/support/resources.aspx?id=2495

NOTE: You must log in to the Dell customer support site to access these pages.

Dell Online Services

To learn more about Dell EqualLogic products and new releases being planned, visit the Dell EqualLogic TechCenter site: Delltechcenter.com/page/EqualLogic. Here you can also see articles, demos, join online discussions, and get more details about the benefits of our product family.

Technical Support and Customer Service

Dell support service is available to answer your questions about PS Series arrays and FS Series appliances.

Contacting Dell

Dell provides several online and telephone-based support and service options. Availability varies by country and product, and some services might not be available in your area. To contact Dell for sales, technical support, or customer service issues, go to Dell.com/support/home.

1

Introduction to HIT/Linux

The Dell EqualLogic Host Integration Tools for Linux (HIT/Linux) provide a collection of applications and utilities to simplify the configuration and administration of Dell EqualLogic PS Series storage arrays. This chapter provides an overview of HIT/Linux features and information about what you can expect when installing, configuring, and using this product.

Topics:

Overview

Dell EqualLogic Host Integration Tools for Linux (HIT/Linux) provides the following features:

- Dell EqualLogic Auto-Snapshot Manager/Linux Edition (ASM/LE)—Creates point-in-time copies of Dell EqualLogic volumes as part of a backup and recovery strategy for your data. From the ASM/LE command-line interface, you can:
  - Create, list, and manage Smart Copies, collections, and schedules
  - Mount and unmount Smart Copies
  - Create and manage group access files to store group names, group IP addresses, optional management IP addresses, user names, and passwords
- Remote Setup Wizard Command Line Interface (RSWCLI)—Use the rswcli setup utility to:
  - Discover and initialize PS Series arrays
  - Configure and manage access to a PS Series group
  - Set configuration parameters for Multipath I/O (MPIO)
  - Manage which subnets are used for MPIO
- Dell EqualLogic iSCSI Multipathing—A combination of user and kernel mode binaries that:
  - Automate iSCSI session management to create and maintain the optimal iSCSI session configuration for high availability and increased performance
  - Improve I/O performance by intelligently routing I/O directly to the correct PS Series group member, thereby reducing intragroup I/O forwarding
- Dell EqualLogic Configuration Utility (eqlconfig)—Lets you interactively configure the components in HIT/Linux, including ASMCLI and MPIO. It is automatically invoked as part of the installation process.
- Dell EqualLogic Host Performance and Configuration Tuning Suite (eqltune)—Verifies that an optimal set of tuning parameters are set for a Linux system, to enhance performance and usability
- Dell EqualLogic Log Gathering Facility (eqllog)—Collects system logs and information about the current running state of your configuration
- Dell EqualLogic Volume Rethinning Facility (eqlvolume)—Manages the available storage on thin-provisioned volumes
- Capability to import EqualLogic storage to PowerStore. For more information see *Importing External Storage to PowerStore Guide*.

Terminology

The following terminology is used by HIT/Linux:

- DM device—A block device created by the Device Mapper that is composed of constituent block devices
- DM target—A kernel module that defines how I/O sent to a DM device is mapped to its constituent block devices. Examples include `dm-multipath`, `dm-linear`, and `dm-switch`.
- Ethernet interface—A physical network interface. Also called Ethernet port or NIC port
- iSCSI interface—An Open-iSCSI concept; a logical interface for creating iSCSI sessions. HIT/Linux creates iSCSI interfaces that map to the physical Ethernet interfaces.
- Multipath device—The Device Mapper devices that create high bandwidth, redundant paths by bundling together connections to each volume

- SD block device—Disk block devices for iSCSI; a direct connection between a host and a PS Series group member
- Smart Copy—A snapshot, clone, or replica of an individual volume or component, or a collection of volumes or components
- Switch device—A top-tier switch device for each volume. The host system uses the dm-switch driver to log in to the device and then route I/O to the appropriate multipath device
- Volume slice—The portion of a volume that is located on a single member

For additional terminology specific to Dell EqualLogic storage products, see the *Dell EqualLogic Master Glossary*.

PS Series Group Network Recommendations

The following table describes the network recommendations for PS Series group operation and for computer access to the group. All the normal networking rules apply; however, network configuration is beyond the scope of this manual.

Table 1. PS Series Group Network Recommendations

Recommendation	Description
Switched Gigabit Ethernet network	Connect arrays and computers to a switched network. Dell recommends that all network connections between computers and arrays are Gigabit Ethernet. (An array can operate at 10 and 100 Mbits, but performance will be significantly degraded.) Network interface cards (NICs) or host bus adapters (HBAs) must operate at a minimum of 1 gigabit per second (1 Gb/sec). NICs that operate at speeds of less than 1 Gb/sec are not used for multipath I/O by default.
Network switch types	Dell recommends using enterprise class switches of a nonblocking design for your iSCSI SAN. If using multiple switches, link the switches to enable communication across different subnets. Use a link of sufficient bandwidth.
Multiple network connections	For increased bandwidth and availability, connect multiple network interfaces on an array to the network (use different switches, if possible). Connect the interfaces from lowest port (port 0) to the highest. At least two network connections are required for multipath I/O. The initial group setup configures only one network interface on an array (port 0). After creating the group, use the Group Manager GUI or CLI to assign an IP address and netmask to other network interfaces.
Access to the group IP address	In a multisubnet group, each configured network interface should have access to the subnet on which the group IP address resides.
Redundant network paths	Multipath I/O provides a highly available network path between computers and arrays.
For replication, a reliable, adequately sized network link	For effective and predictable replication, make sure that the network link between the primary and secondary groups is reliable and provides sufficient bandwidth for copying data.
Flow control enabled on switches and NICs	Enable flow control on each switch port and NIC that handles iSCSI traffic. PS Series arrays will correctly respond to flow control.
Unicast storm control disabled on switches	Disable unicast storm control on each switch that handles iSCSI traffic, if the switch provides this feature. However, the use of broadcast and multicast storm control is encouraged on switches.
Jumbo Frames enabled on switches and NICs	Enable jumbo frames on each switch and NIC that handles iSCSI traffic to obtain performance benefits and for consistency of behavior.
No STP functionality on switch ports that connect end nodes	Do not use spanning-tree (STP) on switch ports that connect end nodes (iSCSI initiators or storage array network interfaces). If you want to use STP or RSTP (preferable to STP), you should enable the port settings available on some switches that let the port immediately transition into STP-forwarding state upon linkup. This can reduce network interruptions that occur when devices restart, and should be enabled only on switch ports that connect end nodes. Dell recommends using spanning-tree for a single-cable connection between switches and trunking for multicable connections between switches.

2

Installing and Configuring HIT/Linux

This chapter guides describes how to install, update, and configure HIT/Linux.

To verify that installation was successful, follow the Verification instructions.

Topics:

Obtain the Installation Kit

Obtain the installation kit from the Dell EqualLogic customer support web site.

To download HIT/Linux, you must have a Dell customer support account. To set up an account, go to https://eqlsupport.dell.com.

Log in to your support account and follow these steps:

1. Select **Downloads** in the navigation bar and select **Host Integration Tools for Linux**.
2. Under Recommended Host Integration Tools for Linux, select the latest version of the toolkit.
3. Select the Download Page link.
4. Select the link for the current version of the tool kit.
5. Accept the EULA and click the **Download** button. Save the ISO installation image to a temporary location on your local computer.
6. Return to the Downloads page and select the **GPG Key** link.
7. Download the GPG Key (`RPM-GPG-KEY-DELLEQL`) and save it to a temporary location on your local computer. The installation requires a public key to authorize your RPM signature and run the installation.

Installation Prerequisites

A successful installation depends on several prerequisites, which include:

- Using supported Linux Distributions
- Running supported versions of PS Series array firmware
- Access to a public key used for installing the RPMs
- All prerequisite Linux packages are available
- When installing or upgrading HIT/Linux, existing EqualLogic volumes to which the host is connected must be online and healthy. If installing on a cluster, ensure the cluster nodes and shared storage are online.
- Your configuration has multiple NICs that are on the same subnet (recommended by Dell for improved performance and higher throughput).

The installation software takes certain actions to ensure your environment is ready for using the product. Atypical installations might require you to change certain defaults prior to installing the kit. Particular cases mentioned in this chapter refer to solutions provided in the Troubleshooting section.

Verify that the prerequisites in the following sections are met before installing the kit.

Installation Script and RPMs

HIT/Linux is packaged as an ISO image, `equallogic-host-tools-`*`version`*`.iso`. This image contains all necessary user-mode and kernel-mode RPM files and an installation script for your particular Linux distribution.

Dell End-User License Agreement (EULA)

When you perform a fresh installation or an update, you must accept the Dell End-User License Agreement (EULA). For information about how to accept the EULA, see Installation Steps.

Supported Linux Distributions

HIT/Linux version 1.7 supports the Linux distributions (platforms) specified in the following table.

Linux Distribution	Version
Red Hat Enterprise Linux (RHEL) Server	7.3, 7.4, 7.5
RHEL Server 32-bit or 64-bit	6.9
SUSE Enterprise Linux Server (SLES)	12 SP2 and 12 SP3
Oracle Enterprise Linux (OEL) with Unbreakable Enterprise Kernel (UEK4)	7.3 and 7.4
Oracle Enterprise Linux UEK4 32-bit or 64-bit	6.8 and 6.9

NOTE: The capability to import EqualLogic Storage to PowerStore is supported only on RHEL 7.5 and SLES 12 SP3.

Array Firmware

- HIT/Linux requires PS Series array firmware version 9.0 or later.
- The capability to import EqualLogic storage to PowerStore requires PS Series array firmware version 10.0.3.

Supported File Systems

ASM/LE supports the following Linux file system types: `ext3`, `ext4`, `xfs`, and `gfs2`.

The following table describes file system restrictions when using ASM/LE.

File System	Restriction
`ext2, btrfs`	Not supported by ASM/LE; displays an error message when you attempt to create a Smart Copy.
`ext3, ext4, xfs`	Read-write on SLES 12 and later. **NOTE: Select `ext4` or `xfs` when installing SLES 12. The default file system installed during SLES 12 installation is `btrfs`, which is not supported by ASM/LE.**
`ocfs2`	ASM/LE does not support the cluster file system `ocfs2`.
`gfs2`	Supported on SLES 12 and later.

NOTE: Unsupported file systems might not allow file-system freeze and thaw operations. Without the ability to freeze all of the data at once, ASM/LE cannot create consistent Smart Copies.

Use of the `multipathd` Daemon

The `multipathd` daemon is not required for EqualLogic devices. Blacklist Dell EqualLogic devices so `multipathd` does not attempt to manage them. For more information, see Migrating to HIT/Linux Multipath I/O. However, the `multipathd` daemon is required to manage PowerStore devices. Therefore, to import EqualLogic storage to PowerStore, make sure that the `multipathd` service is running. For more information see *Importing External Storage to PowerStore Guide*.

Linux Packages and Dependencies

The HIT/Linux installation script attempts to automatically install all required Linux packages (for example, OpenSSL, libpcap, open-iscsi, DKMS, and Python). All packages are part of the RHEL and SLES base repositories. If the packages are not available in a local repository from which the package manager can install them, the installation script notifies you of missing dependencies and you must manually install the required Linux packages.

Installation Procedure

This section describes installing HIT/Linux from the ISO image `(equallogic-host-tools-version.iso)` you downloaded from the Dell support site.

NOTE: If you have installed a previous version of HIT/Linux, see Upgrading HIT/Linux for information about updating an existing installation.

The ISO image contains user-mode and kernel-mode RPM packages, an installation script, and documentation. The installation script uses the host resident package manager to verify and install the Dell EqualLogic host tools and packages and to resolve any dependencies.

As the installation script executes, it installs the HIT/Linux software and guides you through the initial system configuration. The installation script accomplishes the following:

- Examines the installation packages, checks and resolves any dependencies
- Installs a compatible prebuilt kernel module, if available. Otherwise, the installation script loads the Dynamic Kernel Module Support (DKMS) package and compiles the Dell EqualLogic-supplied kernel module.
- Prompts you about entering migration mode if this is a new installation and you have existing connections to Dell EqualLogic PS Series storage volumes using LVM, `multipathd`, or are using the SCSI devices directly. Answering "Yes" to this question behaves identically as if you had run the install script with the `--migrate` option. See Migrating to HIT/Linux Multipath I/O for more details.
- Prompts you with the Import Capability option that is used to enable the capability to import EqualLogic storage to PowerStore. Answering "Yes" to this prompt behaves identically as if you had run the install script with the `--accepted_Migration_Capability` option. If you do not intend to enable the capability to import to PowerStore, then answer "No" to this prompt.
- Invokes the EqualLogic Configuration Utility (eqlconfig) by the installation script after the initial installation has completed. This utility performs several configuration functions, then invokes the EqualLogic Host Performance and Tuning Suite (eqltune):
  - Prompts you to enable Multipath I/O, chooses the addressing protocol (IPv4 or IPv6) and available subnets, then configures MPIO.
  - Configures https parameters if the Import Capability option is selected.

    NOTE: IPv6 is not supported for importing EqualLogic storage to PowerStore.
- Invokes eqltune to verify your Linux system for optimal iSCSI performance and can fix critical system errors. (You can run both eqlconfig and eqltune manually. For more information, see EqualLogic Configuration Utility (eqlconfig) and EqualLogic Host Performance and Tuning Suite (eqltune)).
- Sets the iSCSI initiator daemon to `autostart`.
- Starts the Dell EqualLogic Host Connection Manager daemon (ehcmd) and sets it to `autostart`.
- Prompts you for PS Series group credentials used by ASMCLI. If you choose the default (Yes), the installation script asks you for a group name, user name, and password. If you do not enter credentials at this time, you must provide them to run certain ASMCLI commands. Group credentials are stored in a group access file that you can access and manage from ASMCLI (see Creating and Managing Group Access Files).
- Invokes eqltune to fix all critical errors, if you chose to run the installation in noninteractive mode.

Pre-Installation Steps

Before installing HIT/Linux:

1. Verify that you have met the prerequisites described in Installation Prerequisites.
2. Download the installation kit from the EqualLogic customer support site, per the instructions in Obtain the Installation Kit. There are two downloads, one for the ISO image and another for the public key, `RPM-GPG-KEY-DELLEQL`, that is used for checking the integrity and authenticity of the ISO image and its contents.
3. When you download the key, import the key: `# rpm -import file-name`.

If you are migrating an existing Linux system connected to PS Series storage to use Dell EqualLogic MPIO, see the installation steps in Migrating to HIT/Linux Multipath I/O.

By default, the installation runs interactively. To run the installation in the background, accept all defaults, and accept the Dell EULA, enter the following command:

```
# /media/iso/install --noninteractive --accepted-EULA
```

Failure to provide the `--accepted-EULA` option aborts a noninteractive installation.

To enable the capability to import EqualLogic storage to PowerStore, enter the following command:

```
--accepted_Migration_Capability
```

NOTE: This command is to be used only in conjuction with `-n | --noninteractive` option.

Installation Steps

To install HIT/Linux, perform the following steps.

1. Create a directory in which you will mount the ISO image:

   ```
   # mkdir -p /media/iso
   ```

2. Mount the ISO image in the `/media/iso` directory:

   ```
   # mount -o loop equallogic-host-tools-version.iso /media/iso
   ```

   After mounting, you should see the following files and directories:

   ```
   # ls /media/iso
     EULA    install    LICENSES    packages    README    support    welcome-to-HIT.pdf
   ```

3. The installation script provides several options. To see these options, enter `/media/iso/install --help`. To run the installation script for a fresh install, enter:

   ```
   # /media/iso/install
   ```

For more complicated installations, the following options might also be of help:

- You can replace the configuration defaults with those in a given file. This option applies to multihost deployments; existing settings are automatically merged on update. Use the `--file` or `-f` option on the command line.
- To migrate an existing Linux system connected to PS Series storage, use the `--migrate` or `-M` option. For complete information on migrating an existing system, see Migrating to HIT/Linux Multipath I/O.
- The `--always-use-dkms` option specifies not to use precompiled binary packages, even when available, and to use DKMS packages exclusively. This option is needed only if your system has a mix of supported and unsupported kernels.
- The `--nogpgcheck` option specifies not to check the integrity of packages using GPG. Dell recommends not using this option, unless you are unable to download or import the `RPM-GPG-KEY-DELLEQL` key.

When you perform a fresh installation or an update, you must accept the Dell End User License Agreement (EULA) for the installation script to continue. The EULA contents differ slightly between interactive and noninteractive installations.

During interactive installations and updates, the EULA contents display as scrollable pages for you to view.

1. Press **q**. The following message displays:

   `Please indicate your response to EULA terms and conditions by typing (`**Accept**`/`**Reject**`):`
2. Type `Accept` at the prompt.

Installation Results

Some files are installed only in certain situations or for certain Linux distributions. For a full listing of all files that were installed, run the following commands after a successful installation.

- For user package contents:

  ```
  # rpm -q equallogic-host-tools --filesbypkg
  ```

- (Conditional) On OEL 7.x systems, to view kernel module contents that might have been installed during the installation procedure:

```
# rpm -qa *dm-switch* equallogic-host-tools-kernel --filesbypkg
```

The following table summarizes the installation components.

Table 2. Installation Components

Path	Description
Configuration Files	
/etc/equallogic/eql.conf	Configuration file location for HIT/Linux
/etc/equallogic/eqltune.d/	Eqltune-specific configuration scriptlets
Binary Applications	
/usr/sbin/asmcli	Dell EqualLogic Auto-Snapshot Manager CLI (ASMCLI)
/usr/sbin/ehcmd	EqualLogic Host Connection Manager daemon
/usr/sbin/ehcmcli	EqualLogic Host Connection Manager CLI
/usr/sbin/eqlconfig	EqualLogic Configuration Utility
/usr/sbin/eqllog	EqualLogic Log Gathering Facility
/usr/sbin/eqltune	EqualLogic Host Performance and Configuration Tuning tool
/usr/sbin/eqlvolume	EqualLogic Volume Utility for Thin-Provisioned Volumes
/usr/sbin/rswcli	Remote Setup Wizard CLI (RSWCLI)
Kernel Module	
/lib/modules/kernel-version/ extra/ dm-switch/dm-switch.ko	
/usr/src/dm-switch-1.0	Source and pre-built binaries for dm-switch kernel module; applies to systems that install these kernel components using DKMS
Common Files	
/usr/lib[64]/librswbusiness.so	Shared library that implements much of the RSWCLI functionality
/usr/lib[64]/libpsapi-*.so	Shared library for back-end ASM functions
/usr/lib/libeqlasm.so	Shareable image library to support asmcli and eqlvolume.
/usr/share/equallogic/eql/	Supporting files for eqltune, eqlconfig, and eqllog
/usr/share/equallogic/39-eql_user.rules	Example udev customization file
/usr/share/equallogic/bash_completion.d/	Bash completion support scripts
/etc/bash_completion.d/equallogic	Top-level bash completion script
/etc/profile.d/equallogic.sh	Shell profile integration script
/etc/init.d/ehcmd	Startup script for /usr/sbin/ehcmd
/etc/init.d/scsi_reserve_eql	Cluster SCSI fencing integration startup script (RHEL Server only)
/lib/udev/rules.d/ (all others)	Location of the installed udev rules. For a full explanation of the rules installed, see Using the udev Facility.
/var/lib/equallogic/	HIT/Linux runtime data storage
/var/lib/equallogic/eqltune.backup/	System configuration files backed up by eqltune

Path	Description
Manpages and License	
`/usr/share/man/man8/`	Location of manpages for Dell EqualLogic tools
`/usr/share/docs/equallogic-host-tools-version/`	Location of the End-Users License Agreement (EULA) and license

Uninstalling HIT/Linux

You are not required to manually uninstall the HIT/Linux kit when performing an update. HIT/Linux automatically uninstalls the previous version (see Updating HIT/Linux). However, if you want to manually uninstall the kit, perform these steps:

1. Rerun the installation script as follows:

```
# /media/iso/install -e
```

2. Unmount the ISO image:

```
# umount /media/iso
```

Alternatively, you can manually uninstall the user-mode and kernel-mode packages. Depending on your Linux distribution, use either `yum`, `rpm`, or `zypper`.

The uninstallation procedure removes all binaries but does not stop the iSCSI initiator daemon or modify its autostart settings. The following files and system modifications are also left on the system:

- `/etc/equallogic/eql.conf.rpmsave`
- `/var/lib/equallogic/` and its contents
- ASM backup documents
- ASM credential files
- DKMS, open-iscsi, and any other dependency packages that were installed by the installation script
- Any changes implemented by running `eqltune fix`

NOTE: Subsequent installations will detect any obsolete dependency packages and install the required versions.

Updating HIT/Linux

To update from a previous version of HIT/Linux:

1. Obtain the latest HIT/Linux installation kit.
2. If necessary, update your operating system to a supported version, and reboot if prompted by the operating system update. See Supported Linux Distributions for information on supported operating system versions.
3. Mount the new ISO file and rerun the installation script. The installation script automatically detects an update situation and installs the new RPMs accordingly. For information on installing HIT/Linux, see Installation Procedure.

During an update, the Device Mapper devices are modified to use the `dm-linear` target instead of `dm-switch`, which allows the previous version of the `dm-switch` kernel module to be unloaded. After the new kernel module is installed and loaded, the Device Mapper devices will be changed back to use the updated `dm-switch` target. These changes can occur while the devices are in use but will cause a brief pause in I/O. Any configuration settings that were customized will be preserved across the update.

Configuration

After you install HIT/Linux, perform configuration steps specific to HIT/Linux and ASM/LE:

- Configuring HIT/Linux
- Configuring ASM/LE

Configuring HIT/Linux

To begin using HIT/Linux, review the configuration parameters on your Linux server that impact iSCSI operation.

Step 1: Verify Configurable Parameters with eqltune and eqlconfig

When you run the HIT/Linux installation script, the installation procedure invokes the eqltune utility to evaluate if your configurable parameters conform to the Dell recommended values. After eqltune identifies critical system issues in a summary table, it repairs each category of critical issue. You can manually run eqltune in verbose mode to see more detailed system information. For more information, see EqualLogic Host Performance and Tuning Suite (eqltune).

Similarly, you can run the eqlconfig utility to re-ask the set-up configuration questions that you answered during the initial installation. For information, see EqualLogic Configuration Utility (eqlconfig).

Step 2: Verify Bash Command Line Completion is Configured

HIT/Linux installs third-party compliant Bash command line completion scripts that function independently. After the installation, the `/etc/profile.d/equallogic.sh` script initializes the Bash completion scripts. Existing shells can enable Bash completion with the following command:

```
. /etc/bash_completion.d/equallogic
```

For more information about using Bash command completion, see Bash Command Completion.

Step 3: View and Modify MPIO Configuration Settings

View and modify MPIO configuration settings with the `rswcli --mpio-parameters` command. You can designate the number of connections, the default load balancing policy, the number of I/Os to send down each path, the minimum speed of NICs, whether to use MPIO for snapshots, and whether to use IPv4 or IPv6 addresses.

For information, see Setting MPIO Parameters.

Step 4: Enable or Disable Multipath I_O

If you did not enable MPIO during the initial system configuration, you can enable it later by running:

```
# eqlconfig --mpio
```

You can run `eqlconfig --mpio` to disable MPIO after it has been enabled.

When enabling MPIO for the first time, you might want to add the `--migrate` flag to eqlconfig. If you have existing connections to EqualLogic PS Series storage volumes using LVM, Multipathd, or are using the SCSI devices directly, use this flag to transition to begin using EqualLogic MPIO devices instead. For details about migration see Migrating to HIT/Linux Multipath I/O. You can instruct the eqlconfig utility to begin the procedure by running:

```
# eqlconfig --mpio --migrate
```

For more information about eqlconfig, see EqualLogic Configuration Utility (eqlconfig).

Step 5: Discover iSCSI Targets and Create Sessions to Volumes

After the iSCSI service is running, run the iscsiadm utility to discover iSCSI targets and create iSCSI sessions to each volume of interest. For information, see Discovering Targets.

EqualLogic Configuration Utility (eqlconfig)

The EqualLogic Configuration Utility (eqlconfig) is automatically invoked by the installation script after the initial installation has completed. You can also run the utility anytime after installing and configuring HIT/Linux.

EqualLogic Host Performance and Tuning Suite (eqltune)

The EqualLogic Host Performance and Configuration Tuning Suite (eqltune) verifies a Linux system for optimal iSCSI performance configuration. The suite analyzes a large number of system settings for compliance with Dell's recommended practices, categorizing them into four categories:

- *Critical*—The setting must be changed to avoid serious performance issue.
- *Warning*—The setting should be changed to enhance performance.
- *Suggestion*—The setting could be changed. Depending on the actual patterns of data usage, end-users might want to experiment with a range of values to find a more optimal configuration.
- *OK*—The setting matches Dell's recommendation.

While most messages are informational, you should examine *Critical* and *Warning* messages carefully (see Evaluating Output Information).

The eqltune utility runs in three modes: check mode, list mode, or fix mode. Check mode (`eqltune check`) is the default unless other commands are specified. In this mode, parameters are verified against a recommended range of values. Check mode produces critical, warning, suggestion, or OK messages with recommendations for changing values outside Dell's recommended settings. List mode (`eqltune list`) lists all parameters checked by the eqltune utility but does not perform the verification/recommendation analysis. Fix mode (`eqltune fix`) repairs any issues identified in the analysis phase (running `eqltune check`).

Features

The eqlconfig utility performs a full or partial configuration of HIT/Linux components. It provides the following features:

- Invokes the eqltune utility (see EqualLogic Host Performance and Tuning Suite (eqltune)) to detect any critical configuration issues in the system. If any are found, it asks you if you want to correct the issues before proceeding. The default (Yes) fixes configuration errors and completes the installation.
- Runs through the configuration process for Multipath I/O and ASM/LE. For MPIO configurations, eqlconfig reviews the network protocol and subnets, and presents available subnets or lets you choose individual NICs, and then starts the services. In the process, it merges the configuration files.
- Provides options for you to return to the configuration process anytime and choose to configure MPIO or ASM separately. By default, the eqlconfig utility runs through the configuration of all components.
- When enabling EqualLogic Multipath I/O for the first time, assists you in migrating a live system to start using MPIO devices.
- Lets you select the defaults by simply hitting Enter to all questions or run in noninteractive mode where all defaults are automatically selected as the eqlconfig utility executes.

eqlconfig Command Syntax

The command syntax is as follows:

```
eqlconfig [parameter]
  --version
  --help | -h
  --all
  --migrate | -M
  --mpio
  --asm
  --noninteractive | -n
  --file filename | -f
```

The following table describes the eqlconfig utility's parameters:

Table 3. eqlconfig Parameters

Parameter	Abbr.	Argument	Description
--version	none	none	Shows the version number of the utility and exits
--help	-h	none	Displays command-line parameters for eqlconfig
--all	none	none	Configures all components
--migrate	-M	none	Assists in migrating a live system to start using EqualLogic MPIO devices. For more information, see Migrating to HIT/Linux Multipath I/O.

Parameter	Abbr.	Argument	Description
--mpio	none	none	Configures Multipath I/O component only
--asm	none	none	Configures ASM/LE component only
--noninteractive	-n	none	Runs the eqlconfig utility noninteractively and accepts all configuration defaults
--filename	-f	*filename*	Specifies a configuration file name

eqltune Command Syntax

The command syntax is as follows:

```
eqltune [command] [parameters]

[--version]
       [ --help    | -h]
       [--verbose | -v | -vv]
       [--quiet    | -q]
[check [--verbose | -v | -vv][--quiet | -q]]
[fix   [--verbose | -v][--quiet | -q]]
[list  [--verbose | -v][--quiet | -q]]
```

The following table describes the eqltune utility's commands and parameters.

Table 4. Eqltune Commands and Parameters

Command	Abbr.	Parameter	Description
n/a	none	--version	*Optional.* Shows version number and exits.
n/a	-h	--help	*Optional.* Shows the eqltune utility's options.
n/a	-v	--verbose	*Optional.*
n/a	-q	--quiet	*Optional.*
check	none	-v | -vv | -q	*Default.* Performs all eqltune verifications. Optionally, use with `-v` or `-q` flags. With no verbose flags, shows a summary of all settings. With one verbose flag (`-v`), shows all parameters as they are checked and displays relevant help text. With two verbose flags (`-vv`), shows all settings and all help text.
fix	none	-v | -vv | -q	*Optional.* Repairs any Critical issues identified by the utility.
list	none	-v | -q	*Optional.* Lists all parameters that the eqltune utility surveys but does not perform a verification. For verification, run check mode (`check`). With one verbose flag (`-v`), prints help text for each parameter.

Evaluating Output Information

As the eqltune utility executes, it verifies your installation in several categories.

Block Devices

The eqltune utility looks at every block device that corresponds to an iSCSI session on a PS Series array. Most settings examined by the eqltune utility only provide suggestions because the optimal values are dependent on the expected I/O load on the volume. The critical block device verification is the SCSI device timeout value. HIT/Linux configures the timeout value, based on the EqualLogic `udev` rules installed with HIT/Linux. Example rules are also provided in the installed `udev` file that demonstrate how to set other block device settings.

Sysctl Tunables

Sysctl tunables include Network buffer sizes and Multi-NIC settings. The Network buffer size settings are less critical but you should increase these settings to avoid overflowing transmit and receive buffer queues, which can lead to degraded performance. The Multi-NIC settings are critical because they pertain to the ARP Flux and Return Path Filtering issues.

Ethernet Devices

For every Ethernet device available in the system, the eqltune utility performs a few minor verifications. The most important verification is whether the `Generic Receive Offload` is enabled or disabled. If the Generic Receive Offload setting is enabled, extremely poor read performance results. Dell recommends you disable this setting.

The other Ethernet device settings that the eqltune utility evaluates depend more on the overall network environment than on the host itself. However, you should optimize your network environment settings as suggested by eqltune.

iSCSI Settings

Similar to sysctl tunables, there are a number of settings in the initiator configuration file that you can adjust for performance gains. The eqltune utility verifies that all discovered nodes are synchronized with the main `iscsid.conf` file. The Linux iSCSI initiator allows settings to be set per node. However, because these settings are reset when performing a rediscovery, Dell recommends keeping these settings identical to those in `iscsid.conf`.

Logical Volume Management (LVM)

The LVM device filter expressions must be set up to exclude hidden Dell EqualLogic DM devices. All LVM configuration settings are in `/etc/lvm/lvm.conf`.

EqualLogic Host Tools

This category summarizes whether the running ehcmd process has detected warnings or errors on the running system. For details and an explanation of how to fix these runtime issues, enter `ehcmcli status`.

Examples: eqltune Output

This section includes examples of default and verbose outputs from the `eqltune` command.

Example 1: Default eqltune Output

When you run `eqltune` and specify no options, summary information is provided. This information shows the number of issues detected per category by severity level. For example:

```
# eqltune
```

```
Checking your Linux system for optimal iSCSI performance...
```

```
 Sysctl Tunables        Critical     Warnings      Suggestions  Ok
--------------------------------------------------------------------
ARP Flux                 0            0             0           10
RP Filter                0            0             0            5
Network Buffers          0            2             8            0
Scheduler                0            0             1            0

Ethernet Devices
--------------------------------------------------------------------
eth8                      0           0             2            1
eth5                      0           0             2            1
eth7                      0           0             2            1
eth6                      0           0             2            1
eth9                      0           0             2            1
```

```
iSCSI Settings
-------------------------------------------------------------------
```

```
iscsid.conf defaults      0          4             1            2

External Utility Settings
-------------------------------------------------------------------
Blacklists                0           0            0            1

EqualLogic Host Tools
-------------------------------------------------------------------
Running system checks     0           0             0           2
eqlvolume checks          0           0             0           3

Run in verbose mode (eqltune -v) for more details and instructions on how to adjust your
settings.
```

Example 2: Verbose eqltune Output

The following example shows `eqltune -v` output for the iSCSI settings:

```
===================================================================
iSCSI Settings
===================================================================
The default settings in /etc/iscsi/iscsid.conf are propagated to individual nodes on
discovery or redescovery.

The following command will re-discover all existing nodes (Warning:resets any per-node
settings back to the defaults in iscsid.conf, including 'node.startup'):
iscsiadm -m discovery -t st -p portal

iscsid.conf defaults
--------------------
These settings must be manually edited in /etc/iscsi/iscsid.conf,but will only take effect
for newly-discovered nodes.

node.startup =? manual
  If node.startup is 'automatic', ALL discovered nodes will be logged   in at boot. If by
default this is 'manual', you can designate only those nodes you actually want to auto-login
on a case-by-case basis with iscsiadm:
   iscsiadm -m node -T target -p portal [-I interface]\
         -o update -n node.startup -v automatic
Note that this is the only value that should be set on a per-node basis.
[C] node.session.iscsi.FastAbort =? No
  All EqualLogic products prefer this set to 'No'
node.session.initial_login_retry_max =? 12
  More retries will make it more likely that login will succeed at boot, at the cost of a
slighly longer time to actually fail.
node.conn[0].iscsi.MaxRecvDataSegmentLength =? between 65536 and 524288
  A lower value improves latency at the cost of higher IO throughput
node.session.cmds_max =? 1024
  Maximum number of queued iSCSI commands per session.
  Must be an even power of 2.
node.session.queue_depth =? 128
  The device queue depth
node.conn[0].timeo.noop_out_interval =? any value greater than 1
  iSCSI pings (noops) must not be disabled for proper  error detection and handling
out-of-sync nodes =? 0
  Any nodes that are not in-sync with the defaults in /etc/iscsi/iscsid.conf will have their
settings reset the next time a rediscovery takes place.
```

Configuring ASM/LE

Before you can begin using ASM/LE to create and manage Smart Copies, you must configure your environment as follows:

- Ensure that your network supports PS Series group, discover your PS Series group, and set MPIO configuration parameters, if required.
- Ensure that your PS Series groups meet requirements. See PS Series Group Network Recommendations

- Make sure that you have the correct access controls between the computer and the PS Series group.
- Verify login access to the PS Series group volumes for which Smart Copies will be created.
- Ensure that the correct Linux iSCSI initiators are installed on your computer and that you can connect to the iSCSI targets.

Environment Prerequisites

Ensure that your environment meets the following prerequisites:

- Verify the proper network recommendations are met for PS Series group operation and computer access to the group (see PS Series Group Network Recommendations).
- Optionally, set configuration parameters for MPIO (see Configuring Multipath I/O Devices).

