

SnapDrive™ for UNIX® Installation and Administration Guide (IBM® AIX®, HP-UX®, Linux®, Solaris®)

Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089 USA
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 4-NETAPP
Documentation comments: doccomments@netapp.com
Information Web: <http://www.netapp.com>

Part number 210-00699_D0
July 2005

Copyright and trademark information

Copyright information

Copyright © 1994—2005 Network Appliance, Inc. All rights reserved. Printed in the U.S.A.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted Network Appliance material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETWORK APPLIANCE “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL NETWORK APPLIANCE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Network Appliance reserves the right to change any products described herein at any time, and without notice. Network Appliance assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Network Appliance. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of Network Appliance.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the Network Appliance logo, the bolt design, NetApp—the Network Appliance Company, DataFabric, FAServer, FilerView, MultiStore, NearStore, NetCache, SecureShare, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapVault, SyncMirror, and WAFL are registered trademarks of Network Appliance, Inc. in the United States, and/or other countries. Data ONTAP, gFiler, Network Appliance, SnapCopy, Snapshot, and The Evolution of Storage are trademarks of Network Appliance, Inc. in the United States and/or other countries and registered trademarks in some other countries. ApplianceWatch, BareMetal, Camera-to-Viewer, ComplianceClock, ComplianceJournal, ContentDirector, ContentFabric, EdgeFiler, FlexClone, FlexVol, FPolicy, HyperSAN, InfoFabric, LockVault, Manage ONTAP, NOW, NetApp on the Web, ONTAPI, RAID-DP, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, SharedStorage, Simulate ONTAP, Smart SAN, SnapCache, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapMigrator, SnapSuite, SnapValidator, SohoFiler, vFiler, VFM, Virtual File Manager, VPolicy, and Web Filer are trademarks of Network Appliance, Inc. in the United States and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of Network Appliance, Inc. in the United States. Spinnaker Networks, the Spinnaker Networks logo, SpinAccess, SpinCluster, SpinFS, SpinHA,

SpinMove, and SpinServer are registered trademarks of Spinnaker Networks, LLC in the United States and/or other countries. SpinAV, SpinManager, SpinMirror, SpinRestore, SpinShot, and SpinStor are trademarks of Spinnaker Networks, LLC in the United States and/or other countries.

Apple is a registered trademark and QuickTime is a trademark of Apple Computer, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Network Appliance is a licensee of the CompactFlash and CF Logo trademarks.

Network Appliance NetCache is certified RealSystem compatible.

Table of Contents

	Preface	ix
Chapter 1	Overview of SnapDrive for UNIX	1
	How SnapDrive for UNIX works.	2
	Comparing SnapDrive for UNIX on different host platforms	9
	Where to go for more information	11
Chapter 2	Installing SnapDrive for UNIX.	13
	Verifying that your storage systems are ready	14
	Installation instructions for all hosts	19
	Installing and uninstalling SnapDrive for UNIX on an AIX host	22
	Installing and uninstalling SnapDrive for UNIX on an HP-UX host.	38
	Installing and uninstalling SnapDrive for UNIX on a Linux host	45
	Installing and uninstalling SnapDrive for UNIX on a Solaris host.	50
	Completing the installation	54
	Files installed by SnapDrive for UNIX on the host	55
	Upgrading to a new version of SnapDrive for UNIX	59
Chapter 3	Configuring and Using SnapDrive for UNIX	61
	Setting configuration information	62
	Preparing hosts for adding LUNs.	85
	Setting up audit, recovery, and trace logging.	86
	Setting up autosupport	94
	General steps for executing commands.	97
Chapter 4	Setting Up Security Features	99
	Overview of SnapDrive for UNIX's security features	100
	Setting up access control	101

View the current access control settings	.105
Specifying the current login information for storage systems	.107
Enabling SSL encryption	.111

Chapter 5	Creating and Using Snapshots	.113
	Creating snapshots	.114
	Displaying information about snapshots	.120
	Renaming a snapshot	.127
	Restoring a snapshot	.129
	Connecting to a snapshot	.137
	Disconnecting a snapshot.	.144
	Deleting snapshots	.146

Chapter 6	Provisioning and Managing Storage	.149
	Overview of storage provisioning	.150
	Creating storage.	.152
	Displaying information about storage	.160
	Increasing the size of the storage	.166
	Connecting existing LUNs to the host	.170
	Disconnecting LUN mappings from the host.	.172
	Connecting only the host-side of storage.	.176
	Disconnecting only the host-side of storage	.177
	Deleting storage from the host and storage system.	.179

Chapter 7	Troubleshooting	.183
	Data collection utility.	.184
	General troubleshooting tips for SnapDrive for UNIX	.187
	Understanding error messages	.190
	Common error conditions.	.192
	Standard exit status values	.194

Chapter 8	Checklists for Using SnapDrive Commands.	.207
	Collecting information needed by SnapDrive for UNIX Commands	.208
	Checklist for using the snapshot create command	.209
	Checklist for using the snapshot rename command	.211
	Checklist for using the snapshot show command	.213
	Checklist for using the snapshot restore command.	.215
	Checklist for using the snapshot delete command	.217
	Checklist for using the snapshot connect command	.219
	Checklist for using the snapshot disconnect command.	.223
	Checklist for using the storage create command	.225
	Checklist for using the storage resize command	.229
	Checklist for using the storage show command	.231
	Checklist for using the storage connect command	.234
	Checklist for using the storage disconnect command	.235
	Checklist for using the host connect command.	.237
	Checklist for using the host disconnect command	.238
	Checklist for using the storage delete command	.240
 Chapter 9	 Command Reference	 .243
	Summary of the SnapDrive for UNIX commands	.244
	SnapDrive for UNIX keywords and arguments	.247
	 Glossary	 .259
	 Index	 .263

Preface

About this guide

This document describes how to install, configure, and operate SnapDrive™ software for UNIX® servers. It does not cover basic system or network administration topics, such as IP addressing, routing, and network topology. It also does not cover topics that are handled in the Network Appliance™ FCP Host Attach Kit or iSCSI Support Kit documentation. For example, this document does not include information on setting up multipathing.

Audience

This guide is for system administrators who possess working knowledge of NetApp® storage systems, such as filers.

This guide assumes that you are familiar with the following topics:

- ◆ Fibre Channel Protocol (FCP)
- ◆ iSCSI Protocol
- ◆ Basic network functions and operations
- ◆ UNIX servers
- ◆ UNIX security
- ◆ Data storage array administration concepts
- ◆ NetApp storage system management
- ◆ Logical volume manager on the system you are using

Command conventions

You can enter filer commands on the system console or from any client that can obtain access to the filer using a Telnet session. In examples that illustrate commands executed on a UNIX® workstation, the command syntax and output might differ, depending on your version of UNIX. The examples might be from any of the hosts supported for SnapDrive for UNIX.

Terminology

This guide uses the following terms:

- ◆ *LUN* (Logical Unit Number) refers to a logical unit of storage identified by a number.
- ◆ *LUN ID* refers to the numerical identifier for a LUN.

- ◆ *Storage system* refers to a NetApp® storage system that incorporates the CompactFlash™ unit. Storage systems can support FCP and iSCSI.

Note

NetApp's storage products (filers, FAS appliances, NearStore® systems, and storage appliances) are all storage systems. All these products are also sometimes called filers or storage appliances.

- ◆ Some UNIX operating systems use the term disk group while others use the term volume group. This guide uses the term *disk group* to refer to both disk groups and volume groups.
- ◆ Some operating systems refer to host volumes while others refer to logical volumes. This guide uses the term host volume to refer to both host volumes and logical volumes.

Formatting conventions

The following table lists different character formats used in this guide to set off special information.

Formatting convention	Type of information
<i>Italic type</i>	◆ Words or characters that require special attention. ◆ Placeholders for information you must supply. For example, if the guide requires you to enter the <code>fctest <i>adaptername</i></code> command, you enter the characters “fctest” followed by the actual name of the adapter. ◆ Book titles in cross-references.
Monospaced font	◆ Command and daemon names. ◆ Information displayed on the system console or other computer monitors. ◆ The contents of files.
Bold monospaced font	Words or characters you type. What you type is always shown in lowercase letters, unless your program is case-sensitive and uppercase letters are necessary for it to work properly.

**Keyboard
conventions**

This guide uses capitalization and some abbreviations to refer to the keys on the keyboard. The keys on your keyboard might not be labeled exactly as they are in this guide.

What is in this guide...	What it is used to mean...
hyphen (-)	Separates individual keys. For example, Ctrl-D means holding down the Ctrl key while pressing the D key.
<i>Enter</i>	Refers to the key that generates a carriage return; the key is named Return on some keyboards.
<i>type</i>	Pressing one or more keys on the keyboard.
<i>enter</i>	Pressing one or more keys and then pressing the Enter key.

Special messages

This guide contains special messages that are described as follows:

Note _____
A note contains important information that helps you install or operate the system efficiently.

Caution _____
A caution contains instructions that you must follow to avoid damage to the equipment, a system crash, or loss of data.

WARNING _____
A warning contains instructions that you must follow to avoid personal injury.

About this chapter

This chapter provides information about how SnapDrive for UNIX works and what it does.

Topics in this chapter

This chapter includes the following topics:

- ◆ [“How SnapDrive for UNIX works”](#) on page 2
- ◆ [“Comparing SnapDrive for UNIX on different host platforms”](#) on page 9
- ◆ [“Where to go for more information”](#) on page 11

How SnapDrive for UNIX works

What SnapDrive for UNIX does

SnapDrive for UNIX is a tool that simplifies the backup of data so that you can recover it should it be accidentally deleted or modified. SnapDrive for UNIX uses NetApp Snapshot™ technology to create an image (i.e., a snapshot) of the data on a storage system attached to a UNIX host at a specific point in time. If the need arises later, you can restore the data to the storage system. When you restore a snapshot, it replaces the current data on the storage system with the image of the data in the snapshot.

In addition, SnapDrive for UNIX lets you provision storage on the storage system. SnapDrive for UNIX provides a number of storage features that enable you to manage the entire storage hierarchy, from the host-side application-visible file down through the volume manager to the storage-system-side LUNs providing the actual repository.

With SnapDrive for UNIX installed, you can perform the following tasks:

- ◆ Create a snapshot of one or more volume groups on a storage system. Then you can rename the snapshot, restore it, or delete it. You can also connect a snapshot to a different location on the host or a different host or disconnect it. In addition, SnapDrive for UNIX lets you display information about snapshots that you created with it.
- ◆ Create storage on a storage system in the form of LUNs, file systems, logical volumes, or disk groups. Then you can resize the storage or delete it. You can also connect the storage to a host or disconnect it. In addition, SnapDrive for UNIX lets you display information about storage that you created with it.

Note

SnapDrive for UNIX works only with snapshots it creates. It cannot restore snapshots that it did not create.

How SnapDrive for UNIX works

The SnapDrive for UNIX software interacts with the host operating system and volume manager. It lets you easily and quickly back up and restore data about host volume groups that you stored on a NetApp storage system. You can use it to manage the snapshots you create using it.

SnapDrive for UNIX coordinates the host Logical Volume Manager (LVM) volume groups and file systems to ensure that the host file systems stored on NetApp LUNs have consistent images in the snapshot. This action enables you to restore data from the backup snapshots without requiring significant data recovery steps on the host.

You can also use SnapDrive for UNIX to create and manage storage. The `snapdrive storage` commands work with LVM to let you create LVM objects and file systems that use the storage. They also let you remove the mappings between the storage and the host as well as delete the storage.

SnapDrive for UNIX communicates with the storage system using the host IP interface that you specified when you set up the storage system.

SnapDrive for UNIX and logical volumes

The host LVM combines LUNs from a storage system into disk or volume groups. This storage is then divided into logical volumes, which are used as if they were raw disk devices to hold file systems or raw data.

Note

This documentation refers to logical volumes as host volumes to distinguish them from storage system volumes.

SnapDrive for UNIX integrates with the host LVM to determine which NetApp LUNs make up each disk group, host volume, and file system requested for snapshot. Because data from any given host volume can be distributed across all disks in the disk group, snapshots can be taken and restored only for whole disk groups.

Security considerations

You must be logged in as root to use the SnapDrive for UNIX commands.

To enable SnapDrive for UNIX to access the storage systems connected to the host, you must configure it to use the login names and passwords assigned to the storage systems when you set them up. If you do not provide this information, SnapDrive for UNIX cannot communicate with the storage system. For information on supplying the login names and passwords, see [“Specifying the current login information for storage systems”](#) on page 107.

SnapDrive for UNIX stores this information on the host in an encrypted file.

By default, SnapDrive for UNIX does not encrypt the password information when it sends it out across the network. It communicates using HTTP over the standard IP connection. You can tell SnapDrive for UNIX to use HTTPS (SSL

encryption) to communicate with the storage system by setting the `use-https-to-filer` variable to On in the `snapdrive.conf` configuration file. (For more information on the configuration file, see [“Setting configuration information”](#) on page 62.).

Note

SnapDrive for UNIX does not use rsh, so you do not need to configure rsh.

Using access permissions on a storage system

SnapDrive for UNIX lets you specify access permissions for each host. These permissions indicate whether a host can perform certain snapshot and storage operations. They do not affect any of the show or list operations.

You set up permissions for each host in a file that resides on the storage system. You can control what action SnapDrive for UNIX takes when it does not find a permissions file for a given host by the value you set in the `snapdrive.conf` configuration file for the `all-access-if-rbac-unspecified` variable. The options are to allow all access to that storage system or to disable all access to it.

For more information on using access control, see [“Setting up access control”](#) on page 101. For more information on the configuration file, see [“Setting configuration information”](#) on page 62.

Methods for executing SnapDrive for UNIX

You execute SnapDrive for UNIX from the UNIX host. SnapDrive for UNIX manages snapshots and storage provisioning using a command-line interface. SnapDrive for UNIX lets you:

- ◆ Enter individual commands at the command line prompt.
- ◆ Run scripts you create that contain embedded SnapDrive for UNIX commands.

Prerequisites for using SnapDrive for UNIX

Before installing SnapDrive for UNIX, you must have done the following:

- ◆ Installed the Network Appliance™ FCP Host Attach Kit (attach kit) or iSCSI Host Support Kit (support kit) for your host. SnapDrive for UNIX works with the following kits:
 - ❖ IBM® AIX® FCP attach kit and iSCSI support kit
 - ❖ HP-UX® FCP attach kit and iSCSI support kit
 - ❖ Red Hat® Enterprise Linux® iSCSI support kit
 - ❖ Solaris® FCP attach kit

- ◆ Set up your host and storage systems.

Follow the instructions provided with the kit to set up your storage systems to work with the host.

Note

The most current information about SnapDrive for UNIX and its requirements is in the *SnapDrive for UNIX Compatibility Matrix* on the NOW site (<http://now.netapp.com/NOW/knowledge/docs/docs.shtml>).

SnapDrive for UNIX stack

SnapDrive for UNIX requires the following stack:

- ◆ Host operating system and appropriate patches
- ◆ Volume management for the operating system
- ◆ File system
 - ❖ AIX: JFS and JFS2
 - ❖ HP-UX: VxFS
 - ❖ Linux: Ext3
 - ❖ Solaris: VxFS
- ◆ Multipathing (except on Linux hosts and AIX hosts running iSCSI)
- ◆ IP access between the host and storage system
- ◆ FCP license or iSCSI license, depending on your host platform and its needs
- ◆ SnapRestore® license
- ◆ Attach or support kit software

For example, if you are using SnapDrive for UNIX with an AIX host, you must set up the features required by the attach kit for that host when you are using multipathing. This means you must install the Dot Hill® SANpath® software, which the attach kit uses to handle multipathing.

Note

For this release, AIX does not support multipathing if you have the iSCSI support kit installed.

- ◆ Data ONTAP™ software on your storage system

The most up-to-date information about the requirements for your host operating system are in the online *SnapDrive for UNIX Compatibility Matrix* and the online *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products*. These matrices are on the NOW site.

Supported configurations

SnapDrive for UNIX supports the same host cluster and storage system cluster configurations that the kit supports. These include the following topologies:

- ◆ Nonclustered configuration where a single host is connected to a single storage system
- ◆ Any of the topologies involving NetApp storage system cluster failover
- ◆ Any of the topologies involving NetApp-supported host clusters

See the kit documentation for more information about the recommended configurations for your host and the storage systems you are using.

Note

If you need a SnapDrive for UNIX configuration that is not mentioned in the kit documentation, consult with your NetApp sales representative.

Architectural limitations for SnapDrive for UNIX

This release of SnapDrive for UNIX has the following limitations:

- ◆ A LUN managed by SnapDrive for UNIX cannot serve as either of the following:
 - ❖ A boot disk or a system disk
 - ❖ A location for the system paging file or memory dump files (or swap disk)
- ◆ Support for volume groups spanning multiple storage system volumes or multiple storage systems is limited on Linux hosts because there is not a freeze operation in the software stack.

As a result, SnapDrive for UNIX cannot guarantee the generation of consistent snapshots across storage system volumes (or storage systems) on Linux hosts.

On Linux hosts, you must stop or perform whatever steps are necessary to quiesce the application before taking a snapshot. SnapDrive for UNIX then attempts to take a snapshot without freezing the target storage. It checks after the fact to ensure that the snapshot was consistent. If it was not, SnapDrive for UNIX discards the snapshot and informs the user of the failure. The consistency check is completely robust, so non-consistent snapshots are guaranteed to be noticed. But obtaining a consistent snapshot can take multiple attempts and may not be possible if the target storage is under heavy load.

- ◆ Taking a snapshot where applications are using host volumes as raw storage requires additional steps on HP-UX and Solaris hosts as well as Linux hosts. For each of these hosts, you must stop the application or perform the steps necessary to quiesce it before taking the snapshot.

You do not need stop applications on an AIX host because on that platform, SnapDrive for UNIX can freeze operations at the volume group level.

- ◆ On AIX, HP-UX, and Solaris hosts, support for volume groups spanning multiple storage system volumes or multiple storage systems is limited. In particular, SnapDrive for UNIX has the following limitation:
 - ❖ You cannot create volume groups that span storage systems using the `storage create` command.

Four key commands that SnapDrive for UNIX supports in these cases are:

- ❖ The `snap create` command
 - ❖ The `snap restore` command
 - ❖ The `snap connect` command
 - ❖ The `snap disconnect` command
- ◆ On HP-UX hosts, the default behavior of the LVM is to limit the host to 16 physical volumes per host volume group. This limits the maximum number of LUNs to 16. This limit affects the `snapdrive storage create` and the `snapdrive storage resize` commands. If you attempt to create a host volume with more than 16 LUNs, the `snapdrive storage create` command fails. If you try to resize a volume group so that it contains more than 16 LUNs, the `snapdrive storage resize` command fails.

To have more than 16 LUNs (physical volumes) on an HP-UX host, you must manually create the host volume group using the `vgcreate` command and set the value for `-p max_pv` to the maximum value you expect to need for LUNs. The only time you may set this limit is when you create the host volume group. See the HP-UX `vgcreate(1)` man page for more information,

- ◆ Support for using the `snap connect` command to connect to the originating host is limited to AIX and HP-UX hosts. SnapDrive for UNIX does not support this use of the command on a Solaris or Linux host. The originating-host case of `snap connect` arises when you connect a snapshot at a new location on the same host where it was last connected (or is still connected).
- ◆ In creating a logical volume or file system with `storage create`, you cannot directly specify the size of the entity being created. Instead, you must specify the size of the underlying volume group; any host volume or file system created within it will be marginally smaller than it is.
- ◆ If you want to resize your storage, you must use the `-addlun` option on the `storage resize` command line. SnapDrive for UNIX does not currently support expanding existing LUNs. For more information, see [“Increasing the size of the storage”](#) on page 167.
- ◆ In creating host volumes, SnapDrive for UNIX does not provide any options to control their formatting. It creates only concatenated host volumes.

SnapDrive for UNIX does operate correctly on host volumes of other formats (such as striped volumes) that were created outside of it.

- ◆ SnapDrive for UNIX supports qtrees, but it does not support qtree quotas.
- ◆ You cannot restore a portion of a disk group. SnapDrive for UNIX backs up and restores whole disk groups only.
- ◆ Solaris and Linux hosts have operating system limits on how many LUNs you can create. To avoid having a problem when you create LUNs on these hosts, use the `snapdrive config check luns` command. For more information, see “[Preparing hosts for adding LUNs](#)” on page 85.
- ◆ SnapDrive for UNIX does not support colons in the long forms of the names for LUNS and snapshots. The only place SnapDrive for UNIX accepts colons is between the components of a long snapshot name or between the storage system name and the storage system volume name of a LUN. For example, `toaster:/vol/vol1:snap1` would be a typical long snapshot name, while `toaster:/vol/vol1/lunA` would be a typical long LUN name.