Group Requirements

The PS Series group must meet the following requirements for use with ASM/LE:

- PS Series groups are running the required firmware version.
- Users must have credentials for access rights to volumes in PS Series groups.
- Volumes must have snapshot space reserved to create snapshots. See the *Dell EqualLogic PS Series Group Administrator's Guide*.
- The group and volume must be configured for replication to create volume replicas. See the *Dell EqualLogic PS Series Group Administrator's Guide*.

Access to Volumes and Snapshots

To log in to a PS Series group volume or snapshot (iSCSI target), you must configure the group IP address as the iSCSI target discovery address. A computer must be logged in to the group volumes for which it will create Smart Copies.

To log in to a volume, the computer must match an access control record configured in the group. If access to the volume is being authenticated with CHAP, you must provide a correct CHAP user name and password. After you logged in to a volume, the volume is shown as a regular iSCSI disk.

Run the Remote Setup Wizard Command Line Interface to initialize an array and create or expand a PS Series group or to configure computer access to the group.

ASM/LE can only operate on EqualLogic volumes.

Verification

To verify that HIT/Linux successfully installed, run `ehcmcli status` to see the current status of the EqualLogic Host Connection Manager daemon (ehcmd). You should see an entry for every EqualLogic volume on which you have active iSCSI sessions and no warning messages should appear.

3

Remote Setup Wizard Command Line Interface

The Remote Setup Wizard Command Line Interface (RSWCLI) provides setup commands that you can execute at the Linux command prompt and a command line interface where you can perform all required setup actions.

Topics:

General Command Syntax for the RSWCLI

The general RSWCLI command syntax is as follows:

```
# rswcli --subcommand | -abbreviated_subcommand [--parameter|
--abbreviated_parameter [argument]]
```

The following rules apply:

- When spelling out the full subcommand, prefix with a double dash (`--`). For example, `rswcli --list-group`.
- Prefix abbreviated subcommands with a dash (`-`). For example, `rswcli -l`.
- When specifying parameters, use a double dash with both full and abbreviated parameter names. For example, `--group-name` or `--gn`.
- Use of the equal sign (`=`) to specify a value to parameters is optional. For example, `--member-name=eqlgroup1` and `--member-name eqlgroup1` are equivalent.
- Brackets ([]) indicate that a parameter or argument is optional.
- Several parameters are mutually exclusive and specifying more than one of these parameters causes the command to fail. The command parser ignores any redundant parameters.
- To display all applicable parameters, enter an entire subcommand followed by a tab.
- To display a list of applicable options, enter part of any subcommand or parameter followed by two tabs. For information, see Bash Command Completion.
- To access help on any RSWCLI subcommand, use the `--help` or `-h` parameter after specifying the subcommand. For example, `rswcli --discover --help`.

Configuration — Discovering a PS Series Array

Use the `rswcli --discover` (or `rswcli -d`) command for a list of uninitialized PS Series array serial numbers.

rswcli --discover Command Syntax

The command syntax is as follows:

```
rswcli --discover | -d

      [--detail]
```

The following table describes the optional command parameter.

Table 5. Discover Command Parameter

Parameter	Description
--detail	Specifies whether RSWCLI should output the array MAC address, Host Adapter MAC address, firmware version, and supported RAID levels for each array. Additional information might display, depending on your PS Series Array firmware version.

Example: Discovering Uninitialized PS Series Arrays

The following example shows the detailed output of the `--discover` command.

```
# rswcli --discover --detail | more

Processing discover command...

Discovering uninitialized PS Series arrays that are accessible from
this computer.

(This could take up to a minute)
1 uninitialized PS Series array found.

Service Tag: JJRKWP1
Serial Number: SHU0935411YBF04
Array MAC Address for eth0: 00:09:8A:0A:4F:30
Host Adapter MAC Address: 00:26:B9:3A:2C:35
Firmware Version: X.Y.Z
Supported RAID Levels: 6, 10, 50
Disk Info: 15 931.52GB 7200rpm SATA(HDD) disks
Product Family: PS6000

The discover command succeeded.
```

Initializing a PS Series Array

Use the `rswcli --initialize` (or `rswcli -i`) command to initialize a PS Series array and join a group.

NOTE: This command does not support selecting a pool. Use the PS Group Manager GUI or CLI to select a pool when the member is in a group.

rswcli --initialize Command Syntax

The command syntax is as follows:

```
rswcli --initialize | -i
       --create-group | --cg (or, --jg | --join-group)
       --group-ip IPv4_address | -gip
       --group-name group-name | --gn
       --member-ip-eth0 IPv4_address | --mip
       --member-name member-name | --mn
       --membership password
       --member-subnet subnet_mask | --msub
       --member-gateway IPv4_address | --mgwy
       --service-tag service-tag | --st
```

```
   (or, --serial-number serial_number | --sn)
[--admin password (use with cg | --create-group only)]
[--chap-user username | --cu (use with cg | --create-group only)]
[--chap-secret password | --cs (use with cg | --create-group only)]
[--dcb-vlan-id DCBVlanID | --vid]
[--raid 10|50|6|6-accelerated (use with --cg | --create-group only)]
```

The following table describes the required and optional command parameters.

Table 6. Initialize Command Parameters

Parameter	Abbr.	Argument	Description
Required Parameters			
--create-group or --join-group	--cg or --jg	none	Specifies whether to create a new group for the member using the specified group name and group IP address, or whether the member should join the group with specified group name and group IP address. If the `--create-group` parameter is specified, the `--group-name` and `--group-ip` command parameters are required and must not be in use by any other group.
--group-ip	--gip	*IPv4_address*	Designates the IP Address of the new PS Series group that will be created for the member or the existing group that the member will join. Only IPv4 addresses are supported.
--group-name	--gn	*group-name*	Designates the name of the new PS Series group that will be created for the member or the existing group that the member will join.
--member-ip-eth0	--mip	*IPv4_address*	Designates the IP Address for eth0 of the PS Series array to be initialized. Only IPv4 addresses are supported.
--member-name	--mn	*member-name*	Designates the name of the PS Series array to be initialized.
--membership	none	*password*	Specifies the password on the command line to use for authenticating membership access to the specified group. If you omit this parameter, you will be prompted to enter the group membership password as the command executes.
--member-subnet	--msub	*subnet_mask*	Designates the subnet mask for the subnet controlling access to the PS Series array to be initialized.
--member-gateway	--mgwy	*IPv4_address*	Designates the IP Address of the gateway to use for access to the PS Series array to be initialized. Only IPv4 addresses are supported.
--service-tag or --serial-number	--st or --sn	*service-tag* or *serial number*	Designates the service tag or serial number of the PS Series array to be initialized. The `--service-tag` parameter is not supported on all versions of the PS Series array firmware.
Optional Parameters			
--admin	none	*password*	Specifies the password to use for default group admin user access to the specified group. Required if `--create-group` is specified.
--chap-user	--cu	*username*	Specifies the CHAP username of the member, to only be used with the `--create-group` option. This optional parameter is provided for compatibility with other host-based management tools, but is not stored in the Linux host.
--chap-secret	--cs	*password*	Specifies the CHAP password of the member, to only be used with the `--create-group` option. This optional parameter is provided for compatibility with other host-based management tools, but is not stored in the Linux host.
--dcb-vlan-id	--vid	*DCBVLanID*	Specifies the VLan ID to use for iSCSI traffic if data center bridging is enabled. This parameter is not supported on all versions of the PS Series array firmware.
--raid[1]	none	*10|50|6| 6-accelerated*	Specifies whether the RAID policy for the member in a new group should be RAID10, RAID50, RAID6, or RAID6-accelerated. This option is required with `--create-group`.

[1]RAID5 is no longer considered a best practice for data protection of business critical data on Dell EqualLogic arrays, but it can be configured using the Group Manager CLI.

For more information about RAID policies on PS Series systems, see the Dell Technical Report, *PS Series Storage Arrays: Choosing a Member RAID Policy*, at http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/19861480.

Example: Initializing a PS Series Array

The following example includes the serial number (`--sn`), member name (`--mn`), member IP address (`--mip`), subnet (`--msub`), member gateway (`--mgwy`), create group (`--cg`), group name (`--gn`), group IP address (`--gip`), and RAID policy (`--raid`) parameters. By omitting the membership (`--membership`) parameter from the command line, you are prompted to enter the password as the command executes. Using this method hides the password for added security.

```
# rswcli --initialize --sn VM000c2939fd8b3 --mn farm31 --mip 192.x.x.xx
     --msub 255.255.0.0 --mgwy 10.xxx.x.x --cg --gn farm31-grp
     --gip 192.x.x.xy --raid 50

Processing initialize command...

Validating command line parameters for initialize command...

Membership password: ******
Initialization of array 'VM000c2939fd8b3' in progress...

(This may take a minute or more)

Initialization of array 'VM000c2939fd8b3' as member 'farm31' of group
 'farm31-grp' completed successfully.
Refer to /var/log/equallogic/rswsetup.log for complete details.

The initialize command succeeded.
```

Managing Access to a PS Series Group

The RSWCLI provides several commands for managing group access. With these commands, you can list the available groups, add access to a group, modify group parameters, and remove group access. PS Series arrays initialized using RSWCLI are automatically added to the list of accessible groups.

Adding group access registers the group with iscsiadm and adds it to the discovery portal list. Similarly, removing a group removes the group IP from the list of discovered portals. If a group has been manually discovered with iscsiadm, you do not need to add group access with the `rswcli --add-group-access` command.

This section provides the following information for managing group access:

- Adding Access to a Group
- Listing Accessible Groups
- Modifying Access to a Group
- Removing Host Access to a Group
- Configuring Multipath I/O

Adding Access to a Group

Use the `rswcli --add-group-access` (or `rswcli -a`) command to configure host access to a PS Series group. The `--group-name` and `--group-ip` command parameters are required and must not be in use by any other group.

rswcli --add-group-access Command Syntax

The command syntax is as follows:

```
rswcli    --add-group-access | -a
          --group-name group-name | --gn
          --group-ip IP_address | --gip
         [--noping]
```

The following table describes the required and optional command parameters.

Table 7. Add Group Access Command Parameters

Parameter	Abbr.	Argument	Description
Required Parameters			
--group-name	--gn	*group-name*	Designates the name of the PS Series group to which access is being configured.
--group-ip	--gip	*IP_address* or FQDN	Designates the IP address or fully-qualified domain name (FQDN) to be used for accessing the specified group.
Optional Parameter			
--noping	none	none	Specifies whether RSWCLI should ping the specified IP Address to determine whether it is accessible by the host.

Example: Adding Group Access

In the following example, access is added to the PS series group `farm31-grp`. The command syntax includes the group name (`--gn`) and group IP address (`--gip`).

```
# rswcli --add-group-access --gn farm31-grp --gip 192.x.x.xx

Processing add-group-access command...

Group Name: farm31-grp
Group IP Address: 192.x.x.xx

The add-group-access command succeeded.
```

The following example specifies a FQDN for the group IP:

```
# rswcli --add-group-access --gn farm32-grp --gip testgroup03.mycompany.com

Processing add-group-access command...

Group Name: farm32-grp
Group IP Address: testgroup03.mycompany.com
```

```
The add-group-access command succeeded.
```

Listing Accessible Groups

Use the `rswcli --list-group` (or `rswcli -l`) command to for a list of the names and IP addresses of all PS Series groups to which access is configured for this host.

rswcli --list-group Command Syntax

The command syntax is as follows:

```
rswcli --list-group | -l
      [--group-name group-name] | --gn]
```

The following table describes the optional command parameter.

Table 8. List Group Optional Command Parameter

Parameter	Abbr.	Argument	Description
--group-name	--gn	*group-name*	Displays a list of names and IP addresses of all PS Series groups to which access is configured for this host.

Example: Listing Information for Groups

In the following example, the `list` command produces a list of names and IP addresses for a PS Series group to which access is configured.

```
# rswcli --list-group --gn farm31-grp

Processing list-group command...

Groups accessible from this computer:

Group Name: farm31-grp
Group IP Address: 192.x.x.xx

The list-group command succeeded.
```

Modifying Access to a Group

Use the `rswcli --modify-group-access` (or `rswcli -m`) command to modify the configuration for existing host access to a PS Series group. This command fails if access to the group was not configured previously.

Group access credentials are managed by ASMCLI. For more information, see About Group Access Files.

rswcli --modify-group-access Command Syntax

The command syntax is as follows:

```
rswcli --modify-group-access | -m
       --group-ip IP-address | --gip
       --group-name group-name --gn
      [--new-group-name group-name] | --ngn]
      [--noping]
```

The following table describes the required and optional command parameters.

Table 9. Modify Group Access Command Parameters

Parameter	Abbr.	Argument	Description
Required Parameters			
--group-ip	--gip	*IP-address* or FQDN	Designates a new IP address or fully-qualified domain name (FQDN) to be used for accessing the specified group.
--group-name	--gn	*group-name*	Designates the PS Series group whose host access configuration is being modified.
Optional Parameters			
--new-group-name	--ngn	*group-name*	Designates a new name for an existing group to which access is configured.
--noping	none	none	Directs `rswcli` to not ping the specified IP Address to learn if it is accessible by the host.

Example: Modifying Group Access

In the following example, group access is modified from group `farm31-grp` to group `new31-grp` using the group name (`--gn`) and new group name (`--ngn`) parameters:

```
# rswcli --modify-group-access --gn farm31-grp --ngn
new31-grp

Processing modify-group-access command...
```

```
Group Name: new31-grp
Group IP Address: 192.xx.xx.xx

The modify-group-access command succeeded.
```

Removing Host Access to a Group

Use the `rswcli --remove-group-access` (or `rswcli -r`) command to remove host access to the PS Series group.

rswcli --remove-group-access Command Syntax

The command syntax is as follows:

```
rswcli --remove-group-access | -r
       --group-name group-name | --gn
```

The following table describes the required command parameter.

Table 10. Remove Group Access Command Required Parameter

Parameter	Abbr.	Argument	Description
--group-name	-gn	*group-name*	Designates the name of the PS Series group to which access is to be removed.

Example: Removing Host Access to Group

In the following example, access is removed to the group `farm30-grp`:

```
# rswcli --remove-group-access --gn farm30-grp

Processing remove-group-access command...

The remove-group-access command succeeded.
```

Configuring Multipath I/O

Several configuration parameters control the behavior of the EqualLogic Host Connection Manager daemon (ehcmd). The `rswcli --mpio-parameters` command lets you view and modify the most commonly used multipath I/O parameters. These parameters determine:

- Number of iSCSI sessions to create to each volume
- Minimum adapter speed to use for iSCSI sessions
- Load balancing policy to set
- IP version to use (IPv4 or IPv6)

For complete information on configuring multipath I/O using the `rswcli --mpio-parameters` command parameters, see Configuring Multipath I/O Devices.

4

Configuring Multipath I/O Devices

EqualLogic Multipath I/O (MPIO) in HIT/Linux enables you to configure redundant network paths from a Linux host to PS Series volumes for high availability and high performance SAN data traffic. MPIO accomplishes two goals:

- Creates and manages an optimal number of iSCSI sessions to each SAN volume based on the number of host NICs and the number of group members the volume is using. HIT/Linux monitors the SAN configuration and adds sessions as the configuration changes.
- Enhances performance by directing each I/O to the proper PS Series group, reducing the amount of I/O forwarding that is necessary between members.

This chapter describes how to use HIT/Linux to set up access to EqualLogic PS Series storage volumes, and explains how the EqualLogic Host Connection Manager daemon (ehcmd) sets up a two-tier multipath device structure to accomplish these goals.

Topics:

Accessing Volumes Using Multipath I/O

To access volumes using multipath I/O, you must discover targets (volumes) and then log into at least one iSCSI session to the volume.

By default, the EqualLogic Host Connection Manager (ehcmd) uses the software iSCSI initiator to connect to volumes. For information on enabling and configuring the Broadcom iSCSI offload functionality, see Enabling iSCSI Offload on HIT/Linux.

Discovering Targets

The PS Series group address must be configured as a discovery address and a send-targets discovery executed before any volumes can be accessed. HIT/Linux does this for you when you use the `rswcli --add-group` command. You can view the currently configured discovery addresses with:

```
# iscsiadm -m discoverydb
```

Also, you can view the current list of discovered target IQN names with:

```
# iscsiadm -m node | sort -u
```

If you created additional volumes, removed volumes, or modified access to existing volumes so their visibility to the Linux host has changed, you must run the `discoverydb` command to add any new volume records and remove any old volume records:

```
# iscsiadm -m discoverydb -p group-IP-address -t st -o new delete -D
```

NOTE: Support of the `iscsiadm -m discoverydb` command syntax applies to systems with version 2.0-872 and later of the iSCSI initiator tools. Earlier versions must use this syntax: `iscsiadm -m discovery [-I iface ] -t st -p group-IP-address`. This command syntax will not preserve the settings in existing node records, but will reapply the default settings specified in the `iscsid.conf` file. See the `iscsiadm(8)` manpage for more information.

Log In to an MPIO Volume

To log in to an MPIO volume, specify:

```
# ehcmcli login --target target-name --portal portal
```

where *target-name* indicates the full iSCSI-qualified name or a volume name for the PS Series group target node, and *portal* indicates iSCSI portal (group IP address or name).

NOTE: You must specify the full iSCSI-qualified name in *target-name* if there are multiple volumes with identical names.

The appropriate device to mount is returned by the `ehcmcli login` command. For example:

```
# ehcmcli login --target iqn.2001-05.com.equallogic:0-290c06-52f7c99c6-
   1a2000001aa4f479-myvolume
```

```
Login succeeded. Device to mount:
```

```
/dev/eql/myvolume
```

NOTE: Unlike logging in using `iscsiadm`, you do not need to specify the iface record; ehcmcli selects the interfaces based on your included subnets.

After you create a single session, the EqualLogic Host Connection Manager (ehcmd) will analyze the configuration and create additional iSCSI sessions as appropriate. See Multipath I/O Sessions for more information.

When complete, an iSCSI session is initialized for the appropriate PS Series volume, `/dev/eql/volume-name`, and the volume is available for use. All further iSCSI management by ehcmd is transparent to any application using `/dev/eql/volume-name`. For more details, see Multipath Device Configuration.

Logging In to an MPIO Volume

After you configure multipath I/O and create volumes on the PS Series group, you need to create at least one iSCSI session to the volume.

Mount an MPIO Volume

To use a volume with multipath I/O:

1. Confirm the name of the device to mount by running `ehcmcli status`. Look for the device to mount line:

   ```
   Device to mount: /dev/eql/volume-name
   ```

   For example:

   ```
   Device to mount: /dev/eql/VolA
   ```

2. Build a Linux file system. For example:

   ```
   # mkfs.ext4 /dev/eql/VolA
   ```

   For a list of supported file systems, see Supported File systems.

3. Specify a directory and create a mount point for the new volume. For example

   ```
   # mkdir /mount/VolA
   ```

   ```
   # mount /dev/eql/volA /mount/VolA
   ```

You can now write files to `/mount/VolA`

NOTE:

- EqualLogic storage supports partitions (for example, using fdisk), but recommends using file systems on the full volume instead of partitions on that volume.
- If you want to suppress creating partition devices on a volume, see Setting Volume-Specific Configuration Parameters.
- If you want to mount a mount point on system reboot, add it to your `/etc/fstab` file.

Logging Out of an MPIO Volume

To log off a volume with multiple connections, log out of all iSCSI sessions using the following command:

```
ehcmcli logout --target target-name
```

where *target-name* is the full iSCSI-qualified name or a volume name for the PS Series group target node.

NOTE: You must specify the full iSCSI-qualified name in *target-name* if there are multiple volumes with identical names.

Displaying Connections to a Volume

Issue the `ehcmcli status` command to show the diagnostic state of the EqualLogic Host Connection Manager (ehcmd) and report the iSCSI sessions connected to each volume. For information on using ehcmcli, see Using the ehcmcli Diagnostic Tool.

Managing Automatic Logins

In many scenarios, you want to have iSCSI volumes connected at startup without user intervention. To connect the volumes, set the `node.startup` property to `automatic` for individual iSCSI targets.

Dell recommends that you not set the default `node.startup` property to `automatic` in the `iscsid.conf` file. Setting the property causes the iSCSI initiator to log in to all discovered targets from all configured network interfaces at startup, and your computer will connect to every iSCSI volume that it discovers when booting. In addition, multiple iface records are configured for the server. Logging in through all of these records results in excess iSCSI sessions being created at boot time, which negatively impacts startup performance and unnecessarily strains the SAN resources.

To set a volume to automatically log in at the next system startup, use the `--login-at-boot` option with the `ehcmcli login` command:

```
# ehcmcli login --target target [--portal portal] --login-at-boot
```

To set a volume to not automatically log in at the next system startup, use the `--no-login-at-boot` option with either the `ehcmcli login` or the `ehcmcli logout` command:

```
# ehcmcli login --target target [--portal portal] --no-login-at-boot
# ehcmcli logout --target target --no-login-at-boot
```

It is safe to run `ehcmcli login` with either of these options, even if you are already logged in to the specified target. The currently logged-in sessions are not disturbed, but the automatic login management changes take place.

Likewise, it is safe to run `ehcmcli logout --no-login-at-boot` on a target that is not currently logged in for the purpose of automatic login management changes.

For the current version of the iSCSI initiator tools, ehcmcli and ehcmd detect the leading login facility, `node.leading_login`, and use it to control the number of automatic logins at boot.

NOTE: Rediscovering targets with `iscsiadm -m discovery` resets the node properties to the defaults in `iscsid.conf`, including the node.startup properties. Dell recommends using `iscsiadm -m discoverydb` instead, if available in your version of the iSCSI initiator tools.

Managing iSCSI Interfaces

For iSCSI multipathing, the ehcmd daemon creates iface records based on the MAC addresses that are bound to particular hardware ports. Multiple iface records for each Ethernet port can be created to allow ehcmd to create multiple sessions from the same Ethernet port on the server to the same iSCSI target. With newer versions of the iSCSI initiator tools, ehcmd does not have to create multiple ifaces for each network port.

If you created iface records for specific adapters, ehcmd uses the existing records. If ehcmd creates new records, it uses the naming scheme `eql.eth0_X`. For example:

```
# iscsiadm -m iface | sort
default tcp,<empty>,<empty>,<empty>,<empty>
eql.eth0_0 tcp,00:26:B9:3A:6B:9E,<empty>,<empty>,<empty>
eql.eth0_1 tcp,00:26:B9:3A:6B:9E,<empty>,<empty>,<empty>
eql.eth1_0 tcp,00:26:B9:3A:6B:9F,<empty>,<empty>,<empty>
eql.eth1_1 tcp,00:26:B9:3A:6B:9F,<empty>,<empty>,<empty>
eql.eth2_0 tcp,00:1B:21:46:5E:18,<empty>,<empty>,<empty>
eql.eth2_1 tcp,00:1B:21:46:5E:18,<empty>,<empty>,<empty>
eql.eth3_0 tcp,00:1B:21:46:5E:19,<empty>,<empty>,<empty>
eql.eth3_1 tcp,00:1B:21:46:5E:19,<empty>,<empty>,<empty>
iser iser,<empty>,<empty>,<empty>,<empty>
```

To avoid interrupting access to any iSCSI storage, these additional iSCSI interfaces will be left in place if the HIT/Linux software is ever uninstalled.

Multipath Device Configuration

When you configure multipath device I/O, a mounted volume on a host is shown as a single disk. MPIO configuration uses multiple iSCSI sessions to the SAN volume to provide load balancing and ensure that I/O operations, such as writing data to a volume, succeed even if there is a network error.

Traditional Multipath Device Configurations

Traditional multipath configurations are implemented as a single tier of Device Mapper multipath devices. These configurations spread I/O sent to the multipath device to the constituent SCSI disk (SD) block devices.

Multipath Configuration

A PS Series volume can have data spread across multiple group members. Multipath device configuration in HIT/Linux improves performance over standard Linux multipathing by ensuring the host has iSCSI sessions to each member. Using information about the volume layout, it routes each I/O directly to the destination group member, reducing the amount of I/O forwarding that is necessary between members.

HIT/Linux includes two components that use this configuration:

- The EqualLogic Host Connection Manager daemon (ehcmd) monitors the iSCSI session state and the configuration of the Linux server and PS Series group. Running in the background, ehcmd uses the Open-iSCSI management tool (iscsiadm) to add, remove, or modify iSCSI sessions to maintain optimal iSCSI sessions, and gathers information on the volume layout from the PS Series group.
- A loadable kernel module, `dm-switch`, which implements a new Device Mapper target to the multipath devices. Based on the volume layout on the group members, the `dm-switch` module routes each I/O to the optimal path to the PS Series group.

The result is a two-tier device structure for each volume with multiple multipath devices: one for each member containing a portion of the volume and a switch device on top of the multipath devices.

The following figure shows the relevant components in the storage stack.

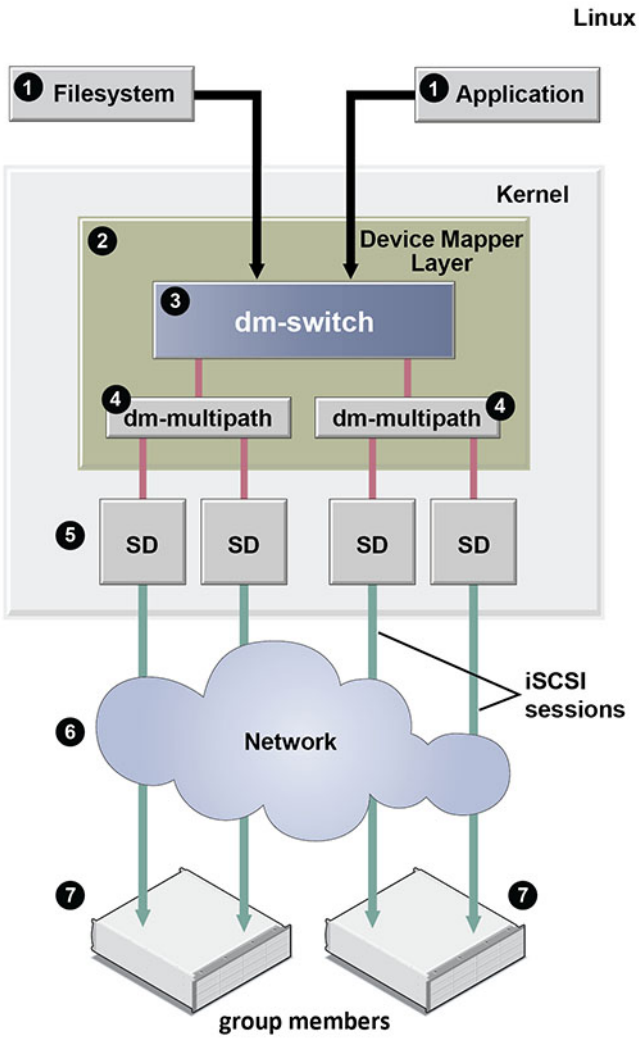

Figure 1. Dell EqualLogic Multipath Device Configuration - Logical View

Table 11. Callouts in Figure 1

Callout Number	Description of Feature
1	Host file system or application using a block device
2	Device Mapper layer—The Device Mapper allows layering block devices to set up a multipath configuration. HIT/Linux uses the Device Mapper to build a two-tier device structure with multiple multipath devices (`dm-multipath`) for each volume and a Device Mapper switch target on top of the multipath devices.
3	Switch device (one per volume)—The ehcmd daemon builds a single top-tier switch device for each volume the host system logs in to with the `dm-switch` device. The `dm-switch` device determines on which group the I/O data resides, routes I/O to the appropriate multipath device and thereby to the appropriate member based on the starting logical block number of the I/O. Because this device is not a multipath device, the standard multipath command line tools do not report it.
4	Multiple dm-multipath devices are built for each volume that the host uses, one for each group member that holds a portion of the volume. Each of the dm-multipath devices lists all SD devices, some in a preferred priority group and some in a non-preferred group. The SD block devices that represent direct paths to that member are placed in a preferred priority group and are used for all I/O in normal conditions; all other SD devices are put in a nonpreferred priority group so they are available for failover of I/O. The ehcmd daemon builds and maintains this configuration, responds to any changes in SAN topology, such as device addition or device removal, and modifies the Device Mapper devices as necessary. Based on load factors, the `dm-multipath` device determines which iSCSI path should be taken to the group member from the multiple paths in the preferred priority group. When available, the `dm-multipath` devices use a

Callout Number	Description of Feature
	queue-length path selector for load balancing. Otherwise, the devices use a round-robin path selector. The `dm-multipath` device also provides error handling capability and on an I/O error will retry I/Os on alternate paths. If all paths in the preferred priority group fail, paths from the nonpreferred priority group are used for I/O. For information on the supported path selection policies, see Conditions for Configuring Multipath Devices.
5	SD block device (one per iSCSI session)—An SD block device represents a direct connection between a host and a group member.
6	iSCSI sessions over Internet/Intranet—While this entire process is executing, the ehcmd daemon runs in the background to provide iSCSI session management. It monitors the iSCSI session state, the configuration of the Linux server, and the PS Series group. It adds, removes, or modifies iSCSI sessions to maintain the optimal iSCSI sessions. The ehcmd daemon uses the Open-iSCSI management tool (iscsiadm) to manage iSCSI sessions. It also gathers information on the volume layout from the PS Series group.
7	PS Series group members

Example Multipath Device Configuration

The following example output from the `dmsetup` table illustrates a typical configuration that results from creating four iSCSI sessions to a single EqualLogic volume located on two group members. Note that the following naming conventions are used:

Device Mapper Device	Naming Convention	Explanation
dm-switch	`eql- PS_volume_id-`*volume_name*	*volume_name* is the name specified by the volume creator.
dm-multipath	`eql- PS_volume_id_`*[abc...]*	*[abc...]* indicates the dm multipath device to define the preferred and nonpreferred path to group members.

```
# dmsetup table
eql-0-8a0906-19d376406-b06003614cc4c053-vol5: 0 41963520 switch 2 30720 253:8 0 253:7 0
eql-0-8a0906-19d376406-b06003614cc4c053_b: 0 41963520 multipath 0 0 2 1 round-robin 0 2 1
65:0 1 65:16 1 round-robin 0 2 1 8:240 1 8:176 1
eql-0-8a0906-19d376406-b06003614cc4c053_a: 0 41963520 multipath 0 0 2 1 round-robin 0 2 1
8:240 1 8:176 1 round-robin 0 2 1 65:0 1 65:16 1
```

In the example, `vol5` is a volume on the PS series group. The file system or application uses the top-tier switch device (`eql-0-8a0906-19d376406-b06003614cc4c053-vol5`), not the lower-tier devices indicated by `eql-PS_volume_id_[abc...]`.

All DM devices are visible in `/dev/eql`, for example:

```
# ls /dev/eql
0 brw-rw---- 1 root disk 253, 9 Oct 19 14:48 vol5
```

When you mount the device, for convenience, the udev rules creates symlinks under `/dev/eql`. In the previous example, you would then mount the device using the following convention:

```
/dev/eql/volume-name
```

For example:

```
/dev/eql/vol5
```

To avoid volume name symlink collisions, Dell recommends that every volume to which you connect have a unique name. For more information, see Using the udev Facility.

Special MPIO Configuration Cases

There are special cases when an alternative device structure is built. However, you always mount `/dev/eql/volume_name`. These cases include:

- Single-member volume with multiple connections—In some cases, a volume might reside in a single member but might have multiple connections to that member (for example, very small volumes only reside on a single member). In this case, the ehcmd daemon builds a single Device Mapper device that uses the Device Mapper multipath target to allow I/O to all available sessions, because all paths are equally good for each I/O. This device is called `eql-PS_volume_id _a`. A top-tier device called `eql-PS_volume_id-`*`volume_name`* is also created that uses the Device Mapper linear target to map to the lower-tier device.
- Single-member / Single-path volume—If you specify that a volume is only accessible by a single network address, you will only get one session established to the volume. In this case, because there is only a single path to use, the ehcmd daemon builds a Device Mapper device with a linear target called `PS_volume_id-`*`volume_name`*.
- Multi-partition volume—The ehcmd daemon builds an additional layer over the `dm-switch` layer consisting of linear targets offsetting into the volume. This method uses the same strategy as kpartx when creating partitions on multipath devices. The same naming convention is used, `eql-PS_volume_id-`*`volume_namepN`*, where *N* identifies the partition number.

For example, if the target volume is `iqn.2001-05.com.equallogic:eql-0-8a0906-8f7f95f06-9f5210a13884c73d-sample`, the entire volume can be mounted as `/dev/eql/sample` and partition 2 can be mounted as `/dev/eql/samplep2`.

```
# ls /dev/eql
brw------- 1 root root 253,  2 Sep 20 20:00 sample
brw------- 1 root root 253, 22 Sep 20 20:00 samplep1
brw------- 1 root root 253, 21 Sep 20 20:00 samplep2
brw------- 1 root root 253, 20 Sep 20 20:00 samplep3
```

NOTE:

- **Do not mount the lower-tier devices `/dev/mapper/eql-0-8a0906-8f7f95f06-9f5210a13884c73d_a` and `/dev/mapper/eql-0-8a0906-8f7f95f06-9f5210a13884c73d_b` or the SD devices used to build them. These devices are used by the `dm-switch` kernel module and should not be directly accessed by users (for example, file system mounting).**
- **EqualLogic storage supports partitions (for example, using fdisk), but recommends using file systems on the full volume instead of partitions on that volume.**
- **If you want to suppress creating partition devices, see Setting Volume-Specific Configuration Parameters.**

Conditions for Configuring Multipath Devices

Certain conditions should be met to successfully configure multipath devices using HIT/Linux.

Use of the `multipathd` Daemon

The `multipathd` daemon is not required for EqualLogic devices. Blacklist Dell EqualLogic devices so `multipathd` does not attempt to manage them. For more information, see Migrating to HIT/Linux Multipath I/O. However, the `multipathd` daemon is required to manage PowerStore devices. Therefore, to import EqualLogic storage to PowerStore, make sure that the `multipathd` service is running. For more information see *Importing External Storage to PowerStore Guide*.

Use of dm-switch Kernel Module

To efficiently route I/O requests to the best path, the Linux host uses the `dm-switch` kernel module, which provides a new Device Mapper (DM) target, `switch`. To verify that the kernel module installed correctly and completed the handshake with the DM environment, use the `dmsetup targets` command to list the known DM targets and include an entry for `switch` in that list. For example:

```
# dmsetup targets
switch v1.0.0
mirror v1.2.0
multipath v1.0.5
raid45 v1.0.0
striped v1.1.0
linear v1.0.2
error v1.0.1
```

If you suspect that the module is not installed, use the `modprobe` command to determine whether the kernel module is installed. If it is not installed, look at the DKMS log to ensure that the module was built successfully and if rebuilding the module was needed.