Comparing SnapDrive for UNIX on different host platforms

About SnapDrive for UNIX on multiple platforms

In the majority of ways, SnapDrive for UNIX executes the same way on multiple host operating system platforms. This section describes the differences in how SnapDrive for UNIX executes across host operating systems. Except for these points, this document describes how to use SnapDrive for UNIX regardless of the host operating system. Differences in SnapDrive for UNIX behavior between the operating systems include the following:

- ◆ Installation instructions
- ◆ Pathnames
- ◆ Terminology
- ◆ Output from examples
- ◆ Cross-volume snapshot support
- ◆ Originating-host support of `snap connect` operation

This document contains separate installation instructions for each operating system. In addition, it provides a list of files that each operating system installs so you can see the pathnames used by that operating system.

In an effort to make this document easier to read, it uses one term for an item regardless of the operating system. For example, the terms *disk group* and *volume group* in this documentation can refer to either a disk group or a volume group. The term *host volume* refers to either a host volume or a logical volume; both are distinguished from storage system volumes.

Wherever possible, this document uses generic examples. In some cases, this document provides multiple examples to show the output from different operating systems. In most cases, the differences in output have to do with the differences in pathnames and terminology on an operating system so they are clear to users familiar with UNIX operating systems. For example, an HP-UX host uses `/dev/rdisk/c22t0d0` as a device name while a Solaris host uses `/dev/vx/rdmp/c1t0d2s2`, an AIX host uses `/dev/hdisk2`, and a Linux host uses `/dev/sdd`. To avoid making the examples overly long, this document sometimes uses one operating system for one example and a different operating system for the next example.

Differences between host platforms

There are three differences in supported SnapDrive for UNIX functionality across different host operating systems (for more information on these differences, see “[Architectural limitations for SnapDrive for UNIX](#)” on page 6).

- ◆ You can take snapshots that cross storage systems and storage system volumes on AIX, HP-UX, and Solaris hosts. Linux, however, does not include a freeze function, so any cross-volume snapshots you take from that host may not be consistent.
- ◆ You can only use the `snap connect` operation to connect to the originating host when that host is running AIX or HP-UX. SnapDrive for UNIX does not support connecting to the originating host when that host is running Solaris or Linux.
- ◆ You must prepare Linux and Solaris hosts before you add LUNs. See “[Preparing hosts for adding LUNs](#)” on page 85.

SnapDrive for UNIX terms for volume managers

The following table summarizes some of the differences in terms when referring to volume managers on the host platforms.

Host	Volume manager	Volume or disk groups	Location of logical volumes	Location of multipathing devices
AIX	Native Logical Volume Manager (LVM)	Volume groups (vg)	/dev/lvol-name Note: All logical volumes share the same namespace.	/dev/hdisk (FCP only) Multipathing not supported with iSCSI
HP-UX	Native Logical Volume Manager (LVM)	Volume groups (vg)	/dev/dg-name /volgroup-name	/dev/rdisk
Linux	Native Logical Volume Manager (LVM)	Volume groups (vg)	/dev/dg-name /lvol-name	Multipathing not supported
Solaris	Veritas Volume Manager (VxVM)	Disk groups (dg)	/dev/vx/dg-name/lvol-name	/dev/rdmp

Where to go for more information

Relevant documentation

This guide provides information about the basic tasks involved in installing SnapDrive for UNIX on your host and working with snapshots. The following documentation might also be useful to you:

- ◆ *SnapDrive for UNIX Release Notes (IBM AIX, HP-UX, Linux, Solaris)*
This document comes with SnapDrive for UNIX. You can also download a copy from <http://now.netapp.com>. It contains any last-minute information that you need to get your configuration up and running smoothly. It also contains late-breaking problems and their workarounds.
- ◆ *SnapDrive for UNIX Quick Start Guide (IBM AIX, HP-UX, Linux, Solaris)*
This document comes with SnapDrive for UNIX. You can also download a copy from <http://now.netapp.com>. It provides high-level steps so that you can quickly start using SnapDrive for UNIX to create snapshots and manage storage.
- ◆ *SnapDrive for UNIX Compatibility Matrix*
This document is available at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>. It is a dynamic, online document that contains the most up-to-date information specific to SnapDrive for UNIX and its platform requirements.
- ◆ *SnapDrive for UNIX man page*
This online document comes with the product. It contains descriptions of the SnapDrive for UNIX commands and covers issues such as using initiator groups and internal name generation.
- ◆ *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products*
This document is available at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/. It is a dynamic, online document that contains the most up-to-date information about the requirements for setting up a system in a NetApp SAN environment. It provides the most current details about storage systems and host platforms, cabling issues, switch issues, and configurations.
- ◆ *The FCP host attach kit documentation*
This documentation comes with the attach kit. It consists of the installation guide for your host and the release notes for that attach kit.

- ◆ The NetApp iSCSI host support kit documentation
This documentation comes with the support kit. It consists of the installation guide for your host and the release notes for that support kit.
- ◆ *System Configuration Guide*
This document is available at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml/>. This is an online document that contains information about the supported storage system models for Data ONTAP.
- ◆ *Data ONTAP Block Access Management Guide*
This document is available from the Data ONTAP library on the NOW site at <http://now.netapp.comNOW/knowledge/docs.shtml/>. It provides information about using Data ONTAP and setting up your storage system to work with Data ONTAP.
- ◆ Host operating system and host bus adapter (HBA) information
NetApp does not provide these documents. See the readme files and other documentation that you received with your host operating system.

About this chapter

This chapter explains the procedures you must follow to install SnapDrive for UNIX.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Verifying that your storage systems are ready”](#) on page 14
- ◆ [“Installation instructions for all hosts”](#) on page 19
- ◆ [“Installing and uninstalling SnapDrive for UNIX on an AIX host”](#) on page 22
- ◆ [“Installing and uninstalling SnapDrive for UNIX on an HP-UX host”](#) on page 38
- ◆ [“Installing and uninstalling SnapDrive for UNIX on a Linux host”](#) on page 45
- ◆ [“Installing and uninstalling SnapDrive for UNIX on a Solaris host”](#) on page 50
- ◆ [“Completing the installation”](#) on page 54
- ◆ [“Files installed by SnapDrive for UNIX on the host”](#) on page 55
- ◆ [“Upgrading to a new version of SnapDrive for UNIX”](#) on page 59

Verifying that your storage systems are ready

Verify storage system readiness

Before installing SnapDrive for UNIX, verify that your the storage systems are ready by making sure that the following conditions are true:

- ◆ The storage systems are online.
- ◆ The storage systems meet the minimum system requirements for SnapDrive for UNIX. See “[Requirements for storage systems](#)” on page 14.
- ◆ The HBAs and network interface cards (NICs) in your storage systems meet the requirements for your host operating system.
- ◆ The hosts and the storage systems can communicate using an IP interface. (You should have set this up when you set up your storage system.)
- ◆ SnapRestore is licensed and running. (You should have set this up when you set up your storage system.)

For detailed information about administering your storage system, see your Data ONTAP *Storage Management Guide*.

Requirements for storage systems

Each storage system in your SnapDrive for UNIX configuration must meet the requirements in the following table.

Component	Minimum requirement
Operating system	Data ONTAP 6.5.2 or later Note —— SnapDrive for UNIX supports flexible volumes. It does not does not take full advantage of all the features associated with flexible volumes. ——

Component	Minimum requirement
storage system setup	You must specify the partner IP address in the storage system cluster that can be used if a storage system failover occurs. Note——— You specify this address when you run the setup program on the storage system. See “Confirm storage system has partner IP address” on page 15.
Licenses	◆ FCP or iSCSI, depending on your host platform Note——— You must have the correct protocol running on the storage system for SnapDrive for UNIX to execute. ◆ SnapRestore software

Note———

For the latest SnapDrive for UNIX requirements, see the online *SnapDrive for UNIX Compatibility Matrix* and the *Compatibility and Configuration Guide for NetApp’s FCP and iSCSI Products*. These matrices are at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml/>.

Confirm storage system has partner IP address

When you ran the setup program on your storage system, it prompted you for an IP address for a partner storage system to use in case of a failover. Make sure that you supplied this address. If you didn’t supply it, SnapDrive for UNIX cannot inquire about the LUNs on a storage system that was taken over.

Example: The following is the portion of the storage system setup script that requests the IP address. This example uses the IP address 10.2.21.35.

```
filer_A> setup
...
Should interface e0 take over a partner IP address during failover?
[n]: y
Please enter the IP address or interface name to be taken over by e0
[]: 10.2.21.35
...
reboot -t 0
```

Cautions for using SnapDrive for UNIX

Net App strongly recommends that you heed the following cautions whenever you use SnapDrive for UNIX:

- ◆ Never disable the space reservation setting for any LUN managed by SnapDrive.
- ◆ Set the snap reserve on the storage system to 0 percent for each volume. (This is the recommended setting for a SAN environment.)
- ◆ Place all LUNs connected to the same host on a dedicated storage system volume accessible by just that host.
- ◆ If you use snapshots, you cannot use the entire space on a storage system volume to store your LUNs.

The storage system volume hosting the LUNs should be at least twice the combined size of all the LUNs on the storage system volume.

- ◆ Do not create any LUNs in /vol/vol0.

This storage system volume is used by Data ONTAP to administer the storage system and should not be used to contain any LUNs.

Preparing a storage system volume

You need to perform the following tasks on the storage system to create a volume that can hold the SnapDrive for UNIX LUNs attached to a single host:

- ◆ Create a filer volume.

Note

You can use either the command-line prompt on the storage system (for example, through a Telnet session) or FilerView to create a storage system volume dedicated to SnapDrive for UNIX.

For more information about the following procedures, see the *Data ONTAP Block Access Management Guide*.

- ◆ Create a qtree (only necessary if you plan to store LUNs at a qtree root, rather than at the dedicated storage system volume root).
- ◆ Reset the snap reserve option to 0 percent on the storage system volume holding all the LUNs attached to the host (optional, but highly recommended in a SAN environment).

Guidelines for creating storage system volumes: When you create a volume on a storage system to hold LUNs, keep the following in mind:

- ◆ You can create multiple LUNs on a storage system volume.

- ◆ A LUN must reside at either the root of a storage system volume or the root of a qtree.

Note

Do not create LUNs on the root filer volume (for example, /vol/vol0).

Recommendations for configuring storage system volumes: You can optimize your storage system volumes in the following ways:

- ◆ When multiple hosts share the same storage system, each host should have its own dedicated storage system volume to hold all the LUNs connected to that host.
- ◆ NetApp recommends that, when multiple LUNs exist on a storage system volume, the dedicated volume on which the LUNs reside contains only the LUNs for a single host. It must not contain any other files or directories.

Resetting the snap reserve option: By default, the `snap reserve` option for Data ONTAP 6.5.x is 20 percent. When you use Data ONTAP in a SAN environment, NetApp strongly recommends that you reset the `snap reserve` option to 0 percent on all storage system volumes holding SnapDrive for UNIX LUNs.

To reset the `snap reserve` option on the storage system, complete the following steps.

Step	Action
1	Access the storage system either by using a command such as <code>telnet</code> from the host or by going to the storage system console.
2	Enter the following command: <pre># snap reserve vol_name 0</pre> <i>vol_name</i> is the name of the volume where you want to set the <code>snap reserve</code> option.

To reset the `snap reserve` option using FilerView, complete the following steps.

Step	Action
1	Open a FilerView session to the storage system holding the volume whose <code>snap reserve</code> setting is to be changed.

Step	Action
2	From the main FilerView menu, navigate to Volumes > Snapshots > Configure.
3	In the Volume field, select the volume whose snap reserve setting is to be changed
4	In the space reservation field, enter “0.”
5	Click Apply at the bottom of the panel.

Installation instructions for all hosts

About installing the attach kit

Before you install SnapDrive for UNIX, you must install the attach kit for your operating system and get it working. Use the documentation that came with the attach kit. It contains information about volume managers, multipathing, and other features you need to set up before you install SnapDrive.

Verify that the hosts are ready

Verify that the hosts are ready by performing the following tasks:

- ◆ Confirm that the host and storage system can communicate.
- ◆ Set up a LUN on the storage system and map it to the host.
- ◆ Confirm that you set up the host and storage system correctly according to the instructions in the attach kit for that host.
- ◆ Verify that the host meets the minimum requirements for SnapDrive for UNIX, including the required operating system patches.

Note

For the latest information about the SnapDrive for UNIX requirements and the attach kits that NetApp provides, see the online *SnapDrive for UNIX Compatibility Matrix* and the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products*. These matrices are at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml/>.

Get a copy of the SnapDrive for UNIX software package

You can obtain the SnapDrive for UNIX software package in two ways:

- ◆ Download the package from <http://now.netapp.com>. The SnapDrive for UNIX software is bundled into a single compressed file. For information about getting this file, see “[Downloading the SnapDrive for UNIX software from NOW](#)” on page 20.
- ◆ Get the package from the SnapDrive for UNIX CD-ROM. For information about getting the software package from the CD-ROM, see “[Getting SnapDrive for UNIX software from the CD-ROM](#)” on page 21.

Note

To verify that you have the latest software package, check the online *SnapDrive for UNIX Compatibility Matrix* and the *Compatibility* and the *Configuration Guide for NetApp's FCP and iSCSI Products*. They are both available at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

Downloading the SnapDrive for UNIX software from NOW

To download the SnapDrive for UNIX software from NOW, complete the following steps.

Step	Action
1	Log in to http://now.netapp.com . In the fields on the right side of the page, complete the following steps: 1. Enter your user name and password. 2. From the Select Start Page list box, click Download Software. 3. Click the Login button.
2	Go to the SnapDrive for UNIX product row of the Software Download table and select your host operating system from the Select Platform drop-down list.
3	Click Go!
4	Follow the prompts to reach the Software Download page.
5	Download the software file to a local directory. If you did not download the file to your host machine, move it to that machine.
6	Go to the instructions for installing the SnapDrive for UNIX software on your host operating system.

Getting SnapDrive for UNIX software from the CD-ROM

To get the SnapDrive for UNIX software from the CD-ROM, complete the following steps.

Step	Action
1	Insert the CD-ROM containing the version of SnapDrive for UNIX for your host operating system into the CD-ROM drive.
2	Change to the directory containing the SnapDrive for UNIX software.

Installing and uninstalling SnapDrive for UNIX on an AIX host

System requirements

The following table lists the minimum requirements for using SnapDrive with an AIX host.

Component	Requirement
NetApp iSCSI Support Kit or FCP AIX Attach Kit	To make sure you have the correct version of the kit, see the online <i>SnapDrive for UNIX Compatibility Matrix</i> and the <i>Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products</i> .
Additional disk space	SnapDrive for UNIX maintains three log files. While SnapDrive for UNIX rotates the files when they reach a maximum size, you should make sure you have sufficient disk space for them. For more information about the log files, see “Setting up audit, recovery, and trace logging” on page 86. Based on the default settings for the audit and trace log files, you need at least 1.1 MB of space. There is no default size for the recovery log because it rotates only after an operation completes, not when it reaches a specific size.

Setting up your host

Set up the host and storage system according to the instructions in the installation and setup guide for the iSCSI or FCP kit. You must do this before you install SnapDrive for UNIX.

Getting the software package

You can obtain the SnapDrive for UNIX software package in two ways:

- ◆ Download the package from <http://now.netapp.com>. See “[Downloading the SnapDrive for UNIX software from NOW](#)” on page 20.
- ◆ Get the package from the SnapDrive for UNIX CD-ROM provided with the media package for your host operating system. See “[Getting SnapDrive for UNIX software from the CD-ROM](#)” on page 21.

Uncompress the software if you downloaded it

If you downloaded the software package from NOW, you can uncompress it by completing the following steps:

Step	Action
1	Uncompress the file and extract the software by entering the following command: <pre># uncompress NetApp.snapdrive_aix_2_0_1.Z</pre> Result: This command extracts the installation package NetApp.snapdrive_aix_2_0_1.
2	You're now ready to install SnapDrive for UNIX. Follow the steps in the section " Installing SnapDrive for UNIX on an AIX host " on page 23.

Installing SnapDrive for UNIX on an AIX host

You use the System Management Interface Tool (SMIT) to install the AIX package. Complete the following steps.

Step	Action
1	Make sure you are logged in as root.

Step	Action	
2	If you are installing...	Then...
	Software from the CD-ROM	1. Insert the CD-ROM into the CD-ROM drive. 2. Change to the root directory, which contains the software package NetApp.snapdrive_aix_2_0_1. Note _____ This package is in a standard AIX Licensed Program Product (LPP) format. _____ 3. Optional: Copy the package from the CD-ROM's root directory to a temporary location on the host. (You can also install the package from the CD-ROM.) Example: If you mount the CD-ROM on /mnt and want to place the package in the /tmp directory, you might enter commands similar to the following: <pre># mount -v cdrfs -o ro /dev/cd0 /mnt # cp /mnt/install /tmp # cd /tmp</pre>
	Software from NOW	Change to the directory where you placed the uncompressed file NetApp.snapdrive_aix_2_0_1.
3	Start SMIT by entering the following command: <pre># smit</pre>	

Step	Action
4	Select the option titled “Software Installation and Maintenance.” Example: When you start SMIT, it displays the following screen. On this screen, the “Software Installation and Maintenance” option is the first menu option. <pre>+-----+ root> # smit System Management Move cursor to desired item and press Enter. Software Installation and Maintenance Software License Management Devices System Storage Management (Physical & Logical Storage) Security & Users Communications Applications and Services Print Spooling Problem Determination Performance & Resource Scheduling System Environments Processes & Subsystems Applications Cluster System Management Using SMIT (information only) F1=Help F2=Refresh F3=Cancel Esc+8=Image Esc+9=Shell Esc+0=Exit Enter=Do +-----+</pre>

Step	Action
5	At the screen that appears, select the Install and Update Software menu option. Example: The following is an example of the Software Installation and Maintenance screen. <pre> +-----+ Software Installation and Maintenance Move cursor to desired item and press Enter. Install and Update Software List Software and Related Information Software Maintenance and Utilities Network Installation Management System Backup Manager F1=Help F2=Refresh F3=Cancel Esc+8=Image Esc+0=Exit Enter=Do Esc+9=Shell +-----+ </pre>
6	At the next screen, select the Install Software menu option. Example: The following is an example of the Install and Update Software screen. <pre> +-----+ Install and Update Software Move cursor to desired item and press Enter. Install Software Update Installed Software to Latest Level (Update All) Install Software Bundle Update Software by Fix (APAR) Install and Update from ALL Available Software F1=Help F2=Refresh F3=Cancel Esc+8=Image Esc+0=Exit Enter=Do Esc+9=Shell +-----+ </pre>

Step	Action
7	At the Install Software screen, specify the location of the software in one of the following ways: ◆ Manually enter the location by providing the following information. ❖ If you're installing from the CD-ROM, enter the CD-ROM drive. ❖ If you're installing from the host machine, enter the path to the software package (for example, /tmp/NetApp.snapdrive_aix_2_0_1). ◆ Press F4 to display a list of options. If you want to use the F4 method, complete the following steps: 1. Press F4. 2. At the prompt for software to install, enter NetApp.snapdrive. 3. At the prompt asking whether you want to continue or cancel, press Enter to complete the installation. The installation process checks the version of AIX. The installation process displays a warning message before it completes if you are not running a version of AIX that SnapDrive for UNIX supports: Example 1: The following is an example of entering the path to the software package when you're at the Install Software screen. <pre> +-----+ Install Software Type or select a value for the entry field. Press Enter AFTER making all desired changes. [Entry Fields] * INPUT device / directory for software [/tmp/aix/NetApp.snapdrive_aix_2_0_1]+ F1=Help F2=Refresh F3=Cancel F4=List Esc+5=Reset Esc+6=Command Esc+7=Edit Esc+8=Image Esc+9=Shell Esc+0=Exit Enter=Do +-----+</pre>

Step	Action
	Example 2: After you enter path to the software package, SMIT displays the following screen. This is the screen where you enter name of the software package, NetApp.snapdrive. <pre> +-----+ Install Software Type or select values in entry fields. Press Enter AFTER making all desired changes. [Entry Fields] * INPUT device / directory for software /tmp/aix/NetApp.snapdrive_aix_2_0_1 * SOFTWARE to install [NetApp.snapdrive] + PREVIEW only? (install operation will NOT occur) no + COMMIT software updates? yes + SAVE replaced files? no + AUTOMATICALLY install requisite software? yes + EXTEND file systems if space needed? yes + OVERWRITE same or newer versions? no + VERIFY install and check file sizes? no + Include corresponding LANGUAGE filesets? yes + DETAILED output? no + Process multiple volumes? yes + ACCEPT new license agreements? no + Preview new LICENSE agreements? no + F1=Help F2=Refresh F3=Cancel F4=List Esc+5=Reset Esc+6=Command Esc+7=Edit Esc+8=Image Esc+9=Shell Esc+0=Exit Enter=Do +-----+ </pre>

Step	Action
	<pre>+-----+ Installing Software... +-----+ installp: APPLYING software for: NetApp.snapdrive 2.0.1.0 << Copyright notice for NetApp.snapdrive >> Copyright (c) 2004 Network Appliance, Inc. All Rights Reserved. << End of copyright notice for NetApp.snapdrive >>. . . . Finished processing all filesets. (Total time: 2 secs). +-----+ Summaries: +-----+ Installation Summary ----- Name Level Part Event Result ----- NetApp.snapdrive 2.0.1.0 USR APPLY SUCCESS</pre>
8	Verify your installation of the package by entering the <code>lslpp -l NetAppsnapdrive</code> command line. It should produce the following output: <pre># lslpp -l NetApp.snapdrive Fileset Level State Description ----- Path: /usr/lib/objrepos NetApp.snapdrive 2.0.1.0 COMMITTED Network Appliance Snapdrive</pre> You can also check the SMIT log file (<code>smit.log</code> and <code>smit.script</code>). These files are in the SMIT log directory (<code>\$HOME</code>).