For additional information troubleshooting problems with the `dm-switch` kernel module, see Assuring Correct Operation of the dm-switch Kernel Module.

Configuring Settings for Multipath I/O

HIT/Linux lets you configure multiple multipath devices as redundant paths from a Linux iSCSI initiator to a PS Series volume. RSWCLI commands for Multipath I/O (MPIO) let you:

- Set MPIO parameters
- List MPIO parameters
- Include a subnet for use by MPIO
- Exclude a subnet for use by MPIO

Multipath I/O Sessions

The number of paths created by HIT/Linux MPIO depends on the topology of your SAN and the MPIO settings on the host computer. Every volume is distributed across one or more members in the PS Series group. The portion of a volume that is located on a single member is referred to as a *volume slice*. The ehcmd utility creates one or more iSCSI sessions to each volume slice. The actual number of sessions is determined through the following actions:

1. Calculating the optimal number of sessions to maximize the bandwidth between the host and the member. This calculation takes into account the speed of available host adapters and network interface cards (NICs) on each member hosting part of the volume. The algorithm does not create unnecessary sessions when there is no performance benefit. Therefore, in configurations with limited numbers of computer and member Ethernet ports, the actual number of sessions created might be fewer than the user limits specified in the next step.

   Example A: You have 2 x 1 Gb host NICs, and are connecting to a volume on members with 4 x 1 Gb eth ports. The optimal number of sessions is 2 per member, because that number saturates the host NICs.

   Example B: You have 1 x 10 Gb host NICs, and are connecting to a volume on members with 4 x 1 Gb eth ports. The optimal number of sessions is 4 per member, because that number saturates the member eth ports.

   Example C: You have 2 x 10 Gb host NICs, and are connecting to a volume on members with 1 x 10 Gb eth ports. The optimal number of sessions is 1 per member, because that number saturates the member eth ports.
2. Applying any user-configured session limits. The default behavior is to limit to 2 sessions per volume slice or 6 sessions per volume, whichever is reached first. However, you can modify these limits, as described in Setting MPIO Parameters.
3. Apply a per-pool throttle to ensure the group remains below 90% of the maximum number of allowed connections per pool. This throttling logic ensures the connections are equitably shared among all the pool members and all the hosts using HIT multipathing to connect to volumes on the group. See the *Dell EqualLogic PS Series Storage Arrays Release Notes* for the current configuration limits.

Setting MPIO Parameters

Use the `rswcli --mpio-parameters` (or `rswcli -P`) command to change the settings used for configuring MPIO. All parameters are optional. If you do not enter parameters for this command, the current values for all parameters are listed.

Listing MPIO Subnets

Use the `rswcli --mpio-list` (or `rswcli -L`) command to list the subnets that are included and excluded for use by MPIO.

rswcli --mpio-parameters Command Syntax

The command syntax is as follows:

```
rswcli --mpio-parameters | -P
      [--max-sessions-per-entire-volume 1-12 | --maxvole]
      [--dlb | --default-load-balancing LQD|RR]
      [--ios | --io-per-path 1-1000]
      [--max-sessions-per-volume-slice 1-4 | --maxvols]
      [--minsp | --min-adapter-speed 100|1000|10000 (as Mbps)]
      [--msnap | --use-mpio-snapshots yes|no]
      [--ip | --useip or --ip6 | --useip6]
```

The following table describes the optional command parameters.

Table 12. MPIO Optional Parameters

Parameter	Abbr.	Argument	Description
--max-sessions-per-volume-slice	maxvols	*1-4*	Specifies the maximum number of connections for MPIO to use per volume per member.
--max-sessions-per-entire-volume	maxvole	*1-12*	Specifies the maximum number of connections for MPIO to use per volume for all members.
--default-load-balancing	dlb	*LQD* or *RR*	Specifies the default load balancing policy for MPIO as Least Queue Depth (LQD) or Round Robin (RR).
--io-per-path	ios	*1 - 1000*	Specifies the number of I/Os to send down each path to a member before the `dm-multipath` target switches to another path. Larger values provide better throughput on sequential I/O workloads. The default is 16.
--min-adapter-speed	minsp	*100* or *1000* or *10000*	Specifies the minimum speed of the NICs MPIO should use, if at least two meet the speed.
--use-mpio-snapshots	msnap	*Yes* or *No*	Specifies whether to use MPIO for snapshots.
--useip or --useip6	ip or ip6	none	Specifies that MPIO will use IPv4 or IPv6 addresses. Only one parameter can be specified.

Examples: Setting MPIO Parameters

This section includes examples of using the `rswcli` command to set parameters.

Example 1: Setting maxvols, maxvole, dlb, and Internet Protocol

In the following example, the command syntax includes the maximum number of connections that MPIO should use per volume (`--maxvols`), maximum sessions per entire volume (`--maxvole`), the default load balancing mechanism (`--dlb`), and Internet Protocol (`--ip6`):

```
# rswcli -P --maxvols 3 --maxvole 12 --dlb RR --ip6

Processing mpio-parameters command...

MPIO Parameters:

Max sessions per volume slice: 3
Max sessions per entire volume: 12
Minimum adapter speed: 1000
Default load balancing policy configuration: Round Robin (RR)
IOs Per Path: 16
Use MPIO for snapshots: Yes
Internet Protocol: IPv6

The mpio-parameters command succeeded.
```

Example 2: Setting I/Os Per Path

The following example shows output after setting the I/Os per path from 16 to 8:

```
# rswcli --mpio-parameters --io-per-path 8

Processing mpio-parameters command...

MPIO Parameters:
Max sessions per volume slice: 2
Max sessions per entire volume: 6
Minimum adapter speed: 100
Default load balancing policy configuration: Round Robin (RR)
IOs Per Path: 8
Use MPIO for snapshots: Yes
Internet Protocol: IPv4
```

```
The mpio-parameters command succeeded.
```

rswcli --mpio-list Command Syntax

The command syntax is as follows:

```
rswcli --mpio-list | -L
      [--include-only | -io or --exclude-only | --eo]
```

The following table describes the optional command parameters.

Table 13. List of Optional MPIO Parameters

Parameter	Abbr.	Description
--include-only	--io	Specifies that the list should only contain the subnets that are included for use by MPIO. If this parameter is specified, the `exclude-only` parameter cannot be specified.
--exclude-only	--eo	Specifies that the list should only contain the subnets that are excluded from use by MPIO. If this parameter is specified, the `include-only` parameter cannot be specified.

Example: Listing Included MPIO Subnets

In the following example, the command syntax lists only included MPIO subnets:

```
# rswcli --mpio-list --io

Processing mpio-list command...

Subnets included for MPIO:
192.x.x.xx - 225.xxx.x.x

Subnets excluded from MPIO: none

The mpio-list command succeeded.
```

Including a Subnet or Adapter for MPIO

Use the `rswcli --mpio-include` (or `rswcli -I`) command to select the subnets or adapters to include for MPIO. You can specify a subnet, an adapter identified by a user-defined string, an IP address, or a MAC address. Dell recommends using a subnet or adapter name based on inclusion/exclusion rules.

rswcli --mpio-include Command Syntax

```
rswcli --mpio-include | -I
      [--network ip-address | --net]
      [--mask subnet-mask ]
      [--adapter=adapter-name ]
      [--mac=mac-address ]
      [--ip-address=ip-address ]
      [--force]
      [--default=include/exclude ]
```

The following table describes the optional command parameters.

Table 14. Include MPIO Subnet Optional Parameters

Parameter	Abbr.	Argument	Description
--network	--net	*ip-address*	Specifies a subnet to include. Used with the `--mask` parameter.

Parameter	Abbr.	Argument	Description
--mask	none	*subnet-mask*	Specifies a subnet mask to include. Used with the `--network` parameter.
--adapter	none	*adapter-name*	Specifies a user-defined string as the name of the adapter to include. For example, eth0, eth1, and so on.
--mac	none	*mac-address*	Specifies a MAC address, or physical address, for an adapter that is to be included. The MAC address can be specified in any of the following formats: `AABBCCDDEEFF`, `AA-BB-CC-DD-EE-FF`, `AA:BB:CC:DD:EE:FF`, or `AA.BB.CC.DD.EE.FF`.
--ip-address	none	*ip-address*	Specifies an IP address to include.
--force	none	none	Forces the included list to be overwritten with all the subnets, adapters, or IP addresses listed. This option must be used when switching from one type of rule to another. For example, when changing from inclusion by subnet to inclusion by adapter name.
--default	none	include|exclude	Specifies to include or exclude by default the adapters not explicitly listed.

Examples: mpio-include Command

The following examples show using the `--mpio-include` command to include a subnet and include an adapter.

Example 1: Including a Subnet

In the following example, the `rswcli --mpio-include` command includes a subnet.

```
# rswcli --mpio-include --network=192.x.x.xx --mask=255.255.0.0 --force

Processing mpio-include command...
Adding '192.x.x.xx - 255.255.0.0' to include list
Subnet Rules:
 Included:
   192.x.x.xx - 255.255.0.0
 Excluded:
   All Other
All other subnets, adapters or ip addresses are excluded by default.

The mpio-include command succeeded.
```

Example 2: Including an Adapter

In this example, the `rswcli --mpio-include` command includes an adapter by a user-defined string.

```
# rswcli --mpio-include --adapter=eth0 --force

Processing mpio-include command...
Adding 'eth0' to include list
Adapter Rules:
 Included:
   eth0
 Excluded:
   All Other
All other subnets, adapters or ip addresses are excluded by default.

The mpio-include command succeeded.
```

Excluding a Subnet or Adapter From MPIO

Use the `rswcli --mpio-exclude` (or `rswcli -E`) command to select the subnets or adapters to exclude from MPIO. You can specify a subnet, an adapter identified by a user-defined string, an IP address, or a MAC address. Dell recommends using a subnet or adapter name based on inclusion/exclusion rules.

Including or Excluding Specific Host NICs

Include or exclude a host NIC with the IP address of the NIC and a netmask of 255.255.255.255:

```
# ip addr show eth1

3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
qlen 1000
    link/ether 08:00:27:31:25:55 brd ff:ff:ff:ff:ff:ff
    inet x.x.x.xx/16 brd xx.xxx.255.255 scope global eth1

# rswcli --mpio-exclude --net x.x.x.xx --mask 255.255.255.255

Processing mpio-exclude command...
Adding x.x.x.xx - 255.255.255.255 to exclude list
The mpio-exclude command succeeded.
```

Configuring I/O Path Selection

HIT/Linux provides configurable parameters that let you establish path selection policy for the middle tier multipath device. These parameters set load balancing policy and determine number of I/Os to send on each path before switching paths.

Load Balancing Policy

The default load balancing parameter (`--mpio-parameters --dlb`) lets you specify the default load balancing policy for MPIO: Round Robin (RR) or Least Queue Depth (LQD). The RR policy rotates among each path to the member sending a fixed number of I/Os to each. The LQD (or queue-length) policy considers the number of outstanding I/Os on that path.

Each load balancing policy is implemented by a path selector kernel module. Availability of path selectors depends on your kernel version. If available, LQD is the default policy. If not, HIT/Linux uses RR as the default policy.

You can set the load balancing parameter by editing the `eql.conf` file or with the following RSWCLI command:

```
#  rswcli --mpio-parameters --dlb LQD|RR
```

I/O Scheduling

The I/O per path parameter lets you specify the number of I/Os to send on each path before switching paths.

I/O scheduling is performed above the multipath layer, at the `dm-switch` layer. The I/O per path parameter (`--mpio-parameters --ios`) defaults to a value of 16. You can adjust this value in the `eql.conf` file or with the following RSWCLI command:

```
#  rswcli --mpio-parameters --ios=1-1000
```

Setting Volume-Specific Configuration Parameters

The configuration parameters in the MPIO Parameters section are global settings applied to all volumes. The `ehcmd` command supports volume-specific settings for the following parameters:

- EnableMPIO
- IOPerPath

- LoadBalanceType
- MaxDevicesPerMPIOSession
- MaxConnectionsPerMember
- SuppressPartitions

To set volume-specific values for any of these values, add a section to the `eql.conf` file name `[MPIO Volume Params volume-name]` and define the volume-specific parameter settings in this section. Wildcards can be used in *volume-name*. For example, to increase the IOPerPath value to 100 for all `LogVol*` volumes, add the following lines to the `eql.conf` file:

```
[MPIO Volume Params LogVol*]
IOPerPath = 100
```

By default, SuppressPartitions is set to `False`, which enables you to use partitions on an EqualLogic device. Some software solutions, such as the XEN server, might not detect volumes correctly if a partitioned device is present. To suppress creating partition devices, set SuppressPartitions to `True`, which is similar to the *no-partitions* option when using Linux native multipath devices. For example, create no partition devices on `LogVol03` when users log in to the volume, add the following line to the `eql.conf` file:

```
[MPIO Volume Params LogVol03]
  SuppressPartitions = True
```

5

Introduction to ASM/LE

Dell EqualLogic Auto-Snapshot Manager/Linux Edition (ASM/LE) is a component of Dell EqualLogic Host Integration Tools for Linux. ASM/LE is installed as part of the HIT/Linux installation and presented as a configuration option.

Topics:

About ASM/LE

ASM/LE creates consistent copies of data stored on one or more PS Series groups to help prevent data loss and facilitate data recovery. You can copy a single volume or multiple volumes. The resulting collection of copied data is called a Smart Copy, which consists of the copy itself, as well as a backup document that describes the Smart Copy. When you configure your PS Series group for replication, you can create Smart Copies from ASM/LE to create a secure storage environment.

While you create Smart Copies, all applications remain online and performance is not impacted significantly. The time required for copying is minimized, and the data is always consistent and usable.

A Smart Copy can be one of several types: a snapshot, clone, thin clone, or replica (see Types of Smart Copies Created).

The Smart Copies created by ASM/LE are application consistent, which means that the data within the Smart Copies is consistent within the scope of the different transaction streams taking place between one or more applications when the Smart Copy was created.

While ASM/LE enhances and supplements your regular backup regimen by providing fast and efficient data recovery, it is not a replacement for a regular and complete backup of your data to long-term media. You can use your backup software to transfer the data in Smart Copies to long-term backup media. Because the applications in your production environment remain online during such transfers, ASM/LE significantly reduces your planned computer downtime.

ASM/LE provides a command-line interface from which you can perform the following operations:

- Create, list, and manage Smart Copies that consist of snapshots, clones, or replicas
- Create, list, and manage Smart Copy collections that contain multiple source mount points
- Create, list, and manage schedules to automatically perform Smart Copy operations at regular intervals
- Create, store, and manage group access files for automatic authentication of PS Series group identification information
- Mount a Smart Copy, log in to a Smart Copy, and unmount and log out of a Smart Copy
- Restore a Smart Copy

For information about ASMCLI command syntax and examples for using the CLI, see Auto-Snapshot Manager Command Line Interface.

About Smart Copies

A Smart Copy is a snapshot, clone, or replica of an individual volume, or a collection of volumes. You can also create Smart Copy thin clones of template volumes.

You use the ASMCLI `create smart-copy` command, along with optional parameters, to create the desired type of Smart Copy. You can create Smart Copies instantly, or you can create a schedule for Smart Copy creation using the `create schedule` command. For information about Smart Copy schedules, see Creating a Schedule.

When you create a Smart Copy, it consists of the following elements:

- One or more snapshots, replicas, or clones, depending on the type of Smart Copy operation and the original object
- A backup document, describing the Smart Copy

The following figure illustrates the Smart Copy creation process.

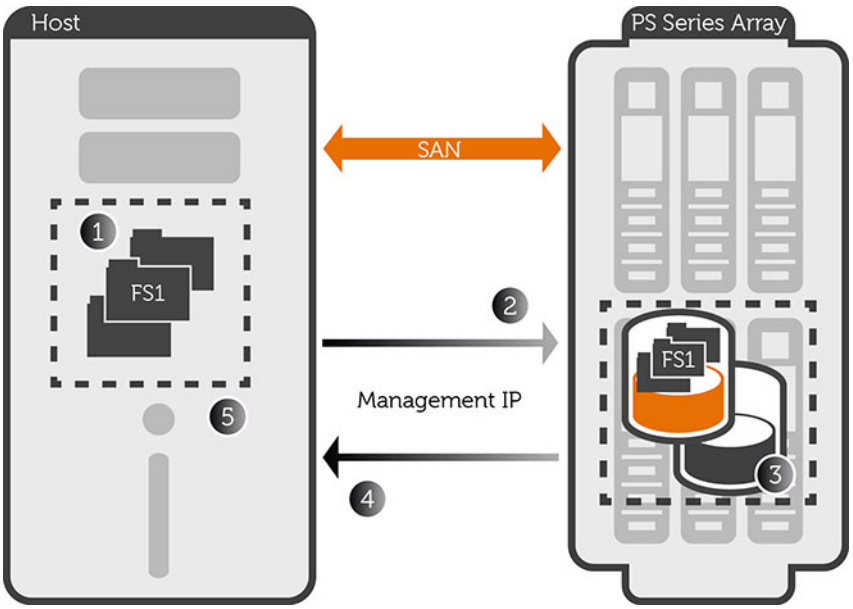

Figure 2. Smart Copy Creation Process

Table 15. Smart Copy Creation Process Callouts

Callout Number	Description
1	Freeze the file system
2	Request made to create a snapshot
3	Snapshot is created
4	Array sends snapshot complete message back to the host
5	File system is thawed

After you create a Smart Copy, you can import and mount it, and then restore the original objects from the Smart Copy.

When you create a Smart Copy with ASM/LE, the Smart Copies are visible in the PS Series Group Manager GUI and CLI. (The Group Manager is a web-based and Java application, easily accessible from Linux systems. It is also accessible as a CLI.) For troubleshooting, it might be necessary to access and manage a Smart Copy from the PS Series group using the Group Manager GUI. However, Dell recommends that you manage Smart Copies from ASM/LE and not the Group Manager GUI.

Types of Smart Copies Created

A Smart Copy can be a snapshot, clone, replica, replica clone, thin clone, or a synchronous replication.

Snapshot Smart Copies

A Snapshot Smart Copy is a point-in-time copy of a PS Series volume. Restoring a snapshot restores the volume to the state represented by the snapshot. By default, ASM/LE creates snapshot Smart Copies, unless you specify otherwise.

ASM/LE creates a snapshot for each volume comprising the original object. For example, if the original object is a volume on a PS Series group, the resulting Smart Copy Set will contain one snapshot. If the original object consists of a collection of two PS Series volumes, the resulting Smart Copy Set will contain two snapshots.

Replica Smart Copies

A Replica Smart Copy is a point-in-time copy of a PS Series volume, where the original volume and the replica are located on different PS Series groups for disaster tolerance. The groups and the volume must be configured for replication.

When replication begins, snapshots are created and copied to the destination replication site. The replication might not complete until much later.

To create a Replica Smart Copy, you must provide credentials for accessing the source volumes. To log in to or mount the completed replicas, you must provide credentials for the replication site partner (that is, the destination).

ASM/LE creates a replica for each volume comprising the original object on the PS Series group configured as a replication partner for the original object. Each volume that is part of the original object must already be configured for replication in the group.

Clone Smart Copies

A Clone Smart Copy is a new, independent volume containing the same data as the original volume at the time the clone was created. A clone is an independent volume that cannot be restored, but data items can be copied from a clone.

ASM/LE creates a new volume for each volume comprising the original object.

Replica Clones

A Replica Clone Smart Copy is created directly from a Replica Smart Copy, without promoting (mounting) the replica set or interrupting inbound replication. The Replica Clone resides in the replication partner and is accessible only on the local host. Replica clones support the same functions as clones (for example, list, mount, unmount, and delete). See Creating a Smart Copy for an example of creating a Replica Clone Smart Copy.

Thin Clone Smart Copies

A Thin Clone Smart Copy is a copy of a template volume, which is a read-only, thin-provisioned volume from which you can create thin clone volumes. A thin clone volume depends on its template volume. Template and thin clone volumes are useful when you need to create multiple volumes with common data. The common data can be written to a volume, which can be converted to a template volume, and any thin clones created from the template volume will have that common data. Then, you can modify each thin clone volume as needed. See the *Dell EqualLogic PS Series Group Manager Administrator's Guide* for more information on thin clones.

Synchronous Replication and Smart Copies

Synchronous replication (SyncRep) simultaneously writes volume data to two storage pools that belong to a PS Series group: an active pool (SyncActive) and an alternate pool (SyncAlternate). You can switch the pools at any time, making the active pool the alternate and the alternate the active pool.

NOTE: Configuring synchronous replication volumes and switching the roles of the two storage pools must be performed through the Dell EqualLogic Group Manager GUI. For more information, see the *Dell EqualLogic PS Series Group Manager Administrator's Guide*.

HIT/Linux Smart Copy snapshots can be created and managed only on the active storage pool. To create a snapshot on the alternate storage pool, switch the roles of the pools, then create the snapshot on the currently active storage pool.

Using the `asmcli` command, you can perform the following functions on the active storage pool: create (snapshots, clones, and thin clones only), delete, mount, restore (from a snapshot on the currently active storage pool), and unmount.

NOTE: To restore a snapshot Smart Copy that was created on the alternate storage pool, you must switch the roles of the pools, or the `asmcli restore` command fails. Smart Copies can be mounted or deleted from either volume, but restore operations are limited to snapshots created on the volume that is currently active.

Select a Smart Copy

If the set of all Smart Copies you created contains dozens or even hundreds of entries, finding the particular Smart Copy you need can become difficult. For example, the commands used to restore data from a Smart Copy require you to select an individual Smart Copy. If the number of available Smart Copies is large, finding and specifying an individual Smart Copy is an important function.

Many ASMCLI commands include search options to specify subsets of the available Smart Copies. Using the search options, you can narrow the set of available Smart Copies to locate Smart Copies:

- Created by a specific source mount point
- Contained in a specific collection
- Created by a specific host
- Created by a specified schedule
- Of a specified type (clone, replica, snapshot, or thin clone)
- That are the oldest or most recently created Smart Copy
- By Object ID, a unique identifier assigned to each Smart Copy and stored in the Smart Copy backup data

These options can be combined, for example, to search for a replica Smart Copy on a particular host.

You can also search by comment string or volume name. ASM/LE also performs a case-insensitive `glob` search and will match any string if you specify multiple strings:

- By the comment string used when the Smart Copy was created
- By the source target volume name

ASM/LE supports the following `glob` syntax:

- Asterisk (*) matches zero or more characters.
- Question mark (?) matches exactly one character. You can combine a question mark with the asterisk (?*) to match any string with one or more character.
- Square brackets ([ ]) specify a group of possible characters.
- Backward slash (\) indicates an escape character to identify a special character (for example,* ? ^ !]) as a normal character.

For more information about `glob` pattern matching, see the `glob(7)` manpage. For ASMCLI search command descriptions and examples, enter:

```
# asmcli help selection
```

See ASMCLI Search Options for a complete list of search options and associated parameters.

Requirements for Creating Smart Copies

The following requirements must be met before you can create a Smart Copy:

- The computer must meet the requirements in Environment Prerequisites.
- The computer must be logged into all iSCSI volumes that comprise the Smart Copy object.
- You must have an administrator identity on the PS Series group. Set up pool or volume administrators to avoid using the group administrator account. To create an administrator identity, see the *Dell EqualLogic Group Manager Administrator's Guide*.
- The Linux mount points must be defined as PS Series volumes and no other device types.
- To create a snapshot Smart Copy, the PS Series volume must have sufficient snapshot space reserved.
- To create a replica Smart Copy, a replication partnership must be configured between PS Series groups, and replication must be enabled on the volumes. The primary group must be available and ready to start replication, or the operation causes an error. ASM/LE does not need access to the secondary group. For information on configuring a replication partnership with the Group Manager, see the *Dell EqualLogic Group Manager Administrator's Guide*.
- To create a clone Smart Copy, there must be free space on the PS Series group equal to the size of the original volumes.
- To create a Smart Copy of a volume or collection that affects cluster resources, you must perform the ASM/LE operation from the node that owns the cluster resources.
- You cannot take cluster snapshots of `ocfs2` file systems.

About Snapshot Space Borrowing

Snapshot space borrowing temporarily increases the available snapshot space for a volume by borrowing from free pool space and the snapshot reserve of other volumes. This increase enables you to prevent the oldest snapshots in your collection from being automatically deleted. If the borrowed space is needed for other functions, such as creating a new volume or increasing the size of existing snapshots, the PS Series group deletes snapshots created using borrowed space.

Snapshot borrowing is enabled through group configuration but can be enabled and disabled for an individual volume. The `asmcli` command detects if snapshot space borrowing is enabled on the volume and if the snapshots are using borrowed space. When snapshots are using borrowed space, `asmcli` notifies you that snapshot space borrowing is occurring, and warn you if the borrowed snapshot space has crossed the defined threshold limits.

Methods of Creating Smart Copies

You can use ASM/LE to create Smart Copies in the following ways:

- On demand—Select a volume, mount point, or collection to immediately create a Smart Copy of it.
- Automated schedule—Select a volume or collection and then create a Smart Copy schedule for that object. A schedule will automate Smart Copy creation and allow you to control the timing, frequency, and number of retained copies.
- Linux Script— Incorporate ASM/LE into a Linux script using a variety of scripting languages.

About Group Access Files

To successfully access a PS Series group, administrators require network access to the group IP address and a valid group identity (that is, administrative user name and password). Different account types provide different privileges. The default account, `grpadmin`, provides all privileges. Credentials can also be set up for pool and volume administrators, restricting access to the entire group. Dell recommends using volume administrator credentials.

When you connect to PS Series groups from the Auto-Snapshot Manager command line interface (ASMCLI) to create Smart Copies, some commands require administrator credentials. This requirement might not be desirable in secure environments where user names and passwords are to be hidden.

ASM/LE lets you create a group access file to store administrative identities for the group with the ASMCLI `create group-access` command (see Creating Access to a PS Series Group). This file contains two-way encrypted user names and passwords for administrative accounts on PS Series groups, the group name, the group IP address, and optionally, the management IP address. When you create a group access file, ASMCLI does not require you to give the group name, administrator name, and password each time you create a Smart Copy.

If you do not want ASM/LE to store a password locally, you can override password encryption with the `--no-password` option to the `create group-access` command. When you create a Smart Copy with the `create smart-copy` command, ASM/LE will prompt you for the password.

ASM/LE decrypts the necessary passwords when making connections to the PS Series group. If a group password has been stored, when you execute an ASMCLI command that requires credentials, ASM/LE locates the group credentials, decrypts the password, and makes the connection.

ASMCLI provides commands to create a group access file, list passwords in the file, modify the group access file, and delete password information. For information and examples of using these commands, see Creating Access to a PS Series Group.

About Collections

Collections are related sets of mount points from which you create Smart Copies. Collections are useful when you want the Smart Copies to be created simultaneously in one set. For example, if two or more file systems (multiple mount points) are being used to support a given application, you could group them into a collection.

To create a collection, use the ASMCLI `create collection` command. You must specify a name for the collection and one or more mount points. For example:

```
# asmcli create collection --name=sample --source=/user1 /user2
```

All Smart Copy operations have a collection as the source. If you create a Smart Copy with only a single volume, ASM/LE regards it as a collection with a single mount point.

When creating a Smart Copy snapshot of the collection, ASM/LE freezes all applicable mount points, creates the necessary snapshots, and then thaws the mount points.

Creating a Smart Copy of a collection creates a set of Smart Copies that freeze and thaw the same mount points. ASM/LE validates the collection defining this Smart Copy set. For example, if you create a Smart Copy set from a collection of two LVM logical volumes, mount it, and then add a third logical volume to the same virtual group, you invalidate the collection. The mount point of the third volume cannot be safely frozen and thawed, so the Smart Copy operation fails.

NOTE: Make sure the elements of a collection support your intended backup and restore plans. Dell recommends creating collections only for objects that are logically related to restore as a group. Avoid creating collections that contain volumes for multiple databases, or that include database volumes and volumes used by other, unrelated applications.

You can define collections for specific tasks. For example, you have a cron job that runs every hour to copy several mount points. The collection for that job might be:

```
# asmcli create collection --name=hourly-user-vols --source="/user1, /user2, /user3"
```

With ASM/LE, you can list all Smart Copies created with a given collection. For example,

```
# asmcli list collection --name daily-vols
```

You can modify a collection by changing its name, by replacing all existing mount points, by adding one or more mount points, or by deleting it. For example,

```
# asmcli update collection --name daily-user-vols --add /source/VolC
```

For information about using ASMCLI to perform these collection operations, see Creating and Managing Collections.

About Smart Copy Schedules

Schedules perform Smart Copy operations at regular intervals. They are predefined Smart Copy creation commands that you create once and use repeatedly. You can also control how many Smart Copies ASM/LE preserves.

To create a Smart Copy schedule, use the `asmcli create schedule` command and specify a schedule name and a list of one or more mount points using the `--source` or `--collection` parameters.

Advantages of Smart Copy Schedules

Using Smart Copy schedules provides several advantages over manual Smart Copy creation:

- You can name commonly used commands for creating a Smart Copy to indicate its purpose.
- You can incorporate schedules into scripts or cron jobs.
- You can include the scheduled operation in a script or cron table, and change schedules without modifying any scripts or cron tables.
- You can find Smart Copies that were created by a particular schedule using the `list smart-copy --schedule` command in addition to other commands.
- You can limit the number of Smart Copies a schedule keeps by using the `--keep` option. By default, ASM/LE keeps 12 Smart Copies. If you specify a value of 10 for `--keep`, when the schedule creates copy 11, ASM/LE either deletes copy 1, if that copy is not in use (logged in or mounted), or the oldest copy that is not in use.

  NOTE: Snapshot and replica Smart Copies are also limited by the snapshot reserve and replica reserve configured on the PS Series group. See the *Dell EqualLogic Group Manager Administrator's Guide* for details.

For information about commands for schedule operations, see Creating and Managing a Smart Copy as a Scheduled Operation.

Constraints of Smart Copy Schedules

There are several constraints to consider when using Smart Copy schedules:

- You can only schedule snapshot or replica Smart Copy operations. You cannot create clones by using a schedule.
- While you can incorporate scheduled operations into scripts or cron jobs and enable them, creating Smart Copies every few minutes could adversely affect performance. ASMCLI enforces a minimum elapsed time limit between using a schedule, and Dell recommends specifying at least 10 minutes for cron jobs that run a Smart Copy schedule.
- By default, ASMCLI keeps 12 schedules. You can use the `--keep` option in the `asmcli create schedule` command to modify the number of schedules retained. If the `--keep` value is set to 12, when the schedule creates copy 13, ASMCLI either deletes copy 1 if that copy is not in use (logged in or mounted), or deletes the oldest copy that is not in use. The deleted Smart Copies result in a broken Smart Copy, because the backup document refers to snapshots or replicas that are no longer on the PS Series group.

Data Restoration

Data restoration involves mounting a Smart Copy, restoring data from it, and then unmounting and logging off the Smart Copy.

How you access or restore data from a Smart Copy depends on the original object (volume or collection) and the result of the Smart Copy operation (snapshot, replica, or clone).

Collections and applications can have multiple components (for example, multiple volumes or databases spanning multiple volumes). Based on the components of the original object, a Smart Copy can include one or more snapshots, replicas, or clones created at the same time,

Options for accessing data include:

- Quickly restore the original object from a Smart Copy. Applies to Smart Copies that contain volume snapshots.
- Mount a Smart Copy as read-only. Applies to replicas, or Smart Copies that contain replicas of volumes.
- Restore the original object in a new location. Applies to Smart Copies that contain snapshots and clones.

If you make changes to a volume layout and have a Smart Copy that predates the layout changes, you cannot use that Smart Copy to recover data. Attempting to recover the data might cause data loss on the production volume.

You can mount or restore data from a Smart Copy of a volume using mount points referenced when you created the Smart Copy. Restored mount points reference the same volumes they did at the time you created the Smart Copy, unless you manually change the restored mount points to reference different volumes, such as another mounted Smart Copy.

About Mounting Smart Copies

Mounting a Smart Copy makes its contents accessible to the computer. If you are restoring certain files or a subset of files contained in a Smart Copy, you must mount it first.

Mounting a Smart Copy supports mount points, physical locations in the partition that server uses as an access point for a new mounted file system. To access data in a snapshot or clone Smart Copy created from volumes or application components, you first mount it.

To mount a Smart Copy, use the `asmcli mount smart-copy` command and specify the directory where the Smart Copy is to be temporarily mounted. You can optionally specify the earliest or latest created Smart Copy, or a list of one or more mount points to mount. After a Smart Copy is mounted, you can access the data that you want to restore on it.

For more information about using ASMCLI to mount a Smart Copy, see Mounting a Smart Copy.

Before you mount a Smart Copy:

- Ensure that the computer has the security credentials to access the Smart Copies in the set.
- Identify the snapshot or Smart Copy set that you want to mount.
- Identify the directory in which to mount the file systems.

When you mount a Smart Copy, ASM/LE performs the following actions:

- Sets the Smart Copy online on the PS Series group, making it an accessible iSCSI target
- Logs in sessions to the targets
- Mounts the file system to the specified location

About Unmounting Smart Copies

When you are finished using a mounted Smart Copy, unmount the Smart Copy using `unmount smart-copy`. For information about `unmount smart-copy`, see Unmounting a Smart Copy.

Before unmounting a Smart Copy:

- Ensure that the computer has the security credentials to access the Smart Copies in the set.
- Identify the Smart Copy or Smart Copy set that you want to unmount.
- Stop using the Smart Copy.

The `unmount smart-copy` command performs the following actions:

- Unmounts the file systems
- Logs off all sessions to the Smart Copy targets
- Sets the Smart Copy targets offline

About Logging In to Smart Copies

Logging in to a Smart Copy lets you access data without mounting a file system, enabling you to access data even when you are unable to mount a file system. The `login smart-copy` command performs the same functions that the `mount smart-copy` command does, without mounting the file system. The Smart Copy data is set online, sessions are logged in to the targets, and devices are made available. For information on `login smart-copy`, see Logging In to a Smart Copy.

Before logging in to a Smart Copy:

- Ensure that the computer has the security credentials to access the Smart Copies in the set.
- Identify the snapshot or Smart Copy Set that you want to log in to.

About Logging Off Smart Copies

When you are finished with a Smart Copy that is logged in, the `logout smart-copy` command logs out all sessions to the Smart Copy targets and sets the targets offline. For information about `logout smart-copy`, see Logging Out of a Smart Copy.

Before you log off:

- Ensure that the computer has the security credentials to access the Smart Copies in the set.
- Identify the snapshot or Smart Copy Set that you want to log off.
- Stop using the Smart Copy.

About Restoring Data From a Smart Copy

If a catastrophic system failure results in data loss, it might be necessary to restore data from a Smart Copy. If the Smart Copy contains a copy of a volume (a mount point or set of mount points) on a PS Series group, this operation will reset the data on the volume to the previous state. If the Smart Copy is a consistent copy of multiple mount points and multiple volumes, the restore operation will restore all of the mount points and all of the volumes to the earlier state.

When you restore a complete volume, it is called an in-place restore. For snapshot Smart Copies of volumes or volume collections, the restore operation replaces the current data in the volumes or volume collection with the data in the Smart Copy. The Smart Copy continues to exist and you can restore from it as often as necessary.

The `asmcli restore smart-copy` command provides options for identifying the exact data to restore from a set of mount points used to create the Smart Copy. You can also conveniently restore from the newest or oldest Smart Copy data. See Restoring a Smart Copy.

Preconditions for Restoring Data

Before you restore data, several preconditions must be met:

- The Smart Copy must be a snapshot type.
- The Smart Copy cannot already have been logged in or mounted by ASM/LE.
- The EqualLogic volumes to be restored can only have sessions logged in by the host doing the restore in place. Other systems cannot have sessions logged in. As part of the restore operation, ASM/LE unmounts all mount points specified by the Smart Copy. If a file system is in use, the unmount and the restore operations fail.
- All of the mount points (directories) in the Smart Copy must exist on the host.
- If a new mount point has been added to the EqualLogic volumes, the mount point must be unmounted before the restore operation. Because the new mount point did not exist before the Smart Copy was created, the new file system will be lost when the restore completes.

Restoring Data Process

When you restore data from a Smart Copy using the `asmcli restore smart-copy` command, ASM/LE automatically performs the following tasks:

1. Unmounts the mount points. If processes are still using any of the file systems, the restore operation fails.
2. Deletes all Device Mapper (DM) devices used to support the EqualLogic volumes
3. Logs out all sessions to the EqualLogic volumes
4. Sets the EqualLogic volumes offline
5. Restores the EqualLogic volumes using the specified snapshots. (The group automatically creates a new snapshot that represents the current state of the volume.)
6. Sets the EqualLogic volumes online
7. Logs in sessions to the EqualLogic volumes
8. Rebuilds the DM devices to support the EqualLogic volumes
9. Mounts the file systems

About Restoring a Soft-Deleted Smart Copy

When a Smart Copy clone is deleted, the clone is soft-deleted and moved to a volume recovery bin, enabling the clone to be restored. As long as the clone is in the recovery bin, a backup document is retained. The `asmcli list smart-copy --check-status` command uses the backup document to identify the clones that have been soft-deleted.

NOTE:

- **Restoring soft-deleted volumes can be performed only for Smart Copy clones running PS Series array firmware version 6.0 or later. Deleting a Smart Copy clone that is running an earlier version of the firmware immediately deletes that clone and does not preserve the backup document for that Smart Copy.**

- **If a volume is running a pre-6.0 version of the firmware, updating that volume to the current of the firmware enables the soft-delete feature on that volume.**
- **The `asmcli list smart-copy --check-status` command can display soft-deleted Smart Copies, but it cannot restore or purge them. To restore or purge a soft-deleted Smart Copy, use the Recover button in the Volume Recovery Bin interface of the EqualLogic Group Manager GUI.**

Deleting a Smart Copy clone stores that clone in the volume recovery bin until one of the following events occur:

- The bin is filled by subsequent Smart Copy clone deletions. If the bin is full, the clone stored first is permanently deleted as each newer clone is stored.
- Seven days have passed since the clone was stored in the bin.
- You manually either restore or purge a clone from the bin.

NOTE: When a clone is purged from the volume recovery bin, either manually or automatically, no confirmation message is given. Use the `asmcli list smart-copy --check-status` command to review the current state of a Smart Copy created by ASMCLI and list that Smart Copy if it is a soft-deleted Smart Copy clone.

If you change credentials so that the volume is no longer accessible, ASM/LE cannot see the volume in the recovery bin. As such, if you enter a `list smart-copy --check-status` command, this action permanently deletes the Smart Copy. This behavior is identical to the way that changing credentials can create unreachable Smart Copies. However, if you delete credentials such that you can no longer connect to the PS Series group, a `list smart-copy --check-status` command does not purge the Smart Copy backup documents. To resolve this, you must define new credentials to connect to the PS Series group; even if those credentials are volume administrator credentials that do not allow access to any volumes. The `list smart-copy --check-status` command can then clean up the old Smart Copies, and the credentials can be revoked.

The size of the volume recovery bin is determined by the free pool space on the system. Smart Copy clones that are larger than the available free pool space are permanently deleted. After a clone is permanently deleted from the volume recovery bin, that clone can no longer be restored.

About Managing Smart Copy Clones After Changing or Deleting Credentials

For groups running PS Series firmware version 6.0 and later, if you create a Smart Copy clone and delete it, the firmware places the deleted clone in a recovery bin. The deleted volume can be restored by the PS Series group or, eventually, it automatically is purged. Depending on which action occurs, entering an `asmcli list smart-copy --check-status` command either restores or permanently deletes the Smart Copy. However, if you change credentials so that the volume is no longer accessible, ASM/LE cannot see the volume in the recovery bin, and when you enter `list smart-copy --check-status`, ASM/LE permanently deletes the Smart Copy. This behavior is identical to the way that changing credentials can create unreachable Smart Copies.

If you delete credentials so that you cannot connect to the PS Series group at all, entering an `asmcli list smart-copy --check-status` command does not purge the Smart Copy backup documents. In that case, ASM/LE cannot clean up the Smart Copy backup documents.

Example

Create a Smart Copy clone of volume `vol2`:

```
# asmcli create smart-copy --source /mnt/vol2 --type clone
Dell EqualLogic Auto-Snapshot Manager CLI Version 1.7.0
Copyright (c) 2010-2020 Dell Inc.

The mount point has been frozen.
Created clone rhel69x-vol2-clone-2017:11:21-11:25:13.375 on eqlgroup
The mount point has been thawed.
Successfully created a Smart Copy from 1 target.
```

Some time later, the Smart Copy clone is deleted. The Smart Copy is not displayed as part of a normal `asmcli list smart-copy` command; it is hidden.

```
[root@rhel69x ~]# asmcli delete smart-copy --latest
Dell EqualLogic Auto-Snapshot Manager CLI Version 1.7.0
Copyright (c) 2010-2020 Dell Inc.

Deleted Smart Copy: 5-82cd1c-c445d3028-eecf8fc826db539a
Successfully removed the Smart Copy.

[root@rhel69x ~]# asmcli list smart-copy --type clone
Dell EqualLogic Auto-Snapshot Manager CLI Version 1.7.0
Copyright (c) 2010-2020 Dell Inc.

============================================================
Smart Copies
============================================================

Source: rhel69x.mycompany.com : /mnt/vol1
  Clone: 25-Sep-2017 10:16:03.150384
  ObjectId: 9-ad0bde-e441bbe0d-da25bffc416638b4
    Rhel63x-vol1--> KevinO-a08-clone-2017:09:25-10:16:03.150 on eqlgroup

1 Smart Copy found.
```

When you enter an `asmcli list smart-copy --type clone --check-status` command, the Smart Copy is displayed and the status of the volume indicates it is in the recovery bin.

```
# asmcli list smart-copy --type clone --check-status
Dell EqualLogic Auto-Snapshot Manager CLI Version 1.7.0
Copyright (c) 2010-2020 Dell Inc.

============================================================
Smart Copies
============================================================
Source: rhel69x.mycompany.com : /mnt/vol1
  Clone: 25-Sep-2017 10:16:03.150384
  ObjectId: 9-ad0bde-e441bbe0d-da25bffc416638b4
     rhel63x-vol1--> rhel69x-vol1-clone-2017:09:25-10:16:03.150
on eqlgroup

Source: rhel69x.mycompany.com : /mnt/vol2
  Clone: 21-Nov-2017 11:25:13.375355
  Status: Warning: smart copy is unavailable.
  ObjectId: 5-82cd1c-c445d3028-eecf8fc826db539a
    rhel69x-vol2 --> rhel63x-vol2-clone-2017:11:21-11:25:13.375 on eqlgroup\
       [recovery bin]

2 Smart Copies found.
```

At this point, you can restore the volume on the PS Series group, which moves the volume out of the recovery bin and back into the list of normal volumes. The next `asmcli list smart-copy --check-status` command you enter restores the Smart Copy and the status of the Smart Copy changes from hidden to a normal state. If you then enter `asmcli list smart-copy`, the Smart Copy is visible and you can mount the Smart Copy.

You can also manually purge the volume from the Group Manager, which deletes it from the recovery bin. (Over time, the PS Series group automatically purges the volume if you do not manually purge it.) After the volume has been removed from the recovery bin, a subsequent `asmcli list smart-copy --check-status` command causes ASM/LE to delete the Smart Copy completely, because the volume no longer exists.

A serious issue can occur when you delete a Smart Copy clone and then enter an `asmcli list smart-copy --check-status` command with different credentials. This situation is further complicated if you use different credentials that prevent access to volumes in the recovery bin.

In this example, the user's default credentials provided volume administrator (voladmin) access to all volumes for the local host (rhel69x). When you enter an `asmcli list smart-copy - -check-status` command and specify a different set of voladmin credentials,

you can connect to the PS Series group, but access to the volumes for rhel69x is not allowed. As a result, the volumes for rhel69x are hidden.

```
# asmcli list smart-copy --type clone --check-status --group-access-file \
  voladmin-sles
  Dell EqualLogic Auto-Snapshot Manager CLI Version 1.7.0
  Copyright (c) 2010-2020 Dell Inc.

============================================================
Smart Copies
============================================================

Source: rhel69x.mycompany.com : /mnt/vol1
Clone: 25-Sep-2017 10:16:03.150384
Status: Error: smart copy is unusable.
ObjectId: 9-ad0bde-e441bbe0d-da25bffc416638b4
    rhel69x-vol1--> rhel69x-vol1-clone-2017:09:25-10:16:03.150 on eqlgroup

1 Smart Copy found.
```

In this case, Smart Copy snapshots appear to be broken and Smart Copy clones appear unreachable. When ASM/LE processes deleted Smart Copy clones, the volumes are not visible in the recovery bin. ASM/LE cannot locate them and assumes that the volumes have been purged. ASM/LE then deletes the Smart Copy completely. You can use the Group Manager to restore the volume in the PS Series group. However, subsequently entering an `asmcli list smart-copy - -check-status` command does not return the Smart Copy to its normal state because the Smart Copy was deleted.

NOTE: Dell recommends that you do not change credentials when using the `asmcli list smart-copy --check-status` command. Credentials used to create a Smart Copy will correctly handle all other operations. If you use a different set of credentials, be careful when checking for broken, unreachable, or deleted Smart Copy clones. For example, if a single Smart Copy storage location is being used to save all of the Smart Copy documents for a facility, a single host can be used to maintain the set of documents. That host might have pool or group admin privileges for its access to the PS Series groups, meaning that the host has access to all volumes.

Best Practices For Using ASM/LE

Dell recommends the following best practices:

- When selecting the elements of a collection, choose only elements that are logically related. Avoid creating collections that include volumes used by other, unrelated applications. Make sure your collections support your intended backup and restore plans.
- Do not implement a large number of schedules.
- When working with schedules, develop them in stages. Begin with an `asmcli create smart-copy` command that you know works, then cut and paste the Smart Copy information to create the corresponding schedule. Test the schedule from the command line using `asmcli do schedule`. Lastly, add the schedule to a crontab or equivalent.
- Use pool and volume administration credentials instead of group administration credentials for each host or related hosts.
- When creating a Smart Copy with the `asmcli create smart-copy` command or mounting a Smart Copy with the `asmcli mount smart-copy` command , use the `--retry-count` option in scheduled operations and scripts. If necessary, you can specify a longer time interval to wait before retrying with the `--retry-wait` option. If an intermittent failure occurs, ASM/LE will try as many times as you indicate, increasing the likelihood of success.
- Use the `asmcli list smart-copy` command to identify a specific Smart Copy before performing other actions, such as deleting a Smart Copy.

6

Auto-Snapshot Manager Command Line Interface

The Auto-Snapshot Manager Command Line Interface (ASMCLI) provides Smart Copy creation and management commands that you execute at the Linux command prompt.

Topics:

Introduction to ASMCLI

ASMCLI presents commands that enable you to create and manage:

- Smart Copies that consist of snapshots, clones, or replicas. These commands also enable you to list, mount/unmount, log in/log out, and restore data from Smart Copies.
- Smart Copy collections that contain snapshots, clones, or replicas
- Group access file for storing PS Series group authentication information
- Smart Copy schedules

For conceptual information about these features, see Introduction to ASM/LE.

Entering ASMCLI Commands

Commands are executed from the Linux command prompt in one of two ways:

- Manually typing the command, using command help to see options that complete a command (see General Command Syntax for ASMCLI and Access Help on ASMCLI Commands)
- Using Bash command completion to interactively complete a command (see Appendix C Bash Command Completion)

General Command Syntax for ASMCLI

The general ASMCLI command syntax is as follows:

```
# asmcli command-verb command-noun [--parameter [argument]]
```

The following rules apply:

- ASMCLI command syntax uses verb-noun pairs to form a complete command. A complete command must include both a command-verb and a command-noun. For example, `asmcli create smart-copy` includes the verb *create* and the noun *smart-copy*.
- When entering a command verb or noun, type the entire command-verb followed by a tab for ASMCLI to display all applicable noun commands. Enter part of any command followed by two tabs and ASMCLI, using Bash command completion, for a list of applicable options. For information, see Appendix C Bash Command Completion.
- When specifying parameters, enter two hyphens followed by the parameter. Enter part of the parameter name followed by two tabs and Bash command completion for a list of applicable parameters.

- Use of the equal sign (=) to specify a value to parameters is optional. For example, `--name=eqlgroup1` and `--name eqlgroup1` are equivalent.
- Not shown in the syntax for an ASMCLI command is the `--verbose` option, which you can use to display more detailed information. For more verbose output information, specify `--verbose 2` or simply `-vv`.
- Use the `--help` option to display arguments to parameters. Use `asmcli help` to display all command noun-verb pairs and options. For more information about command-line help, see Access Help on ASMCLI Commands.
- Brackets ([]) indicate that a parameter or argument is optional.
- Several parameters are mutually exclusive and if more than one parameter is specified, it is treated as an error and causes the command to fail.
- The command parser takes the latest value entered when redundant options are given.

Access Help on ASMCLI Commands

From the Linux command prompt, to access general help on ASMCLI commands, enter:

```
# asmcli help
```

For information on using the ASMCLI to create Smart Copies, using Bash command completion, and creating group access files, enter:

```
# asmcli help begin
```

For information on using ASMCLI to create Smart Copy collections, enter:

```
# asmcli help collection
```

For information on setting up Smart Copy schedules, enter:

```
# asmcli help schedule
```

For information on using ASMCLI to store authentication information (group names, group IP addresses, optional management IP addresses, user names, and passwords) used when creating Smart Copies and collections, enter:

```
# asmcli help group-access
```

For information on search options are available for several ASMCLI commands, enter:

```
# asmcli help selection
```

ASMCLI Search Options

The ASMCLI provides search option parameters that are helpful when you need to locate a Smart Copy among dozens or hundreds. The search options are common to the following ASMCLI commands, unless otherwise noted in Table 16. ASMCLI Search Options. (This table lists the common search options, applicable arguments, and describes how each is used.)

- `delete smart-copy`
- `list smart-copy`
- `login smart-copy`
- `logout smart-copy`
- `mount smart-copy`
- `unmount smart-copy`
- `restore smart-copy`

Syntax for the search options is as follows:

```
# asmcli smart-copy command
  [--collection collection-name | -c]
  [--comment comment-string | -C}
  [--copy-group destination-group-name]
  [--copy-volume destination-volume-name]
  [--group group-name
  [--earliest | -e]
  [--hostname host-name | -H]
  [--in-use ]
  [--latest |-l]
  [--mounted ]
  [--not-in-use ]
  [--not-mounted ]
  [--object object-id | -o]
```

```
[--schedule schedule-name | -S]
[--source mount-point1 [mount-point2...] -s]
[--type volume-type | -t]
[--volume volume-name]
```

Table 16. ASMCLI Search Options

Search Option Parameter	Abbr.	Argument	Description
--collection	-c	*collection-name*	Search by the name of a Smart Copy collection. Do not use with `--source`.
--comment	-C	*comment-string*	Search by the comment string you entered when creating a Smart Copy. The comment is set for the PS Series group and stored in the backup document. By default, ASM/LE performs a case-insensitive `glob` search and matches any string if you enter multiple strings. For more information about `glob` pattern matching, see Select a Smart Copy.
--copy-group	none	*destination group-name*	Search by the destination group name. By default, ASM/LE performs a case-insensitive `glob` search and matches any string if you enter multiple strings.
--copy-volume	none	*destination volume-name*	Search by the destination volume name. By default, ASM/LE performs a case-insensitive `glob` search and matches any string if you enter multiple strings.
--group	none	*group-name*	Search by the PS Series group name. By default, ASM/LE performs a case-insensitive `glob` search and matches any string if you enter multiple strings.
--earliest	-e	none	Search for Smart Copies by the earliest (oldest) Smart Copy.
--hostname	-H	*host-name*	Search for the name of the host that created the Smart Copy.
--in-use	none	none	Search for Smart Copies that are logged in or mounted by ASM. **NOTE: The search option does not apply to the `login smart-copy` command.**
--latest	-l	none	Search for Smart Copies by the latest (newest) Smart Copy.
--logged-in	none	none	Search for Smart Copies that are logged in but not mounted by ASM. **NOTE: The search option does not apply to the `login smart-copy` command.**
--mounted	none	none	Search for Smart Copies that are mounted by ASM. **NOTE: The search option does not apply to the `login smart-copy` or `mount smart-copy` commands.**
--not-in-use	none	none	Search for Smart Copies that are not logged in or mounted by ASM.
--not-logged-in	none	none	Search for Smart Copies that are not logged in by ASM.
--not-mounted	none	none	Search for Smart Copies that are not mounted by ASM. **NOTE: The search option does not apply to the `login smart-copy` or `mount smart-copy` commands.**
--object	-o	*object-id*	Search by an object ID that uniquely identifies a replica Smart Copy. Use the `list smart-copy` command to locate the Object ID.
--schedule	-S	*schedule-name*	Search for the schedule used to create the Smart Copy.
--source	-s	*mount-point*	Specifies one or more mount points to copy when creating the Smart Copy. Do not use with `--collection`.

Search Option Parameter	Abbr.	Argument	Description
--type	-t	*volume-type*	Search by type of Smart Copy: either `clone`, `replica`, `snapshot`, or `thinclone`. **NOTE: The search option does not apply to the `restore smart-copy` command.**
--volume	none	*volume-name*	Search by the source volume name. By default ASM/LE performs a case-insensitive `glob` search and matches any string if you enter multiple strings.

You can combine search options to narrow down a long list of Smart Copies. For example:

```
asmcli list smart-copy --collection accounting --mounted
```

You can also identify the oldest mounted Smart Copy for that collection as follows:

```
asmcli list smart-copy --collection accounting --mounted --earliest
```

Creating and Managing Group Access Files

Authentication files contain two-way encrypted user names and passwords for administrative accounts on PS Series groups. These credentials are used by ASM/LE to connect to PS Series groups when needed to support ASM/LE functions, such as Smart Copy creation. When you create a group access file, ASMCLI does not require that you specify the group name, administrator name, and password each time you create the Smart Copy.

NOTE: When creating an administrator identity on the PS Series group, you can avoid using group administrator credentials by establishing pool or volume administrators. See the *Dell EqualLogic Group Manager Administrators Guide* for information on setting up pool and volume administration credentials.

This section provides the following information:

- Creating Access to a PS Series Group
- Updating Access to a Group
- Listing Group Access Information
- Deleting Group Access Information

Creating Access to a PS Series Group

Use the `asmcli create group-access` command to create a group access file to store your PS Series group administration credentials (group name, group IP address, optional management IP address, user name, and password).

asmcli create group-access Command Syntax

The command syntax is as follows:

```
# asmcli create group-access
       --name group-name | -n
       --user-name user-name | -u
      [--file file-location | -f]
      [--ip-address IP-address | -i]
      [--management-ip mgmt-IP-address | -I]
      [--no-password]
      [--password password | -p]
```

The following table describes the required and optional command parameters.

Table 17. Create Group Access Command Options

Parameter	Abbr.	Argument	Description
Required Parameters			

Parameter	Abbr.	Argument	Description
--name	-n	*group-name*	Specifies the name of the PS Series group.
--user-name	-u	*user-name*	Specifies the user name on the PS Series group.
Optional Parameters			
--file	-f	*file-location*	Specifies the location of the group access file.
--ip-address	-i	*IP-address*	Specifies the permanent IP address of the PS Series group. You must use parameter when creating a new entry for the PS Series group.
--management-ip	-l	*mgmt-IP-address*	Specifies the management IP address for the PS Series group
--no-password	none	none	Specifies not to store a password for this user. If you use this option, you will be prompted to enter a password each time an ASMCLI command is run (for example, `create smart-copy`) and you are not able to run a scheduled operation.
--password	-p	*password*	Specifies the password to be encrypted and stored.

Examples: Creating Group Access

This section includes examples of creating group access records.

Example 1: Creating Group Access for Volumes on Multiple Groups

In this example, a system uses volumes on three different PS Series groups. You need to store credentials for all three PS Series groups to use ASM/LE for all of the volumes, as follows:

```
# asmcli create group-access  --name eqlgroup1 --ip-address
  192.x.x.xx  --user-name j_smith

# asmcli create group-access --name eqlgroup2 --ip-address
  192.x.x.xy --user-name j_smith

# asmcli create group-access --name eqlgroup6 --ip-address
  192.x.x.xz  --user-name John_Smith
```

If the `/user1` and `/user2` file system mount points are built by volumes on eqlgroup1 and if the credentials you provided are valid, you can create a Smart Copy for those file systems without specifying a user name and password for the operation:

```
# asmcli create smart-copy --source /user1 /user2
```

If you try to perform an ASMCLI operation that requires administrator access to a PS Series group without providing valid credentials, you will get an error. For example, if the file system mount points `/AccPay` and `/AccRecv` are on a member that is not included in the group access file, ASMCLI does not have the necessary credentials to create Smart Copies for those volumes:

```
# asmcli create smart-copy --source /AccPay /AccRecv
Error: Cannot connect to PS Array 192.x.x.xx : No group entry for that IP address in the
group access file.
```

Example 2: Specifying a Management IP Address

In this example, the `--management-ip` option is used to create a group access file for PS Series groups that are configured with separate IP addresses for management:

```
asmcli create group-access --name eqlgroup1 --ip-address
  192.x.x.xx --management-ip eqlmgmt1.mycompany.com
  --user-name j_smith
Password for j_smith on eqlgroup4:
```

Example 3: Managing Multiple Identities

It might be necessary to use different credentials for different ASM/LE operations. You can manage multiple identities in two ways:

- Using multiple group access files
- Inserting multiple user names into a group access file and selecting the user name for a given operation on the ASMCLI command line

In this example, volume administrator `user-admin` is an administrator for several user volumes on PS Series group, eqlgroup1. That PS Series group has volumes used by the accounting applications. Those volumes are administered by volume administrator `accnt-admin`. The system administrator creates a group access record in the default group access file, `/etc/equallogic/asm-group-access` for `user-admin` to manage Smart Copies for the user volumes. The system administrator uses the default group access file for all ASMCLI operations for the user volumes, as follows:

```
# asmcli create group-access --name eqlgroup1 --ip-address
  192.x.x.xx --user-name user-admin
  Password for user-admin on eqlgroup1:

# asmcli create smart-copy --source /user1 /user2
```

To manage the accounting application volumes, the system administrator creates a separate group access file for administrator `accnt-admin`, as follows:

```
# asmcli create group-access --name eqlgroup1 --ip-address
  192.x.x.xx \--user-name accnt-admin --group-access
  ~/.asm-group-access-accnt
  Password for accnt-admin on eqlgroup1:

# asmcli create smart-copy --source /acc-recv /acc-pay --group-access
  ~/.asm-group-access-accnt
```

Alternatively, you can insert both sets of credentials into the same file and select which name to use on the ASMCLI command line. In this case, you use the `--user-name` option with the `create smart-copy` command, as follows:

```
# asmcli create group-access --name eqlgroup1 --ip-address
  192.x.x.xx --user-name user-admin
  Password for user-admin on eqlgroup1:

# asmcli create group-access --name eqlgroup1 --user-name accnt-admin
  Password for accnt-admin on eqlgroup1:

# asmcli create smart-copy --source /user1 /user2 --user-name user-admin

# asmcli create smart-copy --source /accnt-recv --user-name accnt-admin
```

NOTE: In the previous example, ASM/LE needs to connect to a PS Series group. If a group access file has multiple user names for the group, but the `create smart-copy` command does not include the `--user-name` option to select the user name to use, the command will fail. A unique identity is required to process the request.

Example 4: Creating a Group Access Record Without Specifying a Password

When you create a group access file and specify the `--no-password` option, ASMCLI does not store and encrypt a password in the group access file. In this example, the group access file for Group3 is created without storing a password.

```
# asmcli list group-access

============================================================
Group Access
============================================================

Group name: Group1
IP address: 192.x.x.xx
User names: user-voladmin
```

```
Group name: Group2
IP address: 192.x.x.xy
User names: user-voladmin

2 Groups defined.
```

```
# asmcli create group-access --name Group3 --ip-address 192.x.x.xz
--user-name user-voladmin --no-password
```

```
Successfully created credentials record in file /etc/equallogic/asm-group-access.

# asmcli list group-access

============================================================
Group Access
============================================================

Group name: Group1
IP address: 192.x.x.xx
User names: user-voladmin

Group name: Group2
IP address: 192.x.x.xy
User names: user-voladmin

Group name: Group3
IP address: 192.x.x.xz
User names: user-voladmin (No stored password)

3 Groups defined.
```

For an example of changing an existing group access record that stores a password to one that does not, see Updating Access to a Group.

Listing Group Access Information

Use the `asmcli list group-access` command to list all user names and passwords in the group access file.

asmcli list group-access Command Syntax

The command syntax is as follows:

```
# asmcli list group-access
  [--name group-name | -n]
  [--file file-location | -f]
```

The following table describes the optional command parameters:

Table 18. List Group-Access Optional Command Parameters

Parameter	Abbr.	Argument	Description
--name	-n	*group-name*	Specifies the name of the PS Series group for which you want to list all user names and passwords. Omit this parameter to show all groups.
--file	-f	*file-location*	Specifies the location of a nondefault group access file. Omit this parameter to use the default group access file.

Examples: Listing Group Access Information

This section includes examples of listing all group access records and listing group access records for a specific group.

Example 1: Listing All Group Access Records

To list all group access records, enter the `list group-access` command without any options:

```
# asmcli list group-access

============================================================
Group Access
============================================================

Group name:     eqlgroup1
IP address:     192.x.x.xx
User names:     j_smith
```

```
Group name:     eqlgroup2
IP address:     192.x.x.xy
User names:     j_smith

Group name:     eqlgroup8
IP address:     192.x.x.zz
User names:     John_Smith

3 Groups defined.
```

Example 2: Listing Group Access Records for a Specific Group

To list group access information for the group eqlgroup1, specify a group name:

```
# asmcli list group-access --name eqlgroup1

Group name:     eqlgroup1
IP address:     192.x.x.xx
User names:     j_smith
```

Updating Access to a Group

Use the `asmcli update group-access` command to modify a group access file and its contents.

asmcli update group-access Command Syntax

The command syntax is as follows:

```
# asmcli update group-access
   --name group-name | -n
  [--file file-location | -f]
  [--ip-address IP-address | -i]
  [--management-ip mng-IP-address | -I]
  [--new-name new-name | G]
  [--no-password]
  [--password password | -p]
  [--user-name user-name | -u]
```

The following table describes the required and optional command parameters.

Table 19. Update Group-Access Command Options

Parameter	Abbr.	Argument	Description
Required Parameter			
--name	-n	*group-name*	Specifies the name of the PS Series group.
Optional Parameters			
--file	-f	*file-location*	Specifies the location of the nondefault group access file. Omit this parameter to use the default group.
--ip-address	-i	*IP-address*	Specifies the permanent IP address of the PS Series group.
--management-ip	-I	*mng-IP-address*	Specifies the management IP address for the PS Series group
--new-name	-G	*new-name*	Specifies the new name for the PS Series Group.
--no-password	none	none	Specifies not to store a password or this user. If you use this option, you are prompted to enter a password each time an ASMCLI command is run (for example, `create smart-copy`) and it will not be possible to run a scheduled operation.
--password	-p	*password*	Specifies the password to be encrypted and stored.

Parameter	Abbr.	Argument	Description
--user-name	-u	*user-name*	Specifies the user name on the PS Series group.

Examples: Updating Group Access Information

This section includes examples of replacing a group access file, changing a group name, updating a group access file for a management IP, and removing a stored password for a group.

Example 1: Replacing Information in a Group Access File

In this example, information in an existing group access file is replaced with new values (ASMCLI prompts for the password):

```
# asmcli update group-access --name eqlgroup1 --ip-address
  192.0x.x.xx
# asmcli update group-access --name eqlgroup1 --user-name j_smith
  Password for j_smith on eqlgroup1:
```

Example 2: Changing the Group Name

This example specifies the `--new-name` option to change the name of group `sample` to `eqlgroup8`:

```
# asmcli update group-access --name sample --new-name eqlgroup8

Successfully renamed group "sample" to "eqlgroup8" in file "/etc/
mycompany/asm-group-access".
```

Example 3: Updating the Group Access File for a Management IP

This example configures a PS Series group with a separate IP address for management. Use the `--management` option to specify the management IP address:

```
# asmcli update group-access --name eqlgroup1 --ip-address
  192.x.x.xx --management 192.x.x.xx
```

You can also use the `update group-access` command to remove the management IP address by specifying the `--management` option without a value, as follows:

```
# asmcli update group-access --name eqlgroup1 --management
```

Example 4: Removing a Stored Password for a Group

In this example, `Group2` has an encrypted password stored in its group access file. The `--no-password` option in the `update group-access` command updates the stored credentials to remove the password. When you create a Smart Copy of `Group2`, ASMCLI will prompt you to enter the password.

```
# asmcli list group-access

==========================================================
Group Access
==========================================================

Group name: Group1
IP address: 192.x.x.xx
User names: user-voladmin

Group name: Group2
IP address: 192.x.x.xy
User names: user-voladmin

Group name: Group3
IP address: 192.x.x.xz
User names: user-voladmin (No stored password)

3 Groups defined.

# asmcli update group-access --name EqlGroup2 --user-name user-voladmin
```

--no-password

```
Successfully updated credentials for user "user-voladmin" in group
"Group2" in file /etc/equallogic/asm-group-access.
```

```
# asmcli list group-access

============================================================
Group Access
============================================================

Group name: Group1
IP address: 192.x.x.xx
User names: user-voladmin

Group name: Group2
IP address: 192.x.x.xy
User names: user-voladmin (No stored password)

Group name: Group3
IP address: 192.x.x.xz
User names: user-voladmin (No stored password)

3 Groups defined.
```

Deleting Group Access Information

Use the `asmcli delete group-access` command to delete a group access file.

asmcli delete group-access Command Syntax

The command syntax is as follows:

```
# asmcli delete group-access
   --name group-name | -n
  [--file file-location | -f]
  [--user-name user-name | -u]
```

The following table describes the required and optional parameters.

Table 20. Delete Group-Access Command Options

Parameter	Abbr.	Argument	Description
Required Parameter			
--name	-n	*group-name*	Specifies the name of the PS Series group.
Optional Parameters			
--file	-f	*file-location*	Specifies the location of the group access file.
--user-name	-u	*user-name*	Specifies the user name on the PS Series group.

Examples: Deleting Group Access Information

This section provides examples of deleting group access information for a specific user and for all users.

Example 1: Deleting Group Access Information for a Specific User

This example deletes credentials for a specific user, j_smith, in group `eqlgroup1`:

```
# asmcli delete group-access --name eqlgroup1 --user-name j_smith
```

Example 2: Deleting Group Access Information for All Users

To delete all credentials and group information for `eqlgroup1`, the `--user-name` option is omitted:

```
# asmcli delete group-access --name eqlgroup1
```

Creating and Managing Smart Copies

ASMCLI provides commands for creating and managing Smart Copies. This section provides the following information:

- Creating a Smart Copy
- Listing Available Smart Copies
- Deleting a Smart Copy

For additional Smart Copy operations, see the following information:

- Logging In to and Out of Smart Copies
- Mounting and Unmounting Smart Copies
- Restoring a Smart Copy

Creating a Smart Copy the First Time

When you create a Smart Copy the first time, you can set up your environment so ASMCLI will not require that you specify an identity each time you create the Smart Copy.

1. Create an administrator identity on the PS Series group. Dell recommends that you restrict unauthorized access by setting up pool or volume administration accounts. See the *Dell EqualLogic PS Series Group Manager Administrator's Guide*.
2. Save one or more administrator user names and passwords in a group access file. If you do not specify the password as an option, ASMCLI will prompt you for the password (see examples in steps 3 and 4). For more information on creating and managing passwords, see Creating and Managing Group Access Files.

   ```
   # asmcli create group-access --name PS-group-name --ip-address  ip-address
     --username user-name
   ```

   If you have not defined the PS group name, specify the IP address. For example:

   ```
   # asmcli create group-access --name -PSGroup2 --ip-address 192.x.x.xx
     --username=sample-voladmin
     Password for sample-voladmin on PSGroup2:
   ```

3. Create a Smart Copy of a given mount point (default type is snapshot) by specifying a user name. If you did not set up a group access file in Step 2, ASMCLI prompts for a password.

   ```
   asmcli create smart-copy --source mount-point --username  user-name
   ```

   NOTE: The interactive installation script asks if you want to enter and save passwords for each PS Series group IP it discovers. If you answer Yes, these login credentials are saved in a group access file. For features provided by the installation script, see Installation Results.

Creating a Smart Copy

Use the `asmcli create smart-copy` command to create a Smart Copy of a volume. For general information about Smart Copies, see About Smart Copies.

asmcli create smart-copy Command Syntax

The command syntax is as follows:

```
# asmcli create smart-copy
    --collection collection-name | -c
    --source mount-point1 [mount-point2...] | -s
   [--comment comment-string | -C]
   [--destination directory | -d]
```

```
    [--force | -F]
    [--freeze-time n]
    [--group-access file-location | -g]
    [--object object-id | -o]
    [--output error | always]
    [--read-only | -R]
    [--retry-count n | -r]
    [--retry-wait n | -w]
    [--type volume-type | -t]
    [--username user-name | -u]
```

The following table describes the required and optional command parameters.