Step	Action
9	Complete your setup by configuring SnapDrive for UNIX for your system. Most of this information is set by default; however, you need to specify the following information: ◆ The login information for the storage system. See “Specifying the current login information for storage systems” on page 107. ◆ The autosupport settings (autosupport is an optional feature, but NetApp recommends you enable it). See “Setting up autosupport” on page 94. ◆ The correct configuration value for the following options based on whether you are using the FCP protocol or the iSCSI protocol: ❖ default-transport. Set this equal to your protocol. The acceptable values are “fcp” and “iscsi”. ❖ multipathing-type. For FCP, set this to “SANPath”. For iSCSI, set this to “none”. Multipathing is not available with iSCSI for this release. Note_____ For general information about configuration settings, see “Setting configuration information” on page 62.

Uninstalling the software from an AIX host

You use SMIT to uninstall SnapDrive for UNIX from an AIX system. Complete the following steps.

Step	Action
1	Ensure that you are logged in as root.
2	Start SMIT by entering the following command: <pre># smit</pre>

Step	Action
3	Select the “Software Installation and Maintenance” menu option. Example: When you start SMIT, it displays the following screen. <pre> +-----+ root> # smit System Management Move cursor to desired item and press Enter. Software Installation and Maintenance Software License Management Devices System Storage Management (Physical & Logical Storage) Security & Users Communications Applications and Services Print Spooling Problem Determination Performance & Resource Scheduling System Environments Processes & Subsystems Applications Cluster System Management Using SMIT (information only) F1=Help F2=Refresh F3=Cancel Esc+8=Image Esc+9=Shell Esc+0=Exit Enter=Do +-----+ </pre>

Step	Action
4	At the screen that appears, select the Software Maintenance and Utilities menu option. Example: The following is an example of the Software Installation and Maintenance screen. <pre> +-----+ Software Installation and Maintenance Move cursor to desired item and press Enter. Install and Update Software List Software and Related Information Software Maintenance and Utilities Network Installation Management System Backup Manager F1=Help F2=Refresh F3=Cancel Esc+8=Image Esc+9=Shell Esc+0=Exit Enter=Do +-----+</pre>

Step	Action
5	At the next screen, select the Remove Installed Software menu option. Example: The following is an example of the Software Maintenance and Utilities screen. <pre>+-----+ Software Maintenance and Utilities Move cursor to desired item and press Enter. Commit Applied Software Updates (Remove Saved Files) Reject Applied Software Updates (Use Previous Version) Remove Installed Software Copy Software to Hard Disk for Future Installation Check Software File Sizes After Installation Verify Software Installation and Requisites Clean Up After Failed or Interrupted Installation F1=Help F2=Refresh F3=Cancel Esc+8=Image Esc+9=Shell Esc+0=Exit Enter=Do +-----+</pre>

Step	Action
6	Remove the software in one of the following ways: ◆ Enter the package name (NetApp.snapdrive). (Make sure that the “Preview only” option is set to no.) ◆ Press F4 to display a list of names. If you use F4, complete the following steps: 1. Press F4. 2. Scroll down the list of names until you reach NetApp.snapdrive. 3. Select NetApp.snapdrive and press Enter. 4. At the prompt asking whether you want to continue or cancel, press Enter to complete the uninstall. The uninstall process places the log file in /tmp/snapdrive_uninstall. Note_____ When you uninstall the software, the log files are not removed. You must go to the /var/snapdrive directory and manually remove them. _____ Example: The following is an example of the Remove Installed Software screen. Note_____ By default PREVIEW only? is set to yes. You must change it to no if you want to uninstall. _____ <pre> +-----+ Remove Installed Software Type or select values in entry fields. Press Enter AFTER making all desired changes. [Entry Fields] * SOFTWARE name [NetApp.snapdrive] + PREVIEW only? (remove operation will NOT occur) no + REMOVE dependent software? no + EXTEND file systems if space needed? no + DETAILED output? no + F1=Help F2=Refresh F3=Cancel F4=List Esc+5=Reset Esc+6=Command Esc+7=Edit Esc+8=Image Esc+9=Shell Esc+0=Exit Enter=Do +-----+ </pre>

Step	Action
	<pre>+-----+ Deinstalling Software... +-----+ installp: DEINSTALLING software for: NetApp.snapdrive 2.0.1.0 Finished processing all filesets. (Total time: 1 secs). +-----+ Summaries: +-----+ Installation Summary ----- Name Level Part Event Result ----- NetApp.snapdrive 2.0.1.0 USR DEINSTALL SUCCESS Example of warning: If you are trying to uninstall a product that is not currently installed, you see messages such as the following: Not Installed ----- No software could be found on the system that could be deinstalled for the following requests: NetAppSnapdrive (The fileset may not be currently installed, or you may have made a typographical error.)</pre>

Installing and uninstalling SnapDrive for UNIX on an HP-UX host

System requirements

The following table lists the minimum requirements for using SnapDrive with an HP-UX host.

Component	Requirement
NetApp FCP HP-UX Attach Kit or iSCSI HP-UX Initiator Support Kit	To make sure you have the correct version of the kit, see the online <i>SnapDrive for UNIX Compatibility Matrix</i> and the <i>Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products</i> .
Additional disk space	SnapDrive for UNIX maintains three log files. While SnapDrive for UNIX rotates the files when they reach a maximum size, you should make sure you have sufficient disk space for them. For more information about the log files, see “Setting up audit, recovery, and trace logging” on page 86. Based on the default settings for the audit and trace log files, you need at least 1.1 MB of space. There is no default size for the recovery log because it rotates only after an operation completes, not when it reaches a specific size.

Setting up your host

Set up the host and filer according to the instructions in the documentation for your attach kit or support kit. You must do this before you install SnapDrive for UNIX.

Getting the software package

You can obtain the SnapDrive for UNIX software package in two ways:

- ◆ Download the package from <http://now.netapp.com>. See “[Downloading the SnapDrive for UNIX software from NOW](#)” on page 20.

- ◆ Get the package from the SnapDrive for UNIX CD-ROM provided with the media package for your host operating system. See [“Getting SnapDrive for UNIX software from the CD-ROM”](#) on page 21.

Moving the downloaded file to a local directory

If you downloaded the file and did not place it on your HP-UX machine, you must move it to that machine. Complete the following steps.

Step	Action
1	Use commands similar to the following ones to move the file you downloaded from the NOW site to your HP-UX machine. <pre># mkdir /tmp/hpux # cd /tmp/hpux # cp /u/hpux/HP-pkgs/NTAPsnapdrive_hpux_2_0_1.depot .</pre> Make sure you include the period (.) at the end of the copy command line.
2	You’re now ready to install SnapDrive for UNIX. Follow the steps in the section “Installing SnapDrive for UNIX on an HP-UX host” on page 39.

Installing SnapDrive for UNIX on an HP-UX host

To install SnapDrive for UNIX on an HP-UX host, complete the following steps.

Step	Action
1	Make sure you are logged in as either root or an authorized SD (software distributor) user for that HP-UX host. Note_____ The installation file is a standard HP-UX depot file, so any user on your system with software installation privileges can perform the installation. For more information on authorized users, see the HP-UX <code>swacl</code> man page. _____

Step	Action	
2	If you are installing...	Then...
	Software from the CD-ROM	1. Mount the CD-ROM on the host system by entering a command similar to the following: <pre># mount /dev/cdrom /mnt/cdrom</pre> 2. Change to the directory where your CD-ROM is mounted. For example, change to /cdrom/cdrom0: <pre>cd /cdrom/cdrom0</pre> Note _____ You can install the software directly from the CD-ROM. NetApp, however, recommends that you copy the file to the local disk before installing. 3. If you want to install the software from a local directory on the host, copy the file to that directory. For example, you might need to enter commands similar to the following: <pre># mkdir /tmp/hpux # mount /dev/cdrom /mnt/cdrom # cp /mnt/cdrom/NTAPsnapdrive_hpux_2_0_1.depot /tmp/hpux # cd /tmp/hpux</pre>
	Software from NOW	Change to the directory where you put the software you downloaded from NOW.

Step	Action
3	Enter the following command to install the software: <pre># swinstall -s <pathname>/NTAPsnapdrive_hpx_2_0_1.depot snapdrive</pre> <i>pathname</i> is the full path to the NTAPsnapdrive file. This path must begin with a slash (/). The file needs to be on a NFS mounted directory. The script installs SnapDrive for UNIX in the /opt/NetApp/snapdrive/ directory. Example: The swinstall command installs the SnapDrive for UNIX software without a problem. It writes installation information to a log file. <pre># swinstall -s /tmp/hpx/NTAPsnapdrive_hpx_2_0_1.depot snapdrive</pre> <pre>===== 03/11/04 15:49:28 EST BEGIN swinstall SESSION (non-interactive) (jobid=netapp2-0168)</pre> <pre> * Session started for user "root@netapp2". * Beginning Selection * Target connection succeeded for "netapp2:". * Source: /tmp/hpx/NTAPsnapdrive_hpx_2_0_1.depot * Targets: netapp2:/ * Software selections: snapdrive.command_parisc,r=1,fr=2.0.1,fa=HP-UX_B.11.00_32/64 snapdrive.snapdrive_diag_scripts,r=1 snapdrive.snapdrive_docs,r=1 * Selection succeeded. * Beginning Analysis and Execution * Session selections have been saved in the file "/.sw/sessions/swinstall.last". * The analysis phase succeeded for "netapp2:". * The execution phase succeeded for "netapp2:". * Analysis and Execution succeeded.</pre> NOTE: More information may be found in the agent logfile using the command "swjob -a log netapp2-0168 @ netapp2:/". <pre>===== 03/11/04 15:49:48 EST END swinstall SESSION (non-interactive) (jobid=netapp2-0168)</pre>

Step	Action
4	Verify your installation using the <code>swverify</code> command. Example: The following example uses the <code>swverify</code> command to confirm that the SnapDrive for UNIX software installed. <pre># swverify snapdrive ===== 03/12/04 11:55:49 EST BEGIN swverify SESSION (non-interactive) (jobid=netapp2-0172) * Session started for user "root@netapp2". * Beginning Selection * Target connection succeeded for "netapp2:/". * Software selections: snapdrive.command_parisc,l=/opt/NetApp/snapdrive,r=2,fr=2.0.1,fa=HP- UX_B.11.00_32/64 snapdrive.snapdrive_diag_scripts,l=/opt/NetApp/snapdrive,r=2.1 snapdrive.snapdrive_docs,l=/opt/NetApp/snapdrive,r=1 * Selection succeeded. * Beginning Analysis * Session selections have been saved in the file "/.sw/sessions/swverify.last". * The analysis phase succeeded for "netapp2:/". * Verification succeeded. NOTE: More information may be found in the agent logfile using the command "swjob -a log netapp2-0172 @ netapp2:/". ===== 03/12/04 11:55:54 EST END swverify SESSION (non-interactive) (jobid=netapp2-0172) #</pre>

Step	Action
5	Complete your setup by configuring SnapDrive for UNIX for your system. Most of this information is set by default; however, you need to specify the following information: ◆ The login information for the storage system. See “Specifying the current login information for storage systems” on page 107. ◆ The autosupport settings (autosupport is an optional feature, but NetApp recommends you enable it). See “Setting up autosupport” on page 94. Note_____ For general information about configuration settings, see “Setting configuration information” on page 62. _____

Uninstalling the software from an HP-UX host

To uninstall SnapDrive for UNIX from an HP-UX system, complete the following steps.

Step	Action
1	Ensure that you are logged in as either root or an authorized SD (software distributor) user for that HP-UX host.
2	Use the <code>swremove</code> command to remove the NetApp software. <pre># swremove snapdrive</pre> Note_____ This command does not remove the log files. You must go to the <code>/var/snapdrive</code> directory and manually remove them. _____

Step	Action
	Example: The following example uninstalls SnapDrive for UNIX on an HP-UX system. <pre> # swremove snapdrive ===== 03/11/04 16:22:44 EST BEGIN swremove SESSION (non-interactive) (jobid=netapp2-0170) * Session started for user "root@netapp2". * Beginning Selection * Target connection succeeded for "netapp2:/". * Software selections: snapdrive.command_parisc,l=/opt/NetApp/snapdrive,r=2,fr=2.0.1,fa=HP- UX_B.11.00_32/64 snapdrive.snapdrive_diag_scripts,l=/opt/NetApp/snapdrive,r=1 snapdrive.snapdrive_docs,l=/opt/NetApp/snapdrive,r=1 * Selection succeeded. * Beginning Analysis * Session selections have been saved in the file "/.sw/sessions/swremove.last". * The analysis phase succeeded for "netapp2:/". * Analysis succeeded. * Beginning Execution * The execution phase succeeded for "netapp2:/". * Execution succeeded. NOTE: More information may be found in the agent logfile using the command "swjob -a log netapp2-0170 @ netapp2:/". ===== 03/11/04 16:22:51 EST END swremove SESSION (non-interactive) (jobid=netapp2-0170) # </pre>

Installing and uninstalling SnapDrive for UNIX on a Linux host

System requirements

The following table lists the minimum requirements for using SnapDrive with a Red Hat Enterprise Linux host using iSCSI.

Component	Requirement
NetApp iSCSI Red Hat Linux Initiator Support Kit	To make sure you have the correct version of the kit, see the online <i>SnapDrive for UNIX Compatibility Matrix</i> and the <i>Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products</i> .
Additional disk space	SnapDrive for UNIX maintains three log files. While SnapDrive for UNIX rotates the files when they reach a maximum size, you should make sure you have sufficient disk space for them. For more information about the log files, see “Setting up audit, recovery, and trace logging” on page 86. Based on the default settings for the audit and trace log files, you need at least 1.1 MB of space. There is no default size for the recovery log because it rotates only after an operation completes, not when it reaches a specific size.

Setting up your host

Set up the host and filer according to the instructions in the installation and setup guide for the iSCSI Red Hat Linux Initiator support kit. You must do this before you install SnapDrive for UNIX.

Getting the software package

You can obtain the SnapDrive for UNIX software package in two ways:

- ◆ Download the package from <http://now.netapp.com>. See “[Downloading the SnapDrive for UNIX software from NOW](#)” on page 20.
- ◆ Get the package from the SnapDrive for UNIX CD-ROM provided with the media package for your host operating system. See “[Getting SnapDrive for UNIX software from the CD-ROM](#)” on page 21.

Moving the downloaded file to a local directory

If you downloaded the file and did not place it on your Linux host, you must move it to that host. Complete the following steps.

Step	Action
1	Copy the downloaded file to your Linux host. You can place it in any directory on the host. Example: You can use commands similar to the following ones to move the file you downloaded from the NOW site to your Linux machine. <pre># mkdir /tmp/linux # cd /tmp/linux # cp /u/linux/netapp.snapdrive.linux_2_0_1.rpm .</pre> Make sure you include the period (.) at the end of the copy command line.
2	You're now ready to install SnapDrive for UNIX. Follow the steps in the section “Installing SnapDrive for UNIX on a Linux host” on page 46.

Installing SnapDrive for UNIX on an Linux host

To install SnapDrive for UNIX on a Linux host, complete the following steps.

Step	Action
1	Make sure you are logged in as root. If you are executing this file remotely and your system configuration does not allow you to log in as root, use the <code>su</code> command to become root. Note_____ The installation file is a standard Linux .rpm file. _____

Step	Action	
2	If you are installing...	Then...
	Software from the CD-ROM	1. Mount the CD-ROM on the host system by entering a command similar to the following: <pre># mount /dev/cdrom /mnt/cdrom</pre> 2. Change to the directory where your CD-ROM is mounted. For example, change to /cdrom/cdrom0: <pre>cd /cdrom/cdrom0</pre> Note _____ You can install the software directly from the CD-ROM or you can copy it to a local disk. 3. If you want to install the software from a local directory on the host, copy the file to that directory. For example, you might need to enter commands similar to the following: <pre># mkdir /tmp/linux # mount /dev/cdrom /mnt/cdrom # cp /mnt/cdrom/netapp.snapdrive.linux_2_0_1.rpm /tmp/linux # cd /tmp/linux</pre>
	Software from NOW	Change to the directory on your Linux host where you put the software you downloaded from the NOW site.
3	Use the rpm command to install the software. <pre># rpm -U -v <pathname>/netapp.snapdrive.linux_2_0_1.rpm</pre> Example 1: The rpm command installs the SnapDrive for UNIX software without a problem. It writes installation information to a log file. <pre># rpm -U -v netapp.snapdrive.linux_2_0_1.rpm Preparing packages for installation... netapp.snapdrive-2.0.1 #</pre>	

Step	Action
4	Verify your installation. Example: The following example uses the <code>rpm</code> and <code>ls</code> commands to verify that the SnapDrive for UNIX software installed. <pre># rpm -qa netapp.snapdrive netapp.snapdrive-2.0.1 # ls -R /opt/NetApp /opt/NetApp: snapdrive /opt/NetApp/snapdrive: bin diag docs snapdrive.conf /opt/NetApp/snapdrive/bin: snapdrive /opt/NetApp/snapdrive/diag: filer_info linux_info SHsupport.pm snapdrive.dc Telnets.pm /opt/NetApp/snapdrive/docs: man1 /opt/NetApp/snapdrive/docs/man1: brocade_info.1 filer_info.1 linux_info.1 snapdrive.1 snapdrive.dc.1 # ls -l /usr/sbin/snapdrive -rwxr-xr-x 1 root root ../../opt/NetApp/snapdrive/bin/snapdrive #</pre>
5	Complete your setup by configuring SnapDrive for UNIX for your system. Most of this information is set by default; however, you need to specify the following information: ◆ The login information for the storage system. See “Specifying the current login information for storage systems” on page 107. ◆ The autosupport settings (autosupport is an optional feature, but NetApp recommends you enable it). See “Setting up autosupport” on page 94. Note _____ For general information about configuration settings, see “Setting configuration information” on page 62. _____

Uninstalling the software from a Linux host

To uninstall SnapDrive for UNIX from a Linux system, complete the following steps.

Step	Action
1	Ensure that you are logged in as root.
2	Use the <code>rpm</code> command to remove the NetApp software. Example: The following example uses the <code>rpm</code> command to uninstall the SnapDrive for UNIX software. <pre># rpm -e netapp.snapdrive-2.0.1 #</pre> Note_____ This command does not remove the log files. You must go to the <code>/var/log</code> directory and manually remove them. _____
3	Verify that the package uninstalled. Example: The following example verifies that SnapDrive for UNIX is no longer installed. <pre># rpm -qa netapp.snapdrive #</pre>

Installing and uninstalling SnapDrive for UNIX on a Solaris host

System requirements

The following table lists the minimum requirements for using SnapDrive with a Solaris host.

Component	Requirement
NetApp FCP Solaris Attach Kit	To make sure you have the correct version of the kit, see the online <i>SnapDrive for UNIX Compatibility Matrix</i> and the <i>Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products</i> .
Additional disk space	SnapDrive for UNIX maintains three log files. While SnapDrive for UNIX rotates the files when they reach a maximum size, you should make sure you have sufficient disk space for them. For more information about the log files, see “Setting up audit, recovery, and trace logging” on page 86. Based on the default settings for the audit and trace log files, you need at least 1.1 MB of space. There is no default size for the recovery log because it rotates only after an operation completes, not when it reaches a specific size.

Setting up your host

Set up the host and filer according to the instructions in the installation and setup guide for the Solaris attach kit. You must do this before you install SnapDrive for UNIX.

Getting the software package

You can obtain the SnapDrive for UNIX software package in two ways:

- ◆ Download the package from <http://now.netapp.com>. See “[Downloading the SnapDrive for UNIX software from NOW](#)” on page 20.
- ◆ Get the package from the SnapDrive for UNIX CD-ROM provided with the media package for your host operating system. See “[Getting SnapDrive for UNIX software from the CD-ROM](#)” on page 21.

Uncompressing the software package you downloaded

If you downloaded the software package from NOW, you can uncompress it by completing the following steps:

Step	Action
1	Change to the directory where you downloaded the compressed file. For example, change to the /tmp directory: <pre># cd /tmp</pre>
2	Enter the following command to uncompress the file: <pre># uncompress NTAPsnapdrive.tar.Z</pre>
3	Enter the following command to extract the software: <pre># tar -xvf NTAPsnapdrive.tar</pre>
4	You're now ready to install SnapDrive for UNIX. Follow the steps in the section " Installing SnapDrive for UNIX on a Solaris host " on page 51.

Installing SnapDrive for UNIX on a Solaris host

To install SnapDrive for UNIX on a Solaris host, complete the following steps.

Step	Action	
1	Ensure that you are logged in as root.	
2	If you are installing...	Then...
	Software from the CD-ROM	1. Mount the CD-ROM on the host system. 2. Change to the directory where your CD-ROM is mounted. For example, change to /cdrom/cdrom0: <pre>cd /cdrom/cdrom0</pre>
	Extracted software from NOW	Change to the directory where you extracted the software.

Step	Action
3	Enter the following command to install the software: <pre># ./install</pre> NetApp recommends that you have both the attach kit and the ASL installed before you install SnapDrive for UNIX. Example: The script installs the SnapDrive for UNIX software without a problem. It writes installation information to a log file. <pre># ./install Installing NTAPsnapdrive now ... NTAPsnapdrive install completed successfully. snapdrive Installation complete. Log is in /tmp/snapdrive_install_log.23752. #</pre>
4	Complete your setup by configuring SnapDrive for UNIX for your system. Most of this information is set by default; however, you need to specify the following information: ◆ The login information for the storage system. See “Specifying the current login information for storage systems” on page 107. ◆ The autosupport settings (autosupport is an optional feature, but NetApp recommends you enable it). See “Setting up autosupport” on page 94. Note_____ For general information about configuration settings, see “Setting configuration information” on page 62. _____

Uninstalling the software on a Solaris host

To uninstall SnapDrive for UNIX from a Solaris system, complete the following steps.

Step	Action
1	Ensure that you are logged in as root.

Step	Action
2	Use the <code>uninstall</code> command to remove the NetApp software. Answer <code>y</code> (or <code>yes</code>) when the command asks whether you are sure you want to remove the attach kit. If you answer <code>n</code> (or <code>no</code>), the command does not uninstall the attach kit. <pre># ./uninstall Are you sure you want to uninstall the snapdrive package? [y,n,?,q] y Removing NTAPsnapdrive ... NTAPsnapdrive successfully removed. Remove process complete. Log is in /tmp/snapdrive_remove_log.23789 #</pre> To avoid getting the confirmation query, enter the <code>./uninstall</code> command with the <code>-f</code> option, which prevents the uninstall script from running in interactive mode. Note The uninstall process might fail if the package has been partially uninstalled. If that happens, you can uninstall SnapDrive for UNIX by inserting the CD-ROM and executing the <code>./uninstall</code> command from the top-level directory.

Completing the installation

Completing the installation

Once you've installed SnapDrive for UNIX on your host, you need to complete the installation by performing the following tasks:

- ◆ Verify that the installation program installed all the necessary files on your host. For more information, see [“Files installed by SnapDrive for UNIX on the host”](#) on page 55.
- ◆ Confirm that the configuration variables in the snapdrive.conf have the correct settings. For the majority of these variables, the default values should be correct. In a few cases, you may need to modify these values. See [“Setting configuration information”](#) on page 62.
- ◆ Supply SnapDrive for UNIX with the current storage system login information. When you set up your storage system, you supplied a user login for it. SnapDrive for UNIX needs this login information in order to work with the storage system,. For information on supplying the login information, see [“Specifying the current login information for storage systems”](#) on page 107.
- ◆ Specify the access control permissions for each host on each storage system. If you do not specify access control permissions, SnapDrive for UNIX checks the value of the variable `all-access-if-rbac-unspecified` in the snapdrive.conf file. If the variable is set to on (the default), it allows the host to perform all SnapDrive for UNIX operations on that storage system. If the value is set to off, it prevents the host from performing any operations other than the show/list operations. For information on setting permissions, see [“Setting up access control”](#) on page 101.

Files installed by SnapDrive for UNIX on the host

Files installed by the software

When you install SnapDrive for UNIX, it installs a number of files on the host. These files serve different purposes.

Note

There are slight differences in the files based on your host operating system. In most cases, these differences have to do with the path names used by that host. Also, some hosts use symbolic links (symlinks) to files while other hosts simply install the file in two locations.

Executables: These are the executables SnapDrive for UNIX uses.

AIX

- ◆ /opt/NetApp/snapdrive/bin/snapdrive

HP-UX

- ◆ /opt/NetApp/snapdrive/bin/snapdrive
- ◆ /opt/NetApp/snapdrive/bin/snapdrive_parisc

Linux

- ◆ /opt/NetApp/snapdrive/bin/snapdrive

Solaris

- ◆ /opt/NTAPsnapdrive/bin/snapdrive

Configuration file: SnapDrive for UNIX stores configuration information in this file. You should modify it for your system. If you upgrade your version of SnapDrive for UNIX, it maintains your current snapdrive.conf file. For more information on setting configuration variables see “[Setting configuration information](#)” on page 62. For information on upgrading your version of SnapDrive for UNIX and how it handles the snapdrive.conf file, see “[Upgrading to a new version of SnapDrive for UNIX](#)” on page 59.

AIX, HP-UX, and Linux

- ◆ /opt/NetApp/snapdrive/snapdrive.conf

Solaris

- ◆ /opt/NTAPsnapdrive/snapdrive.conf

Uninstall files: If you decide to remove SnapDrive for UNIX, it uses these files. These files are also in the top-level directory on the SnapDrive for UNIX CD-ROM. See “[Installing SnapDrive for UNIX on a Solaris host](#)” on page 51 for more information.

AIX

You use SMIT to uninstall SnapDrive for UNIX from an AIX host so there is not an uninstall file.

HP-UX

You use swremove to uninstall SnapDrive for UNIX from a HP-UX host so there is not an uninstall file.

Linux

You use rpm to uninstall SnapDrive for UNIX from a Linux host so there is not an uninstall file.

Solaris

- ◆ /opt/NTAPsnapdrive/bin/snapdrive.admin
- ◆ /opt/NTAPsnapdrive/bin/uninstall
- ◆

Diagnostic files: You can run the following files if you have a problem with SnapDrive. For more information about these files, see “[Data collection utility](#)” on page 184.

AIX and HP-UX

- ◆ /opt/NetApp/snapdrive/diag/snapdrive.dc
- ◆ /opt/NetApp/snapdrive/diag/filer_info
- ◆ /opt/NetApp/snapdrive/diag/brocade_info
- ◆ /opt/NetApp/snapdrive/diag/cisco_info
- ◆ /opt/NetApp/snapdrive/diag/mcdata_info
- ◆ /opt/NetApp/snapdrive/diag/SHsupport.pm
- ◆ /opt/NetApp/snapdrive/diag/Telnet.pm

AIX only

- ◆ /opt/NetApp/snapdrive/diag/aix_info

HP-UX only

- ◆ /opt/NetApp/snapdrive/diag/hpux_info

Linux

- ◆ /opt/NetApp/snapdrive/diag/snapdrive.dc
- ◆ /opt/NetApp/snapdrive/diag/filer_info
- ◆ /opt/NetApp/snapdrive/diag/linux_info
- ◆ /opt/NetApp/snapdrive/diag/SHsupport.pm
- ◆ /opt/NetApp/snapdrive/diag/Telnet.pm

Solaris

- ◆ /opt/NTAPsnapdrive/diag/snapdrive.dc
- ◆ /opt/NTAPsnapdrive/diag/solaris_info
- ◆ /opt/NTAPsnapdrive/diag/filer_info
- ◆ /opt/NTAPsnapdrive/diag/brocade_info
- ◆ /opt/NTAPsnapdrive/diag/cisco_info
- ◆ /opt/NTAPsnapdrive/diag/mcdata_info
- ◆ /opt/NTAPsnapdrive/diag/SHsupport.pm
- ◆ /opt/NTAPsnapdrive/diag/Telnet.pm

Log files: SnapDrive for UNIX creates several log files.

AIX and HP-UX

- ◆ /var/snapdrive/sd-audit.log
- ◆ /var/snapdrive/sd-recovery.log
- ◆ /var/snapdrive/sd-trace.log

Linux and Solaris

- ◆ /var/log/sd-audit.log
- ◆ /var/log/sd-recovery.log
- ◆ /var/log/sd-trace.log

Man pages: SnapDrive for UNIX provides man pages in several formats:

AIX and HP-UX

- ◆ /opt/NetApp/snapdrive/man/man1/snapdrive.dc.1
- ◆ /opt/NetApp/snapdrive/man/man1/snapdrive.1
- ◆ /opt/NetApp/snapdrive/docs/snapdrive.1.cat
- ◆ /opt/NetApp/snapdrive/docs/snapdrive.1.html
- ◆ /opt/NetApp/snapdrive/docs/snapdrive.1.pdf
- ◆ /opt/NetApp/snapdrive/docs/snapdrive.1.txt
- ◆ /opt/NetApp/snapdrive/docs/brocade_info.1

- ◆ /opt/NetApp/snapdrive/docs/mcdata_info.1
- ◆ /opt/NetApp/snapdrive/docs/cisco_info.1
- ◆ /opt/NetApp/snapdrive/docs/filer_info.1
- ◆ /opt/NetApp/snapdrive/docs/aix_info.1 (AIX host only)
- ◆ /opt/NetApp/snapdrive/docs/hpux_info.1 (HP-UX host only)

Linux

- ◆ /opt/NetApp/snapdrive/docs/man1/snapdrive.dc.1
- ◆ /opt/NetApp/snapdrive/docs/man1/snapdrive.1
- ◆ /opt/NetApp/snapdrive/docs/man1/filer_info.1
- ◆ /opt/NetApp/snapdrive/docs/man1/linux_info.1

Solaris

- ◆ /usr/share/man/man1/snapdrive.1
- ◆ /opt/NTAPsnapdrive/docs/snapdrive.1.man
- ◆ /opt/NTAPsnapdrive/docs/snapdrive.1.cat
- ◆ /opt/NTAPsnapdrive/docs/snapdrive.1.txt
- ◆ /opt/NTAPsnapdrive/docs/snapdrive.1.pdf
- ◆ /opt/NTAPsnapdrive/docs/snapdrive.1.html
- ◆ /opt/NTAPsnapdrive/docs/solaris_info.1
- ◆ /opt/NTAPsnapdrive/docs/brocade_info.1
- ◆ /opt/NTAPsnapdrive/docs/mcdata_info.1
- ◆ /opt/NTAPsnapdrive/docs/cisco_info.1
- ◆ /opt/NTAPsnapdrive/docs/filer_info.1

Upgrading to a new version of SnapDrive for UNIX

Upgrading your version of SnapDrive for UNIX

SnapDrive for UNIX makes it easy for you to upgrade from an earlier version, such as 1.1, to the current version.

You do not have to uninstall SnapDrive for UNIX. Instead, install the latest version of the software on top of the current version.

When you install a new version, SnapDrive for UNIX checks to see if you already have a version installed. If you do, it preserves the current `snapdrive.conf` file and renames the version of the file it is installing to `snapdrive.conf.2.0`. This way it avoids overwriting your existing `snapdrive.conf` file, so you do not lose any changes you made if you customized your settings in the file.

By default, SnapDrive for UNIX comments out the variables in the `snapdrive.conf` file. This means it automatically uses the default values for all variables except the ones you customize. As a result, SnapDrive for UNIX will use the default values for variables that are new even if you have an earlier version of the `snapdrive.conf` file. If you want to change these values, you must add the variables to your current `snapdrive.conf` file and specify the values you want.

For more information on the `snapdrive.conf` file and its variables, see [“Setting configuration information”](#) on page 62.

New variables in the `snapdrive.conf` file

You can use the `snapdrive config show` command to display the information in the `snapdrive.conf` file. A sample 2.x `snapdrive.conf` file is included in the section [“Setting configuration information”](#) on page 62. This section also includes details on the variables and their defaults.

Some of the variables that are new in this release include the following:

- ◆ `all-access-if-rbac-unspecified=on`. This variable permits the host to perform all `snap` and `storage` operations on a storage system if the access permissions file is missing.
- ◆ `available-lun-reserve=8`. This variable is used by host platforms that require preparation before they can add LUNs.
- ◆ `enable-implicit-host-preparation=on`. This variable is used by host platforms that require preparation before they can add LUNs.

- ◆ `default-noprompt=off`. This variable enables or disables the `-noprompt` option in the command line.
- ◆ `default-transport="iscsi"`. This variable specifies the transport type to use for storage provisioning, if a decision is needed.
- ◆ `multipathing-type="SANPath"` This variable specifies the multipathing software to use when more than one multipathing solution is available.
- ◆ `prefix-filer-lun=" "`. This variable adds a prefix to all internally-generated filer LUN names.
- ◆ `prepare-lun-count=16`. This variable is used by host platforms that require preparation before they can add LUNs.
- ◆ `snaprestore-snapmirror-check=on`. This variable enables the `snap restore` command to check the SnapMirror destination volume.
- ◆ `space-reservations-enabled=on`. This variable enables space reservations when you use SnapDrive for UNIX to create new LUNs.
- ◆ `use-https-to-filer=off`. This variable specifies that all communication with the storage system be done via HTTPS instead of HTTP.

About this chapter

This chapter provides details about setting your SnapDrive for UNIX configuration options and general information about using SnapDrive for UNIX.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Setting configuration information”](#) on page 62
- ◆ [“Preparing hosts for adding LUNs”](#) on page 85
- ◆ [“Setting up audit, recovery, and trace logging”](#) on page 86
- ◆ [“Setting up autosupport”](#) on page 94
- ◆ [“General steps for executing commands”](#) on page 97

Setting configuration information

Using the snapdrive.conf file

The `snapdrive.conf` file controls the configurable variables available in SnapDrive for UNIX. This file contains a name/value pair for each configurable variable.

You use a text editor to modify this file. SnapDrive for UNIX automatically checks the information in this file each time it starts.

The `snapdrive.conf` file is in the SnapDrive for UNIX installation directory (see the installation instructions for your operating system for the complete path to this directory).

SnapDrive for UNIX also provides some commands you can use to work with this file, such as the `snapdrive config show` command, which displays this file.

Verify the settings in snapdrive.conf

The `snapdrive.conf` file comes with most of the variables set to default values. In most cases, these are the values you should use when you run SnapDrive for UNIX. To complete your SnapDrive for UNIX installation, though, you should verify the default values in the `snapdrive.conf` file. In particular, you should check the values for the following variables:

- ◆ Logging
- ◆ Autosupport
- ◆ Accurate path information for the system

To view the current settings, execute the `snapdrive config show` command.

Some operations available with snapdrive config command

The `snapdrive config` command has several formats. These commands allow you to perform a number of configuration tasks, including working with the `snapdrive.conf` file, setting the user/password information SnapDrive for UNIX needs to access the storage systems, setting access control permissions, and determining LUN configuration information for your system. The following is a brief overview of the `snapdrive config` command formats:

- ◆ `snapdrive config show [host_file_name]`

This command displays the `snapdrive.conf` current name/value pairs and their defaults. If you include the `host_file_name` argument, it writes the information to the file you specify on the host.

- ◆ `snapdrive config list`
`snapdrive config set user_name filename [filename ...]`
`snapdrive config delete filename [filename ...]`

These versions of the `config` command handle user logins for the storage systems.

You can use the `snapdrive config list` command to display the user names that have been configured in SnapDrive for UNIX for accessing storage systems.

The `set` version of the command lets you tell SnapDrive for UNIX which user name/password pair to use to access a storage system while the `delete` version removes the specified user name/password pair from SnapDrive for UNIX.

For more information, see “[Specifying the current login information for storage systems](#)” on page 107.

- ◆ `snapdrive config access {list | show} filename`

This operation lets you display the access control settings for a host trying to access a storage system.

For more information, see “[Setting up access control](#)” on page 101.

- ◆ `snapdrive config prepare luns -count count`
`snapdrive config check luns`

These commands let you prepare the host for creating a specific number of LUNs as well as to determine how many LUNs you could create.

For more information, see “[Preparing hosts for adding LUNs](#)” on page 85.

Note

You only need these commands on Linux and Solaris hosts. AIX and HP-UX hosts do not require any preparation prior to creating LUNs.

Determining options and their default values

The supported configurable items and their default settings can vary across host operating systems and the different versions of SnapDrive for UNIX. An easy way to determine current configurable items and their settings is to execute the `snapdrive config show` command. The following example shows output that

the `config show` option produces when you execute it on a Solaris host. You should run the `snapdrive config show` command to get a current copy of the file for your host platform.

Note

If you are running SnapDrive for UNIX on a different host operating system, some of the defaults may be different. For example, on an AIX or HP-UX host, the default path for a log file is `/var/snapdrive/...` instead of the Linux or Solaris default path of `/var/log/...`.

```
# snapdrive config show
# Snapdrive Configuration
#   file: /opt/NTAPsnapdrive/snapdrive.conf
#   Version 2.0.1 (Change 400099 Built Tue July 12 18:26:25 EST 2005)
#
#
# Default values are shown by lines which are commented-out in this file.
# If there is no un-commented-out line in this file relating to a particular value, then
# the default value represented in the commented-out line is what SnapDrive will use.
#
# To change a value:
#
#   -- copy the line that is commented out to another line
#   -- Leave the commented-out line
#   -- Modify the new line to remove the '#' and to set the new value.
#   -- Save the file and exit
audit-log-file="/var/log/sd-audit.log" # audit log file
trace-log-file="/var/log/sd-trace.log" # trace log file
recovery-log-file="/var/log/sd-recovery.log" # recovery log file
autosupport-enabled=off # Enable autosupport (requires autosupport-filer be set)
autosupport-filer="" # Filer to use for autosupport (filer must be configured for
autosupport)
audit-log-max-size=20480 # Maximum size (in bytes) of audit log file
audit-log-save=2 # Number of old copies of audit log file to save
available-lun-reserve=8 # Number of LUNs for which to reserve host resources
device-retries=3 # Number of retries on a NetApp filer LUN device inquiry
device-retry-sleep-secs=1 # Number of seconds between NetApp filer LUN device inquiry
retries
enable-implicit-host-preparation="on" # Enable implicit host preparation for LUN
creation
filer-restore-retries=140 # Number of retries while doing lun restore
filer-restore-retry-sleep-secs=15 # Number of secs between retries while restoring lun
filesystem-freeze-timeout-secs=300 # File system freeze timeout in seconds
default-noprompt=off # A default value for -noprompt option in the command line
```

```

mgmt-retries=3 # Number of retries on ManageONTAP control channel
mgmt-retry-sleep-secs=2 # Number of seconds between retries on ManageONTAP control
channel
mgmt-retry-sleep-long-secs=90 # Number of seconds between retries on ManageONTAP
control channel (failover error)
prepare-lun-count=16 # Number of LUNs for which to request host preparation
multipathing-type="" # Multipathing software to use when more than one multipathing
solution is available.
PATH="/sbin:/usr/sbin:/bin:/usr/lib/vxvm/bin:/usr/bin" # toolset search path
password-file="/opt/NTAPsnapdrive/.pwfile" # location of password file
prefix-filer-lun="SD_INT" # Add prefix to all internally-generated filer LUN names
recovery-log-save=20 # Number of old copies of recovery log file to save
snapcreate-consistency-retries=3 # Number of retries on snapshot consistency check
failure
snapcreate-must-make-snapinfo-on-qtree=off # snap create must be able to create
snapinfo on qtree
snapdelete-delete-rollback-with-snap=off # Delete all rollback snapshots related to
specified snapshot
snaprestore-delete-rollback-after-restore=on # Delete rollback snapshot after a
successful restore
snaprestore-snapmirror-check=on # Enable snapmirror destination volume check in snap
restore
snaprestore-make-rollback=on # Create snap rollback before restore
snaprestore-must-make-rollback=on # Do not continue 'snap restore' if rollback
creation fails
space-reservations-enabled=on # Enable space reservations when creating new luns
default-transport="iscsi" # Transport type to use for storage provisioning, when a
decision is needed
trace-enabled=on # Enable trace
trace-level=5 # Trace levels: 1=FatalError; 2=AdminError; 3=CommandError; 4=warning,
5=info, 6=verbose
trace-log-max-size=204800 # Maximum size of trace log file in bytes
trace-log-save=5 # Number of old copies of trace log file to save
all-access-if-rbac-unspecified=on # Allow all access if the RBAC permissions file is
missing
use-https-to-filer=off # Communication with filer done via HTTPS instead of HTTP

```

The following table describes the configuration variables in this file.

Option	Description
all-access-if-rbac-unspecified=on	Set this value to either On or Off where ◆ On tells SnapDrive for UNIX to enable all access permissions if there is no access control permissions file on the storage system. This is the default value. ◆ Off specifies that the storage system allow the host only the permissions specified in the access control permissions file. You can specify access control permissions for each host where SnapDrive for UNIX runs by entering the permission string in an access control file. The string you specify controls the which SnapDrive for UNIX snapshot and storage operations that host may perform on a given storage system. (These access permissions do not affect the show/list operations.) If you provide an access control file, this option has no effect. For more information on access control and how to set it up, see “Setting up access control” on page 101.
audit-log-file="/var/log/sd-audit.log"	Specify the location where SnapDrive for UNIX writes the audit log file. The default value depends on your host operating system. The path shown in this example is the default path for a Linux or Solaris host. For an AIX or HP-UX host, the default path is /var/snapdrive/sd-audit.log. See “Changing the defaults for audit logs” on page 90 for more information.