Table 21. Create Smart-Copy Command Options

Parameter	Abbr.	Argument	Description
Required Parameters			
--collection	-c	*collection-name*	Specifies the Smart Copy collection name. All Smart Copies, even a single one, are stored in a collection. Do not use with `--source`.
--source	-s	*mount-point*	Specifies one or more mount points to copy when creating the Smart Copy. Do not use with `--collection .`
Optional Parameters			
--comment	-C	*comment-string*	Provides a comment string to describe the Smart Copy you are creating. The comment is set for the PS Series group and stored in the backup document.
--destination	-d	*directory*	Specifies where to mount the Smart Copy. If you specify `mnt/temp` for Smart Copy volume `/x`, ASM/LE mounts the volume at `/mnt/temp/x`. If you omit this option, mount the Smart Copy later with the `mount smart-copy` command.
--force	-F	none	Creates the Smart Copy even if all mount points cannot be frozen. When creating a Smart Copy, ASM/LE attempts to freeze all mount points. This option creates the Smart Copy, but the Smart Copy might be inconsistent.
--freeze-time		*n*	Specifies the maximum number of seconds to freeze the file systems when creating a Smart Copy. If this value is exceeded, the operation times out and a consistent Smart Copy is not created. For most configurations, the default time is sufficient. When creating Smart Copies with a large number of mount points or a heavy I/O load, ASM/LE might not have enough time to freeze all of the mount points and build the Smart Copy. Use this option to increase the time limit.
--group-access-file	-g	*file-location*	Specifies the location of the group access file that stores the user names and encrypted passwords.
--object	-o	*object-id*	Specifies an object ID to uniquely identify the replica Smart Copy being created. Use this option only for replica clones (see Example 5). For all other Smart Copy types, use `--source` or `--collection`.
--output		error | always	Specifies when to display output to the console. Use `error` to display output only when an operation fails, such as with a scheduled operation. Use `always` (default) to display all output from the operation.
--read-only	-R	none	Mounts the file system from the Smart Copy as read-only. Use this parameter only with the `--destination` option.
--retry-count	-r	*n*	Specifies the number of times to retry a failed operation.

Parameter	Abbr.	Argument	Description
--retry-wait	-w	*n*	Specifies the number of seconds to wait before retrying a failed operation.
--type	-t	*volume-type* clone|replica| snapshot| thinclone	Specifies the Smart Copy type to create. If omitted, creates `snapshot`. ASMCLI reports an error if an incorrect type is applied, such as `thinclone` for a volume that is not a template volume.
--user-name	-u	*user-name*	Specifies a particular user identity, when multiple user names have been stored in the group access file.

Examples: Creating Smart Copies

This section includes the following examples:

Example 1: Creating a Smart Copy of a Given Mount Point

This example creates a Smart Copy of the mount point `/source/VolA`. The default Smart Copy type is snapshot.

```
# asmcli create smart-copy --source /source/VolA

The mount point has been frozen.
Created snapshot vola-2017-01-05-14:13:50.220.1 on glengrp
The mount point has been thawed.
Successfully created a Smart Copy from 1 target.
```

Example 2: Creating a Smart Copy and Mounting in a Single Operation

In this example, a Smart Copy is created on a given mount point with the `--destination` option.

```
# asmcli create smart-copy --source /source/VolC --destination
/snapshots/VolC
```

```
The mount point has been frozen.
Created snapshot volc-2017-01-05-14:13:50.220.1 on glengrp
The mount point has been thawed.
Successfully created a Smart Copy from 1 target.
```

```
Logging in 1 Smart Copy target.
Logged in volc-2017-01-05-14:13:50.220.1 on glengrp
Mounted /snapshots/VolC/source/VolC
Successfully mounted 1 file system from 1 Smart Copy target.
  Successfully created 1 smart copy.
```

Example 3: Creating a Smart Copy with Multiple Volumes

In this example, multiple sets of mount points are specified for creating Smart Copies. The `--retry-count` option is specified to continue trying to connect to targets if they are temporarily unavailable. In this example, all targets are available.

```
# asmcli create smart-copy --source /source/VolA /source/VolB
/source/VolC \
--retry-count 2

All mount points (3) have been frozen.
```

```
Created snapshot volb-2017-01-05-14:39:14.221.1 on glengrp
Created snapshot volc-2017-01-05-14:39:14.222.1 on glengrp
Created snapshot vola-2017-01-05-14:39:14.223.1 on glengrp
All mount points (3) have been thawed.
Successfully created a Smart Copy from 3 targets.
```

Example 4: Retrying Creation of Smart Copy with Multiple Volumes

In this example, ASMCLI fails to create a consistent Smart Copy for a large number of source mount points on the first attempt. Because the `--retry` option was used, the creation succeeded on the second attempt.

```
# asmcli create smart-copy --source /vol1 /vol2 /vol3
/vol4 /vol5 /vol6 \
```

```
/vol7a /vol7b /vol7c /vol8 /vol9 /vol10 /vol11 /vol12
--retry-count 10

All mount points (14) have been frozen.
Created snapshot vola03-2017-01-09-13:42:13.16776 on eqlgroup1
Created snapshot vola02-2017-01-09-13:42:13.16777 on eqlgroup1
Created snapshot vola01-2017-01-09-13:42:13.16778 on eqlgroup1
Created snapshot vola11-2017-01-09-13:42:13.16779 on eqlgroup1
Created snapshot vola05-2017-01-09-13:42:13.16780 on eqlgroup1
Created snapshot vola08-2017-01-09-13:42:13.16781 on eqlgroup1
Created snapshot vola10-2017-01-09-13:42:14.16782 on eqlgroup1
Created snapshot vola12-2017-01-09-13:42:14.16783 on eqlgroup1
Created snapshot vola04-2017-01-09-13:42:14.16784 on eqlgroup1
Created snapshot vola07-2017-01-09-13:42:14.16785 on eqlgroup1
Created snapshot vola06-2017-01-09-13:42:14.16786 on eqlgroup1
Created snapshot vola09-2017-01-09-13:42:14.16787 on eqlgroup1
Timeout limit (10 seconds) reached: Thawing all mount points early.
All mount points (14) have been thawed.
Error: Cannot create a consistent Smart Copy.
Sleeping for 10 seconds and then will retry the operation.

All mount points (14) have been frozen.
Created snapshot vola03-2017-01-09-13:42:25.16788 on eqlgroup1
Created snapshot vola02-2017-01-09-13:42:25.16789 on eqlgroup1
Created snapshot vola01-2017-01-09-13:42:26.16790 on eqlgroup1
Created snapshot vola11-2017-01-09-13:42:26.16791 on eqlgroup1
Created snapshot vola05-2017-01-09-13:42:26.16792 on eqlgroup1
Created snapshot vola08-2017-01-09-13:42:26.16793 on eqlgroup1
Created snapshot vola10-2017-01-09-13:42:26.16794 on eqlgroup1
Created snapshot vola12-2017-01-09-13:42:26.16795 on eqlgroup1
Created snapshot vola04-2017-01-09-13:42:26.16796 on eqlgroup1
Created snapshot vola07-2017-01-09-13:42:26.16797 on eqlgroup1
Created snapshot vola06-2017-01-09-13:42:26.16798 on eqlgroup1
Created snapshot vola09-2017-01-09-13:42:26.16799 on eqlgroup1
All mount points (14) have been thawed.
Successfully created a Smart Copy from 12 targets.
```

Example 5: Creating a Replica Clone Smart Copy

This example creates a replica Smart Copy (`--type replica`), lists the result, and creates a replica clone Smart Copy. The replica Smart Copy is identified by object ID (`--object`), the clone by `--type clone`.

```
# asmcli create smart-copy --source /mnt/asm-acl --type replica

The mount point has been frozen.
The replica was started successfully.
The mount point has been thawed.
Getting remote replica information, please wait ...
Created replica asm-acl.1-2017-02-02-01:50:58.533.1 on HitStore4
Successfully created a Smart Copy from 1 target.
```

```
# asmcli list smart-copy

==========================================================
```

```
Smart Copies
============================================================
Source: myhost.mycompany.com : /mnt/asm-acl
 Replica: 02-Feb-2017 14:58:37.692088
 ObjectId: b-6d91c1-b5427dfda-ccfc656a321b56aa
  asm-acl on windev33-grp --> asm-acl.1-2017-02-02-01:50:58.533.1 on HitStore4

1 Smart Copy found.

# asmcli create smart-copy --object b-6d91c1-b5427dfda-ccfc656a321b56aa
--type clone

Created clone asm-acl-clone-2017:02:02-14:59:38.424 on HitStore4

# asmcli list smart-copy --type clone

============================================================
Smart Copies
============================================================
Source: myhost.mycompany.com : /mnt/asm-acl
 Clone: 02-Feb-2017 14:59:38.424281
 ObjectId: 0-498326-5e4f5afec-e9b4a8185760f3b8
  asm-acl on windev33-grp --> asm-acl-clone-2017:02:02-14:59:38.424
   on HitStore4

1 Smart Copy found.
```

Example 6: Creating a Smart Copy When a Password Is Required

In the first part of this example, the group access file for Group2 does not contain a password and ASMCLI prompts for one when you create a Smart Copy. If the group access file contained multiple identities, you could specify one with the `--user-name` option.

```
# asmcli create smart-copy --source /test10 /test11 /test12

Password for user user-voladmin on 1.2.3.120: ******
All mount points (3) have been frozen.
Created snapshot User-test-12-2017-02-20-11:03:17.25257 on Group2
Created snapshot User-test-10-2017-02-20-11:03:17.25258 on Group2
Created snapshot User-test-11-2017-02-20-11:03:17.25259 on Group2
All mount points (3) have been thawed.

Successfully created a Smart Copy from 3 targets.
```

The next part of this example shows the `list smart-copy` command, which does not require group access and ASMCLI does not prompt you for a password.

```
# asmcli list smart-copy

============================================================
Smart Copies
============================================================

Source: User-voladmin.mycompany.com : /test10 /test11 /test12
  Snapshot: 16-Feb-2017 13:05:05.279150
  ObjectId: 2-4acff9-7641b679e-063072c486b4578b
    User-test-10 --> User-test-10-2017-02-16-13:03:39.25255 on Group2
    User-test-11 --> User-test-11-2017-02-16-13:03:39.25256 on Group2
    User-test-12 --> User-test-12-2017-02-16-13:03:39.25254 on Group2
```

```
  Snapshot: 20-Feb-2017 11:04:46.292034
  ObjectId: c-f5e7be-8b47cb3db-9c6d744d93d25f84
    User-test-10 --> User-test-10-2017-02-20-11:03:17.25258 on Group2
    User-test-11 --> User-test-11-2017-02-20-11:03:17.25259 on Group2
    User-test-12 --> User-test-12-2017-02-20-11:03:17.25257 on Group2

2 Smart Copies found.
```

The following part shows the **--broken** option, which searches for broken Smart Copies. Because this requires ASMCLI to access the group, you are prompted for a password,

```
# asmcli list smart-copy --broken

Password for user user-voladmin on 192.0.2.20: ******

0 Smart Copies found.
```

If you delete the access credentials for a group, the Smart Copy creation operation will fail. In the last part of this example, group access records are deleted for Group3, and the subsequent attempt to create a Smart Copy fails.

```
# asmcli delete group-access --name Group3 --user-name
grpadmin

Successfully deleted credentials for user "grpadmin" in group "Group3"
in file /etc/equallogic/asm-group-access.

# asmcli list group-access

============================================================
Group Access
============================================================

Group name: Group2
IP address: 192.x.x.xx
User names: grpadmin

Group name: Group3
IP address: 192.x.x.xy
User names:

Group name: Group4
IP address: 192.x.x.yy
in (No stored password)

3 Groups defined.

# asmcli create smart-copy --source /test10

Error: Cannot connect to Group3 (192.x.x.xy) : No credentials found
```

Example 7: Creating a Smart Copy Using Snapshot Borrowing

When snapshot borrowing is enabled and creating a Smart Copy snapshot exceeds the available snapshot reserve space, the `asmcli` command attempts to borrow free space. A message displays, either informing you that snapshot borrowing is being used or warning you that borrowed space has crossed the threshold limits.

```
# asmcli create smart-copy --source /user1

Warning: Volume myvol-user1 on group9 (192.x.x.xy) is borrowing snapshot space
Suggestion: Please increase snapshot reserve or delete unnecessary snapshots.
  For scheduled operations, consider using "keep count" (see "asmcli help schedule").
The mount point has been frozen.
Created snapshot myvol-user1-2017-08-16-09:44:33.462.1 on group9
The mount point has been thawed.
Successfully created a Smart Copy from 1 target.
```

Example 8: Creating a Smart Copy That Exceeds Threshold Limits

If creating a Smart Copy snapshot triggers snapshot borrowing that crosses the free pool space threshold and then another snapshot is created, older snapshots might be deleted to recover space.

```
# asmcli create smart-copy --source /user1

Warning: Volume myvol-user1 on windev19grp (192.x.x.xy) is borrowing snapshot space
Warning: Free space for snapshot borrowing on group group9 (192.x.x.xy)
  is running low (99.4% of free space is in use.)
Additional snapshots may trigger the snapshot space recovery policy for this volume.
```

```
Suggestion: Please increase snapshot reserve or delete unnecessary snapshots.
  For scheduled operations, consider using "keep count" (see "asmcli help schedule").
The mount point has been frozen.
Created snapshot myvol-user1-2017-08-16-09:49:20.465.1 on group9
The mount point has been thawed.
Successfully created a Smart Copy from 1 target.
```

Listing Available Smart Copies

Use the `asmcli list smart-copy` command to list all available Smart Copies.

asmcli list smart-copy Command Syntax

The command syntax is as follows:

```
# asmcli list smart-copy
    [--broken | -B]
    [--check-status]
    [--group-access-file file-location | -g]
    [--output always | error]
    [--unreachable | -U]
    [--user-name user-name | -u]
```

Search option parameters also apply to this command. For information, see ASMCLI Search Options.

The following table describes the optional command parameters.

Table 22. List Smart Copy Optional Command Parameters

Parameter	Abbr.	Argument	Description
--broken	-B	none	Locates broken Smart Copies. A broken Smart Copy contains a snapshot or replica snapshot that is missing while the source volume is still accessible. Broken Smart Copies can occur when snapshots are deleted on the PS Series group without using ASM to delete them. This option attempts to access the group access file for a PS Series group. For information on group access files, see About Group Access Files.
--check-status	none	none	Displays the status of existing Smart Copies that are unreachable, broken, or for replica Smart Copies, in progress and promoted. This parameter takes precedence if used with `--broken` or `--unreachable`. Use with `--type replica` to indicate status for all replica Smart Copies.
--group-access-file	-g	*file-location*	Specifies the location of the group access file that contains the user names and encrypted passwords.
--output		always | error	Specifies whether or not to suppress output. Specify `--output always` (default) to display all output. Specify `--output error` to suppress output unless an operation fails, which is useful for scheduled operations.
--unreachable	-U	none	Locates unreachable Smart Copies that fail to connect to the PS Series group or contain a target on the PS Series group that cannot be reached.
--user-name	-u	*user-name*	Specifies a particular user identity when multiple user names have been stored in the group access file.

Examples: Listing Smart Copies

This section includes examples of listing Smart Copies and listing the status of Smart Copies.

Example 1: Listing All Smart Copies

The `list smart-copy` command with no qualifiers lists all Smart Copies. For each Smart Copy, the output identifies the location, the type, the name, the date created, and the status (for example, mounted). The following example shows four Smart Copies, one of which contains multiple mount points.

```
# asmcli list smart-copy

============================================================
Smart Copies
============================================================

Source: hostname.mycompany.com : /source/VolA
  Snapshot: 14-Dec-2016 16:51:46.956404
  Status: Mounted as /snapshots/VolA
    donvola --> donvola-2015-12-14-16:49:30.131.1 on glengrp

Source: hostname.mycompany.com : /source/VolA /source/VolB /source/VolC
  Snapshot: 05-Jan-2017 14:42:10.794530
  donvola --> donvola-2017-01-05-14:39:14.223.1 on glengrp
  donvolb --> donvolb-2017-01-05-14:39:14.221.1 on glengr
  donvolc --> donvolc-2017-01-05-14:39:14.222.1 on glengrp

Source: hostname.mycompany.com : /source/VolB
  Snapshot: 05-Jan-2017 16:24:29.074310
  Status: Mounted as /snapshots/VolB
    donvolb --> donvolb-2017-01-05-16:21:34.215.1 on glengrp

Source: hostname.mycompany.com : /source/VolC
  Snapshot: 05-Jan-2017 19:16:46.801189
  Status: Mounted as /snapshots/VolC
    donvolc --> donvolc-2017-01-05-19:13:50.220.1 on glengrp

4 Smart Copies found.
```

You can display additional detail using the verbose flags. For example:

```
# asmcli list smart-copy --verbose 2

============================================================
Smart Copies
============================================================
Source: hostname.mycompany.com : /source/volA
   Snapshot: 14-Dec-2016 10:46:44.839798
   Status: Mounted as /snapshots/VolA
   ObjectId: 5-392e06-8244f99a7-88d5834731b41690
   Directory: /var/lib/equallogic/asm/smart-copies/asm-le-hosta.com/source-a-
     5d5e2d-874d8ac30-0d2df1dc89f8932c/
   File: 20111214-154644.839798-5-392e06-8244f99a7-88d5834731b41690.bcd
     vola --> vola-2016-12-14-10:46:48.1421.1 on new1-group

Source: hostname.mycompany.com : /source/VolB
   Snapshot: 05-Jan-2017 16:24:29.074310
   Status: Mounted as /snapshots/VolB
   ObjectId: c-01e943-574012772-74bc9cfd88221283
   Directory: /var/lib/equallogic/asm/smart-copies/asm-le-hosta.com/source-1-
   21a09d-9ce520990-c75a20bb9b1eb86f/
   File: 20120104-212429.074310-c-01e943-574012772-74bc9cfd88221283.bcd
      volb --> volb-2017-01-05-16:21:34.215.1 on new1-group

2 Smart Copies found.
```

Example 2: Listing Smart Copy Status

In this example, the `--check-status` option used with the `--type replica` and `--latest` options shows the status of the newest replica Smart Copy.

```
# asmcli list smart-copy --type replica --latest --check-status
-vv

Logged into PS Series group group3.mycompany.com (group3) as
```

```
username grpadmin (GroupAdmin)
=========================================================
Smart Copies
=========================================================
Source: hostname.mycompany.com : /mnt/mnt1 /mnt/mnt3 /mnt/mnt4
   Replica: 28-Feb-2017 10:42:54.077954
   ObjectId: f-ceea53-2f409c86e-6b29668578db1cac
   Directory: /mnt/asm/asm-le-node202.equallogic.com/source-e-794700-
    706f6e2b2-c25ce4ca6fda4f46/
   File: 20120222-154254.077954-f-ceea53-2f409c86e-6b29668578db1cac.bcd
vol1 on group1 --> vol1.3-2017-02-22-11:05:13.1039.1 on group3
vol3 on group4 --> vol3.2-2017-02-22-11:05:13.1041.1 on group3
vol4 on group4 --> vol4.2-2017-02-22-11:05:13.1040.1 on group3
 [in progress ]

1 Smart Copy found.
```

Access Backup Documents

To access backup documents, use the `asmcli list smart-copy` command to display backup document information for all Smart Copies you created. Narrow the list by specifying the name of the collection, indicating earliest or latest created Smart Copy, or using Smart Copy object ID. For information on these options, see Listing Available Smart Copies.

For example, specifying the object ID in the `asmcli list smart-copy` command lists one Smart Copy backup document:

```
# asmcli list smart-copy --object 5-392e06-8244f99a7-88d5834731b41690 --verbose 2
=========================================================
Smart Copies
=========================================================
Source:     HostA: /source/VolA
  Snapshot: 14-Dec-2016 10:46:44.839798
  Status:   Mounted as /snapshots/VolA
  ObjectId: 5-392e06-8244f99a7-88d5834731b41690
  Directory: /var/lib/equallogic/asm/smart-copies/asm-le-hosta.com/source-a-
    5d5e2d-874d8ac30-0d2df1dc89f8932c/
  File: 20161214-154644.839798-5-392e06-8244f99a7-88d5834731b41690.bcd
    vola --> vola-2016-12-14-10:46:48.1421.1 on new1-group
Source: hostname.mycompany.com : /source/VolB
  Snapshot: 04-Jan-2017 16:24:29.074310
  Status: Mounted as /snapshots/VolB
  ObjectId: c-01e943-574012772-74bc9cfd88221283
  Directory: /var/lib/equallogic/asm/smart-copies/asm-le-hosta.com/source-1-
21a09d-9ce520990-c75a20bb9b1eb86f/
  File: 20170104-212429.074310-c-01e943-574012772-74bc9cfd88221283.bcd
    volb --> volb-2017-01-04-16:21:34.215.1 on new1-group
2 Smart Copies found.
```

Deleting a Smart Copy

Use the `asmcli delete smart-copy` command to delete a Smart Copy by name.

NOTE: The same search options `asmcli list smart-copy` uses to select an individual Smart Copy or a set of Smart Copies are also used by the `asmcli delete smart-copy` command.

asmcli delete smart-copy Command Syntax

The command syntax is as follows:

```
# asmcli delete smart-copy
   [--broken | -B] |
   [--force | -F]
   [--group-access-file file-location | g]
   [--limit number | -L]
   [--output always | error]
   [--unreachable]
   [--user-name user-name | -u]
```

Search option parameters also apply to this command. For information, see ASMCLI Search Options.

The following table describes the optional command parameters.

Table 23. Delete Smart Copy Optional Command Parameters

Parameter	Abbr.	Argument	Description
--broken	-B	none	Deletes broken Smart Copies and attempts to access the group access file for a PS Series group. A broken Smart Copy contains a snapshot or replica snapshot that is missing while the source volume is still accessible. Broken Smart Copies can occur when snapshots are deleted on the PS Series group without using ASM. For information about group access files, see About Group Access Files.
--force	-F	none	Deletes the Smart Copy, bypassing a failure if a Smart Copy is broken or unreachable. If a Smart Copy is the last replica and no snapshot exists to ensure consistency of the replicated data on both replication sites, deletes the Smart Copy and as much copied data as possible.
--group-access-file	-g	*file-location*	Specifies the location of the group access file that contains the user names and encrypted passwords.
--limit	-L	*number*	Specifies the maximum number of Smart Copies to delete at one time. The default is 1. Increase this value cautiously. If you specify a large value without sufficiently restrictive search criteria, you might irrevocably delete all of your Smart Copies. Use the `list smart-copy` command before deleting Smart Copies.
--output	none	always | error	Specifies when to display output. Specify `always` (default) to display all output. Specify `error` to suppress output unless an operation fails, which is useful for scheduled operations.
--unreachable	-U	none	Locates Smart Copies that are unreachable because they fail to connect to the PS Series group or contain a target on the PS Series group that cannot be reached. **CAUTION:** **Use extreme caution when using this option. Any resulting changes in access credentials can affect your ability to connect to a group or to see target volumes on a group. If this situation occurs, the Smart Copies might appear unreachable even though the source volume and its copies are on the PS Series group. If you delete the Smart Copies, ASM cannot locate the source volumes or the copies. This option attempts to access the PS Series group.** See About Group Access Files for information about group access files.
--user-name	-u	*user-name*	Specifies a user identity when multiple user names are stored in the group access file.

Examples: Deleting Smart Copies

This section provides examples of deleting the latest Smart Copy and deleting a Smart Copy by Object ID.

Example 1: Deleting the Latest Smart Copy

This example deletes the most recently created Smart Copy in the specified source. The Object ID of the Smart Copy is displayed in the results.

```
# asmcli delete smart-copy --source /source/VolC --latest

Deleted Smart Copy: a-900d01-d34602f55-6290da94e24343ae

Successfully removed the Smart Copy.
```

Example 2: Deleting a Smart Copy Identified By Object ID

To ensure you are deleting the correct Smart Copy, list Smart Copies in very verbose mode to display full information for the Smart Copies.

```
# asmcli list smart-copy --source /source/VolB -verbose

=========================================================
Smart Copies
=========================================================
Source: hostname.mycompany.com : /source/VolB
Snapshot: 04-Jan-2017 16:24:29.074310
  Status: Mounted as /snapshots/VolB
  ObjectId: c-01e943-574012772-74bc9cfd88221283
    volb --> volb-2017-01-04-16:21:34.215.1 on glengrp

  Snapshot: 05-Jan-2017 15:26:17.183220
  Description: Hourly snapshot of user space
  Schedule: user-hourly
  ObjectId: e-bd6416-7e4e51801-ab61e09f2557cf81
    volb --> volb-2017-01-05-15:23:20.224.1 on glengrp

  Snapshot: 05-Jan-2017 17:05:57.846661
  Description: Hourly snapshot of user space
  Schedule: user-hourly
  ObjectId: 8-b8b03d-7444e1f3e-626856b008bde08e
    volb --> volb-2017-01-05-17:03:01.225.1 on glengrp

3 Smart Copies found.
```

Specify the Smart Copy to delete using the `--object` option.

```
# asmcli delete smart-copy --object c-01e943-574012772-74bc9cfd88221283

Error: Failed to delete Smart Copy c-01e943-574012772-74bc9cfd88221283
because it is still mounted at /snapshots/VolB.
Did not remove the Smart Copy.
```

To delete a Smart Copy, you must first unmount and log off the Smart Copy. The `ummount smart-copy` command performs both operations.

```
# asmcli unmount smart-copy --object c-01e943-574012772-74bc9cfd88221283

Unmounted /snapshots/VolB/source/VolB
Removing sessions to 1 Smart Copy target.
Logged out volb-2017-01-04-16:21:34.215.1 on glengrp

Successfully unmounted 1 mount point and logged out from 1 target.
```

Delete the unmounted Smart Copy:

```
# asmcli delete smart-copy --object c-01e943-574012772-74bc9cfd88221283

Deleted Smart Copy: c-01e943-574012772-74bc9cfd88221283

Successfully removed the Smart Copy.
```

Restoring a Smart Copy

Use the `asmcli restore smart-copy` command to restore data from a Smart Copy. For information about data restoration, see About Restoring Data From a Smart Copy.

asmcli restore smart-copy Command Syntax

The command syntax is as follows:

```
# asmcli restore smart-copy
   [--force | -F]
   [--group-access-file file-location | -g]
   [--user-name user-name | -u]
```

Search option parameters also apply to this command. For information, see ASMCLI Search Options.

The following table describes the optional command parameters.

Table 24. Restore Smart Copy Optional Command Parameters

Parameter	Abbr.	Argument	Description
--force	-F	none	Performs the restore operation without first checking for connections to volumes on PS Series groups from other hosts. Without this parameter, the `restore smart-copy` command performs the restore only if no connections exist to volumes on PS Series groups that come from other hosts. **CAUTION: When using this option, you must understand the possible consequences of other systems accessing the data and doing a restore in-place. You assume the responsibility for this action. For more information, see About Restoring Data From a Smart Copy.**
--group-access-file	-g	*file-location*	Specifies the location of the group access file that stores user names and encrypted passwords.
--user-name	-u	*user-name*	Specifies a particular user identity when multiple user names have been stored in the group access file.

Example: Restoring Smart Copies

The following example restores a Smart Copy from a source mount point.

```
# asmcli restore smart-copy --source /mnt/eql/std-vol-06
Dell EqualLogic Auto-Snapshot Manager CLI Version 1.7.0
Copyright (c) 2010-2020 Inc.

Unmounted /mnt/eql/std-vol-06
Logged out iqn.2001-05.com.equallogic:0-8a0906-35f6e3e0d-6c4004d234f57299-sles11sp4-clus1-std-
vol01
Volume sles11sp4-clus1-std-vol01.1 on hitlinux-3-02 restored from sles11sp4-clus1-std-
vol01.1-2017-07-21-01:30:39.267.1
Logged in iqn.2001-05.com.equallogic:0-8a0906-35f6e3e0d-6c4004d234f57299-sles11sp4-clus1-std-
vol01 on 10.11.57.31
Mounted /mnt/eql/std-vol-06
Restore operation completed successfully.

# ls -l /mnt/eql/std-vol-06
total 512008
-rw-r--r-- 1 root root 524288000 May 12 01:28 file.txt
-rw-r--r-- 1 root root       100 May 10 19:43 file1.txt
-rw-r--r-- 1 root root       300 May 10 19:43 file2.txt
-rw-r--r-- 1 root root         0 Jul 20 18:26 file3.txt
-rw-r--r-- 1 root root         0 Jul 20 18:26 file4.txt
```

Soft-Deleted Smart Copies

This section includes information on how to manage soft-deleted Smart Copies. It provides examples of listing the contents of the volume recovery bin. For more information about the soft-deleted Smart Copies, see About Restoring a Soft-Deleted Smart Copy.

Managing Soft-Deleted Smart Copies

Smart Copy clones that have been soft-deleted are held temporarily in the volume recovery bin. You can use the following commands with Smart Copy clones:

- Identify clones that have been soft-deleted in the volume recovery bin using the `asmcli list smart-copy --check-status` command.
- Access the Volume Recovery Bin interface through the EqualLogic Group Manager GUI.

Examples: asmcli list Command

The **asmcli list --check-status** command reports the status of Smart Copy clones. This section provides the following examples:

Example 1: Listing Smart Copies in a Collection

The following example shows the output of the `asmcli list` command which includes the status of clones.

```
# asmcli list smart-copy --type clone --check-status -vv
```

```
Logged into PS Series group 192.x.x.xx (earth39-grp) as username grpadmin (GroupAdmin)
============================================================
Smart Copies
============================================================

Source: mysystem.equallogic.com : /mnt/ns-gl1 /mnt/ns-gl2
  Clone: 18-May-2017 18:47:38.101719
  ObjectId: d-ae90fe-d14693ed0-4b94835698ac308b
  Directory: /tmp/asm3/asm-le-ns.equallogic.com/source-e-86ca5b-
    e4f2bf9ea-3660dd7d9ae03f4b/
  File: 20170518-224738.101719-d-ae90fe-d14693ed0-4b94835698ac308b.bcd
    ns-gl1 --> ns-gl1-clone-2017:05:18-18:47:38.112 on earth39-grp
    ns-gl2 --> ns-gl2-clone-2017:05:18-18:47:38.371 on earth39-grp

1 Smart Copy found.
```

Example 2: Listing Deleted Smart Copies

The following example shows the status reported by the `asmcli list` command when a Smart Copy clone has been deleted.

```
# asmcli delete smart-copy --type clone --latest

Deleted Smart Copy: d-ae90fe-d14693ed0-4b94835698ac308b
Successfully removed the Smart Copy.

# asmcli list smart-copy --type clone -check-status -vv
Logged into PS Series group 192.x.x.xx (earth39-grp) as username
grpadmin (GroupAdmin)
============================================================
Smart Copies
============================================================

Source: mysystem.equallogic.com : /mnt/ns-gl1 /mnt/ns-gl2
  Clone: 18-May-2017 18:47:38.101719
  ObjectId: d-ae90fe-d14693ed0-4b94835698ac308b
  Directory: /tmp/asm3/asm-le-ns.equallogic.com/source-e-86ca5b-
    e4f2bf9ea-3660dd7d9ae03f4b/
  File: .20170518-224738.101719-d-ae90fe-d14693ed0-4b94835698ac308b.bcd
    ns-gl1 --> ns-gl1-clone-2017:05:18-18:47:38.112 on earth39-grp [ recovery bin ]
    ns-gl2 --> ns-gl2-clone-2017:05:18-18:47:38.371 on earth39-grp [ recovery bin ]

1 Smart Copy found.
```

Example 3: Status of Smart Copy After Recovery

The following example shows the status of the Smart Copies after one clone has been recovered and one clone remains in the volume recovery bin.

```
# asmcli list smart-copy --type clone --check-status -vv

Logged into PS Series group 192.x.x.xx (earth39-grp) as username grpadmin (GroupAdmin)
==========================================================
Smart Copies
==========================================================

Source: mysystem.equallogic.com : /mnt/ns-gl1 /mnt/ns-gl2
  Clone: 18-May-2017 18:47:38.101719
  ObjectId: d-ae90fe-d14693ed0-4b94835698ac308b
  Directory: /tmp/asm3/asm-le-ns.equallogic.com/source-e-86ca5b-
    e4f2bf9ea-3660dd7d9ae03f4b/
  File: .20170518-224738.101719-d-ae90fe-d14693ed0-4b94835698ac308b.bcd
    ns-gl1 --> ns-gl1-clone-2017:05:18-18:47:38.112 on earth39-grp
    ns-gl2 --> ns-gl2-clone-2017:05:18-18:47:38.371 on earth39-grp [ recovery bin ]

1 Smart Copy found.
```

Example 4: Status of Restored Smart Copies

The following example shows the status reported by the `asmcli list` command after all deleted Smart Copy clones have been restored.

```
# asmcli list smart-copy --type clone --check-status -vv
Logged into PS Series group 192.x.x.xx (earth39-grp) as username
grpadmin (GroupAdmin)
==========================================================
Smart Copies
==========================================================

Source: mysystem.equallogic.com : /mnt/ns-gl1 /mnt/ns-gl2
  Clone: 18-May-2017 18:47:38.101719
  ObjectId: d-ae90fe-d14693ed0-4b94835698ac308b
  Directory: /tmp/asm3/asm-le.equallogic.com/source-e-86ca5b-
    e4f2bf9ea-3660dd7d9ae03f4b/
  File: 20170518-224738.101719-d-ae90fe-d14693ed0-4b94835698ac308b.bcd
    ns-gl1 --> ns-gl1-clone-2017:05:18-18:47:38.112 on earth39-grp
    ns-gl2 --> ns-gl2-clone-2017:05:18-18:47:38.371 on earth39-grp

1 Smart Copy found.
```

Managing Soft-Deleted Smart Copy Clones Without Credentials

For groups running PS Series firmware version 6.0 and higher, if you create a Smart Copy clone and delete it, the firmware places the deletion in a recovery bin. The deleted volume can be restored by the PS Series group or eventually, it will be purged. Depending on which action occurs, entering an `asmcli list smart-copy --check-status` command will either restore or permanently delete the Smart Copy. However, if you change credentials so that the volume is no longer accessible, ASMCLI cannot see the volume in the recovery bin, and entering `list smart-copy --check-status` will permanently delete the Smart Copy. This behavior is identical to the way that changing credentials can create unreachable Smart Copies.

However, if you delete credentials so that there is no way to connect to the PS Series group at all, entering an `asmcli list smart-copy --check-status` command will not purge the Smart Copy backup documents. In that case, ASM/LE cannot clean up the Smart Copy backup documents.

Smart Copy Operations for Synchronous Replication Volumes

The `asmcli` command provides synchronous replication volumes a subset of the functions that are available from the Group Manager interface.

NOTE: To configure synchronous replication volumes or to switch from the active volume (SyncActive) to the alternate volume (SyncAlternate), use the PS Series firmware array.

The following table describes ASMCLI commands that are available for synchronous replication volumes.

Table 25. ASMCLI Commands Available for Synchronous Replication Volumes

Command	Description
--create	Creates a Smart Copy of a synchronous replication volume. Supports snapshots and clones only.
--mount	Mounts a synchronous replication volume. Applies to active volumes and to alternate volumes after successfully switching the roles of the pools.
--unmount	Unmounts a synchronous replication volume. Applies to active volumes and to alternate volumes after successfully switching the roles of the pools.
--delete	Deletes a synchronous replication volume. Applies to active volumes and to alternate volumes after successfully switching the roles of the pools.
--restore	Restores a synchronous replication volume from a snapshot of the active volume. A `--restore smart-copy` attempted from a snapshot created on the alternate volume will be denied. For multiple mount points, all synchronous restoration volume restores in the set must be from a snapshot of the active volume.

Example: Attempting to Restore a Snapshot Created on the Alternate Volume

The following example shows a listing of the status of all Smart Copies and an attempt to restore a snapshot that was created on the alternate volume (SyncAlternate).

```
# asmcli list smart-copy --check-status

============================================================
Smart Copies
============================================================

Source: node21 : /mnt/mnt1 /mnt/mnt2 /mnt/mnt3
  Snapshot: 29-May-2017 09:46:41.702126
>>  Status: Warning: Snapshot cannot be used for restore.
  Description: test schedule
  Schedule: schedule1
  Collection: collection1
  ObjectId: 5-c1ab42-654308f1f-67fa8e047c0b6599
>>     vol1 --> vol1-2017-05-29-09:49:26.1077.1 on system1 [SyncAlternate]
    vol2 --> vol2-2017-05-29-09:49:26.1078.1 on system1
    vol3 --> vol3-2017-05-29-09:49:26.1076.1 on system1

1 Smart Copy found.
```

```
# asmcli restore smart-copy --verbose --verbose

Directory:
/var/lib/equallogic/asm/smart-copies/asm-le/schedule-4-6a28b1-5a4b83564-a83888fd894bd78a/
Document: 20170529-134641.702126-5-c1ab42-654308f1f-67fa8e047c0b6599.bcd
Connecting to destination PS Groups.
Logged into PS Series group group1.equallogic.com (hitqavinod1)as
username grpadmin (GroupAdmin)
Error: Cannot proceed with snapshot restore.
  The snapshot for the following SyncRep volume is in the SyncAlternate pool.
>>     vol1 on system1 {system1.equallogic.com}

Suggestion: Switch to SyncAlternate using the group manager UI.
```

Creating and Managing Smart Copy Schedules

Schedules are predefined `create smart-copy` commands that can be created once and used repeatedly. Using schedules, you can:

- Name commonly used commands for creating Smart Copies to indicate their purpose (see Creating a Schedule)
- Incorporate schedules into scripts or cron jobs

- Change the schedule without modifying any scripts or cron tables. When the scheduled operation is included in a script or cron table, you can use ASMCLI to change the options to the `create smart-copy` command without making other changes (see Creating a Smart Copy).
- Locate Smart Copies created by a given schedule by specifying the `--schedule` option with the `list smart-copy` command (see Listing Schedules)
- Limit the number of Smart Copies kept for a schedule using the `--keep` option. For example, if you specify `--keep 12`, then only 12 Smart Copies are saved for that schedule. If 12 Smart Copies exist for that schedule and if the schedule is used to create a new Smart Copy, then the oldest Smart Copy that is not logged in or mounted is deleted (see Running a Schedule).
- Modify an existing schedule by adding or replacing an option (see Updating a Schedule)
- Delete a schedule (see Deleting a Schedule)

Creating a Schedule

Use the `asmcli create schedule` command to create a Smart Copy schedule by providing a name and list of one or more mount points. For general information about schedules, see About Smart Copy Schedules.

asmcli create schedule Command Syntax

The command syntax is as follows:

```
# asmcli create schedule
    --collection=collection-name | -c
    --name=schedule-name | -n
    --source=mount-point1 [mount-point2...] | -s
   [--comment=comment-string | -c]
   [--retry-wait=n | -w]
   [--force | -F]
   [--freeze-time n]
   [--group-access-file=file-location]
   [--keep=[n] | -K]
   [--output=always|error]
   [--retry-count=n | -r]
   [--type=volume-type | -t]
   [--user-name=user-name | -u]
```

The following table describes the required and optional command parameters.

Table 26. Create Schedule Command Options

Option	Abbr.	Argument	Description
Required Parameters			
--collection	-c	*collection-name*	Specifies the Smart Copy collection name. All Smart Copies, even a single one, are stored in a collection. Do not use with `--source`.
--name	-n	*schedule-name*	Specifies a Smart Copy schedule name.
--source	-s	*mount-point*	Specifies a list of one or more mount points to be copied when creating the Smart Copy. Do not use with `--collection`.
Optional Parameters			
--comment	-C	*comment-string*	Provides a comment string describing the Smart Copy being created. The comment is set for the PS Series group and is stored in the backup document.
--force	-F	none	Creates the Smart Copy even if all mount points cannot be frozen. When creating a Smart Copy, ASM/LE attempts to freeze all mount points. This option creates the Smart Copy, but the Smart Copy might be inconsistent.
--freeze-time	none	*n*	Specifies the maximum number of seconds to freeze the file systems when creating a Smart Copy. If this value is exceeded,

Option	Abbr.	Argument	Description
			the operation times out and a consistent Smart Copy is not created. For most configurations, the default time is sufficient. When creating Smart Copies with a large number of mount points or a heavy I/O load, ASM/LE might not have enough time to freeze all of the mount points and build the Smart Copy. Use this option to increase the time limit.
--group-access-file	-g	*file-location*	Specifies the location of the group access file that stores the user names and encrypted passwords.
--keep	-K	n	Specifies the maximum number of Smart Copies to keep for the scheduled task. By default, this value is 12. This option prevents too many Smart Copies from accumulating. If more than this number are created, the oldest is deleted until the keep count is reached. If ASMCLI created snapshots or replicas, the deleted Smart Copies result in a broken Smart Copy because the backup document refers to snapshots or replicas no longer in the group.
--output	none	always|error	Specifies when to display output to the console. Use `error` to display output only when an operation fails, such as with a scheduled operation. Use `always` (default) to display all output from the operation.
--retry-count	-r	*n*	Specifies the number of times to retry a failed operation.
--retry-wait	-w	*n*	Specifies the number of seconds to wait before retrying a failed operation.
--type	-t	*volume type* replica|snapshot	Specifies the Smart Copy type to create. If omitted, creates `snapshot`. ASMCLI reports an error if an incorrect type is applied, such as `thinclone` for a volume that is not a template volume.
--user-name	-u	*user-name*	Specifies a particular user identity when multiple user names have been stored in the group access file.

Example: Creating Schedules

The following example creates a schedule named `user-hourly` to capture user activity over a 24-hour period. After creating a schedule, enable it using the `asmcli do schedule` command. For an example, see Running a Schedule.

```
# asmcli create schedule --source /source/VolB --comment
"Hourly snapshot of user space" --keep 24 --name user-hourly

Successfully created schedule "user-hourly".
```

Listing Schedules

Use the `asmcli list schedule` command to list Smart Copies created by a given schedule. The command option lets you specify all schedules or a specific schedule with the name that you supplied when creating the schedule (see Creating a Schedule).

asmcli list schedule Command Syntax

The command syntax is as follows:

```
# asmcli list schedule
    [--name=schedule-name | -n]
```

The following table describes the optional command parameter.

Table 27. List Schedule Optional Command Parameter

Parameter	Argument	Description
--name	*schedule-name*	Specifies the name of the scheduled operation.

Examples: Listing Schedules

This section includes examples of listing schedules with verbose output and listing schedules that are enabled in a cron job.

Example 1: Listing Schedules with Verbose Output

In addition to listing all enabled schedules and showing the command that created the schedule, verbose2 mode also displays the Object ID that uniquely identifies the Smart Copy set.

```
# asmcli list schedule -vv

==========================================================
Schedules
==========================================================

Name:     hourly-schedule
Source:      /source/VolB
Type:     snapshot
Count:     0 (limit 24)
ObjectId: 9-930983-8f4e82980-90bc0ea549ea2096
Command:  asmcli create schedule --keep 24 --type snapshot --source
           /source/VolB

Name:      user-hourly-schedule
Source:    /source/VolA
Type:      snapshot
Count:     0 (limit 12)
ObjectId:  8-7340fd-7d4c571c2-c33a43ffcda4df82
Command:   asmcli create schedule --keep 12 --type snapshot --source
          /source/VolA

2 schedules found.
```

Example 2: Listing Multiple Schedules Enabled in a Cron Job

ASMCLI scheduled operations can be easily incorporated into scripts or cron jobs and enabled. In this example, the schedule named `15 min` retains four Smart Copies before deleting the oldest.

```
# asmcli list schedule

==========================================================
Schedules
==========================================================
Name:        15min
Command:     asmcli create schedule --keep 4 --output error --retry-count 12
--comment "Test schedule cron: 15 minutes" --type snapshot --collection test1

Name:       hourly
Command:    asmcli create schedule --keep 12 --output error --retry-count 12
--comment "Test schedule cron: hourly" --type snapshot --collection test1

Name:      test2
Command:   asmcli create schedule --retry-wait 60 --comment "Test schedule test2" --type
snapshot --collection test2
```

```
Name:     test3
Command:  asmcli create schedule --keep 0 --type snapshot --collection
test3
```

```
4 schedules found.

# crontab -l
PATH=/usr/local/sbin:/usr/sbin:/sbin:/usr/local/bin:/usr/bin:/bin:/root/bin
```

```
# Run this scheduled operation every 15 minutes
*/15 * * * * asmcli do schedule --name 15min >/dev/null

# Run this scheduled operation every hour
0 * * * * asmcli do schedule --name hourly >/dev/null
```

```
# asmcli list smart-copy

============================================================
Smart Copies
============================================================

Source:          hostname.mycompany.com : /vol1
  Snapshot:     10-Jan-2017 10:15:01.459071
  Description:  Test schedule cron: 15 minutes
  Schedule:    15min
  Collection:   test1
    vola01 --> vola01-2017-01-10-10:14:07.16863 on windev11-grp

Snapshot:      10-Jan-2017 10:30:02.077439
  Description: Test schedule cron: 15 minutes
  Schedule:    15min
  Collection: test1
    vola01 --> vola01-2017-01-10-10:29:07.16864 on windev11-grp

Snapshot:      10-Jan-2017 10:45:01.485456
  Description: Test schedule cron: 15 minutes
  Schedule:    15min
  Collection: test1
    vola01 --> vola01-2017-01-10-10:44:07.16865 on windev11-grp

Snapshot:      10-Jan-2017 11:00:01.853489
  Description: Test schedule cron: hourly
  Schedule:    hourly
  Collection:  test1
    vola01 --> vola01-2017-01-10-10:59:07.16866 on windev11-grp

Snapshot:      10-Jan-2017 11:00:01.853586
  Description: Test schedule cron: 15 minutes
  Schedule:   15min
  Collection: test1
   vola01 --> vola01-2017-01-10-10:59:07.16867 on windev11-grp

5 Smart Copies found.
```

At the scheduled interval, the keep count removed one Smart Copy when a new one was created.

```
# asmcli list smart-copy

============================================================
Smart Copies
============================================================

Source:          hostname.mycompany.com : /vol1
  Snapshot:    10-Jan-2017 10:30:02.077439
  Description: Test schedule cron: 15 minutes
  Schedule:    15min
  Collection:   test1
    vola01 --> vola01-2017-01-10-10:29:07.16864 on windev11-grp

  Snapshot:    10-Jan-2017 10:45:01.485456
  Description: Test schedule cron: 15 minutes
  Schedule: 15min
  Collection: test1
     vola01 --> vola01-2017-01-10-10:44:07.16865 on windev11-grp

  Snapshot:       10-Jan-2017 11:00:01.853489
  Description: Test schedule cron: hourly
  Schedule:    hourly
  Collection:  test1
```

```
    vola01 --> vola01-2017-01-10-10:59:07.16866 on windev11-grp

  Snapshot:           10-Jan-2017 11:00:01.853586
  Description: Test schedule cron: 15 minutes
  Schedule:       15min
  Collection:  test1
    vola01 --> vola01-2017-01-10-10:59:07.16867 on windev11-grp

  Snapshot:       10-Jan-2017 11:15:01.881889
  Description: Test schedule cron: 15 minutes
  Schedule:    15min
  Collection:  test1
     vola01 --> vola01-2017-01-10-11:14:07.16868 on windev11-grp

5 Smart Copies found.
```

Running a Schedule

Use the `asmcli do schedule` command to run a schedule to create Smart Copies. The command option lets you specify the schedule name you supplied when creating the schedule (see Creating a Schedule).

asmcli do schedule Command Syntax

The command syntax is as follows.

```
# asmcli do schedule
   --name=schedule-name | -n
```

The following table describes the required command parameter:

Table 28. Do Schedule Required Command Option

Parameter	Argument	Description
--name	*schedule-name*	Specifies the name of the scheduled operation.

Examples: Running Schedules

This section contains examples of running a schedule from the ASMCLI and running a schedule in a cron job.

Example 1: Running a Schedule

This example shows the `user-hourly` schedule. When executing the `do schedule` command, it runs the schedule and creates the first Smart Copy. Use the `asmcli list schedule` command to verify schedule details prior to running the schedule.

```
# asmcli list schedule --name user-hourly

Name:       user-hourly
Command:    asmcli create schedule --keep 24 --comment "Hourly snapshot
of user space: --type snapshot --source /source/VolB

# asmcli do schedule --name user-hourly

The mount point has been frozen.
Created snapshot donvolb-2017-01-05-15:23:20.224.1 on glengrp
The mount point has been thawed.
Successfully created a Smart Copy from 1 target.
```

Example 2: Running a Schedule in a Cron Job

In this example, the scheduled operation is included in a cron table. You can modify the schedule with ASMCLI (see Updating a Schedule) without modifying any scripts or cron tables. That is, you can change the options on the `create smart-copy` command at the command line through ASM without making other changes.

```
# asmcli list schedule

==========================================================
Schedules
==========================================================
Name:        15min
Command:     asmcli create schedule --keep 4 --output error
  --retry-count 12 --comment "Test schedule cron: 15 minutes"
  --type snapshot --collection test1

Name:       hourly
Command:    asmcli create schedule --keep 12 --output error
 --retry-count 12 --comment "Test schedule cron: hourly"
 --type snapshot --collection test1

Name:      test2
Command:   asmcli create schedule --retry-wait 60 --comment "Test
schedule test2" --type snapshot --collection test2

Name:     test3
Command:  asmcli create schedule --keep 0 --type snapshot --collection test3

4 schedules found.
```

```
# crontab -l
PATH=/usr/local/sbin:/usr/sbin:/sbin:/usr/local/bin:/usr/bin:/bin:/root/bin

# Run this scheduled operation every 15 minutes
*/15 * * * * asmcli do schedule --name 15min >/dev/null

# Run this scheduled operation every hour
0 * * * * asmcli do schedule --name hourly >/dev/null
```

Updating a Schedule

Use the `asmcli update schedule` command to modify and update a Smart Copy schedule by providing a name and list of one or more mount points. Use this command to change existing schedules by adding or replacing existing options. For example, you can modify a schedule to create Smart Copy replicas instead of snapshots.

asmcli update schedule Command Syntax

The command syntax is as follows:

```
# asmcli update schedule
    --name schedule-name | -n
   [--collection collection-name | -c]
   [--comment comment-string | -c]
   [--delete option1 [option2]
   [--force | -F]
   [--freeze-time n]
   [--group-access-file=file-location |-g]
   [--keep n| -N]
   [--new-name new-schedule-name | -N]
   [--output error | always]
   [--retry-count n | -r]
   [--retry-wait n | -w]
   [--source mount-point1 [mount-point2...] | -s]
   [--type volume-type | -t]
   [--user-name user-name | -u]
```

The following table describes the required and optional command syntax.

Table 29. Update Schedule Command Options

Parameter	Abbr.	Argument	Description
Required Parameter			
--name	-n	*schedule-name*	Specifies the name of the schedule.
Optional Parameters			
--collection	-c	*collection-name*	Specifies the name of a Smart Copy collection. All Smart Copies, even a single one, are stored in a collection. Do not use with `--source`.
--comment	-C	*comment-string*	Provides a comment string to describe the Smart Copy that you are updating. The comment is set for the PS Series group and is stored in the backup document.
--delete	-x	*option*	Specifies any options to be removed from the schedule. For example, if a comment string is no longer needed, enter `--delete comment`.
--force	-F	none	Updates the Smart Copy even if one or more mount points cannot be frozen. When updating a Smart Copy, ASM/LE attempts to freeze all mount points. If the freeze operation fails, such as if a file system type does not allow freezing and thawing mount points, this option updates the Smart Copy and updates the Smart Copies, but they might be inconsistent.
--freeze-time	none	*n*	Sets the maximum number of seconds to freeze the file systems when updating a Smart Copy. If this value is exceeded, the operation times out and a consistent Smart Copy is not updated. For most configurations, the default time is sufficient. If you are updating Smart Copies with a significant number of mount points or a heavy I/O load, ASM/LE might not be able to freeze all of the mount points and build the Smart Copy in the default time. Use this option to increase the time limit.
--group-access-file	-g	file-location	Specifies the location of the group access file that stores the user names and encrypted passwords.
--keep	-k	*n*	Specifies the maximum number of Smart Copies to keep for the scheduled task. By default, this value is 12. Use this option to prevent too many Smart Copies from accumulating. If more than this number are created, the oldest are deleted until the keep count is reached. If snapshots or replicas were created by ASMCLI, the deleted Smart Copies cause a broken Smart Copy because the backup document refers to snapshots or replicas no longer in the group.
--new-name	-N	*new-schedule-name*	New name for the schedule.
--output	none	error | always	Specifies when to display output on the console. Use `error` to display output only when an operation fails, such as with a scheduled operation. Use `always` (default) to display all output from the operation.
--retry-count	-r	*n*	Specifies the number of times to retry a failed operation.
--retry-wait	-w	*n*	Specifies the number of seconds to wait before retrying a failed operation.
--source	-s	*mount-point*	Lists one or more mount points to be copied. Do not use with `--collection`.
--type	-t	*volume type* clone| replica|snapshot| thinclone	Indicates the type of Smart Copy to create. If omitted, defaults to `snapshot`. ASMCLI displays an error when applying an incorrect type, such as specifying `thinclone` for a volume other than a template volume.

Parameter	Abbr.	Argument	Description
--user-name	-u	*user-name*	Specifies a particular user identity, when multiple user names have been stored in the group access file.

Examples: Updating Schedules

This section includes examples of updating a schedule to change the Smart Copy type, to change a schedule name, to delete an option in a schedule, and to create a collection.

Example 1: Updating a Schedule to Change the Smart Copy Type

In this example, you change a schedule so that it creates replicas instead of snapshots using the `--type` option:

```
asmcli update schedule --name user-hourly --type replica
--comment "Schedule to create replica smart copies"

Successfully updated schedule "user-hourly".

# asmcli list schedule

============================================================
Schedules
============================================================

Name:      user-hourly
Command:    asmcli create schedule --keep 24 --comment "Schedule
to create replica smart copies" --type replica --source /user1 /user2
/user3

1 schedule found.
```

Example 2: Changing the Schedule Name

This example uses the `--new-name` option to change the name of the schedule.

```
# asmcli update schedule --name user-hourly --new-name
user-hourly-replicas

Successfully updated schedule "user-hourly".

# asmcli list schedule

============================================================
Schedules
============================================================


Name:      user-hourly
Command:    asmcli create schedule --keep 24 --comment "Schedule to create replica
smartcopies" --type replica --source /user1 /user2 /user3

1 schedule found.
```

Example 3: Deleting an Option in a Schedule

In this example, the `--delete` option is used to remove a comment string.

```
# asmcli list schedule

============================================================
Schedules
============================================================

Name:      user-hourly
Command:    asmcli create schedule --keep 24 --comment "Schedule to update user
information on an hourly basis" --type snapshot --source /user1 /user2 /user3
```

```
1 schedule found.

# asmcli update schedule --name user-hourly --delete comment

Successfully updated schedule "user-hourly".

# asmcli list schedule

=========================================================
Schedules
=========================================================

Name:    user-hourly
Command: asmcli create schedule --keep 24 --type snapshot --source
/user1 /user2 /user3

1 schedule found.
```

Example 4: Updating a Schedule to Create a Collection

If you have a schedule that explicitly defines the source mount points using the `--source` option and if you want to use a collection instead, you cannot add the `--collection` option to the schedule without first deleting the `--source` option. For example:

```
# asmcli create schedule --name user-hourly --source /user1 /user2 /user3
  --keep 24

Successfully created schedule "user-hourly".

# asmcli update schedule --name user-hourly --collection all-user-space

Error: Cannot specify --collection and --source at the same time.

# asmcli create schedule --name user-hourly --collection all-user-space
  --delete source

Successfully updated schedule "user-hourly".
```

Deleting a Schedule

Use the `asmcli delete schedule` command to delete Smart Copy schedules. The command option lets you specify the schedule name you supplied when creating the schedule (see Creating a Schedule).

asmcli delete schedule Command Syntax

The command syntax is as follows:

```
# asmcli delete schedule
   --name schedule-name | -n
```

The following table describes the required command parameter.

Table 30. Delete Schedule Required Command Option

Parameter	Argument	Description
--name	*collection-name*	Specifies the name of the collection.

Example: Deleting a Schedule

This example lists all schedules and then deletes a specific schedule.

```
# asmcli list schedule

=========================================================
Schedules
=========================================================
```

```
Name:      user-hourly
Command:   asmcli create schedule --keep 24 --comment "Hourly snapshot of user space"
--type snapshot --source /source/VolB

Name:      user-hourly-all
Command: asmcli create schedule --keep 24 --type snapshot --source /source/VolA /source/VolB

2 schedules found.

# asmcli delete schedule --name user-hourly-all

Successfully deleted schedule "user-hourly-all".

# asmcli list schedule

==========================================================
Schedules
==========================================================
Name:      user-hourly
Command:   asmcli create schedule --keep 24 --comment "Hourly snapshot of user space"
--type snapshot --source /source/VolB

1 schedule found.
```

Mounting and Unmounting Smart Copies

Mounting a Smart Copy makes the file system contents available.

This section provides the following information:

- Mounting a Smart Copy
- Unmounting a Smart Copy

Mounting a Smart Copy

Use the `asmcli mount smart-copy` command to mount a Smart Copy. For general information and prerequisites for mounting Smart Copies, see About Mounting Smart Copies.

Mounting Read-Only Smart Copies

ASM/LE creates a consistent Smart Copy by freezing all file systems at the same time. When a journaling file system is frozen, the journals are locked in a consistent state. In most cases, emptying the journal by performing all of the updates would take too long. Because the journals are an integral part of the file system, the frozen file systems are consistent as long as the changes in their journals can be replayed.

If a `mount` command is used when the media is read-write and the file system is mounted with read-only access, the journals are replayed and the changes are made to on-disk structures, but no further changes can be made to the file system.

A problem might occur if a file system journal contains critical updates but the media is read-only. For example, if a journal needs to be replayed but the iSCSI target has been switched to read-only access, the journal cannot be loaded because the updates to the on-disk structures cannot be made, and the mount command fails.

On PS Series groups, this situation can be created by changing the access type of a volume from read-write to read-only or by changing a volume from a standard volume to a template volume.

Some file systems offer the `norecovery` or `noload` option to mount the file system without replaying the journals.

```
# mount -t xfs -r -o norecovery -o nouuid
  /dev/eql/temporary/1-e81da0-a541acdd9-a23a0a2be971b6b3+sample-xfs /mnt/temp
```

This option can create inconsistencies if the modifications contained in the journals are needed to correctly reflect the state of the file system at the time that the target was copied. Changes to the file system that are only reflected in the journal are ignored. To avoid this situation, use the following process:

1. Mount the file system with read-write media.
2. Let the file system replay the journal.

3. Unmount the file system.

asmcli mount smart-copy Command Syntax

The command syntax is as follows:

```
# asmcli mount smart-copy
    --destination directory-name | -d
   [--group-access-file file-location | -g]
   [--output always | error]
   [--read-only | -R]
   [--retry-count n | -r]
   [--retry-wait n | -w]
   [--user-name user-name | -u]
```

Search option parameters also apply to this command. For information, see ASMCLI Search Options.

The following table describes the required and optional command parameters.

Table 31. Mount Smart Copy Command Options

Parameter	Abbr.	Argument	Description
Required Parameter			
--destination	-d	*directory-name*	Directory where the Smart Copies are temporarily mounted.
Optional Parameters			
--group-access-file	-g	*file-location*	Specifies the location of the group access file that stores user names and encrypted passwords.
--output	none	always | error	Specifies when to show output. For scheduled operations it is often useful to suppress output to the console unless the operation fails (`--output error`). The default is `--output always`, indicating that output is always shown.
--read-only	-R	none	Specifies to mount as read-only the file systems from the Smart Copy.
--retry-count	-r	*n*	Number of times to retry a failed operation.
--retry-wait	-w	*n*	Number of seconds to wait before retrying a failed operation.
--user-name	-u	*user-name*	Specifies a particular user identity when multiple user names have been stored in the group access file.

Example: Mounting a Smart Copy

This example mounts a snapshot Smart Copy by specifying the source location and specific Object ID.

```
# asmcli mount smart-copy --source /source/VolC --object
   3-7f1425-994a453e5-075236f1b70ee888 --destination /mount/VolC

Logging in 1 Smart Copy target.
Logged in donvolc-2017-01-05-14:13:50.220.1 on glengrp
Mounted /mount/VolC/source/VolC

Successfully mounted 1 file system from 1 Smart Copy target.
```

Unmounting a Smart Copy

Use the `asmcli unmount smart-copy` command to unmount a Smart Copy. Unmounting automatically logs out of the volume, unless overridden by the `--nologout` option.

asmcli unmount smart-copy Command Syntax

The command syntax is as follows:

```
# asmcli unmount smart-copy
   [--group-access-file file-location | -g]
   [--limit number ]
   [--nologout]
   [--output always | error]
   [--retry-count n | -r]
   [--retry-wait n | -w]
   [--user-name user-name | -u]
```

Search option parameters also apply to this command. For information, see ASMCLI Search Options.

The following table describes the optional command parameters.

Table 32. Unmount Smart Copy Optional Command Parameters

Parameter	Abbr.	Argument	Description
--group-access-file	-g	*file-location*	Specifies the location of the group access file that contains the user names and encrypted passwords.
--limit		*number*	Limits the number of Smart Copies ASM will unmount to this number. **CAUTION: Use caution when limiting a large number of Smart Copies. You might inadvertently unmount all Smart Copies by specifying too large a number.**
--nologout	none	none	Keeps sessions logged in after the volume is unmounted.
--output	none	always | error	Specifies whether or not to suppress output. For scheduled operations, it is often useful to suppress output to the console unless the operation fails (`--output error`). The default is `--output always`, indicating that output is always shown.
--retry-count	-w	*n*	Specifies the number of times to retry a failed operation. For example, if you are creating a Smart Copy snapshot of two separate volumes and one is offline, the operation will attempt to connect every *n* seconds
--retry-wait	-r	*n*	Specifies the number of seconds to wait before retrying a failed operation.
--user-name	-u	*user-name*	Specifies a particular user identity when multiple user names have been stored in the group access file.

Example: Unmounting a Smart Copy

To unmount a Smart Copy, you need to specify the Object ID.

1. Locate the object ID (`Objectid`) of the Smart Copy by listing the Smart Copies in very verbose mode:

```
# asmcli list smart-copy --verbose 2

==========================================================
Smart Copies
==========================================================
Source: hostname.mycompany.com : /source/VolA
Snapshot: 14-Dec-2016 10:46:44.839798
Status: Mounted as /snapshots/VolA
ObjectId:5-392e06-8244f99a7-88d5834731b41690
Directory: /var/lib/equallogic/asm/smart-copies/asm-le-hosta./source-
a-5d5e2d-874d8ac30-0d2df1dc89f8932c/
File: 20161214-154644.839798-5-392e06-8244f99a7-88d5834731b41690.bcd
vola --> vola-2016-12-14-10:46:48.1421.1 on GroupA
1 Smart Copy found.
```

2. Specify the object ID in the unmount command:

```
# asmcli unmount smart-copy --object 5-392e06-8244f99a7-88d5834731b41690
```

If you know the Smart Copy is the latest created, you can bypass using the Object ID and unmount using the `--latest` option. Similarly, you can unmount the earliest-created Smart Copy using the `--earliest` option. For example:

```
 # asmcli unmount smart-copy --latest --source /source/volA

Unmounted /snapshots/vola/source/volA
Removing sessions to 1 Smart Copy target.
Logged out VolA-2016-12-14-10:46:48.1421.1 on Group A
Successfully unmounted 1 mount point and logged out from 1 target.
```

NOTE: To unmount, I/O cannot be in progress to the volume and the volume cannot have any open files.

Logging In to and Out of Smart Copies

This section presents the following information:

- Logging In to a Smart Copy
- Logging Out of a Smart Copy

Logging In to a Smart Copy

Use the `asmcli login smart-copy` command to log in to a Smart Copy. For conceptual information about logging in to Smart Copies, see About Logging In to Smart Copies.

asmcli login smart-copy Command Syntax

The command syntax is as follows:

```
# asmcli login smart-copy
   [--group-access-file file-location | g]
   [--output always | error]
   [--retry-count n | -r]
   [--retry-wait n | -w]
   [--username user-name | -u]
```

Search option parameters also apply to this command. For information, see ASMCLI Search Options.

The following table describes the optional command parameters.

Table 33. Login Smart-Copy Optional Command Parameters

Parameter	Abbr.	Argument	Description
--group-access-file	-g	*file-location*	Specifies the location of the group access file where the user names and encrypted passwords are stored.
--output	none	always | error	Specifies whether or not to suppress output. For scheduled operations it is often useful to suppress output to the console unless the operation fails (`--output error`). The default is `--output always`, which always shows output.
--retry-count	-r	*n*	Specifies the number of times to retry a failed operation.
--retry-wait	-w	*n*	Specifies the number of seconds to wait before retrying a failed operation.
--user-name	-u	*user-name*	Specifies a particular user identity when multiple user names have been stored in the group access file.

Example: Logging In to Smart Copies

The following example logs in to the latest Smart Copy created by a schedule.

```
# asmcli login smart-copy --schedule user-hourly --source/source/VolB
  --latest

Logging in 1 Smart Copy target.
Logged in donvolb-2017-01-05-17:03:01.225.1 on glengrp
Created device /dev/mapper/eql-8-b8b03d-7444e1f3e-626856b008bde08e+donvolb

Successfully logged in 1 file system from 1 Smart Copy target.
```

Listing the Smart Copy shows it is logged in:

```
# asmcli list smart-copy --verbose 2

============================================================
Smart Copies
============================================================

Source: hostname.mycompany.com : /source/VolB
  Snapshot: 05-Jan-2017 17:05:57.846661
  Status: Logged-in
  Description: Hourly snapshot of user space
  Schedule: user-hourly
  ObjectId: 8-b8b03d-7444e1f3e-626856b008bde08e
  Directory: /var/lib/equallogic/asm/smart-copies/asm-le-dk-test2.com
  /source-1-21a09d-9ce520990-c75a20bb9b1eb86f/
  File: 20170105-220557.846661-8-b8b03d-7444e1f3e-626856b008bde08e.bcd
    donvolb --> donvolb-2017-01-05-17:03:01.225.1 on glengrp
```

Logging Out of a Smart Copy

Use the `asmcli logout smart-copy` command to log out of a Smart Copy. You must unmount a Smart Copy before you can log out (see Unmounting a Smart Copy).

asmcli logout smart-copy Command Syntax

The command syntax is as follows:

```
# asmcli logout smart-copy
   [--group-access-file file-location | -g]
   [--limit number]
   [--output=always | error]
   [--retry-count n | -r]
   [--retry-wait n | -w]
   [--user-name user-name | -u]
```

Search option parameters also apply to this command. For information, see ASMCLI Search Options.

The following table describes the optional command parameters.

Table 34. Logout Smart Copy Optional Command Parameters

Parameter	Abbr.	Argument	Description
--group-access-file	-g	*file-location*	Specifies the location of the group access file that contains the user names and encrypted passwords.
--limit		*number*	Limit the number of Smart Copies ASM will log out to this number. **NOTE: Use caution when limiting a large number of Smart Copies; you might inadvertently log off all Smart Copies by specifying too large a number.**

Parameter	Abbr.	Argument	Description
--output		always | error	Specifies when to display output to the console. Use `error` to display output only when an operation fails. Use `always` (default) to display all output from the operation.
--retry-count	-w	*n*	Specifies the number of times to retry a failed operation. For example, if you are creating a Smart Copy snapshot of two separate volumes and one is offline, the operation will attempt to connect every *n* seconds.
--retry-wait	-r	*n*	Specifies the number of seconds to wait before retrying a failed operation.
--user-name	-u	*user-name*	Specifies a particular user identity when multiple user names have been stored in the group access file.