Option	Description
audit-log-max-size=20480	Specify the maximum size, in bytes, of the audit log file. When the file reaches this size, SnapDrive for UNIX renames it and starts a new audit log. The default value is 20,480 bytes. Because SnapDrive for UNIX never starts a new log file in the middle of an operation, the actual size of the file could vary slightly from the value specified here. Note NetApp recommends that you use the default value. If you decide to change the default, keep in mind that having a large number of large log files can take up space on your disk and might eventually affect performance.
audit-log-save=2	Specify how many old audit log files SnapDrive for UNIX should save. After this limit is reached, SnapDrive for UNIX discards the oldest file when it creates a new one. SnapDrive for UNIX rotates this file based on the value you specify for the option audit-log-max-size. (For more information on log rotation, see “Settings affecting log file rotation” on page 88.) The default value is 2. Note NetApp recommends that you use the default value. If you decide to change the default, keep in mind that having a large number of large log files can take up space on your disk and might eventually affect performance.

Option	Description
autosupport-enabled=off	Set this value to either On to enable autosupport or Off to disable it. By default this option is off because you need to perform additional setup steps to enable it, including specifying a value for the autosupport-filer option. NetApp recommends that you enable this option. See “Setting up autosupport” on page 94 for details about how to enable this option.
autosupport-filer=" "	Enter the name or IP address of the storage system that autosupport should use to send the message. To disable this option, leave the storage system name blank. This option is disabled by default because it requires information specific to your setup. NetApp recommends that you enable this option. To enable autosupport, you must enter a value here and also set autosupport-enabled to On. See “Setting up autosupport” on page 94 for details about how to enable this option.

Option	Description
available-lun-reserve=8	Enter the number of LUNs that the host must be prepared to create when the current SnapDrive for UNIX operation completes. If too few operating system resources are available to create the number of LUNs specified, SnapDrive for UNIX requests additional resources, based on the value supplied with the enable-implicit-host-preparation variable. The default value is 8. Note_____ This variable applies only to systems where host preparation is needed before you can create LUNs. Currently, only Linux and Solaris hosts require that preparation.
default-noprompt=Off	Specify whether you want the <code>-noprompt</code> option to be available or not. By default, this option is not available (Off). If you turn this option On, then SnapDrive for UNXI does not prompt you to confirm an action requested by <code>-force</code>.
default-transport="iscsi"	Specify which protocol SnapDrive for UNIX uses as the transport type when creating storage, if a decision is required. The acceptable values are <code>iscsi</code> or <code>fc</code>. The default is <code>iscsi</code>. Note_____ On AIX hosts, you must also make sure the <code>multipathing-type</code> option is set correctly. If you specify <code>"fc"</code> here, you must set <code>multipathing-type</code> to <code>"SANPath"</code>. If you specify <code>"iscsi"</code> here, you must specify <code>"none"</code> as the value for <code>multipathing-type</code>.

Option	Description
device-retries=3	Specify the number of times SnapDrive for UNIX attempts to inquire about the device where the LUN is. The default value is 3. In normal circumstances, the default value should be adequate. Under other circumstances, LUN queries for a <code>snap create</code> operation could fail simply because the storage system is exceptionally busy. If the LUN queries keep failing even though the LUNs are online and correctly configured, you might want to increase the number of retries.
device-retry-sleep-secs=1	Specify the number of seconds SnapDrive for UNIX waits between inquiries about the device where the LUN resides. The default value is 1 second. In normal circumstances, the default value should be adequate. Under other circumstances, LUN queries for a <code>snap create</code> operation could fail simply because the storage system is exceptionally busy. If the LUN queries keep failing even though the LUNs are online and correctly configured, you might want to increase the number of seconds between retries.

Option	Description
enable-implicit-host-preparation="on"	Specify whether SnapDrive for UNIX implicitly requests host preparation for LUNs or simply notifies you that it is required and exits. When this value is set to On, SnapDrive for UNIX implicitly requests the host to make more resources available if there will not be enough resources to create the correct number of LUNs once the current command completes. The number of LUNs being created is specified in the available-lun-reserve variable. When this variable is set to Off, SnapDrive for UNIX informs you if additional host preparation is necessary for the LUN creation and exits. You can then perform the operations necessary to free up resources needed for the LUN creation. For example, you can execute the <code>snapdrive config prepare luns</code> command. Once the preparation is complete, you can re-enter the current SnapDrive for UNIX command. By default, this variable is set to On. Note This variable applies only to systems where host preparation is needed before you can create LUNs. Currently, only Linux and Solaris hosts require that preparation.

Option	Description
filer-restore-retries=140	Specify the number of times SnapDrive for UNIX attempts to restore a snapshot on a storage system if a failure occurs during the restore. The default value is 140. In normal circumstances, the default value should be adequate. Under other circumstances, this operation could fail simply because the storage system is exceptionally busy. If it keeps failing even though the LUNs are online and correctly configured, you might want to increase the number of retries.
filer-restore-retry-sleep-secs=15	Specify the number of seconds SnapDrive for UNIX waits between attempts to restore a snapshot. The default value is 15 seconds. In normal circumstances, the default value should be adequate. Under other circumstances, this operation could fail simply because the storage system is exceptionally busy. If it keeps failing even though the LUNs are online and correctly configured, you might want to increase the number of seconds between retries.
filesystem-freeze-timeout-secs=300	Specify the amount of time, in seconds, that SnapDrive for UNIX should wait when it can't access the file system before trying again. The default value is 300 seconds (5 minutes).
mgmt-retries=3	Specify the number of times SnapDrive for UNIX retries an operation on the ManageONTAP control channel after a failure occurs. The default value is 3.
mgmt-retry-sleep-secs=2	Specify the number of seconds SnapDrive for UNIX waits before retrying an operation on the ManageONTAP control channel. The default value is 2 seconds.

Option	Description
mgmt-retry-sleep-long-secs=90	Specify the number of seconds SnapDrive for UNIX waits before retrying an operation on the ManageONTAP control channel after a failover error occurs. The default value is 90 seconds.
multipathing-type=" "	Specify the multipathing software to use. This option applies only if there is more than one multipathing solution available. If a platform supports only one type of multipathing, it does not use this option. For this release, the AIX platform is the only one where you must set this value. Solaris and HP-UX platforms provide support for only one type of multipathing each. The Linux platform does not support multipathing, so the default value there is “none”. The value you set for AIX depends on which protocol you are using. If you are using AIX with ◆ FCP, set this value to “SANPath”. In addition, set the default-transport option to “fcp”. ◆ iSCSI, set this value to “none”. In addition set the default-transport option to “iscsi”. Note _____ For this release, mulitpathing is not supported on AIX hosts running iSCSI. _____

Option	Description
<code>path="/sbin:/usr/sbin:/bin:/usr/lib/ vxvm/diag.d:/usr/bin:/usr/local/bin"</code>	Specify the search path the system uses to look for tools. NetApp recommends that you verify that this is the correct path for your system. If it is incorrect, change it to the correct path. The default value may vary depending on your operating system. In this case, this is the default path for a Solaris system or an HP-UX system. AIX systems don't use this variable because they process the commands differently.
<code>password-file="/opt/NTAPsnapdrive/.pwfile"</code>	Specify the location of password file for the user logins for the storage systems. The default value depends on your operating system. For example, this is the default path for a Solaris system. Other hosts use a default path of <code>/opt/NetApp/snapdrive/.pwfile</code>.
<code>prefix-filer-lun=" "</code>	Specify the prefix that SnapDrive for UNIX applies to all LUN names it generates internally. The default value for this prefix is the empty string. This variable allows the names of all LUNs created from the current host, but not explicitly named on a SnapDrive for UNIX command line, to share an initial string.

Option	Description
prepare-lun-count=16	Specify how many LUNs SnapDrive for UNIX should prepare to create. SnapDrive for UNIX checks this value when there is a request to prepare the host to create additional LUNs. By default, this variable is set to 16. That means the system will be able to create 16 additional LUNs once the preparation is complete. Note This variable applies only to systems where host preparation is needed before you can create LUNs. Currently, only Linux and Solaris hosts require that preparation.
recovery-log-file="/var/log/sd-recovery.log"	Specify the location where SnapDrive for UNIX writes the recovery log file. The default value depends on your host operating system. The path shown in this example is the default path for a Linux or Solaris host. For an AIX or HP-UX host, the default path is /var/snapdrive/sd-recovery.log. See “Contents of a recovery log” on page 90 for more information.

Option	Description
recovery-log-save=20	Specify how many old recovery log files SnapDrive for UNIX should save. After this limit is reached, SnapDrive for UNIX discards the oldest file when it creates a new one. SnapDrive for UNIX rotates this log file each time it starts a new operation. (For more information on log rotation, see “Settings affecting log file rotation” on page 88.) The default value is 20. Note NetApp recommends that you use the default value. If you decide to change the default, keep in mind that having a large number of large log files can take up space on your disk and might eventually affect performance.
snapcreate-must-make-snapinfo-on-qtree=off	Set this value to On to enable the <code>snapdrive snap create</code> operation to create snapshot information on a qtree. Set it to Off to disable this feature. This option is disabled by default. SnapDrive for UNIX always attempts to write <code>snapinfo</code> at the root of a qtree if the LUNs being snapped are at the root of that qtree. Setting this option to On means that SnapDrive for UNIX fails the <code>snap create</code> operation if it can't write this data. You should only turn this option on if you are mirroring snapshots using Qtree SnapMirror® (QSM) software. Note Snapshots of qtrees work the same way snapshots of volumes do.

Option	Description
snapcreate-consistency-retries=3	Specify the number of times SnapDrive for UNIX attempts a consistency check on a snapshot after it receives a message that a consistency check failed. This option is particularly useful on host platforms that do not include a freeze function. Currently, Linux is the only platform without that function. For more information, see “Architectural limitations for SnapDrive for UNIX” on page 6. The default value for this option is 3.
snapdelete-delete-rollback-with-snap=off	Set this value to On to delete all rollback snapshots related to a specified snapshot. Set it to Off to disable this feature. This option is disabled by default. This option only takes effect during a <code>snap delete</code> operation. is used by the recovery log file if you encounter a problem with an operation. NetApp recommends that you accept the default setting.
snaprestore-delete-rollback-after-restore=on	Set this value to On to delete all rollback snapshots after a successful <code>snap restore</code> operation. Set it to Off to disable this feature. This option is enabled by default. This option is used by the recovery log file if you encounter a problem with an operation. NetApp recommends that you accept the default setting.

Option	Description
snaprestore-make-rollback=on	Set this value to either On to create a rollback snapshot or Off to disable this feature. This option is enabled by default. A rollback is a copy of the data on the storage system that SnapDrive for UNIX makes before it begins a <code>snap restore</code> operation. That way, if a problem occurs during the <code>snap restore</code> operation, you can use the rollback snapshot to restore the data to the state it was in before the operation began. If you do not want the extra security of a rollback snapshot at restore time, set this option to Off. If you like the rollback, but not enough to fail your <code>snap restore</code> operation if you can't make one, set the option <code>snaprestore-must-make-rollback</code> to Off. This option is used by the recovery log file, which you send to technical support if you encounter a problem. See “Contents of a recovery log” on page 90. NetApp recommends that you accept the default setting.

Option	Description
snaprestore-must-make-rollback=on	Set this value to On to fail a <code>snap restore</code> operation if the rollback creation fails or Off to disable this feature. The option is enabled by default. When you set this option to On, SnapDrive for UNIX attempts to make a rollback copy of the data on the storage system before it begins the <code>snap restore</code> operation. If it can't make a rollback copy of the data, SnapDrive for UNIX halts the <code>snap restore</code> operation. If you want the extra security of a rollback snapshot at restore time, but not enough to fail your <code>snap restore</code> operation if you can't make one, set this option to Off. This option is used by the recovery log file if you encounter a problem with an operation. NetApp recommends that you accept the default setting.
snaprestore-snapmirror-check=on	Set this value to on to enable the <code>snap restore</code> command to check the SnapMirror destination volume. If it is Off, the <code>snap restore</code> command is unable to check the destination volume. The default value is On.
space-reservations-enabled=on	Set this value to On to tell SnapDrive for UNIX to enable space reservations when it creates new LUNs. Set it to off if you do not want space reservations enabled. The default value is On.

Option	Description
trace-enabled=on	Set this value to On to enable the trace log file or Off to disable it. This file is used by NetApp. The default value is On. Enabling this file does not affect performance. For more information, see “About the trace log file” on page 92.
trace-level=5	Specifies the types of messages SnapDrive for UNIX writes to the trace log file. This option accepts the following values: ◆ 1—Record fatal errors. ◆ 2—Record admin errors. ◆ 3—Record command errors. ◆ 4—Record warnings. ◆ 5—Record information messages. ◆ 6—Record in verbose mode. The default value is 5.
trace-log-file="/var/log/sd-trace.log"	Specify the location where SnapDrive for UNIX writes the trace log file. The default value depends on your host operating system. The path shown in this example is the default path for a Linux or Solaris host. For an AIX or HP-UX host, the default path is /var/snapdrive/sd-trace.log. See “About the trace log file” on page 92 for more information.

Option	Description
trace-max-size=204800	Specify the maximum size of the trace log file in bytes. When the file reaches this size, SnapDrive for UNIX renames it and starts a new trace log. The default value is 204,800. Because SnapDrive for UNIX never starts a new log file in the middle of an operation, the actual size of the file could vary slightly from the value specified here. Note_____ NetApp recommends that you use the default value. If you decide to change the default, keep in mind that having a large number of large log files can take up space on your disk and might eventually affect performance. _____
trace-log-save=5	Specify how many old trace log files SnapDrive for UNIX should save. After this limit is reached, SnapDrive for UNIX discards the oldest file when it creates a new one. SnapDrive for UNIX rotates this file based on the value you specify for the trace-max-size option. (For more information on log rotation, see “Settings affecting log file rotation” on page 88.) The default value is 5. Note_____ NetApp recommends that you use the default value. If you decide to change the default, keep in mind that having a large number of large log files can take up space on your disk and might eventually affect performance. _____

Option	Description
use-https-to-filer=off	Specify whether you want SnapDrive for UNIX to use SSL encryption (HTTPS) when it communicates with the storage system. By default, this value is set to off. Note_____ If you are using a version of Data ONTAP prior to 7.0, you might see slower performance with HTTPS enabled. This is not an issue if you are running Data ONTAP 7.0. _____ For more information on the security features available with SnapDrive for UNIX, see the chapter “Setting Up Security Features” on page 99.

Setting values in snapdrive.conf

To change the values in the snapdrive.conf file or add new name/value pairs, complete the following steps.

Step	Action
1	Log on as root.
2	Make a backup of the snapdrive.conf file.
3	Open the snapdrive.conf file in a text editor.

Step	Action
4	Make the changes to the file. To add a name/value pair, use the following format: <pre>config-option-name=value # optional comment</pre> <i>config-option-name</i> is the name of the option you want to configure; for example, audit-log-file. <i>value</i> is the value you want to assign to this option. If you want to include a comment with the name/value pair, precede the comment with a pound sign (#). You should enter only one name/value pair per line. If the name or the value uses a string, enclose the string in either single (') or double (") quotes. You can place the quotes around either the entire name/value pair or just the value. The following are three examples of how you can use quotes and comments with name/value pairs: <pre>"config-option-one=string with white space" # double quotes around the pair config-option-two="string with white space" # double quotes around the value config-option-2B='string with white space' # single quotes around the value</pre>
5	To modify an existing name/value pair, replace the current value with the new value. The best way to do this is to perform the following steps; 1. Comment out the line you want to modify. 2. Copy the commented-out line. 3. Un-comment the copy by removing the pound sign. 4. Modify the value. This way you always have a record of the default value in the file. If you want to specify a blank value (for example, you want to disable the audit log file), enter a pair of quotation marks (" ").
6	Save the file after you make your changes. SnapDrive for UNIX automatically checks this file each time it start. Your changes take effect the next time it starts.

Checking your version of SnapDrive

To determine which version of SnapDrive you’re using, complete the following step.

Step	Action
1	Enter the following command: <pre>snapdrive version</pre> Example: SnapDrive for UNIX displays its version information when you enter this command. <pre># snapdrive version Version 2.0 #</pre> Note_____ The only argument this command accepts is -v (verbose), which displays additional version details. If you include additional arguments, SnapDrive for UNIX displays a warning and then the version number. _____

Preparing hosts for adding LUNs

Checking host information

If your operating system requires that you prepare it before you create new LUNs, you can use the `snapdrive config` command. This command lets you check information about how many LUNs can be created on a storage system that is mapped to your host.

Not all host platform operating systems require that you prepare the host. Currently, these commands are required only on Linux and Solaris hosts.

Determine how many LUNs could be created

SnapDrive for UNIX lets you determine how many LUNs could be created on the host without exceeding a host-local limit. Use the following command to determine this value:

```
snapdrive config check luns
```

On a Linux host, this command checks the existing `/dev/sg` files to determine how many are unused.

On a Solaris host, this command scans `/kernel/drv/sd.conf` to determine how many unused entries it has that would be suitable for LUNs.

Add host entries for new LUNs

You can also make sure the host is prepared for the creation of a specific number of new LUNs. These LUNs will reside on a storage system that is currently mapped to the host. Use the following command:

```
snapdrive config prepare luns -count count
```

count is the number of new LUNs for which you want the host to be prepared.

On Linux, this command adds a new `/dev/sg` device file for each potential LUN for which a device file is not currently available.

On Solaris, this command adds entries to the file `/kernel/drv/sd.conf`, if necessary, for each potential new LUN that doesn't have an entry. It also generates an entry for each SCSI target to which the storage system is mapped. On Solaris 8, you must reboot the host after adding `sd.conf` entries. This command displays a warning whenever a reboot is necessary.

Setting up audit, recovery, and trace logging

Supported logs

SnapDrive for UNIX supports the following log files:

- ◆ Audit log. SnapDrive for UNIX logs all `snapdrive` commands and their return codes to an audit log. SnapDrive for UNIX makes an entry when you initiate a command and another when the command is complete. That entry includes both the status of the command and the completion time.
- ◆ Recovery log. Some SnapDrive for UNIX operations have the potential to leave the system in an inconsistent or less usable state if interrupted. This could happen if a user kills the program, or if the host crashes in the middle of an operation. The recovery log contains the steps of a `snap restore` operation. It documents the steps that were taken and the progress made. This way, technical support can assist you with the manual recovery process.
- ◆ Trace log. SnapDrive for UNIX reports information useful for diagnosing problems. If you have a problem, technical support might request this log file.

Enabling and disabling log files

To enable or disable each of these log files in the `snapdrive.conf` file, complete the following steps.

Step	Action
1	Log in as root.
2	Open the <code>snapdrive.conf</code> file in a text editor.

Step	Action	
3	If you want to ...	Then ...
	Enable a log file	Specify a file name as the value in the name/value pair of the log file you want to enable. SnapDrive for UNIX only writes log files if it has the name of a file to write to. The default names for the log files are as follows: ◆ Audit log: sd-audit.log ◆ Recovery log: sd-recovery.log ◆ Trace log: sd-trace.log Note_____ The path to these files may vary depending on your host operating system. For example, on AIX or HP-UX host, the default path for a log file is /var/snapdrive/...; on a Linux or Solaris host the default path is /var/log/... .
	Disable a log file	Do not enter a value for the log file name parameter. If you do not supply a value, there is no file name to which SnapDrive for UNIX can write the log information. Example: This example disables the audit log file. <pre>audit-log-file=" "</pre> Note_____ NetApp strongly recommends that you leave all log files enabled.
4	Save the file after you make all your changes. SnapDrive for UNIX automatically checks this file each time it starts, so your changes take effect the next time it starts.	

Settings affecting log file rotation

The values you specify in the `snapdrive.conf` file enable automatic log file rotations. You can change these values, if necessary, by editing the `snapdrive.conf` options. The following options affect log file rotation:

- ◆ `audit-log-max-size`
- ◆ `audit-log-save`
- ◆ `trace-max-size`
- ◆ `trace-log-save`
- ◆ `recovery-log-save`

Note

For information about the default values for these options, see [“Determining options and their default values”](#) on page 63.

With automatic log rotation, SnapDrive for UNIX keeps old log files until it reaches the limit specified in the `<audit/trace/recovery>-log-save` option. Then it deletes the oldest log file.