Example: Specifying an Object ID with Logout

The following listing shows that a snapshot Smart Copy created by schedule `user-hourly` is currently logged in:

```
# asmcli list smart-copy -vv

==========================================================
Smart Copies
==========================================================

Source: myhost.mycompany.com : /source/VolB
  Snapshot: 05-Jan-2017 17:05:57.846661
  Status: Logged-in
  Description: Hourly snapshot of user space
  Schedule: user-hourly
  ObjectId: 8-b8b03d-7444e1f3e-626856b008bde08e
  Directory: /var/lib/equallogic/asm/smart-copies/asm-le-dk-test2.com
  /source-1-21a09d-9ce520990-c75a20bb9b1eb86f/
  File: 20170105-220557.846661-8-b8b03d-7444e1f3e-626856b008bde08e.bcd
    volb --> volb-2017-01-05-17:03:01.225.1 on glengrp
```

The `logout` command is entered with an object ID:

```
# asmcli logout smart-copy --schedule user-hourly --source /source/VolB
   --object 8-b8b03d-7444e1f3e-626856b008bde08e

Removing sessions to 1 Smart Copy target.
Logged out volb-2017-01-05-17:03:01.225.1 on glengrp

Successfully logged out from 1 target.
```

Creating and Managing Collections

Collections are related sets of mount points from which you create Smart Copies. ASMCLI provides several commands for creating and managing collections:

- Creating a Collection
- Listing Collections
- Updating a Collection
- Deleting a Collection

Creating a Collection

Use the `asmcli create collection` command to create a collection by providing a name and list of one or more mount points. For information about collections, see About Collections.

asmcli create collection Command Syntax

The command syntax is as follows:

```
# asmcli create collection
   --name collection-name
   --source mount-point1 [mount-point2...]
```

The following table describes the required command parameters.

Table 35. Create Collection Required Command Options

Parameter	Argument	Description
--name	*collection-name*	Specifies the name of the collection.
--source	*mount-point*	Specifies a list of one or more mount points.

Example: Creating Collections

The following example shows how to use separate collections to save copies of the same volumes:

```
# asmcli create collection --name=daily-user-vols yes--source=/source/VolA
   --source=/source/VolB --source=/source/VolC
```

You could also specify the same command as follows:

```
# asmcli create collection --name daily-user-vols --source /source/VolA
  /source/VolB /source/VolC
```

The output specifies a successfully created collection of the specified name (for example, `daily-user-vols`). Use the `asmcli list collection` command to show the contents of the collection:

```
# asmcli list collection --name daily-user-vols

Name:          daily-user-vols
Mount point:   /source/VolA
Mount point:   /source/VolB
Mount point:   /source/VolC
```

Listing Collections

Use the `asmcli list collection` command to list all collections or a specific collection by a name. For information about collections, see About Collections.

asmcli list collection Command Syntax

The command syntax is as follows:

```
# asmcli list collection
   --name collection-name
```

The following table describes the optional command parameter.

Table 36. List Collection Optional Command Parameter

Parameter	Argument	Description
--name	*collection-name*	Specifies the name of the collection.

Example: Listing Collections

This example specifies the name `daily-user-vols` and displays a collection containing three volumes.

```
# asmcli list collection --name daily-user-vols

============================================
Collections
============================================

Name:           daily-user-vols
Mount point:    /source/VolA
Mount point:    /source/VolB
Mount point:    /source/VolC

1 collection found.
```

Updating a Collection

Use the `asmcli update collection` command to modify a collection by adding to the collection, deleting from the collection, providing a new collection name, and providing new mount points.

asmcli update collection Command Syntax

The command syntax is as follows.

```
# asmcli update collection
    --name collection-name | -n
   [--add mount-point1 [mount-point2...] | -a]
   [--delete mount-point1 [mount-point2...] | -x]
   [--new-name new-name | G]
   [--user-name user-name | -u]
```

The following table describes the required and optional command parameters.

Table 37. Update Collection Command Options

Parameter	Abbr.	Argument	Description
Required Parameter			
--name	-n	*collection-name*	Specifies the name of the collection.
Optional Parameters			
--add	-a	*mount-point*	Specifies a list of mount points to add to the collection.
--delete	-x	*mount-point*	Specifies a list of mount points to remove from the collection.
--new-name	-N	*new-name*	Specifies a new name for the collection.
--source	-s	*mount-point*	Specifies the list of mount points for volumes in the collection.

Examples: Updating Collections

This section provides examples of adding a volume to a collection and deleting a volume from a collection.

Example 1: Adding to a Collection

This example adds the volume `VolC` to the collection `daily-user-vols`.

```
#asmcli update collection --name daily-user-vols --add
/source/VolC

Successfully updated collection "daily-user-vols".
```

Example 2: Deleting from a Collection

This example deletes the volume `VolC` from the collection.

```
# asmcli list collection --name daily-user-vols

==============================================
Collections
==============================================

Name:           daily-user-vols
Mount point:    /source/VolA
Mount point:    /source/VolB
Mount point:    /source/VolC

1 collection found.

#asmcli update collection --name daily-user-vols --delete /source/VolC

Successfully updated collection "daily-user-vols".

# asmcli list collection --name daily-user-vols

==============================================
Collections
==============================================

Name:           daily-user-vols
Mount point:    /source/VolA
Mount point:    /source/VolB

1 collection found.
```

Deleting a Collection

Use the `asmcli delete collection` command to delete a collection by name. For information about collections, see About Collections.

asmcli delete collection Command Syntax

The command syntax is as follows:

```
# asmcli delete collection
  --name=collection-name
```

The following table describes the required command parameter.

Table 38. Delete Collection Command Required Option

Parameter	Argument	Description
--name	*collection-name*	Specifies the name of the collection.

Example: Deleting Collections

This example deletes the collection `daily-user-vols` and all its mount points.

```
# asmcli delete collection --name daily-user-vols
Successfully deleted collection "daily-user-vols"
```

7

Managing Thin-Provisioned Volumes

Thin-Provisioned volumes enable you to recover disk space that has become available and make more efficient use of disk capacity in storage volumes.

Using the eqlvolume tool, you can show information on mount points and the EqualLogic volumes that support them. You can also display a list of volumes recommended for thin-provisioning (also called rethinning) or that are mounted with the `-o discard` option, and specify which volumes are to be thin-provisioned.

NOTE: See the *Dell EqualLogic PS Series Group Manager Administrator's Guide* for considerations and limitations regarding thin-provisioned volumes.

Topics:

About Thin-Provisioning

Thin-provisioning informs a PS Series group that disk space has been deallocated, such as when a file is removed or truncated. The host sends an iSCSI `UNMAP` command to the PS Series group.

NOTE:

Thin-provisioning requires all volumes in the target to be EqualLogic volumes that are on a PS Series group running a supported firmware version.

NOTE:

- **Before you use the `eqlvolume rethin` command, use the `eqlvolume show` command to display information about the mounted volumes. See Show Mount Point Information.**
- **For additional information about the volumes, run the `eqltune -v` command. See EqualLogic Host Performance and Tuning Suite (eqltune).**
- **For considerations regarding thin-provisioning, see Considerations for Thin-Provisioning.**

Thin-provisioning is of two types:

- Realtime Thin-Provisioning
- On-Demand Thin-Provisioning

Realtime Thin-Provisioning

Realtime thin-provisioning occurs when the host automatically deallocates disk space whenever a file is removed or truncated. Realtime volume thin-provisioning is enabled when the file system is mounted with the `-o discard` option, which enables the file system to send SCSI `UNMAP` commands to the SCSI array when sufficient disk space is deallocated.

Older Linux kernels and file systems might not support the realtime discard feature, and this feature is not optimized. Deleting and writing small files repeatedly might cause unnecessary discards to be sent for regions of the target volume that will be overwritten. Before enabling realtime, discard support and investigate the performance impact the discard feature will have on your system.

Some mount points should not have realtime discard support enabled, due to the nature of the file systems and the characteristics of the target volumes. To determine if any mount points are incorrectly mounted with `-o discard`, use `eqlvolume rethin --discard-usage`.

The option is supported in these Linux distributions:

- RHEL 6.8 Server and higher: `ext4`, `gfs2`, and `xfs`.
- SLES 12 and higher: `ext3`, `ext4`, and `xfs`.

NOTE: By default, the file system when installing SLES 12 and later is `btrfs`, which is not supported by ASM/LE. Select `ext4` or `xfs` when installing SLES 12.

On-Demand Thin-Provisioning

On-demand thin-provisioning enables you to disable realtime volume thin-provisioning and to run volume thin-provisioning either manually or as a scheduled maintenance activity. On-demand volume thin-provisioning ensures that storage previously allocated for files that have been deleted is marked as free, and that the volume is as thin-provisioned as possible. This provisioning is useful for recovering free space that was lost in the past and is not likely to be recovered soon. It can also be used, instead of realtime thin-provisioning , as part of periodic maintenance.

To perform on-demand thin-provisioning , the **eqlvolume** utility first attempts to issue a `FITRIM ioctl` to each mount point. The `FITRIM` call prompts the file system to examine its on-disk structure, identify what regions of disk space are not in use, and send iSCSI `UNMAP` commands for every unallocated region of space.

These Linux distributions support `FITRIM ioctl`:

- RHEL 6.8 Server and later: Support for `ext4` and `xfs`
- SLES 12 and later: Support for `ext4` and `xfs`

If `FITRIM` fails, **eqlvolume** remounts the file system with the `-o discard` option. This creates several temporary files to take up most of the free space, then deletes the temporary files, which prompts the file system to send `discard` commands to recover the free space occupied by the temporary files.

The result of a successful on-demand thin-provisioning operation would look as follows:

```
# eqlvolume rethin
Rethinning /lv21a ... Initiated successfully
Rethinning /lv21b ... Initiated successfully
```

As with the `eqlvolume show` command, the `eqlvolume rethin --include` specifies a list of mount points to thin-provision instead of all mount points supported by EqualLogic volumes. Similarly, the `eqlvolume rethin --exclude` command specifies mount points to drop from the list of mount points to be thin-provisioned.

To see what volumes could be thin-provisioned before doing any rethinning operations, use the `eqlvolume rethin --recommend` command.

The `eqlvolume rethin --verbose` command returns the reasons why mount points were not considered for thin-provisioning.

For more examples, see:

- Show Mount Point Information
- Thin-Provisioning a Volume

Thin-Provisioning as Scheduled Maintenance

The `eqlvolume rethin` command can be periodically run using crontab or a comparable facility. For example, the following crontab entry requests eqlvolume thin-provisioning at 2:00 a.m. once a week:

```
0 2 * * 1 /usr/sbin/eqlvolume rethin --output error >/dev/null
```

The full path specification is needed to run ASMCLI commands through crontab. You could also redefine the variable *PATH* inside crontab to include `/usr/sbin`:

```
PATH=/usr/local/sbin:/usr/sbin:/sbin:/usr/local/bin:/usr/bin:/bin:/root/bin
0 2 * * 1 eqlvolume rethin --output error >/dev/null
```

NOTE:

- **If you use the `eqlvolume rethin --output error` command in crontab, output is sent to stderr only if an error is encountered.**
- **If stdout text is piped to `/dev/null`, the result is that email is sent only if an error occurs.**

Group Access Files

Access to the PS Series groups by eqlvolume is controlled by entries in the group access file. These entries are built and maintained by the asmcli utility.

For more information about group access files, use the `asmcli help group-access` command.

Considerations for Thin-Provisioning

Several characteristics are considered when eqlvolume determines whether or not a volume can be thin-provisioned. See the following sections for more information:

- File System and Volume Characteristics
- Multiple File Systems

File System and Volume Characteristics

The eqlvolume tool uses characteristics of the file system and the EqualLogic volume, as well as defined configuration parameters to determine if a volume can be rethinned. For information on the defined configuration parameters, see Volume Thin-Provisioning Configuration Parameters.

File system characteristics that prevent eqlvolume from recommending rethinning include:

- Read-only file system—A file system that is mounted as read-only cannot be updated, nor can space be allocated or deallocated.
- Lack of `FITRIM ioctl()` or `-o discard` support—Thin-provisioning requires either `FITRIM ioctl()` or `-o discard` support. If a file system was not mounted with the `-o discard` option, eqlvolume might need to remount that file system using the `-o discard` option, perform on-demand thin-provisioning, and remount the file system with the `-o nodiscard` option. A file system that does not support `-o nodiscard` cannot be restored to its original state after thin-provisioning, and eqlvolume will not recommend that file system for on-demand thin-provisioning.
- A `gfs2` file system without FITRIM support—Thin-provisioning a `gfs2` file system by allocating most of the disk space and deleting the temporary files is very slow, so eqlvolume does not recommend thin-provisioning a `gfs2` files system without FITRIM.
- Insufficient free space—Without sufficient free space in the file system, eqlvolume is unlikely to recover any significant free space on the volume.

EqualLogic volume characteristics that prevent eqlvolume from recommending thin-provisioning include:

- Target is a snapshot—Snapshots only store changes to the volume, so they are already space efficient.
- Volume is not t1in-provisioned—If the volume is not thin-provisioned, space is preallocated to the volume, and free space cannot be recovered.
- Firmware—Firmware in the PS Series group does not support unmaps, or discards.
- Replication or synchronous replication is configured—Replicated volumes do not support thin-provisioning.
- Volume is a replica failover volume—The product of a replicated volume is not a good candidate for thin-provisioning.

Multiple File Systems

If all of the characteristics of the file system and the volume indicate that on-demand thin-provisioning is possible, eqlvolume recommends thin-provisioning based on a comparison of the space used by the file systems and the volumes that support them.

When eqlvolume evaluates multiple file systems built by multiple logical volumes from one or more target volumes or file systems built by partitions, the function compares the total amount of space used by all of the related file systems to the total amount of space used by the target volumes. If the space used by the target volumes is significantly greater than the space used by the file systems, all of the file systems are recommended for thin-provisioning because it is impossible to predict what free space was originally part of which file system.

If a target volume supports multiple file systems through logical volumes or partitions, and one or more of those file systems are not mounted, eqlvolume cannot include the amount of space used by the unmounted file systems. In this circumstance, eqlvolume will likely recommend that the remaining file systems be thin-provisioned.

For example, if two partitions (file systems) each use 100GB, and the target volume has 300 GB of space in use, eqlvolume will recommend that the two partitions be thin-provisioned. However, if a third partition is using 100 GB of space but is not mounted, the space used by the third (unmounted) partition can cause eqlvolume to recommend thin-provisioning, although the thin-provisioning will likely not recover any free space.

Volume Thin-Provisioning Configuration Parameters

Configuration parameters in `/etc/equallogic/eql.conf` control the behavior of the `eqlvolume rethin` command. The eqlvolume tool uses these parameters, with file system and EqualLogic volume characteristics, to determine if a volume can be thin-provisioned. For information on these characteristics, see Considerations for Thin-Provisioning.

The following table describes the volume thin-provisioning configuration parameters.

Table 39. Volume Thin-Provisioning Configuration Parameters

Parameter	Default Value	Description
ReThinRecommend	3	Specifies the minimum percentage of available space required. If the percentage of space used by the target minus the percentage of space used by the file system is greater than or equal to the value specified in this parameter, the file system is recommended for volume thin-provisioning. Use this parameter to configure volume thin-provisioning on small storage volumes.
ReThinRecommendMB	2048	Specifies the minimum size (MB) of available space required. This value must be larger than 15. If the space used by the target minus the space used by the file system is greater than or equal to the value specified in this parameter, the file system is recommended for volume thin-provisioning. Use this parameter to configure volume thin-provisioning on large storage volumes.
ReThinHeadroom	2	Specifies the minimum percentage of free space that must exist. This value is also the minimum percentage of free space that must be left unallocated if a file system is thin-provisioned. This value is not used on systems where `FITRIM ioctl` is supported. On systems where `FITRIM ioctl` is not supported, this parameter and `ReThinHeadroomMB` ensure that free space is reserved for users during thin-provisioning. Use this parameter to manage free space available on the storage volume. **NOTE: If the file system is particularly active and allocation failures occur for user applications while thin-provisioning is taking place, increase this parameter and/or `ReThinHeadroomMB`.**
ReThinHeadroomMB	15	Specifies the minimum amount of free space, as a fixed limit in MB, that must exist for a volume to be thin-provisioned. This value is also the minimum amount of free space to be left unallocated if a file system is thin-provisioned. This value is not used on systems where `FITRIM ioctl` is supported. On systems where `FITRIM ioctl` is not supported, this parameter and `ReThinHeadroom` ensure that free space is reserved for users while thin-provisioning takes place. **NOTE: If the file system is particularly active and allocation failures occur for user applications while thin-provisioning is taking place, increase this parameter and/or `ReThinHeadroom`.**

Show Mount Point Information

Use the `eqlvolume show` command to locate the file systems, or mount points, supported by EqualLogic volumes and to display characteristics of those mount points and volumes.

eqlvolume show Command Syntax

The command syntax is as follows:

```
# eqlvolume show
   [--include mount-point | -i]
   [--exclude mount-point | -e]
   [--group-access-file file-location | -g]
   [--user-name user-name | -u]
```

The following table describes the optional command parameters.

Table 40. eqlvolume show Optional Command Parameters

Parameter	Abbr.	Argument	Description
--include	-i	*mount-point*	Specifies the mount points for which to display information. If not used, displays information on all mount points using EqualLogic volumes.
--exclude	-e	*mount-point*	Specifies the mount points for which information is not displayed.
--group-access-file	-g	*file-location*	Specifies the location of the group access file where the user names and encrypted passwords are kept.
--user-name	-u	*user-name*	If multiple user names have been stored in the group access file, specifies a particular user identity.

If you are able to connect to the PS Series group and get information about the volume, the characteristics of each volume are also shown.

The space used by each file system is calculated by subtracting the free space available from the total capacity. This number reflects the overhead created by each file system's on-disk structures. The percentage used is calculated from the file system space used and the capacity. These numbers are not the same as those reported by the Linux disk free (df) utility.

The minimum allocation sizes also differs. The file system layout determines whether space is actually needed on the array. For example, the file system reserves areas on the target for future on-disk structures and does not need to allocate space for those areas until they are written.

The space used by the target volumes and the capacity are acquired from the PS Series array. These numbers will differ from the file system numbers because the file systems might have been built using a subset of the target space to leave room for expansion.

Examples: Showing Mount Point Information

This section includes examples of using the `eqlvolume show` command to show mount point information.

The `eqlvolume show` command locates file systems, or mount points, that are supported by EqualLogic volumes to display the characteristics of the mount points and the volumes that support those mount points.

Example 1: eqlvolume show Command

This example shows the output of the `eqlvolume show` command with no options. In this output:

- Mount point `/lv11a` is a thin-provisioned volume with replication enabled. Mount point `/mnt/temp/lv11a` is a temporary mount of a replica mounted by Auto Snapshot Manager (ASM).
- Each mount point is represented in a single line above the volume name and group name of the EqualLogic volume used to support it.
- Relationships are shown by grouping the mount points above the volume on the PS Series group. In this example, the LVM Logical Volumes `/lv12a`, `/lv12b`, and `/lv12c` are built from a single volume group that is built from one Physical Volume. The mount points are grouped above the volume `vg12` on the PS Series group `eqlgroup1`.
- Each mount point is shown with its file-system type and the mount attributes.

```
# eqlvolume show
============================================================
File Systems
============================================================
File system or Volume Used Size Used% Type Attributes

/lv11a                150 MB     9.84 GB   1.5%    ext4    rw
```

```
  vg11 on eqlgroup1  1.71 GB   100.00 GB   1.7%    thin   replication

/lv12a                150 MB     9.84 GB   1.5%    ext4   rw
/lv12b                150 MB     9.84 GB   1.5%    ext4   rw
/lv12c                150 MB     9.84 GB   1.5%    ext4   rw
  vg12 on eqlgroup1  5.07 GB   100.00 GB   5.1%    thin

/VolA                5.18 GB    98.43 GB   0.2%    ext3   rw
  vola on eqlgroup2  2.65 GB   100.00 GB   2.7%    thin

/mnt/temp/lv11a       662 MB     9.84 GB   1.5%    ext4   rw
  vg11-2016-07-20-08:46:30.26820.1 on eqlgroup3 replica

/VolB                 534 MB     9.85 GB   0.3%    ext3   rw
  volb on eqlgroup2  1.29 GB    10.00 GB  12.9%    thin

7 file systems found.
```

Example 2: Including Specific Mount Points

To display the characteristics of certain mount points, use the `--include` option to specify the mount points to include.

```
# eqlvolume show --include /vg12a /vg12b /vg12c
============================================================
File Systems
============================================================
Filesystem or Volume    Used       Size    Used%     Type   Attributes

/lv12a                  150 MB     9.84 GB   1.5%     ext4   rw
/lv12b                  150 MB     9.84 GB   1.5%     ext4   rw
/lv12c                  150 MB     9.84 GB   1.5%     ext4   rw
  vg12 on eqlgroup1     5.07 GB  100.00 GB   5.1%     thin

3 file systems found.
```

Example 3: Excluding Specific Mount Points

To exclude certain mount points, use the `--exclude` option to specify the mount points to be eliminated from the list.

```
# eqlvolume show --exclude /mnt/temp/lv11a
============================================================
File Systems
============================================================
Filesystem or Volume    Used       Size    Used%   Type  Attributes

/lv11a                150 MB     9.84 GB   1.5%   ext4  rw
  vg11 on eqlgroup1  1.71 GB   100.00 GB   1.7%   thin  replication

/lv12a                150 MB     9.84 GB   1.5%   ext4  rw
/lv12b                150 MB     9.84 GB   1.5%   ext4  rw
/lv12c                150 MB     9.84 GB   1.5%   ext4  rw
  vg12 on eqlgroup1  5.07 GB   100.00 GB   5.1%   thin

/VolA                5.18 GB    98.43 GB   0.2%   ext3  rw
  vola on eqlgroup2  2.65 GB   100.00 GB   2.7%   thin

/VolB                 534 MB     9.85 GB   0.3%   ext3  rw
  volb on eqlgroup2  1.29 GB    10.00 GB  12.9%   thin

6 file systems found.
```

Thin-Provisioning a Volume

Use the `eqlvolume rethin` command to perform on-demand thin-provisioning or to get recommendations about thin-provisioning.

eqlvolume rethin Command Syntax

The command syntax is as follows:

```
# eqlvolume rethin
   [--recommend | -R]
   [--discard-usage | -D]
   [--include mount-point | -i]
   [--exclude mount-point | -e]
   [--force mount-point | -F]
   [--group-access-file file-location | -g]
   [--user-name user-name | -u]
```

The following table describes the optional command parameters.

Table 41. eqlvolume rethin Command Options

Parameter	Abbr	Argument	Description
--recommend	-R	none	Generates a list of volumes that appear to be good candidates for thin-provisioning, but does not rethin any volumes.
--discard-usage	-D	none	Displays a list of file systems that are incorrectly mounted with `-o discard`, but does not thin-provision any volumes.
--include	-i	*mount-point*	Specifies the mount points to be thin-provisioned. If this parameter is omitted, all mount points using EqualLogic volumes will be considered for thin-provisioning.
--exclude	-e	*mount-point*	Specifies a list of the mount points that are not to be thin-provisioned. If a mount point is listed with both `--include` and `--exclude` parameters, the mount point is excluded.
--force	-F	*mount-point*	Causes the specified volumes to be thin-provisioned, even if they do not meet the thin-provisioning criteria defined by the `ReThinRecommend` and `ReThinRecommendMB` configuration parameters.
--group-access-file	-g	*file-location*	Specifies the location of the group access file where the user names and encrypted passwords are kept.
--user-name	-u	*user-name*	If multiple user names have been stored in the group access file, specifies a particular user identity.

Examples: Analyzing Thin—Provisioning Recommendations

This section provides examples showing how modifying an **eqlvolume rethin** configuration parameter affects the volumes that are recommended for thin-provisioning. It also shows the verbose output of the`--recommend` parameter showing the reasons that volumes were not recommended.

Example 1: Modifying an eqlvolume rethin Configuration Parameter

This example recommends a file system for volume thin-provisioning, modifies the `--recommend` parameter to change the minimum percentage of space required to recommend thin-provisioning, and displays a different recommendation based on the new parameter value. For more information on these parameters, see Volume Thin-Provisioning Configuration Parameters.

```
 # eqlvolume rethin --recommend --verbose

Locating all mount points supported by EqualLogic targets.
Logged into PS Series group 192.x.x.xy (eqlgroup2) as username grpadmin  (GroupAdmin)
There is 1 mount point recommended for rethinning:
   /test04

The following 2 mount points were not recommended for rethinning:
   /test05 : File system space used (2.65 GB) is greater than space used by
```

```
      target (2.36 GB)
   /test06 : File system space used (3.17 GB) is greater than space used by
      target (2.52 GB)

# vi /etc/equallogic/eql.conf
# tail -n 4 /etc/equallogic/eql.conf

[EQLVOLUME Params]
ReThinRecommend = 50

# eqlvolume rethin --recommend --verbose

Locating all mount points supported by EqualLogic targets.
Logged into PS Series group 192.x.x.xy (eqlgroup2) as username grpadmin (GroupAdmin)
No mount points recommended for rethinning.

The following 3 mount points were not recommended for rethinning:
   /test04 : Potential recoverable space of 1.89 GB is less than
      2.00 GB fixed limit
   /test05 : File system space used (2.65 GB) is greater than space used by
      target (2.36 GB)
   /test0 : File system space used (3.17 GB) is greater than space used by
      target (2.52 GB)
```

Example 2: Volumes Recommended and Not Recommended for Thin-Provisioning

The following example uses the `--recommend` and `--verbose` parameters to display which volumes are recommended for thin-provisioning, and which volumes are not recommended for thin-provisioning, and the reasons those volumes are not recommended.

```
 # eqlvolume rethin --verbose --recommend

Locating all mount points supported by EqualLogic targets.
Logged into PS Series group 192.x.x.xy (eqlgroup2) as username linux-server
(VolumeAdmin) Logged into PS Series group 192.x.x.xx (eqlgroup1)
as username linux-server (VolumeAdmin) There are 2 mount points
recommended for re-thinning:
  /user2
  /user3

The following 6 mount points aren't recommended for re-thinning:
  /lv11a  : Volume user1-vg11 on eqlgroup2 (192.x.x.xx) is not thin
      provisioned
  /partp1 : Potential recoverable space of 287 MB is less than 307 MB
      (3% capacity)
  /partp2 : Potential recoverable space of 287 MB is less than 307 MB
      (3% capacity)
  /partp3 : Potential recoverable space of 287 MB is less than 307 MB
      (3% capacity)
  /user1  : Volume user1-rhel6i-1 on eqlgroup1 (192.x.x.xy) does not
      support unmap (discards)
  /user4  : Readonly file system
```

8

Troubleshooting

This chapter provides troubleshooting information for HIT/Linux.

Topics:

General Troubleshooting Guidelines

The following general troubleshooting guidelines apply when using HIT/Linux.

Ensuring Correct Operation of the dm-switch Kernel Module

If the `dm-switch` kernel module appears to be installed but fails to operate correctly, examine the `/var/log/messages` file because the `dm-switch` kernel module writes status messages using the standard kernel logging facility.

The log file might reveal a problem (for example, a system resource issue). If that occurs, contact your system administrator to adjust the environment to resolve the issue.

The ehcmd daemon writes debugging information into a log file in the `/var/log/equallogic/` directory. Error messages in the log file might provide important clues about the interaction between the HIT/Linux components and the operating system environment. For example, if Device Mapper commands (that is, the `dmsetup` commands) are failing, the commands are written in the log along with the errors returned by Linux. The errors might point to environmental issues that need to be resolved by the system administrator before multipath I/O can operate at peak efficiency.

Establishing iSCSI Sessions

If the ehcmd daemon is having trouble establishing the desired number of sessions, look at the event log on the PS Series group. In particular, you can look at the existing sessions and access control lists on the volume and then verify that all the Ethernet ports are allowed to access the volume if IP ACLs are being used. If multiple Linux initiators are connecting to the same volume, set the option `Allow simultaneous connections from initiators with different IQN names`.

NOTE: To prevent data corruption, set this option only when you have a cluster-aware file system or application using the volumes.

If the ehcmd daemon is still unable to connect from all Ethernet ports, a host or networking issue might be causing the problem. To narrow the scope of the problem, use iscsiadm to manually attempt to log in to a volume from each Ethernet port.

In addition, you can run eqltune and look at diagnostic information to configure an optimal Linux multipath I/O environment. For information on using eqltune, see EqualLogic Host Performance and Tuning Suite (eqltune).

Increasing the Volume Size on a PS Series Group

You can increase the size of a volume on the PS Series group. The ehcmd daemon periodically rescans all devices for the volumes that it manages for multipath I/O. If ehcmd detects the change in volume size, it attempts to rebuild the two-tier device configuration using the newer, larger capacity.

A critical part of this process is for ehcmd to issue an `iscsiadm -m node -R -T iqn` command to direct the Open iSCSI software to rescan all of the block-mode devices so that the attributes (especially the size) for those devices are updated. If this command is slow, you might see ehcmcli diagnostic information similar to the following output:

```
============================================================
Volume list
============================================================
 Volume: user2
   Target name: iqn.2001-05.com:0-8a0906-fde687707-60f1eeafe8a4d260-user2
   Device to mount: /dev/mapper/eql-0-8a0906-fde687707-60f1eeafe8a4d260-user2
   Status: Unable to build Device Mapper tables with the correct
   capacity
   Resolution: If this situation persists, use 'iscsiadm -m session --rescan',
   and restart the ehcmd daemon.
   Session: 4060 /dev/sdc 192.x.x.xv -> 192.x.x.xx 00:00:12
   Session: 4061 /dev/sde 192.x.x.zz -> 192.x.x.xx 00:00:12
   Session: 4064 /dev/sdh 192.x.x.zz -> 192.x.x.xy 00:00:12
   Session: 4063 /dev/sdg 192.x.x.xv -> 192.x.x.yy 00:00:12
   Session: 4053 /dev/sdb 192.x.x.zz -> 192.x.x.yy 00:00:12
   Session: 4052 /dev/sdd 192.x.x.xv -> 192.x.x.yz 00:00:12
```

This information indicates that ehcmd continues trying to rebuild the devices using the larger size.

However, if ehcmd has tried repeatedly and failed, you might see ehcmcli diagnostic information like the following:

```
============================================================
Volume list
============================================================
 Volume: user2
   Target name: iqn.2001-05.com:0-8a0906-fde687707-60f1eeafe8a4d260-user2
   Device to mount: /dev/mapper/eql-0-8a0906-fde687707-60f1eeafe8a4d260-user2
   Status: WARNING: Unable to build Device Mapper tables with the correct
   capacity
   Resolution: Use 'iscsiadm -m session --rescan', and restart the ehcmd
   daemon.
   Session: 4060 /dev/sdc 192.x.x.xv -> 192.x.x.xx 00:00:12
   Session: 4061 /dev/sde 192.x.x.zz -> 192.x.x.xx 00:00:12
   Session: 4064 /dev/sdh 192.x.x.zz -> 192.x.x.xy 00:00:12
   Session: 4063 /dev/sdg 192.x.x.xv -> 192.x.x.yy 00:00:12
   Session: 4053 /dev/sdb 192.x.x.zz -> 192.x.x.yy 00:00:12
   Session: 4052 /dev/sdd 192.x.x.xv -> 192.x.x.yz 00:00:12
```

Should this condition occur, try issuing the `iscsiadm -m node -R -T iqn` command (where `iqn` is the full target string). Alternatively, try issuing the `iscsiadm -m session -R` command to rescan all sessions that are currently logged in. If the iscsiadm rescan operation succeeds, restart the ehcmd daemon. If restarting the daemon does not correct the problem, look at the iscsiadm error diagnostics.

Logical Volume Manager (LVM) Considerations

If you are using Logical Volume Manager (LVM) to manage the disk space on your EqualLogic volumes, the following considerations can apply. eqltune will automatically configure LVM to ignore all lower-tier EqualLogic devices. See EqualLogic Host Performance and Tuning Suite (eqltune) for additional details.

LVM Filtering Rules for EqualLogic Devices

The EqualLogic Host Connection Manager daemon (ehcmd) and Auto-Snapshot Manager CLI (ASMCLI) create a number of DM devices, some of which are "visible" and safe for LVM to access and use (for example, user-facing top-tier MPIO devices) and some of which are "hidden" and should not be accessed by LVM (for example, ASMCLI-mounted snapshot Smart Copies and lower-tier MPIO devices).

Verify and Fix LVM Filters with eqltune

To confirm the proper LVM exclusion filter is set, run eqltune, which will review the filters set up in `lvm.conf` and detect a `critical`error for all hidden devices that are not ignored (that is, removed) by the current LVM device filter. Run `eqltune fix` to edit `lvm.conf` and add a rule to ensure all HIT/Linux hidden devices are properly ignored.

Linux Distribution-Specific udev Rule Considerations

SLES11 ships with a udev rules file (`64-lvm2.rules`) that may activate LVM volumes. Because this rule bypasses `lvm.conf` and the special filter that eqltune added, HIT/Linux installs a udev rules file named `64-eqllvm.rules` to prevent `64-lvm2.rules` from activating LVM on EqualLogic hidden devices.The `64-eqllvm.rules` file does not affect any other devices, allowing the SLES `64-lvm.rules` to continue to discover and act upon EqualLogic visible devices and any other non-EqualLogic LVM devices that appear. None of the other udev rules installed by HIT/Linux affects LVM.

Single-path SD devices configured with LVM may be configured to be activated as they appear on the system. The `64-eqllvm.rules` file disables this single-path LVM auto-activation, but only for EqualLogic volumes. Activation will still occur for devices using EqualLogic MPIO or Linux native MPIO, but not for single-path volumes with MPIO disabled. To activate LVM on non-MPIO EqualLogic volumes, run `vgchange -ay vg-name`.

Considerations when Using the Auto-Snapshot Manager CLI

When ASMCLI creates a snapshot of an LVM device, it records the entire LVM configuration as part of the backup document so that it can be recreated at a later time.

When ASMCLI performs a login or mount of a Smart Copy with a backup document containing LVM information, it recreates the stored LVM configuration using hidden devices, so an existing LVM configuration can exist alongside a temporary login or mount. If the devices were not hidden, LVM might improperly detect that the mounted snapshot belongs to another volume group and incorrectly route I/O. When ASMCLI performs a restore operation, it recreates the stored LVM configuration using visible devices. File system I/O and multipathing will function properly, but because LVM no longer owns the configuration, commands such as `lvchange` cannot resize or perform other operations on the logical volumes. If your distribution activates LVM devices, and `lvscan` lists the devices as `active`, no further steps are necessary. If `lvscan` lists the devices as `inactive` and you want to use certain LVM operations, such as resizing or changing logical volumes or volume groups, LVM must reactivate the configuration. To cause LVM to reactivate the configuration:

1. Unmount the file system.
2. Remove the DM device.
   a. To list all devices for a given volume group:

      ```
      # dmsetup ls | grep  volume-group
      ```

   b. To remove a device:

      ```
      # dmsetup remove device-name
      ```

3. Activate the LVM devices:

   ```
   # vgchange -ay volume-group
   ```

4. Remount the file system(s).

Alternatively, you can perform a reboot to return ownership of the devices to LVM.