SnapDrive for UNIX tracks which file is oldest by assigning the file the number 0 when it creates the file. Each time it creates a new file, it increments by 1 the number assigned to each of the existing log files. When a log file’s number reaches the save value, SnapDrive for UNIX deletes that file

Example: This example uses the `ls` command to produce information about the log files on the system. The `snapdrive.conf` file has the following settings:

```
audit-log-max-size=20480 # Maximum size of audit log file
audit-log-save=2 # Number of old copies of audit log file to save
trace-log-max-size=204800 # Maximum size of trace log file
trace-log-save=5 # Number of old copies of trace log file to save
recovery-log-save=20 # Number of old copies of recovery log file
to save
```

Based on those settings, you would see the following information on log files:

```
# ls -l /var/log/sd*
-rw-r--r-- 1 root other 12247 Mar 13 13:09 /var/log/sd-audit.log
-rw-r--r-- 1 root other 20489 Mar 12 16:57 /var/log/sd-audit.log.0
-rw-r--r-- 1 root other 20536 Mar 12 03:13 /var/log/sd-audit.log.1
-rw-r--r-- 1 root other 3250 Mar 12 18:38 /var/log/sd-recovery.log.1
-rw-r--r-- 1 root other 6250 Mar 12 18:36 /var/log/sd-recovery.log.2
-rw-r--r-- 1 root other 6238 Mar 12 18:33 /var/log/sd-recovery.log.3
-rw-r--r-- 1 root other 191704 Mar 13 13:09 /var/log/sd-trace.log
-rw-r--r-- 1 root other 227929 Mar 12 16:57 /var/log/sd-trace.log.0
-rw-r--r-- 1 root other 213970 Mar 12 15:14 /var/log/sd-trace.log.1
-rw-r--r-- 1 root other 261697 Mar 12 14:16 /var/log/sd-trace.log.2
-rw-r--r-- 1 root other 232904 Mar 12 14:15 /var/log/sd-trace.log.3
-rw-r--r-- 1 root other 206905 Mar 12 14:14 /var/log/sd-trace.log.4
```

As shown in this example, the maximum size for each log indicates roughly what size it will be when SnapDrive for UNIX rotates it. SnapDrive for UNIX never rotates a log in the middle of an operation, so the sizes may vary from the default sizes.

Contents of an audit log

The audit log shows information about commands you issued with SnapDrive for UNIX. It maintains a history of the following information:

- ◆ The commands issued.
- ◆ The return value from those commands.
- ◆ The user ID of the user who invoked the command.
- ◆ A timestamp indicating when the command started (with no return code) and another timestamp indicating when the command finished (with a return code). The audit log record shows only information about snapdrive usage (issued commands).

An audit log file contains the following information.

Field	Description
uid	user ID
gid	group ID
msgText	message text
returnCode	return code from a command

Changing the defaults for audit logs

The snapdrive.conf file enables you to set the following parameters for audit logging.

Note

See “[Determining options and their default values](#)” on page 63 for more information about the default values for these options.

- ◆ The name of the file containing the audit log.
- ◆ The maximum size of the audit log file. The default size is 20K. After the file size reaches the value specified here, SnapDrive for UNIX renames the current audit log file by adding an arbitrary number to the name. Then it starts a new audit file using the name specified by the audit-log-file value.
- ◆ The maximum number of old audit files that Snapdrive saves. The default is 2.

Example of an audit log file

The following example shows an audit log file:

```
2501: Begin uid=0 gid=1 15:35:02 03/12/04 snapdrv snap create -dg
rdg -snapname snap_rdg1
2501: Status=0 15:35:07 03/12/04
2562: Begin uid=0 gid=1 15:35:16 03/12/04 snapdrv snap create -dg
rdg -snapname snap_rdg1
2562: FAILED Status=4 15:35:19 03/12/04
```

The first pair of lines in this example show an operation that succeeded, as indicated by the "Status=0" line.

The second pair of lines indicates an operation that failed. The return code of 4 means “already exists.” If you look at the two command lines, you can see that the first created a snapshot called snap_rdg1. The second line attempted to do the same, but the name already existed, so the operation failed.

Contents of a recovery log

If SnapDrive is halted using the Ctrl-C key sequence to kill the program, or if the host or storage system crashes in the middle of an operation, the system might not be able to recover automatically.

So, during any operation that, if interrupted, could leave the system in an inconsistent state, SnapDrive for UNIX writes information to a recovery log file. If a problem occurs, you can send this file to technical support so they can assist you in recovering the system’s state.

The recovery log utility records the commands that were issued in the process of the operation. Each is marked with an `operation_index` (a number that uniquely identifies the operation being executed), followed by the date/time stamp and the message text).

Changing the defaults for the recovery logs

The `snapdrive.conf` file enables you to set the following parameters for recovery logging.

Note

See “[Determining options and their default values](#)” on page 63 for more information about the default values for these options.

- ◆ The name of the file containing the recovery log, such as `recoverylog`.
- ◆ The maximum number of old recovery files that Snapdrive for UNIX saves. The default is 20. SnapDrive for UNIX keeps this number of recovery logs in case the problem with the process isn’t immediately discovered. SnapDrive for UNIX starts a new recovery log file each time it completes an operation. It renames the previous one by adding an arbitrary number to the name, such as `recoverylog.0`, `recoverylog.1` and so on.

Note

The size of the recovery log file depends on the operation being performed. Each recovery log contains information about a single operation. When that operation is complete, SnapDrive for UNIX starts a new recovery log, regardless of how large the previous file was. As a result, there is no maximum size for a recovery log file.

Example of a recovery log file

The following is an example of entries in a recovery log where SnapDrive for UNIX has restored two snapshots before the operations halted. At this point, you would send this recovery log file to technical support for assistance in restoring the remaining snapshots.

```
6719: BEGIN 15:52:21 03/09/04 snapdrive snap restore -dg jssdg -
snapname natasha:/vol/vol0:abort_snap_restore
6719: BEGIN 15:52:27 03/09/04 create rollback snapshot:
natasha:/vol/vol0:abort_snap_restore.RESTORE_ROLLBACK_03092004_155
225
6719: END 15:52:29 03/09/04 create rollback snapshot:
natasha:/vol/vol0:abort_snap_restore.RESTORE_ROLLBACK_03092004_155
225 successful
6719: BEGIN 15:52:29 03/09/04 deactivate disk group: jssdg
6719: BEGIN 15:52:29 03/09/04 stop host volume:
/dev/vx/dsk/jssdg/jvol_1
6719: END 15:52:30 03/09/04 stop host volume:
/dev/vx/dsk/jssdg/jvol_1 successful
6719: BEGIN 15:52:30 03/09/04 unmount file system: /mnt/demo_fs
6719: END 15:52:30 03/09/04 unmount file system: /mnt/demo_fs
successful
6719: BEGIN 15:52:30 03/09/04 stop host volume:
/dev/vx/dsk/jssdg/jvol_2
6719: END 15:52:30 03/09/04 stop host volume:
/dev/vx/dsk/jssdg/jvol_2 successful
6719: BEGIN 15:52:30 03/09/04 deport disk group: jssdg
6719: END 15:52:30 03/09/04 deport disk group: jssdg successful
6719: END 15:52:30 03/09/04 deactivate disk group: jssdg
successful
6719: BEGIN 15:52:31 03/09/04 SFSR of LUN: /vol/vol0/lun1 from
snapshot: abort_snap_restore
6719: END 15:52:31 03/09/04 SFSR of LUN: /vol/vol0/lun1 from
snapshot: abort_snap_restore successful
6719: BEGIN 15:52:47 03/09/04 SFSR of LUN: /vol/vol0/lun2 from
snapshot: abort_snap_restore
6719: END 15:52:47 03/09/04 SFSR of LUN: /vol/vol0/lun2 from
snapshot: abort_snap_restore successful
```

About the trace log file

This file is for NetApp's use in cases where there is a problem that needs debugging. Enabling it does not affect system performance. By default, this file is enabled. You can disable it by setting the snapdrive.conf option trace-enabled to off.

Changing the defaults for the trace logs

The `snapdrive.conf` file enables you to set the following parameters for trace logging.

Note

See “[Determining options and their default values](#)” on page 63 for more information about the default values for these options.

- ◆ The name of the file containing the trace log.
- ◆ The maximum size of the trace log file. The default size is 204,800 bytes. After the file size reaches the value specified here, SnapDrive for UNIX renames the current trace log file by adding an arbitrary number to the name. Then it starts a new trace file using the name specified by the `trace-log-file` value.
- ◆ The maximum number of old trace files that SnapDrive for UNIX saves. The default is 5.
- ◆ The types of messages that SnapDrive for UNIX writes to the trace log file. By default, the trace log file contains fatal errors, admin errors, command errors, warnings and information messages.

Setting up autosupport

Understanding autosupport

NetApp provides autosupport with its storage systems as a way to provide better service to you, should you have a problem with your system. With autosupport, you can configure your storage system to send an e-mail message to technical support when an error occurs. Then, if you call in with an issue, technical support has information about your filers and configuration and can more quickly help you to solve the problem.

No secure information is ever sent using autosupport.

How SnapDrive for UNIX uses autosupport

For this release, SnapDrive for UNIX only sends an autosupport message from the storage system to NetApp the first time you execute it after a system reboot. It sends one message for each host reboot from the host that rebooted. At this time, it does not send a message when an error condition occurs.

Note

To use this feature, you must have a user login configured for the storage system and you must enable autosupport in the `snapdrive.conf` file. See [“Enabling autosupport”](#) on page 95

The autosupport feature in SnapDrive for UNIX logs into the storage system you configured for autosupport in the `snapdrive.conf` file. It uses that storage system to send an autosupport message to NetApp. This message specifies the following information:

- ◆ SnapDrive for UNIX version
- ◆ Message status: 3 for a warning or information message
- ◆ Host name
- ◆ Host operating system
- ◆ Host operating system release number
- ◆ Host operating system version

If the storage system specified for autosupport in the `snapdrive.conf` file cannot send an autosupport message to NetApp, SnapDrive for UNIX does not log an error to `syslog`. This information is only written to the internal SnapDrive for UNIX log files.

Examples of an autosupport message

The substance of the autosupport message is essentially the same regardless of your operating system. The details of the message, such as information on your operating system, vary according to your system setup.

HP-UX example: The following example is a message sent from a host named DBserver that is running release 5B.11.22 of HP-UX. This is an informational message, as indicated by the number 3 in parentheses (3).

```
snapdrive: 2.0.1 (3)
general: host_name = DBserver, host_os=HP-UX,
host_os_release=B.11.22, host_os_version=U
```

Linux example: The following example is a message sent from a host named DBserver that is running Linux. This is an informational message, as indicated by the number 3 in parentheses (3).

```
snapdrive: 2.0.1 (3)
general: host_name = DBserver, host_os=Linux,
host_os_release=2.4.21-9.ELsmp, host_os_version=#1 SMP Thu Jan 8
17:08:56 EST 2004.
```

Solaris example: The following example is a message sent from a host named DBserver that is running release 5.9 of Solaris. This is an informational message, as indicated by the number 3 in parentheses (3).

```
snapdrive: 2.0.1 (3)
general: host_name = DBserver, host_os=SunOS, host_os_release=5.9,
host_os_version=Generic
```

Enabling autosupport

To enable autosupport, complete the following steps.

Step	Action
1	Log in as root.
2	Set user logins for the storage systems you want SnapDrive for UNIX to access. See “ Specifying the current login information for storage systems ” on page 107.

Step	Action
3	Modify the snapdrive.conf file to enable autosupport. 1. Set the autosupport-enabled parameter to 1. It should read as follows: <code>autosupport-enabled=1</code> 2. Specify the name or IP address of a storage system that you want autosupport to contact and use for sending the messages. For example, if you want autosupport to use the storage system named toaster, the name/value pair would read as follows: <code>autosupport-filer=toaster</code> 3. Save the snapdrive.conf file. SnapDrive for UNIX automatically checks this file each time it starts, so your changes take effect the next time it starts.

General steps for executing commands

What you can do with the commands

SnapDrive for UNIX supports snapshot management (`snap` commands) and storage provisioning (`storage` and `host` commands). In addition, you can use the `config` command to monitor or control some of SnapDrive for UNIX's configuration details.

Running SnapDrive for UNIX from the command line

You must be logged in as root to run SnapDrive for UNIX. You execute SnapDrive for UNIX from the CLI, either directly from a shell prompt or from scripts. When you execute SnapDrive for UNIX from the command line, keep the following in mind:

- ◆ Enter each command on a separate line.
- ◆ If you enter a command without arguments, SnapDrive for UNIX gives you a usage line with examples of commands.
- ◆ Each SnapDrive for UNIX command must use the following format:

```
snapdrive type_name operation_name [<keyword/options>  
<arguments>]
```

type_name specifies the type of task you want to perform. You can perform the following types of tasks:

- ❖ Snapshot management. Use the type name `snap` with the `snapdrive` command to indicate you are performing a snapshot management operation.

```
snapdrive snap operation <arguments>
```
- ❖ Display storage provisioning and management information. Use the type name `storage` with the `snapdrive` command to indicate you are performing an operation dealing with storage on the filer.

```
snapdrive storage operation <arguments>
```
- ❖ Connect or disconnect the host. Use the type name `host` with the `snapdrive` command to indicate you are performing an operation dealing only with the host-side entities. The host commands allow you to connect storage on the storage system to the host or disconnect that storage from the host. They do not affect the storage on the storage system. They affect only the visibility of the storage on the host.

```
snapdrive host operation <arguments>
```
- ❖ Specify or view behavioral options. Use the type name `config` with the `snapdrive` command to indicate you want to create, modify, or delete a

user login for a storage system. You also use the config type name to get SnapDrive for UNIX to view the snapdrive.conf file, view access permissions, or prepare the host for adding LUNs.

```
snapdrive config operation <arguments>
```

- ❖ Determine the version number of SnapDrive for UNIX.

```
snapdrive version <arguments>
```

operation_name specifies the operation you want to perform. Each command *type* has certain operations you can perform. The snapshot management operations are create, rename, restore, show, connect, disconnect, and delete. The storage operations are create, connect, disconnect, resize, show, and delete. The configuration operations are access, show, list, set, delete, and prepare and check luns. And the host operations are connect and disconnect.

<*keyword/options*> are the keywords that you can use to specify information corresponding to the host and storage system objects with which you are working, as well as options you can specify to get information such as verbose output. (See “[SnapDrive for UNIX keywords and arguments](#)” on page 247 for more information.)

<*arguments*> are the values you supply with the keywords that specify the target. (See “[SnapDrive for UNIX keywords and arguments](#)” on page 247 for a list of the acceptable arguments and the value you can supply for them.)

Frequently used command-line options

Some frequently used command-line options are the following:

- ◆ `-force (-f)`. This option forces SnapDrive for UNIX to attempt an operation that it ordinarily would not undertake. SnapDrive for UNIX prompts you to confirm that you want it to perform this operation before it attempts the operation.
- ◆ `-noprompt`. This option keeps SnapDrive for UNIX from prompting you to confirm the operation. You should include this option if you’re running SnapDrive for UNIX in a scripted environment where there is no person to confirm the operation.
- ◆ `-verbose (-v)`. This option displays detailed, or verbose, output.

These options appear in the command examples in the following chapters.

About this chapter

This chapter contains information about the security features available in SnapDrive for UNIX and how to access them.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Setting up access control”](#) on page 101
- ◆ [“View the current access control settings”](#) on page 105
- ◆ [“Specifying the current login information for storage systems”](#) on page 107
- ◆ [“Enabling SSL encryption”](#) on page 111

Overview of SnapDrive for UNIX's security features

Security features provided by SnapDrive for UNIX

SnapDrive for UNIX provides certain features to allow you to work with it in a more secure fashion. These features include letting you:

- ◆ Set up access control permissions
- ◆ Specify login information for the storage systems
- ◆ Specify that SnapDrive for UNIX use HTTPS (secure socket layer)

The access control feature lets you specify which operations a host running SnapDrive for UNIX can perform on a storage system. You set these permissions individually for each host. For more information, see “[Setting up access control](#)” on page 101.

In addition, to allow SnapDrive for UNIX to access a storage system, you must supply the login name and password for that storage system. For more information, see “[Specifying the current login information for storage systems](#)” on page 107.

The HTTPS feature lets you specify that you use it for all interactions with the storage system via the ManageONTAP interface, including sending the passwords.

Together these features give you more control over which users can perform operations on a storage system and from which host.

Setting up access control

Access control in SnapDrive for UNIX

SnapDrive for UNIX lets you control the level of access that each host has to each storage system to which the host is connected. This access level indicates which operations the host is allowed to perform when it targets a given storage system.

With the exception of the `show` and `list` operations, the access control permissions can affect all `snap` and `storage` operations.

Available access control levels

You can set the following access levels:

- ◆ NONE – The host has no access to the storage system.
- ◆ SNAP CREATE – The host can create snapshots.
- ◆ SNAP USE – The host can delete and rename snapshots.
- ◆ SNAP ALL – The host can create, restore, delete, and rename snapshots.
- ◆ STORAGE CREATE DELETE – The host can create, resize and delete storage.
- ◆ STORAGE USE – The host can connect and disconnect storage.
- ◆ STORAGE ALL – The host can create, delete, connect, and disconnect storage.
- ◆ ALL ACCESS– The host has access to all these SnapDrive for UNIX operations.

Each level is distinct. If you specify permission for only certain operations, then SnapDrive for UNIX can execute only those operations. For example, if you specify STORAGE USE, the host can use SnapDrive for UNIX to connect and disconnect to storage, but it cannot perform any other operations governed by access control permissions.

How access control works

SnapDrive for UNIX determines the access control permissions by checking the storage system for a permissions file in the root volume of the storage system. This file is in the directory `sdprbac` (SnapDrive Permissions Roles-Based Access Control). The file name is `sdhost-name.prbac`, where *host-name* is the name of the host to which the permissions apply. You can have a permissions file for each host attached to the storage system.

Note

To determine the name of host, run the `hostname` command.

If the file is empty, unreadable, or has an invalid format, SnapDrive for UNIX does not grant the host access permission to any of the operations.

If the file is missing, SnapDrive for UNIX checks the variable `all-access-if-rbac-unspecified` in the `snapdrive.conf` file. If this variable is set to on (the default), it allows all hosts complete access to all these operations on that storage system. If it is set to off, SnapDrive for UNIX denies the host permission to perform any operations governed by access control on that storage system.

Setting access control permissions

To set the access control permissions for a storage system, complete the following steps.

Step	Action
1	Log in as root on the host.
2	On the storage system, create the directory <code>sdprbac</code> in the root volume of the target storage system. Note— One way to make the root volume accessible is to mount the volume using NFS.
3	In this directory, create the permissions file. Make sure the following is true: ◆ The file must be named <code>sdhost-name.prbac</code> where <i>host-name</i> is the name of the host for which you are specifying access permissions. ◆ The file must be read-only. This ensures that SnapDrive for UNIX can read it, but that it cannot be modified. Example: To give a host named <code>dev-sun1</code> access permission, you would create the following file on the storage system. <pre>/vol/vol0/sdprbac/sddev-sun1.prbac</pre>

Step	Action
4	Set the permissions in the file for that host. You must use the following format for the file. ◆ You can specify only one level of permissions. To give the host full access to all these operations, enter the string ALL ACCESS. ◆ The permission string must be the first thing in the file. The file format is invalid if the permission string is not on the first line. ◆ Permission strings are case insensitive. ◆ No white space can precede the permission string. ◆ No comments are allowed. The valid permission strings are: <pre data-bbox="541 673 817 994"> NONE SNAP CREATE SNAP USE SNAP ALL STORAGE CREATE DELETE STORAGE USE STORAGE ALL ALL ACCESS </pre> These strings allow the following access: ◆ NONE – The host has no access to the storage system. ◆ CREATE SNAP – The host can create snapshots. ◆ SNAP USE – The host can delete and rename snapshots. ◆ SNAP ALL– The host can create, restore, delete, and rename snapshots. ◆ STORAGE CREATE DELETE – The host can create, resize, and delete storage. ◆ STORAGE USE – The host can connect and disconnect storage. ◆ STORAGE ALL – The host can create, delete, connect, and disconnect storage. ◆ ALL ACCESS– The host has access to all the SnapDrive for UNIX operations.

Step	Action
	Each of these permission strings is discrete. If you specify SNAP USE, the host can delete or rename snapshots, but it cannot create snapshots or restore snapshots or perform any storage provisioning operations. Regardless of the permissions you set, the host can perform show and list operations.
5	Verify the access permissions by entering the following command: <pre>snapdrive config access show <i>filer-name</i></pre>

View the current access control settings

About viewing the access control permissions

You use the `snapdrive config access` command to display information about the permissions available for a host on a specific storage system.

Viewing the access control permissions

To view the access control permissions, complete the following step.

Step	Action
1	Execute the <code>snapdrive config access show</code> command. This command has the following format: <code>snapdrive config access {show list} <i>filer_name</i></code> You can use the same arguments regardless of whether you enter the <code>show</code> or <code>list</code> version of the command. Note For details about using the options and arguments in this command line, see the section “SnapDrive for UNIX keywords and arguments” on page 247. Result: This operation displays information about the access permissions available for that host. Example 1: This command line checks the storage system <code>toaster</code> to determine which permissions the host has. Based on the permissions shown, the permissions for the host on this storage system are <code>SNAP ALL</code>. <pre># snapdrive config access show toaster This host has the following access permission to filer, toaster: SNAP ALL Commands allowed: snap create snap restore snap delete snap rename #</pre>

Step	Action
	Example 2: In this example, the permissions file is not on the storage system, so SnapDrive for UNIX checks the variable <code>all-access-if-rbac-unspecified</code> in the <code>snapdrive.conf</code> file to determine which permissions the host has. This variable is set to <code>on</code>, which is equivalent to creating a permissions file with the access level set to <code>ALL ACCESS</code>. <pre># snapdrive config access list toaster This host has the following access permission to filer, toaster: ALL ACCESS Commands allowed: snap create snap restore snap delete snap rename storage create storage resize snap connect storage connect storage delete snap disconnect storage disconnect #</pre> Example 3: This example shows the kind of message you would receive if there is not a permissions file on the storage system toaster and the variable <code>all-access-if-rbac-unspecified</code> in the <code>snapdrive.conf</code> file is set to <code>off</code>. <pre># snapdrive config access list toaster Unable to read the access permission file on filer, toaster. Verify that the file is present. Granting no permissions to filer, toaster. #</pre>

Specifying the current login information for storage systems

About using current logins for storage systems

A user name/password enables SnapDrive for UNIX to access each storage system. It also provides a level of security because, in addition to being logged in as root, the person running SnapDrive for UNIX must supply the correct user name/password when prompted for it. If a login is compromised, you can delete it and set a new user login.

You created the user login for each storage system when you set it up. For SnapDrive for UNIX to work with the storage system, you must supply it with this login information. Depending on what you specified when you set up the storage systems, each storage system could use either the same login or a unique login.

SnapDrive for UNIX stores these logins and passwords in encrypted form on each host. You can specify that it encrypt this information when you send it across the wire by setting the snapdrive.conf variable `use-https-to-filer=on`. For more information, see “[Enabling SSL encryption](#)” on page 111.

Specifying login information

To specify the user login information for a storage system, complete the following steps.

Note—
Depending on what you specified when you set up the storage system, each storage system could use either the same user name/password or a unique user name/password. If all the storage systems use the same user name/password information, you only need to perform the following steps once. If not, repeat these steps for each storage system.

Step	Action
1	Log in as root.

Step	Action
2	Enter the following command: <pre>snapdrive config set user_name filename [filename ...]</pre> <i>user_name</i> is the user name that was specified for that storage system when you first set it up. If you don't include a value for it, SnapDrive for UNIX prompts you for one. <i>filename</i> is the name of the storage system. You can enter multiple storage system names on one command line if they all have the same user login/password. You must enter the name of at least one storage system.
3	At the prompt, enter the password, if there is one. Note If no password was set, press Enter (i.e., the null value) when prompted for a password. Example 1: This example sets up a user called root for a storage system called toaster. <pre># snapdrive config set root toaster Password for root: Retype Password: Access to filer toaster has been configured for user root</pre> Example 2: This example sets up one user called root for three storage systems. <pre># snapdrive config set root toaster oven broiler Password for root: Retype Password: Access to filers toaster, oven, and broiler has been configured for user root</pre>
4	If you have another storage system with a different user name/password, repeat these steps.

Verifying storage system user names associated with SnapDrive for UNIX

You can verify which user name SnapDrive for UNIX has associated with a storage system by executing the `snapdrive config list` command.

Note—
This command does not query the storage system to determine whether addition user names have been configured for it. Nor does it display the password associated with a storage system.

To execute `config list`, complete the following step.

Step	Action
1	Log in as root.
2	Enter the following command: <code>snapdrive config list</code> This command displays the user name/storage system pairs for all systems that have users specified within SnapDrive for UNIX. It does not display the passwords for the storage systems. Example: This example displays the users associated with the storage systems <code>rapunzel</code> and <code>mediumfiler</code>. <pre># snapdrive config list user name filer name ----- rumplestiltskins rapunzel longuser mediumfiler</pre>

Deleting a user login for a storage system

To delete a user login for one or more storage systems, complete the following steps.

Step	Action
1	Log in as root.

Step	Action
2	Enter the following command: <pre data-bbox="494 288 1126 314">snapdrive config delete <i>filename</i> [<i>filename</i> ...]</pre> <i>filename</i> is the name of the storage system for which you want to delete the user login information. Result: SnapDrive for UNIX removes the user name/password login information for the storage systems you specify. Note_____ To enable SnapDrive for UNIX to access the storage system, you must specify a new user login.

Enabling SSL encryption

Enabling SnapDrive for UNIX to use HTTPS

You can set up SnapDrive for UNIX to use HTTPS (secure socket layer encryption) when it sends information about storage system user names and passwords across the network.

SnapDrive for UNIX uses HTTP instead of HTTPS as its default when it sends storage system login information across the network. By changing the `use-https-to-filer` variable in the `snapdrive.conf` file, you can tell SnapDrive for UNIX to use HTTPS (SSL encryption) to communicate with the storage system.

Note Using HTTPS might result in slower performance with versions of Data ONTAP prior to 7.0. It does not affect performance with Data ONTAP. 7.0.

Setting up HTTPS

To set up SnapDrive for UNIX to use HTTPS, complete the following steps:

Note For more information on the configuration file, see [“Setting configuration information”](#) on page 62.

Step	Action
1	Log on as root.
2	Make a backup of the <code>snapdrive.conf</code> file.
3	Open the <code>snapdrive.conf</code> file in a text editor.

Step	Action
4	Change the value for the use-https-to-filer variable to on. <pre>use-https-to-filer variable=on</pre> Tip: A good practice any time you modify the snapdrive.conf file is to perform the following steps: 1. Comment out the line you want to modify. 2. Copy the commented-out line. 3. Un-comment the copy by removing the pound sign. 4. Modify the value. This way you always have a record of the default value in the file.
5	Save the file after you make your changes. SnapDrive for UNIX automatically checks this file each time it starts, so your changes take effect the next time it starts.

About this chapter

This chapter provides details about creating and using SnapDrive for UNIX snapshots.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Creating snapshots”](#) on page 114
- ◆ [“Displaying information about snapshots”](#) on page 120
- ◆ [“Renaming a snapshot”](#) on page 127
- ◆ [“Restoring a snapshot”](#) on page 129
- ◆ [“Connecting to a snapshot”](#) on page 137
- ◆ [“Disconnecting a snapshot”](#) on page 144
- ◆ [“Deleting snapshots”](#) on page 146

Creating snapshots

About creating snapshots

You use the `snapdrive snap create` command to create snapshots, which are point-in-time, read-only images of data on storage system volumes. The `snap create` operation ensures that you have backed up your LUNs. You can use the snapshot you create to restore your databases if you encounter data corruption or other problems.

As part of SnapDrive for UNIX's snapshot creation process, it flushes the file system to disk. This means that the disk image in the snapshot is in a consistent state. While SnapDrive for UNIX ensures this consistency, other methods of creating snapshots may not unless you specifically coordinate the data with the host

A snapshot contains images of all the LUNs in the storage system volumes that are part of the disk group you specify. Before the snapshot is taken, though, SnapDrive for UNIX makes consistent only those LUNs requested on the `snap create` command line. This means that LUNs being used outside those specified by the `snap create` command's *filespec* argument (disk groups or file systems on raw devices) and LUNs that might be in use by other hosts do not have consistent images in the snapshot. As a result, SnapDrive for UNIX only enables you to restore entities specified by the *filespec* argument that are marked as consistent in the snapshot.

Depending on your host operating system and LVM, host volumes that are being used as raw storage are not made consistent in the disk group snapshot. For example, a host volume being used as raw storage by a database is not made consistent.

On a single storage system, SnapDrive for UNIX uses Data ONTAP to guarantee a crash-consistent snapshot. In other words, the data in the snapshot is the same as it would be after a system failure or power outage on the host.

If the data spans multiple storage systems, SnapDrive for UNIX provides a crash-consistent snapshot by making sure no read or write I/O operations are going through while snapshot is taken.

Note

SnapDrive for UNIX has additional requirements for taking a snapshot on Linux hosts because there is not a freeze operation in the software stack. There are also additional requirements on HP-UX, Linux, and Solaris hosts if the host volumes are being used as raw storage. For more information, see “[Architectural limitations for SnapDrive for UNIX](#)” on page 6.

To ensure that a snapshot is “application-consistent,” you may need to stop or do whatever steps are required to quiesce the application before taking the snapshot. If the application has not completed its transactions and written data to the storage system, the resulting snapshot might not be consistent.

Note

If your application can recover from a crash consistent snapshot, you may not need to stop it. Consult the documentation for your application. For more information on taking application consistent snapshots, see “[Architectural limitations for SnapDrive for UNIX](#)” on page 6

You should take a new snapshot whenever you add or remove a host volume, add or remove a LUN, or resize host volumes or file systems. This ensures that you have a consistent copy of the newly configured disk group that you can use if you need to restore the disk group.

Snapshots that span storage systems, volumes

SnapDrive allows you to take snapshots that span multiple storage system volumes or multiple storage systems. These volumes can reside on the same storage system or different storage systems.

When you use the `snap create` command to take a snapshot that spans multiple volumes, you do not need to name the volumes on the command line. SnapDrive for UNIX gets this information from the *file_spec* arguments on the command line. It translates these arguments into disk groups on the host. Then it freezes the file systems on the volumes in the disk groups and takes a snapshot of the storage systems containing one or more LUNs in the disk groups.

In addition to using a *file_spec* argument that is built on LUNs from multiple storage systems and storage system volumes, you can also use a combination of *file_spec* arguments where each value is based on single storage system or storage system volume. For example, suppose you have a setup where the disk group `dg1` spans the storage systems `filer2` and `filer3`, `dg2` is on `filer2`, and `dg3` is on `filer3`. In this case, any of the following command lines would be correct:

```
snap create -dg dg1 -snapname snapdg1
```

```
snap create -dg dg2 dg3 -snapname snapdg23
snap create -dg dg1 dg2 dg3 -snapname snapdg123
```

Something to keep in mind when creating snapshots that span storage systems and volumes is that SnapDrive for UNIX creates the snapshot on each storage system volume using a short name. It does not include the full pathname in the name, even if the volumes are on different storage system. This means that if you later rename the snapshot, you must go to each storage system and volume and rename it there as well.

Note

There are limitations for taking snapshots that span multiple volumes on Linux hosts because there is not a freeze operation anywhere in the software stack. In addition, there are limitations for taking a snapshot where an application is running. For more information, see [“Architectural limitations for SnapDrive for UNIX”](#) on page 6.

Restrictions on snapshot creation

The following restrictions apply to using snapshots:

- ◆ You can keep a maximum of 255 snapshots per storage system volume. This limit is set by the storage system. The total number can vary depending on whether other tools use these snapshots.
When the number of snapshots has reached the maximum limit, the `snapshot create` operation fails. You must delete some of the old snapshots before you can use SnapDrive for UNIX to take any more.
- ◆ SnapDrive for UNIX does not support snapshots that it does not create. For example, it does not support snapshots that are created from the storage system console, because such a practice can lead to inconsistencies within the file system.
- ◆ You cannot use SnapDrive for UNIX to create snapshots of the following:
 - ❖ Root disk groups. The `snap create` operation fails when you try to take a snapshot of a root disk group for a logical volume manager.
 - ❖ Boot device or swap device. SnapDrive for UNIX does not take a snapshot of a system boot device or a system swap device.
 - ❖ Shared disk group. SnapDrive for UNIX does not create snapshots of disk groups that are shared with a host cluster partner.

Creating a snapshot To create a snapshot using SnapDrive for UNIX, complete the following steps.

Note SnapDrive for UNIX can only restore snapshots it takes.

Step	Action
1	Execute the <code>snapdrive snap create</code> command. This command has the following format: <pre>snapdrive snap create {-dg -fs -hostvol } file_spec [file_spec ...] [{-dg -fs -hostvol } file_spec [file_spec...]] -snapname snap_name [-force [-noprompt]]</pre> Note The <code>-dg</code> and <code>-vg</code> options are synonyms that reflect the fact that some operating systems refer to disk groups and others refer to volume groups. In addition, <code>-lvol</code> and <code>-hostvol</code> are also synonyms. This guide uses <code>-dg</code> to refer to both disk groups and volume groups and <code>-hostvol</code> to refer to both logical volumes and host volumes. For details about using all the options and arguments available with this command, see “SnapDrive for UNIX keywords and arguments” on page 247. Examples of possible command lines are included at the end of this table. Keep the following in mind as you create snapshots: ◆ The value you supply for the <code>file_spec</code> argument can span multiple storage systems and storage system volumes. ◆ If you want SnapDrive for UNIX to overwrite an existing snapshot, you must include the <code>-force</code> option on the command line. Otherwise, the <code>create</code> operation fails. (If the old snapshot is in use, this operation fails even with the <code>-force</code> option.) The <code>-force</code> option first asks whether you want to overwrite the existing snapshot. To avoid getting this prompt, include the <code>-noprompt</code> option on the command. When you include this option with the <code>-force</code> option, SnapDrive for UNIX performs the operation without checking with you.

Step	Action
	◆ If you want to use the same <i>file_spec</i> argument multiple times in the same command line, enter the option once and follow it by multiple arguments. For example, if you want to specify multiple disk groups, enter <code>-dg</code> once and list all disk groups after it. So you have the argument <code>-dg dg1 dg2</code>. ◆ If you use the <code>-hostvol</code> option, make sure that it specifies both the requested host volume and the disk group containing it, such as <code>-hostvol dg3/acct_volume</code>. ◆ All the LUNs in the disk group you are taking a snapshot of must be online and available. ◆ Whenever you add or remove a host volume, add or remove a LUN, or resize host volumes or file systems, you should take a new snapshot. This ensures that you have a consistent copy of the newly configured disk group to use for the <code>restore</code> operation when you need it. Result: The <i>file_spec</i>, <code>-lvol</code>, and <code>-hostvol</code> arguments represent a set of LUNs on one or more storage systems. The <code>snap create</code> operation takes a snapshot of the storage system volume containing those LUNs and gives it the name specified in the <i>snap_name</i> argument.

Step	Action
	Examples of command lines: These examples use a host called DBserver. The disk group dg1 has host volumes myvol1 and myvol2. The host volume dg1/myvol2 has a file system mounted on /fs2. The disk group dg1 has three LUNs in it: toaster:/vol/vol1/lun0, toaster:/vol/vol1/lun1, and toaster: /vol/vol1/lun2. The following command lines are all valid, and back up the same data. These command lines fall into three general categories. The first command line uses a file system; the next to refer to disk/volume groups, and the last two refer to host volumes/logical volumes. All the command lines create a snapshot called toaster:/vol/vol1:snap1. <pre> snapdrive snap create -fs /fs2 -snapname snap1 snapdrive snap create -dg dg1 -snapname snap1 snapdrive snap create -vg dg1 -snapname snap1 snapdrive snap create -hostvol dg1/myvol1 -snapname snap1 snapdrive snap create -lvol dg1/myvol2 -fs /fs2 -snapname snap1 </pre> Example 2: This example creates a multivolume snapshot for a Linux host. The snapshot contains the disk group vgmultipvol, which contains the host volumes lv01 and lv02. <pre> # snapdrive snap create -vg vgmultipvol -snapname snapmultipvol Successfully created snapshot snapmultipvol on 2 filer volumes: toaster:/vol/vol1 toaster:/vol/vol2 snapshot snapmultipvol contains: disk group vgmultipvol containing host volumes lv01 lv02 # </pre>
2	Repeat Step 1 to create another snapshot.

Displaying information about snapshots

Command to use to display snapshot information

You can use the `snapdrive snap show` (or `list`) command to display information about each snapshot taken by SnapDrive for UNIX.

You can use this command to display information on the following:

- ◆ Storage systems
- ◆ Volumes on storage systems
- ◆ Host-side entities such as volume groups, disk groups, file systems, logical volumes, and host volumes
- ◆ Snapshots

Displaying information about a snapshot

To display snapshot information, complete the following step.

Step	Action
1	Execute the <code>snapdrive snap show</code> command. This command has the following formats: <pre>snapdrive snap { show list } -filer <i>filename</i> [<i>filename...</i>] [-verbose] snapdrive snap { show list } -filervol <i>filervol</i> [<i>filervol...</i>] [-verbose] snapdrive snap { show list } {-dg -fs -hostvol } <i>file_spec</i> [<i>file_spec ...</i>] [-verbose] snapdrive snap { show list } [-snapname] <i>long_snap_name</i> [<i>snap_name ...</i>] [-verbose]</pre> Note_____ You can use the same arguments regardless of whether you enter <code>snapdrive snap show</code> or <code>snapdrive snap list</code> as the command. These commands are synonyms. In addition, the <code>-dg</code> and <code>-vg</code> options are synonyms that reflect the fact that some operating systems refer to disk groups and others refer to volume groups. In addition, <code>-lvol</code> and <code>-hostvol</code> are also synonyms. This guide uses <code>-dg</code> to refer to both disk groups and volume groups and <code>-hostvol</code> to refer to both logical volumes and host volumes. For details about using all the options and arguments available with this command, see “SnapDrive for UNIX keywords and arguments” on page 247. ◆ It might take a few minutes for SnapDrive for UNIX to display information. This is because it must gather information about the snapshot.

Step	Action
	◆ You can use this command to display information about a storage system entity, such as storage system (-filer), storage system volume (-filervolume), or snapshot (-snapshot). You can also combine these options with host-side options as specified by the filespec argument. For example, you could find all snapshots that have disk group dg1 on the storage system comet even though dg1 doesn't exist on your host by running one command: <pre>snap show -dg dg1 -filer comet</pre> ◆ You can use the wildcard (*) character in snapshot names. The <code>snap show</code> operation lets you use the wildcard character to show all snapshot names that match a certain pattern or all the snapshots on a particular volume. The following rules apply to using wildcards in snapshot names: ❖ You can use a wildcard at the end of the name only. You cannot use the wildcard at the beginning or the middle of a snapshot name. ❖ You cannot use the wildcard in the storage system or storage system volume fields of a snapshot name. See the examples that follow for a sample command line that includes a wildcard. Result: This operation displays, at a minimum, the following information about the snapshot: ◆ The name of the storage system where the snapshot was taken ◆ The name of the host that took the snapshot ◆ The path to the LUNs on the storage system ◆ The date and time the snapshot was taken ◆ The name of the snapshot ◆ The names of the disk groups included in the snapshot You can also use this command to list of all snapshots on specific objects, including storage systems and their volumes, disk groups, host volume groups, file systems, host volumes, and logical volumes. When you include the verbose option, SnapDrive for UNIX displays more information. If you enter a <code>snap show</code> command and SnapDrive for UNIX doesn't locate any snapshots, it displays the message "no matching snapshots." If you specify arguments on the command line, and some portion of them do not exist, SnapDrive for UNIX returns a partial listing of those for which snapshots are found. It also lists the arguments that were invalid.