Logical Volume Manager and Read-Only Media

Before setting EqualLogic volumes read-only in LVM configurations, set any contained Logical Volumes to read-only with the following command:

```
# lvchange -p r volume-group
```

By setting the Logical Volumes read-only, any file systems mounted on the volume will also be mounted as read-only. This change is stored on the metadata LVM stores on the media. You must set LVM volumes read-only before setting EqualLogic volumes read only.

In non-LVM configurations, file system media is detected as read-only when mounted directly on the multipath device.

Logical Volume Manager and Promoted Replica Volumes

When you configure LVM on EqualLogic volumes, enable replication, and then promote a replica volume, Dell recommends that you retain the original target volume name for the promoted replica volume because LVM only recognizes the original volume name and not a new volume name.

Using the eqllog Facility

To gather information for EqualLogic support, you can run the EqualLogic Log Gathering Facility (eqllog) to collect system logs and information about the current running state of your configuration.

The eqllog facility gathers log files, configuration files, and command output, and bundles this information in a compressed tar file for easy storage and transmission. The name of the file is *eqllog-YYYYMMDDHHMMSS.tar.bz2*, where *YYYYMMDDHHMMSS* indicates the date and time eqllog was run. If no `-dir` argument is used to specify a target directory, the file is created in the default `/tmp` directory.

Command Syntax

The command syntax is as follows:

```
eqllog [parameters]
  --version
  --help | -h
  --dir directory | -d
  --output location | -o
  --timeout integer | -t
  --notar | -n
  --no-syslog
  --skip module1 [,module2...] | -s
```

The following table describes the optional eqllog parameters.

Table 42. eqllog Parameters

Parameter	Abbr.	Argument	Description
--version	none	none	Displays the installed version of eqllog.
--help	-h	none	Displays the eqllog command syntax.
--dir	-d	*directory*	Specifies the staging directory for all information. The default is `/tmp`.
--output	-o	*location*	Overrides the default output location. Normally, the output location (directory and tarball) are autogenerated based on the current date and time. You must specify the full path (ignores `--dir`).
--timeout	-t	*integer*	Specifies the maximum time to wait for each command to complete. The default is 30 seconds.
--notar	-n	none	Collects the log information in the staging area but does not create a compressed tarball.
--no-syslog	none	none	Skips syslog collections.
--skip	-s	*module*	Skips the listed module or modules. Available modules are equallogic, system, network, iscsi, block, and cluster.

To run the facility, enter the `eqllog` command alone or with optional arguments. When `eqllog` executes, ignore output errors, because some of the captured information is not relevant for all system configurations. The final line of the command output lists the name of the file that contains your system information.

NOTE: Although `eqllog` does not actively seek to record confidential or personal data, you might want to review the contents of the tar file to ensure it is free of such information.

Example: eqllog Output

The following shows the output of the **eqllog** command without any options.

```
# eqllog

Capturing EqualLogic configuration, logs, and state
------------------------------------------------------------
```

```
Copy directory /var/log/equallogic as logs
Copy directory /etc/equallogic as etc
Capture 'ehcmcli -dv' to 'ehcmcli.dump'
  Return code: 0
Capture 'eqltune -v' to 'eqltune.dump'
  Return code: 1
Capture 'dmsetup status' to 'dmsetup-status.dump'
  Return code: 0
Capture 'dmsetup table' to 'dmsetup-table.dump'
  Return code: 0
Capture 'dkms status' to 'dkms/status.dump'
  Return code: 0
Capturing OS logs, state, and core files
------------------------------------------------------------
Copy file /var/log/messages to var-log
Copy file /var/log/dmesg.old to var-log
Copy file /var/log/messages-20160102 to var-log
Copy file /var/log/dmesg to var-log
Copy file /var/log/messages-20160109 to var-log
Copy file /core.2610
Copy file /core.2759
Copy file /core.2437
Copy file /core.2656
Copy file /core.2653
Copy file /core.2471
Copy file /etc/redhat-release
Copy file /core.2612
Copy file /etc/system-release
Capture 'uname -a' to 'uname.dump'
  Return code: 0
Capture 'lsmod' to 'lsmod.dump'
  Return code: 0
Capture 'lspci' to 'lspci.dump'
  Return code: 0
Capture 'hostname -a' to 'hostname.dump'
  Return code: 0
Capturing Network interfaces and routes
------------------------------------------------------------
Capture 'ip addr' to 'ip-addr.dump'
  Return code: 0
Capture 'ip route' to 'ip-route.dump'
  Return code: 0
Capturing iSCSI configuration and state
------------------------------------------------------------
Copy directory /etc/iscsi as etc
Copy directory /var/lib/iscsi as var-lib
Capture 'iscsiadm --version' to 'iscsiadm.version'
  Return code: 0
Capture 'iscsiadm -m discovery -P1' to 'discovery.dump'
  Return code: 0
Capture 'iscsiadm -m node -P1' to 'node.dump'
  Return code: 0
Capture 'iscsiadm -m iface -P1' to 'iface.dump
  Return code: 0
Capture 'iscsiadm -m session -P3' to 'session.dump'
  Return code: 0
Compressing tarfile, please wait...
Eqllog complete. The data has been collected in:
  /tmp/eqllog-20160112140941.tar.bz2
```

Using the udev Facility

From the udev facility, Linux system administrators can specify attributes on devices. Linux systems typically have a `udevd` process constantly running. The `udevd` process is notified whenever a device is created or removed by the Linux kernel. It consults its set of rules, finds the appropriate rules for the given device (if any), and then returns with instructions to the kernel concerning the device. For more information, enter `man udev` on a Linux system.

HIT/Linux automatically installs a set of udev rule files on your system.These files are installed in `/lib/udev/rules.d/`. Caution: Do not alter these files in any way.

All distributions receive the following rule files:

- `11-eqlenv.rules`

  This udev rule file sets up important environment variables and symlinks used by ehcmd and in other udev rule files.
- `38-equallogic.rules`

  These udev rules should serve as a guide for setting up systemwide udev rules for EqualLogic storage devices, and demonstrate how this setup can be done by setting the readahead value for all EqualLogic devices to 1024kb. You can copy this file to `/etc/udev/rules.d/39-eql_user.rules` for customization.
- `99-eqlsd.rules`

  Set the SCSI subsystem device timeout to 60 seconds for every Dell EqualLogic iSCSI session. This rule file applies only to the individual SD devices. Any other block-layer tuning done here will be bypassed by the top-level DM devices built by ehcmd. To tune those devices, see the information for the `38-equallogic.rules` file.

A number of distribution-specific files might be required to work with other distribution-specific udev rules, such as `64-eqllvm.rules` or `89-eqlfinal.rules`. These files are identified by a `-eql` prefix (`??-eql*.rules`), and contain internal comments that explain their function and purpose.

Using the ehcmcli Diagnostic Tool

The ehcmcli diagnostic tool reports the diagnostic state of the EqualLogic Host Connection Manager (ehcmd) and multipath device information. When entering ehcmcli commands, Bash command completion lets you type part of a command followed by a single horizontal tab, which completes the command. For more information, see Appendix C Bash Command Completion.

ehcmcli Command Syntax

The command syntax is as follows:

```
ehcmcli [subcommand-1] [subcommand-2] [options] [argument]

help [subcommand]
dm create
dm cleanup
dm unload switch
login --target target | -T
     [--portal IP-address-or-name | -p]
     [--hidden]
     [--login-at-boot | --no-login-at-boot]
logout --target target | -T]
     [--no-login-at-boot]
pr init --key key-name
pr remove --key key-name
pr status [--verbose]
status [--verbose]
```

The following table describes the commands and options for the ehcmcli tool.

Table 43. ehcmcli Commands and Options

Subcommands	Options	Argument	Description
Required			
login	--target | -T --portal | -p --hidden --login-at-boot --no-login-at-boot	*target IP address* or *name*	Specifies `--target` to log in to one session to the specified PS Series group. *Optional.* Use `--portal` to provide the iSCSI portal IP address or name to log in if `target` is not already discovered on the system. *Optional.* Use `--hidden` to treat the login as a hidden ASM device. *Optional.* Use either `--login-at-boot` to set the `node.startup` values for the target to perform the login at the next reboot, or `--no-login-at-boot` to set the values to not perform the login at the next reboot.

Subcommands	Options	Argument	Description
Optional			
help	none	*subcommand*	Displays a list of the ehcmcli commands and syntax. Specify a subcommand to list the syntax (for example, `ehcmcli help pr`).
dm create	none	none	Builds the initial DM devices for all current EqualLogic devices.
dm cleanup	none	none	Removes any unmanaged DM devices (`/dev/mapper` devices created by ehcmd) that are no longer needed.
dm unload_switch	none	none	Replaces all `switch` DM targets and unloads the `dm-switch` kernel module.
logout	--target --no-login-at-boot	none	Logs out and removes the given device target. Specify `--no-login-at-boot` to set the `node.startup` values for the given target to not perform the login at the next reboot.
status	--verbose	none	Instructs ehcmcli to gather diagnostic data and report the current multipathing status. With `--verbose`, enables verbose output.

The following table lists the optional persistent reservation (PR) management commands.

Table 44. Persistent Reservation (PR) Management Commands

Subcommands	Options	Argument	Description
pr init	--key	*key-name*	Manages the given PR key on all volumes that have the key registered on at least one path.
pr remove	--key	*key-name*	Removes the given PR key from all volumes.
pr status	none	none	Lists persistent status reservation information.

The following table lists the commands related to enabling the capability to import EqualLogic storage to PowerStore.

Table 45. Enabling Import Capability to PowerStore Commands

Subcommands	Options	Argument	Description
log_show	none	none	Displays Log Level and Log Location.
server_port	--change	*server_port no*	Displays server port that is used for communication and gives option to change to desired port.
mhd_status	none	none	Displays microhttpd status.
restart_webServer	none	none	Restarts web server in case of microhhtpd failure.
recover_device	--device	*source_wwid*	Recovers specified device from hung state.

ehcmcli Diagnostic Output

The ehcmcli diagnostic tool provides an option to gather and display diagnostic information about the status of MPIO sessions. The output includes a list of the adapters for iSCSI sessions and a list of all volumes for which multipathing is configured by ehcmd. If the list does not contain all the adapters you expect, some might be excluded from MPIO usage. Common causes for adapters to be excluded are subnet, adapter speed, or IP protocol version. Reporting verbose diagnostic information (`ehcmcli status --verbose`) will include the list of excluded adapters and the reason for exclusion. To modify the adapter exclusion parameters, use `rswcli -P` (see Setting MPIO Parameters).

The ehcmcli diagnostic output includes a report of all the volumes the ehcmd daemon is managing and a list of iSCSI sessions to each volume. The output for each iSCSI session contains the session ID used by iscsiadm, the SD device, the source and target IP addresses, and the uptime. For example, the following output shows the status of the `linux-hyb-mpio5` volume:

```
Volume: linux-hyb-mpio5
Target name: iqn.2001-05.com:0-8a0906-19d376406-b06003614cc4c053-linux-hyb-mpio5
Device to mount: /dev/eql/linux-hyb-mpio5
Status: Normal
Session: 195 /dev/sdv 192.x.x.xx -> 192.x.x.vv 00:00:09
Session: 194 /dev/sdu 192.x.x.xy -> 192.x.x.vx 00:00:09
Session: 193 /dev/sdt 192.x.x.xz -> 192.x.x.wv 00:00:09
Session: 173 /dev/sdi 192.x.x.zz -> 192.x.x.wx 00:00:09
```

If the volume has a problem, an explanatory message will display in the status field. The possible warnings include:

- The ehcmd daemon is unable to create the desired number of iSCSI sessions due to iSCSI login errors. If this problem occurs, verify that all Ethernet ports are properly configured on the host and that the access control record on the volume allows iSCSI logins from all Ethernet ports.
- The daemon cannot communicate with the `dm-switch` module. If this problem occurs, see Assuring Correct Operation of the dm-switch Kernel Module.

If `ehcmcli status` returns the following message, verify that all of the following conditions have been addressed:

```
Could not communicate with ehcmd (the process may not
be running): 111 (Connection refused)
```

- All `critical` issues reported by `eqltune -v` have been resolved.
- The required services are running; otherwise, start the service:

```
# service iscsid start
# service ehcmd start
```

- The required services are set to auto-start:

```
# chkconfig --list iscsid
# chkconfig --list ehcmd
```

Running the ehcmd Daemon

The EqualLogic Host Connection Manager (ehcmd) daemon runs in the background to manage multipath devices and create additional iSCSI sessions based on the SAN configuration. It manages the iSCSI session by downloading the page table from the group, delivering the page table to the kernel module, and then building and maintaining the device mapper tables. Use the `--foreground` command to run ehcmd in the foreground.

The `ehcmd` command syntax is as follows:

```
ehcmd command [parameters]
  -h | --help
  -f | --foreground [-q | --quiet]
```

The following table describes the optional commands and parameters for the ehcmd daemon.

Table 46. ehcmd Optional Command Parameters

Command	Abbr.	Parameters	Description
--help	-h	none	Displays command help.
--foreground	-f	--quiet | -q	Runs the daemon in the foreground (`-f`). Directs trace messages to a log file (`-q`).

Considerations When Partitioning iSCSI Volumes

As a general rule, partitioning an iSCSI volume adds complexity and overhead to the I/O path. For example, you could create one iSCSI volume and partition it into three pieces, or create three separate volumes. When you need to create partitions, HIT/Linux accommodates partitioned environments.

By partitioning a volume so the starting relative block number on a partition is an odd number, the Linux I/O subsystem often selects a small block size for its I/O operations. When you partition a disk, use a multiple of 8 sectors to ensure that Linux uses a large block size for I/O.

If you create, delete, or modify the partition table on a volume, use a partition management utility, such as fdisk, against the top-tier DM device.

```
/dev/eql/volume_namep#
```

where *volume_name* is the name specified at volume creation and *p#* is the partition number (for example, `p1`, `p2`, and so on.).

When the partition table is written, the partition management utility displays a warning similar to the following fdisk message:

```
WARNING: Re-reading the partition table failed with error
22: Invalid argument.
The kernel still uses the old table.
The new table will be used at the next reboot.
```

The partition management utility tries to use an `ioctl()` function to request that the Linux kernel re-read the partition information. The operation is not allowed because the device characteristics are incompatible with that operation.

After changing the partition table on an EqualLogic storage volume device, log out and log in to that volume to refresh the partition information. For example, if you alter the partition table on a volume named testvolume, unmount any file systems on any of the old or existing partitions and run the following commands:

```
ehcmcli logout --target testvolume
ehcmcli login --target testvolume
```

NOTE: If you alter the partition table, do not run `partprobe` or `kpartx` without specifying a disk device.

If you want to suppress creating partition devices on a volume, see Setting Volume-Specific Configuration Parameters.

A

Migrating to HIT Linux Multipath I/O

The following sections describe how to migrate an existing Linux system connected to Dell EqualLogic PS Series storage to use HIT/Linux multipath I/O (MPIO). If you are installing HIT/Linux in an environment without previously configured storage, follow the basic installation steps in Installing and Configuring HIT/Linux.

NOTE: This migration is not required if you are updating from a previous version of HIT/Linux.

Topics:

Before You Begin

For Logical Volume Manager (LVM)-based configurations, including Clustered Logical Volume Manager (CLVM), live migration can occur without any interruption of service. However, as best practice, remove the node from the cluster before migrating to HIT/Linux. See Enabling Storage Cluster Support for more information.

For other services, such as directly mounted file systems, some service interruption is inevitable.

Be aware of limitations associated with SAN Boot Volumes. If you have questions about migrating these volumes, contact Dell Support for assistance.

During the migration procedure, do not manually add or remove iSCSI sessions using the iscsiadm utility.

When you have completed the procedure, see Accessing Volumes Using Multipath I/O for information about managing iSCSI logins and logouts after HIT/Linux is running.

Step 1: Install HIT Linux in Migration Mode

The HIT/Linux installation normally starts the EqualLogic Host Connection Manager daemon (ehcmd), which immediately begins actively managing your iSCSI devices and block devices. This behavior might interfere with current users of EqualLogic storage volumes because the iSCSI sessions could be terminated without warning. By running the installation script in Migrate Mode, the system provides a window of time in which you can manually migrate existing storage users to the new EqualLogic-managed devices without any unexpected iSCSI session management occurring.

After completing the installation Pprerequisites, enter the following command line to install HIT/Linux in migration mode:

```
# /mnt/iso/install --migrate
```

If you do not use the `--migrate` command option and existing EqualLogic storage is detected on your system, the install process might ask you the following question:

```
Existing EqualLogic devices were found on your system.

Enabling migrate mode can help you migrate your existing devices
to using actively managed MPIO.

Would you like to migrate your existing devices (Yes/No) [Yes]?
```

Answering Yes runs the installation as if you had specified the `--migrate` command option.

When the installation completes, configure the software as in Configuring HIT/Linux.

Step 2: Migrate Existing Storage-Using Applications

After the installation script has completed collecting your configuration responses, it displays the following output:

```
========================================================================
MPIO Device Migration
========================================================================

Starting service: ehcmd
Starting service: scsi_reserve_eql

Creating initial DM devices

Pausing now to allow for manual block device migration.

In another shell, check the output of 'ehcmcli status' for the device
names that have been created, and migrate all services and applications
to use the new devices before proceeding.

Consult the user guide for further instructions.

Are you ready to continue and enable active MPIO and iSCSI session
management (Yes/No) [No]?
```

Switch to a new shell instance and enter the following commands to list all currently detected EqualLogic storage volumes:

```
# ehcmcli status
```

The `Device to mount:` line lists the newly created top-tier device node for each volume. For example:

```
Volume: replicatest
 Target name: iqn.2001-05.com.equallogic:0-8a0906-b30e26209-563ccc969074edcc-
   volname
 Device to mount: /dev/eql/volname
 Status: Normal: MPIO is disabled
```

In the example, you would use `/dev/eql/volname` as the block device to mount.

Logical Volume Manager Considerations

If you use Logical Volume Manager (LVM) to manage the disk space on your EqualLogic volumes, ensure that LVM points to the top-tier multipath device with the following steps:

1. As part of the installation, eqltune added filtering rules to your LVM configuration file (`/etc/lvm/lvm.conf`) to exclude all EqualLogic lower-tier devices on most simple setups. If your system uses a more complicated filtering setup, you might need to manually ensure that all devices under `/dev/ignore_eql/*` are excluded from consideration by LVM. Run `pvscan` to verify the filter rules and ensure proper device names are shown for each physical volume.

   If duplicate devices are detected, add rules to `/etc/lvm/lvm.conf` to remove duplicates and ensure that only the MPIO devices of the format `/dev/eql/`*`volume-name`* are shown.

   NOTE: If you were using Linux native MPIO (`multipathd`) and LVM continues to use the mpath devices created by `multipathd`, add the following rule to your `lvm.conf` filter to exclude all multipathd devices: `r|/dev/mpath/.*|` Do not add this rule if the device is a boot volume, as these filters may not work properly.

2. After `pvscan` shows the correct devices, run `vgchange --refresh` to repoint all active logical volumes to the new EqualLogic devices without service interruption.

Directly Mounted File Systems

If you are using directly mounted file systems instead of LVM, you must manually unmount and remount them using the top-tier multipath device names. This action will temporarily interrupt any users (including applications and services) of these file systems.

1. Unmount the file systems associated with the EqualLogic volumes.

2. If you are using a single path (using the SD device directly without LVM or multipathd), you must manually unmount the file system and run `ehcmcli login --target` *target* to recreate the EqualLogic DM devices after unmounting the file system. You can also run `ehcmcli dm create` to recreate all DM devices instead of `ehcmcli login`, which only affects the devices for a single volume.
3. Remount the file systems using the new top-tier device names (`/dev/eql/volume-name`).
4. Using the new top-tier device names, edit `/etc/fstab` to mount the file systems on the next reboot.

Block Device Access

For applications or services that use block devices directly, consult the documentation for those applications or services, then reconfigure them to use the appropriate new top-tier devices (`/dev/eql/volume-name`). This reconfiguration might require restarting or interrupting the application or service and running `ehcmcli login --target target` or `ehcmcli dm create` to refresh the EqualLogic MPIO DM devices before restarting the application or service. For more details, see step 3 in the procedure in Directly Mounted File Systems.

Excluding Specific Volumes

If a block device cannot be migrated to use EqualLogic MPIO (such as a SAN boot volume), add the following section to the end of `eql.conf` to exclude the device from MPIO:

```
[MPIO Volume Params volume-name]
EnableMPIO = false
```

For more details, see Setting Volume-Specific Configuration Parameters.

Step 3: Reconfigure multipathd

The Linux multipath tools provide facilities to create and manage multiple connections from the Linux host to the target group member. These connections are combined on the Linux host in the form of a Device Mapper (DM) device that provides fault-tolerance and improved performance.

NOTE: If you are using the Linux multipathing tools, HIT/Linux provides automated facilities that without the multipathd service create and maintain the necessary multipath devices for connections to volumes on PS series group. Reconfigure it so that it does not interfere with HIT/Linux.

The multipathd daemon is required to manage PowerStore devices. Therefore, if you want to import existing EqualLogic storage to PowerStore, make sure that the multipathd service is running. For more information about importing to PowerStore, see the *Importing External Storage to PowerStore Guide*.

You must blacklist all EqualLogic devices in the multipath configuration file (`multipath.conf`). As part of the installation, `eqltune fix` introduces blacklist entries that should correctly blacklist all EqualLogic devices.

Run `multipath -r` to ensure multipathd has completely stopped managing all EqualLogic devices before proceeding to the next step.

NOTE:

By default, SLES 12 distributions do not contain the multipath.conf file in `/etc`. Copy the file from the `/usr/share/doc/packages/multipath-tools/` directory using the following command:

```
$cp /usr/share/doc/packages/multipath-tools/multipath.conf.synthetic /etc/multipath.conf
```

If you have already installed HIT/Linux, run `eqltune fix` to blacklist the EqualLogic devices. If you have already created volumes, log out of the volumes and then log back in. (This action is not required if you have not yet installed HIT/Linux.)

Step 4: Complete the Installation

After configuring all software to use EqualLogic top-tier multipath devices, answer `Yes` to the following question from the installation script to begin active management of the MPIO and iSCSI sessions:

```
Are you ready to continue and enable active MPIO and
iSCSI session management (Yes/No) [No]?
```

Any application using an EqualLogic top-tier device will not require any further intervention.

Disabling Active Management

If you answered `No` in step 4, active management remains disabled until you explicitly enable it. To enable active management at this point, run `eqlconfig --mpio` and accept the default to enable:

```
==============================================================================
Active Migration in Progress
==============================================================================

You can choose to end migration and enable active MPIO and iSCSI
session management, abort active migration and leave MPIO disabled,
or just quit eqlconfig immediately, making no changes.

Consult the user guide for further instructions.
```

```
Choose action (Enable/Abort/Quit) [Enable]:
```

- If you choose the default action (`Enable`), active management of MPIO and iSCSI sessions will begin.
- If you choose `Abort`, MPIO will be disabled and active migration will be canceled.
- If you choose `Quit`, no change will be made. Running `eqlconfig --mpio` will present this same dialog again.

Enabling Storage Cluster Support

To enable storage cluster support, HIT/Linux must be compatible with the High Availability Add-On for RHEL 6 and 7 and the SLES 12 for storage clusters. There is not much interaction between the clustering software and HIT/Linux.

Cluster Logical Volume Manager

The Cluster Logical Volume Manager (CLVM), part of the Resilient Storage Add-On, is supported. If you followed all LVM setup instructions, you do not need to take special action when performing a fresh installation or provisioning new storage.

To migrate an existing CLVM node to start using HIT/Linux, follow the migration instructions for LVM in Logical Volume Manager Considerations.

NOTE: Although not strictly required, Dell recommends temporarily fencing the node during the migration process to reduce the possibility of affecting the running cluster on other nodes.

B

Enabling iSCSI Offload on HIT/Linux

Before enabling iSCSI offload functionality, make sure that HIT/Linux is installed on your system and that ehcmd is running and using the (default) TCP initiator.

Follow the steps in this chapter to enable Broadcom iSCSI offload functionality for HIT/Linux.

Topics:

Step 1: Set Up and Verify Your Configuration

For onboard network interfaces, you might need to use a physical key to enable the functionality. For expansion cards, verify that the functionality is enabled by firmware.

If the Broadcom iSCSI Offload Engine (ISOE) functionality is not enabled, messages similar to the following will appear in the `/var/log/messages` file:

```
Jan 14 17:38:01 rh6-910 kernel: Broadcom NetXtreme II
iSCSI Driver bnx2i v2.1.1 (Mar 24, 2017)
Jan 14 17:38:01 rh6-910 kernel: iscsi: registered transport (bnx2i)
Jan 14 17:38:01 rh6-910 kernel: bnx2i: iSCSI not supported, dev=eth3
Jan 14 17:38:03 rh6-910 kernel: bnx2i: iSCSI not supported, dev=eth2
Jan 14 17:38:05 rh6-910 kernel: bnx2i: iSCSI not supported, dev=eth1
Jan 14 17:38:07 rh6-910 kernel: bnx2i: iSCSI not supported, dev=eth0
```

When enabled, messages similar to the following will display:

```
Jan 14 18:13:55 rh6-910 kernel: Broadcom NetXtreme II
iSCSI Driver bnx2i v2.1.1 (Mar 24, 2017)
Jan 14 18:13:55 rh6-910 kernel: iscsi: registered transport (bnx2i)
Jan 14 18:13:55 rh6-910 kernel: bnx2: eth0: using MSIX
Jan 14 18:13:55 rh6-910 kernel: bnx2i [01:00.00]: ISCSI_INIT passed
```

Step 2: Configure the Network Interfaces

Network interfaces are created by **iscsiadm** for each Broadcom iSCSI interface. For example:

```
# iscsiadm -m iface
default tcp,<empty>,<empty>,<empty>,<empty>
iser iser,<empty>,<empty>,<empty>,<empty>
bnx2i.00:00:00:00:00:00 bnx2i,00:00:00:00:00:00,<empty>,<empty>,<empty>
bnx2i.a4:ba:db:38:8b:42 bnx2i,a4:ba:db:38:8b:42,<empty>,<empty>,<empty>
bnx2i.a4:ba:db:38:8b:40 bnx2i,a4:ba:db:38:8b:40,<empty>,<empty>,<empty
```

Broadcom iSCSI offload engines have two MAC addresses per physical network interface: one for TCP-IP and one for ISOE. In the following examples, the system was configured with two Broadcom network interfaces whose TCP-IP MAC addresses were listed by `ifconfig`.

```
#ifconfig
eth2 Link encap:Ethernet HWaddr A4:BA:DB:38:8B:3F
inet addr:x.x.x.x Bcast:x.x.x.z Mask:255.255.254.0
```

```
eth3 Link encap:Ethernet HWaddr A4:BA:DB:38:8B:41
inet addr:x.x.x.y Bcast:x.x.x.z Mask:255.255.254.0
```

To determine the correct ISOE MAC address:

1. Add 1 to the hexadecimal TCP-IP MAC address reported by `ifconfig`.
2. Update the appropriate interface (iface) records in iscsiadm to associate each ISOE MAC address with a new, unique IP address.

For example:

```
# iscsiadm -m iface -I bnx2i.a4:ba:db:38:8b:40 -o update -n iface.ipaddress
   -v 1 xx.xx.x.x
bnx2i.a4:ba:db:38:8b:40 updated.

# iscsiadm -m iface -I bnx2i.a4:ba:db:38:8b:42 -o update -n iface.ipaddress
  -v xx.xx.x.y
bnx2i.a4:ba:db:38:8b:42 updated.
```

After updating the iface records, you should see output similar to the following display:

```
# iscsiadm -m iface
default tcp,<empty>,<empty>,<empty>,<empty>
iser iser,<empty>,<empty>,<empty>,<empty>
bnx2i.a4:ba:db:38:8b:42 bnx2i,a4:ba:db:38:8b:42,1.2.3.5,<empty>,<empty>
bnx2i.a4:ba:db:38:8b:40 bnx2i,a4:ba:db:38:8b:40,1.2.3.4,<empty>,<empty>
```

Step 3: Configure ehcmd to Use the Offload iSCSI Initiator

By default, ehcmd uses the software initiator to connect to iSCSI targets. To specify using the Broadcom iSCSI offload initiator, edit the `/etc/equallogic/eql.conf` file and set the `IscsiInitiator` parameter:

```
IscsiInitiator = tcp
```

to:

```
IscsiInitiator = bnx2i
```

Restart the ehcmd daemon with the following command:

```
# service ehcmd restart
```

In version 2.0-872 and later of the Linux iSCSI initiator package, the iscsid service can be started on an as-needed basis. If iscsid is not running (that is, typically, because there are no currently active iSCSI sessions), ehcmd will not recognize any bnx2i adapters, even if they are properly configured with an IP address. To manually start the iscsid service, you can either log in to at least one iSCSI session with the `iscsiadm` command, or you can run the following command:

```
 $ service iscsid force-start
```

Step 4: Discover iSCSI Targets

After configuring the new interfaces and restarting ehcmd, rediscover the targets on the Broadcom interfaces. The iSCSI initiator initiates discovery on all interfaces using the TCP transport by default, so in this scenario it is necessary to explicitly list the bnx2i interfaces:

```
# iscsiadm -m discovery -t st -p group-IP-address
-I iface -I iface
```

For example:

```
# iscsiadm -m discover -t st -p xxx.x.x.xx -I bnx2i.a4:ba:db:38:8b:40
-I bnx2i.a4:ba:db:38:8b:42
```

After rediscovering the targets, you will see that the isciadm interface records have been updated by ehcmd. For example:

```
# iscsiadm -m iface
default tcp,<empty>,<empty>,<empty>,<empty>
iser iser,<empty>,<empty>,<empty>,<empty>

bnx2i.a4:ba:db:38:8b:40 bnx2i,a4:ba:db:38:8b:40,1.2.3.4,<empty>,<empty>
bnx2i.a4:ba:db:38:8b:42 bnx2i,a4:ba:db:38:8b:42,1.2.3.5,<empty>,<empty>
```

Step 5: Logging In to Additional Targets

If you do not explicitly provide an interface, the iSCSI initiator attempts to log in through all interfaces using the TCP transport. Because you are not using the TCP transport in this scenario, it is necessary to explicitly list one of the bnx2i interfaces:

```
# iscsiadm -m node -T target-name -I iface1
```

It is sufficient to specify a single iface to create a single session and let ehcmd create additional sessions as necessary.

C

Bash Command Completion

Command-line interfaces in HIT/Linux are integrated with a Bash programmable command completion facility. The Bash command-completion script, installed as `/etc/bash_completion.d/equallogic`, integrates with the following CLIs:

- Remote Setup Wizard Command Line Interface (rswcli)
- Auto-Snapshot ManagerCommand Line Interface (asmcli)
- EqualLogic Host Connection Manager Command Line Interface (ehcmcli)
- EqualLogic Configuration Utility (eqlconfig)
- EqualLogic Host Performance and Configuration Tuning Suite (eqltune)
- EqualLogic Log Gathering Facility (eqllog)

Using Bash command completion, you can type part of a command-line element, such as a command, and then enter a single tab. If your input is sufficiently unique, Bash will complete the element. If the input is not sufficient for Bash to complete the entire element, it completes the element up to the point where it needs more input to make it unique. In this case, enter a second tab and Bash will respond with a list of valid inputs.

For example, for ASMCLI Bash completes the command-verb, command-noun, and any options (see General Command Syntax for ASMCLI):

```
# asmcli <TAB><TAB>
```

Bash provides a valid list of command-verbs:

```
# asmcli
create   delete   do   help  list  login  logout   mount  restore  unmount  update
```

If you enter `asmcli c` followed by a tab, Bash completes the command-verb:

```
#asmcli create
```

Now, when you enter two tabs, Bash lists valid command-nouns for the `create` command-verb:

```
#asmcli
collection  create    password     schedule     smart-copy
```

If you enter `s` followed by a tab, Bash shows two options:

```
# asmcli create s<TAB>
schedule  smart-copy
```

To be sufficiently unique, entering `sm` followed by a tab will pertain only to the `smart-copy` option. In this case, Bash completes the entire command:

```
# asmcli create smart-copy
```

If you enter another tab, Bash responds with two hyphens, indicating it expects a parameter:

```
# asmcli create smart-copy --
```

Enter two more tabs and Bash provides a list of acceptable parameters:

```
# asmcli create smart-copy --
--collection --destination --freeze-time --output --retry-count --source
--user-name --comment  --force  --group-access-file --read-only --retry-wait
--type --verbose
```

If you enter `s` followed by a tab, Bash completes the command with the `source` parameter:

```
# asmcli create smart-copy --s<TAB>
# asmcli create smart-copy --source
```

As a general rule, type as much as you think will enable Bash to complete the command. If you are unsure of the options, enter a command, followed by two tabs, and then select from the options.

D

DKMS Considerations for Red Hat Linux

Dynamic Kernel Module Support (DKMS) is required on systems where a precompiled binary kernel module is not available, such as on a system running a custom Linux kernel.

The HIT/Linux installation script will attempt to install all required Linux packages. However, DKMS is not available in the default Red Hat repositories; extra steps are required to obtain and install DKMS. All other prerequisite packages are installed.

Note the following considerations:

- If your system is configured to use the Extra Packages for Enterprise Linux (epel) repository. DKMS is automatically installed by the HIT/Linux installation script. No further action is needed.
- To install DKMS manually, download the package from https://github.com/dell/dkms. Dell recommends using version 2.2.0.1 or later.
- To verify if your system is using the correct DKMS version or the prebuilt binary version, enter the following command:

```
# rpm -q equallogic-host-tools-kernel
```

- If the result shows that the equallogic-host-tools-kernel package is installed, then you are using a DKMS-based installation.
- If DKMS is not installed, entering the `ehcmcli status` command will produce the following error message:

```
ERROR: dm-switch kernel module is not present Resolution:
Ensure the kernel module package appropriate for your system is
installed, and use 'modprobe -v dm-switch' to verify the module can
be loaded.
```

- If DKMS is loaded but contains incompatible module versions (for example, the HIT/Linux `dm-switch` module is still resident in the kernel), entering the `ehcmcli status` command will produce the following error message:

```
ERROR: Incompatible version of dm-switch kernel module
detected Resolution: Ensure the kernel module package appropriate
for your system is up to date, and that any conflicting versions
of the 'dm-switch' module are removed.
```

DKMS requires several other packages: `kernel-devel`, `glibc-headers`, `glibc-devel`, `libgomp`, and `gcc`. The installation script automatically installs these packages. No further action is required.

Index

F

G

H

I

K

L

M

N

P

R

S

T

U

V