Step	Action
	Example 1: The following are examples of <code>snapdrive snap show</code> commands: <pre> snapdrive snap show -snapname toaster:/vol/vol2:snapA snapX snapY snapdrive snap show -verbose toaster:/vol/vol2:snapA /vol/vol3:snapB snapC snapdrive snap show toaster:/vol/vol2:snapA snapdrive snap list -dg dg1 dg2 </pre> Example 2: This example uses a wildcard to display information about the available snapshots on a particular volume: <pre> # snapdrive snap show toaster:/vol/vol0:* snap name host date snapped ----- toaster:/vol/vol0:sss1 DBserver Mar 12 16:19 dg1 toaster:/vol/vol0:testdg DBserver Mar 12 15:35 dg1 toaster:/vol/vol0:t1 DBserver Mar 10 18:06 dg1 toaster:/vol/vol0:hp_1 HPserver Mar 8 19:01 vg01 toaster:/vol/vol0:r3 DBserver Mar 8 13:39 rdg1 toaster:/vol/vol0:r1 DBserver Mar 8 13:20 rdg1 toaster:/vol/vol0:snap2 DBserver Mar 8 11:51 rdg1toaster:/vol/vol0:snap_p1 DBserver Mar 8 10:18 rdg1 </pre> Example 3: In this example, the verbose option is included on the command line. <pre> # snapdrive snap show -v toaster:/vol/vol0:snap_p1 snap name host date snapped ----- toaster:/vol/vol0:snap_p1 DBserver Mar 8 10:18 rdg1 host OS: SunOS 5.8 Generic_108528-23 snapshot name: snap_p1 Volume Manager: vxvm 3.5 disk group: rdg host volume: nvoll file system: nvoll mountpoint: /mnt/rdg_nvoll1 lun path ----- toaster:/vol/vol0/rds0 lun path ----- myfiler:/vol0/lun-1 myfiler:/vol0/lun-1 myfiler:/vol0/lun-5 myfiler:/vol0/lun-5 </pre>

Step	Action																																														
	Example 4: This example on an AIX host includes messages about snapshots that were not created with SnapDrive for UNIX. <pre># snapdrive snap show -filer toaster</pre> <table><tr><th>snap name</th><th>host</th><th>date</th><th>snapped</th></tr><tr><td colspan="4">-----</td></tr><tr><td>toaster:/vol/vol0:hourly.0</td><td>non-snapdrive</td><td>snapshot</td><td></td></tr><tr><td>toaster:/vol/vol1:hourly.0</td><td>non-snapdrive</td><td>snapshot</td><td></td></tr><tr><td>toaster:/vol/vol1:snap1</td><td>DBserver1</td><td>Oct 01 13:42</td><td>dg1 dg2</td></tr><tr><td>toaster:/vol/vol1:snap2</td><td>DBserver2</td><td>Oct 10 13:40</td><td>DBdg/hvol1</td></tr><tr><td>toaster:/vol/vol1:snap3</td><td>DBserver3</td><td>Oct 31 13:45</td><td>DBdg</td></tr></table> Example 5: This example uses the verbose option with the snap show command on an AIX host. <pre># snapdrive snap show -v -fs /mnt/qa/dg1/vol1</pre> <table><tr><th>snap name</th><th>host</th><th>date</th><th>snapped</th></tr><tr><td colspan="4">-----</td></tr><tr><td>blip:/vol/vol1:dg1snapshot</td><td>tester-aix5</td><td>Jun 21 10:51</td><td>dg1</td></tr></table> host OS: AIX 1 5 snapshot name: dg1snapshot Volume Manager: aixlvm 5.1 disk group: dg1 host volume: dg1_log host volume: dg1_vol1 file system: dg1_vol1 file system type: jfs2 mountpoint: /mnt/qa/dg1/vol1 <table><tr><th>lun path</th><th>dev paths</th></tr><tr><td colspan="2">-----</td></tr><tr><td>blip:/vol/vol1/tester_tester-aix5_lun0</td><td>hdisk98,hdisk2</td></tr></table>	snap name	host	date	snapped	-----				toaster:/vol/vol0:hourly.0	non-snapdrive	snapshot		toaster:/vol/vol1:hourly.0	non-snapdrive	snapshot		toaster:/vol/vol1:snap1	DBserver1	Oct 01 13:42	dg1 dg2	toaster:/vol/vol1:snap2	DBserver2	Oct 10 13:40	DBdg/hvol1	toaster:/vol/vol1:snap3	DBserver3	Oct 31 13:45	DBdg	snap name	host	date	snapped	-----				blip:/vol/vol1:dg1snapshot	tester-aix5	Jun 21 10:51	dg1	lun path	dev paths	-----		blip:/vol/vol1/tester_tester-aix5_lun0	hdisk98,hdisk2
snap name	host	date	snapped																																												

toaster:/vol/vol0:hourly.0	non-snapdrive	snapshot																																													
toaster:/vol/vol1:hourly.0	non-snapdrive	snapshot																																													
toaster:/vol/vol1:snap1	DBserver1	Oct 01 13:42	dg1 dg2																																												
toaster:/vol/vol1:snap2	DBserver2	Oct 10 13:40	DBdg/hvol1																																												
toaster:/vol/vol1:snap3	DBserver3	Oct 31 13:45	DBdg																																												
snap name	host	date	snapped																																												

blip:/vol/vol1:dg1snapshot	tester-aix5	Jun 21 10:51	dg1																																												
lun path	dev paths																																														

blip:/vol/vol1/tester_tester-aix5_lun0	hdisk98,hdisk2																																														

Step	Action																																																																																
	Example 6: This example uses the <code>snap list</code> version of the <code>snap show</code> command. <pre># snapdrive snap list -filer toaster</pre> <table><tr><th>snap name</th><th>host</th><th>date</th><th>snapped</th></tr><tr><td colspan="4">-----</td></tr><tr><td>toaster:/vol/vol0:hourly.0</td><td>non-snapdrive</td><td>snapshot</td><td></td></tr><tr><td>toaster:/vol/vol0:hourly.1</td><td>non-snapdrive</td><td>snapshot</td><td></td></tr><tr><td>toaster:/vol/vol0:nightly.0</td><td>non-snapdrive</td><td>snapshot</td><td></td></tr><tr><td>toaster:/vol/vol1:hourly.0</td><td>non-snapdrive</td><td>snapshot</td><td></td></tr><tr><td>toaster:/vol/vol1:hourly.1</td><td>non-snapdrive</td><td>snapshot</td><td></td></tr><tr><td>toaster:/vol/vol1:nightly.1</td><td>non-snapdrive</td><td>snapshot</td><td></td></tr><tr><td>toaster:/vol/vol1:dg12</td><td>hush</td><td>May 14 14:36</td><td>dg1 dg2</td></tr><tr><td>toaster:/vol/vol1:dg1a</td><td>hush</td><td>Mar 31 15:36</td><td>dg1</td></tr><tr><td>toaster:/vol/vol1:dg143s</td><td>hush</td><td>Mar 31 15:20</td><td>dg1 dg3</td></tr><tr><td>toaster:/vol/vol1:dg3_snap</td><td>hush</td><td>Mar 31 15:20</td><td>dg3</td></tr><tr><td>toaster:/vol/vol1:dg1</td><td>hush</td><td>Mar 13 17:57</td><td>dg1</td></tr><tr><td>toaster:/vol/vol1:dg123</td><td>hush</td><td>Mar 11 20:01</td><td>dg2 dg3 dg1</td></tr><tr><td>toaster:/vol/vol1:dg3-b</td><td>hush</td><td>Mar 11 16:03</td><td>dg3</td></tr><tr><td>toaster:/vol/vol1:dg3-a</td><td>hush</td><td>Mar 11 15:52</td><td>dg3</td></tr><tr><td>toaster:/vol/vol1:dg11fs_snap</td><td>hush</td><td>Mar 3 15:30</td><td>dg11</td></tr><tr><td>toaster:/vol/vol1:dg13hvols_snap</td><td>hush</td><td>Mar 3 15:29</td><td>dg13</td></tr><tr><td>toaster:/vol/vol1:dg1_snap</td><td>hush</td><td>Mar 3 15:27</td><td>dg1</td></tr><tr><td>toaster:/vol/vol1:dg2dg3-1</td><td>hush</td><td>Mar 3 15:22</td><td>dg2 dg3</td></tr></table>	snap name	host	date	snapped	-----				toaster:/vol/vol0:hourly.0	non-snapdrive	snapshot		toaster:/vol/vol0:hourly.1	non-snapdrive	snapshot		toaster:/vol/vol0:nightly.0	non-snapdrive	snapshot		toaster:/vol/vol1:hourly.0	non-snapdrive	snapshot		toaster:/vol/vol1:hourly.1	non-snapdrive	snapshot		toaster:/vol/vol1:nightly.1	non-snapdrive	snapshot		toaster:/vol/vol1:dg12	hush	May 14 14:36	dg1 dg2	toaster:/vol/vol1:dg1a	hush	Mar 31 15:36	dg1	toaster:/vol/vol1:dg143s	hush	Mar 31 15:20	dg1 dg3	toaster:/vol/vol1:dg3_snap	hush	Mar 31 15:20	dg3	toaster:/vol/vol1:dg1	hush	Mar 13 17:57	dg1	toaster:/vol/vol1:dg123	hush	Mar 11 20:01	dg2 dg3 dg1	toaster:/vol/vol1:dg3-b	hush	Mar 11 16:03	dg3	toaster:/vol/vol1:dg3-a	hush	Mar 11 15:52	dg3	toaster:/vol/vol1:dg11fs_snap	hush	Mar 3 15:30	dg11	toaster:/vol/vol1:dg13hvols_snap	hush	Mar 3 15:29	dg13	toaster:/vol/vol1:dg1_snap	hush	Mar 3 15:27	dg1	toaster:/vol/vol1:dg2dg3-1	hush	Mar 3 15:22	dg2 dg3
snap name	host	date	snapped																																																																														

toaster:/vol/vol0:hourly.0	non-snapdrive	snapshot																																																																															
toaster:/vol/vol0:hourly.1	non-snapdrive	snapshot																																																																															
toaster:/vol/vol0:nightly.0	non-snapdrive	snapshot																																																																															
toaster:/vol/vol1:hourly.0	non-snapdrive	snapshot																																																																															
toaster:/vol/vol1:hourly.1	non-snapdrive	snapshot																																																																															
toaster:/vol/vol1:nightly.1	non-snapdrive	snapshot																																																																															
toaster:/vol/vol1:dg12	hush	May 14 14:36	dg1 dg2																																																																														
toaster:/vol/vol1:dg1a	hush	Mar 31 15:36	dg1																																																																														
toaster:/vol/vol1:dg143s	hush	Mar 31 15:20	dg1 dg3																																																																														
toaster:/vol/vol1:dg3_snap	hush	Mar 31 15:20	dg3																																																																														
toaster:/vol/vol1:dg1	hush	Mar 13 17:57	dg1																																																																														
toaster:/vol/vol1:dg123	hush	Mar 11 20:01	dg2 dg3 dg1																																																																														
toaster:/vol/vol1:dg3-b	hush	Mar 11 16:03	dg3																																																																														
toaster:/vol/vol1:dg3-a	hush	Mar 11 15:52	dg3																																																																														
toaster:/vol/vol1:dg11fs_snap	hush	Mar 3 15:30	dg11																																																																														
toaster:/vol/vol1:dg13hvols_snap	hush	Mar 3 15:29	dg13																																																																														
toaster:/vol/vol1:dg1_snap	hush	Mar 3 15:27	dg1																																																																														
toaster:/vol/vol1:dg2dg3-1	hush	Mar 3 15:22	dg2 dg3																																																																														

Step	Action																														
	Example 7: This example executes the <code>snap show</code> command on a Linux host. <pre># snapdrive snap show -snapname surf:/vol/vol1:swzldg5snapped</pre> <table><tr><th>snap name</th><th>host</th><th>date</th><th>snapped</th></tr><tr><td colspan="4">-----</td></tr><tr><td>surf:/vol/vol1:bagel5snapped</td><td>pons</td><td>Aug 18 20:06</td><td>dg5</td></tr></table> <pre># # ./linux/ix86/snapdrive snap show -v -snapname surf:/vol/vol1:bagel5snapped ></pre> <table><tr><th>snap name</th><th>host</th><th>date</th><th>snapped</th></tr><tr><td colspan="4">-----</td></tr><tr><td>surf:/vol/vol1:bagel5snapped</td><td>pons</td><td>Aug 18 20:06</td><td>dg5</td></tr></table> host OS: Linux 2.4.21-9.ELsmp #1 SMP Thu Jan 8 17:08:56 EST 2004 snapshot name: bagel5snapped Volume Manager: linuxlvm 1.0.3 disk group: dg5 host volume: vol1 host volume: vol2 host volume: vol3 <table><tr><th>lun path</th><th>dev paths</th></tr><tr><td colspan="2">-----</td></tr><tr><td>surf:/vol/vol1/glk19</td><td>/dev/sdu</td></tr></table> <pre>#</pre> Examples with wildcards: The following examples use wildcards: <pre>snapdrive snap show toaster:/vol/volX:* snapdrive snap show -v toaster:/vol/volX:DB1* filer1:/vol/volY:DB2* snapdrive snap show toaster:/vol/vol2:mynsnap* /vol/vol2:yoursnap* hersnap* snapdrive snap show toaster:/vol/volX:*</pre> Example of an invalid command: The following use of a wildcard is invalid because the wildcard is in the middle of the name instead of at the end: <pre>snap show toaster:/vol/vol1:my*snap</pre>	snap name	host	date	snapped	-----				surf:/vol/vol1:bagel5snapped	pons	Aug 18 20:06	dg5	snap name	host	date	snapped	-----				surf:/vol/vol1:bagel5snapped	pons	Aug 18 20:06	dg5	lun path	dev paths	-----		surf:/vol/vol1/glk19	/dev/sdu
snap name	host	date	snapped																												

surf:/vol/vol1:bagel5snapped	pons	Aug 18 20:06	dg5																												
snap name	host	date	snapped																												

surf:/vol/vol1:bagel5snapped	pons	Aug 18 20:06	dg5																												
lun path	dev paths																														

surf:/vol/vol1/glk19	/dev/sdu																														

Other ways to get snapshot names

Another way to get a snapshot name is to log on to the storage system and use the `snap list` command there. This command displays the names of the snapshots.

Note

The `snapdrive snap show` command is equivalent to the storage system `snap list` command.

Renaming a snapshot

Command to use to rename a snapshot

You can use the `snapshot snap rename` command to change the name of an existing snapshot.

Renaming a snapshot that spans systems, volumes

You can also use this command to rename a snapshot that is across multiple storage systems or multiple storage system volumes.

If you rename one of these snapshots, you must also rename all the related snapshots using the same name. This is because SnapDrive for UNIX uses a short name when it creates the snapshot, even though it spans multiple storage systems or volumes. The rename command changes the name of the current snapshot but it does not change the name of the related snapshots in the other locations.

Changing a snapshot name

To change the name of a snapshot, complete the following step.

Step	Action
1	Execute the <code>snapdrive snap rename</code> command. This command has the following format: <pre>snapdrive snap rename [-snapname] <i>old_long_snap_name</i> <i>new_snap_name</i> [-force [-noprompt]]</pre> <i>old_long_snap_name</i> is the current name of the snapshot. It must use the long format of the name where you enter the storage system name, volume, and snapshot name. <i>new_snap_name</i> is the name you want to give the snapshot. You can use the short form of the name for the <i>new_snap_name</i> because it must exist on the same storage system volume as the original snapshot. An error occurs if you try to rename a snapshot do a different storage system. If you created a snapshot that spans multiple storage systems or storage system volumes, make sure you rename each related snapshot with the same name.

Step	Action
	If the new snapshot name exists, you must use the <code>-force</code> option on the command line to force SnapDrive to use this name. Otherwise, the operation fails. This option prompts you to confirm that you want to use the existing name. To prevent SnapDrive for UNIX from displaying the prompt, include the <code>-noprompt</code> option also. Note_____ The section “SnapDrive for UNIX keywords and arguments” on page 247 provides details about using the options and arguments in this command line. Result: The <code>snap rename</code> operation changes the name of the source snapshot source to the name specified by the target argument. Example: The following are examples of the <code>snapdrive snap rename</code> command. The first command line includes the <code>-force</code> option because a snapshot named <code>newsnapshot</code> already exists. In the second example, both snapshot names use the long form of the name, but they both resolve to the same storage system volume. <pre>snapdrive snap rename -force filer1:/vol/vol1:oldsnap newsnapshot snapdrive snap rename filer1:/vol/vol1:FridaySnap filer1:/vol/vol1:Snap040130</pre>

Restoring a snapshot

Command to use to restore snapshots

The `snapdrive snap restore` command restores data from the snapshot you specify on the command line to the storage system. This operation replaces the contents of the *file_spec* arguments (disk groups and file systems on raw devices) that you specified on the `snap restore` command line in with the contents of the *file_spec* arguments found in the specified snapshot.

You can also restore snapshots for non-existent *file_spec* arguments. This happens when the value you specify no longer exists on the host, but existed when you took the snapshot. For example, it might be a file system that you have now unmounted or a disk group that you have removed.

Normally, you restore snapshots from the host where you took the snapshots (in other words, the originating host). You can also restore snapshots using a different, or non-originating, host (see [“Restoring a snapshot from a different host”](#) on page 136).

Restrictions on restoring snapshots

SnapDrive for UNIX has the following restrictions with regard to restoring snapshots:

- ◆ SnapDrive for UNIX restores only whole disk groups. If you specify part of a disk group, it automatically restores the entire disk group. An error occurs if you enter only a subset of the host volumes and/or file systems in each disk group on the command line. You can include the `-force` option to override this error; however, SnapDrive for UNIX then restores the entire disk group.
- ◆ All applications using the host volumes and file systems you want to restore must be closed. The application is likely to fail if it has open connections to these volumes and file systems when you create the snapshot.
- ◆ If the configuration of the disk group you’re trying to restore has changed since the snapshot was taken, the restore operation fails. If you have added or removed a host volume, file system, or LUNs, changed the way your data is striped, or resized any volume manager entity above the disk group level, you can override and restore an older snapshot by including the `-force` option.

Note

You should always take a new snapshot whenever a LUN has been added to or removed from a disk group.

- ◆ You cannot restore snapshots of
 - ❖ Root disk groups
 - ❖ Shared disk group
 - ❖ Boot device
 - ❖ Swap device

Restoring a snapshot

The restore operation can take several minutes, depending on the type and amount of data being restored.

Caution

Do not interrupt the restore operation by entering Ctrl-C. Doing so could leave your system in an unusable configuration. If that happens, you will need to consult with technical support in order to recover.

To restore a snapshot, complete the following steps.

Step	Action
1	Make sure you are not in any directory on a file system that you want to restore. You can perform the <code>snap restore</code> command from any directory except the one on a file system to which you want to restore the information.

Step	Action
2	Execute the <code>snapdrive snap restore</code> command. This command has the following format: <pre>snapdrive snap restore -snapname snap_name {-dg -fs -hostvol } file_spec [file_spec ...] [{-dg -fs -hostvol } file_spec [file_spec ...] ...] [-force [-noprompt]]</pre> Note For details about using these options and arguments, see the section “SnapDrive for UNIX keywords and arguments” on page 247. ◆ <i>snap_name</i> is the name of the snapshot. It can be either a short name, such as <code>mysnap1</code>, or a long name that includes the storage system name, volume, and snapshot name Generally, NetApp recommends that you use the short name, unless the <i>file_spec</i> argument is non-existent (i.e., no longer exists on the host, see the explanation of the <i>file_spec</i> argument). Then you must use the long name for the snapshot. ◆ By default, SnapDrive for UNIX determines the storage system and storage system volume based on the value you supply for the <i>file_spec</i> argument. If you use a long name for the snapshot and the pathname does not match the storage system/storage volume information on the command line, SnapDrive for UNIX fails. ◆ <i>file_spec</i> is the name of the host entity you specify. For example, if you use the <code>-dg</code> option, <i>file_spec</i> is the name of that disk group. If the snapshot contains host entities that span multiple storage system volumes, you can restore any of the entities in that snapshot. In some cases the value supplied with the <i>file_spec</i> argument may not exist on the host. For example, if you had unmounted a file system or removed a disk group by exporting it, deporting it, or destroying it, that file system or disk group could still be a value for the <i>file_spec</i> argument. It would, however, be considered a non-existent value. SnapDrive for UNIX can restore snapshots for such a non-existent <i>file_spec</i>, but you must use the long snapshot name.

Step	Action
	When you restore snapshots that span multiple storage systems and volumes and contain an non-existent <i>file_spec</i> argument, SnapDrive for UNIX permits an inconsistency in the command line that it does not allow for existing <i>file_spec</i> arguments. If you want to restore only one host entity from a multiple storage system snapshot, the snapshot you specify does not need to be on the same storage system as the storage system containing the host entity. For example, a snapshot might contain dg2, which was created on filer2, and dg3, which was created on filer3. To restore dg2 only, you can specify either the snapshot with the pathname to filer2 or the pathname to filer 3. For more information, see “Example of restoring to a non-existent disk group” on page 135. ◆ If you supply a non-existent <i>file_spec</i> argument that, at the time the snapshot was created, spanned multiple storage system volumes, you can specify any of the snapshots in which the values for the <i>file_spec</i> argument are stored. For example, suppose disk group dg1 spans filer1:/vol/vol1 and filer2:/vol/vol1. When you performed the <code>snap create</code> operation <code>snap create -dg dg1 -snapname snapdg1</code>, SnapDrive for UNIX would have created two snapshots: <code>filer1:/vol/vol1:snapdg1</code> and <code>filer2:/vol/vol1:snapdg1</code> If you want to restore dg1 and it currently exists on the host (i.e., it is <i>live</i>), you only need to supply the short name of the snapshot on the command line and SnapDrive for UNIX finds both snapshots and restores them: <pre>snap restore -dg dg1 -snapname snapdg1</pre> If you exported disk group dg1, it becomes a non-existent entity on the host. You can still restore the snapshot, but now you must specify the long snapshot name that contains the storage system and volume information. The <code>snap restore</code> command only accepts one snapshot name, so you can supply the name of either snapshot that was created by the <code>snap create</code> command. Enter one of the following command lines: <pre>snap restore -dg dg1 -snapname filer1:/vol/vol1:snapdg1 snap restore -dg dg1 -snapname filer2:/vol/vol1:snapdg1</pre>

Step	Action
	◆ You must include the <code>-force</code> option on the command line if you attempt to restore a disk group where the configuration has changed since the last snapshot. For example, if you changed the way data is striped on the disks since you took a snapshot, you would need to include the <code>-force</code> option. Without the <code>-force</code> option, this operation fails. This option ask you to confirm that you want to continue the operation unless you include the <code>-noprompt</code> option with it. Note _____ If you added or deleted a LUN, the restore operation fails, even if you include the <code>-force</code> option. Result: SnapDrive for UNIX replaces the existing contents of the LUNs you specify in the <code>snap restore</code> command line with the contents of the LUNs in the snapshot you specify. This operation could take several minutes. When the operation is complete, SnapDrive for UNIX displays a message similar to the following: <pre>Snap restore <filespec list> succeeded.</pre> Examples of command lines: The following are examples of the <code>snapdrive snap restore</code> command: <pre> snapdrive snap restore -fs /mnt/dir -snapname filer1:/vol/vol1:NewSnap33 snapdrive snap restore -vg OracleVg1 Datavg -snapname Monday snapdrive snap restore -dg dg1 dg2 -snapname Tuesday </pre>

Step	Action
	This example uses a host called DBserver and a disk group dg1 that has host volumes myvol1 and myvol2. The volume dg1/myvol2 has a file system mounted on /fs2. The disk group dg1 has three LUNs in it: toaster:/vol/vol1/lun0, toaster:/vol/vol1/lun1, and toaster:/vol/vol1/lun2. You could use any of the following commands to restore the same data: <pre> snapdrive snap restore -dg dg1 -snapname snap1 snapdrive snap restore -v dg1 -snapname snap1 snapdrive snap restore -hostvol dg1/myvol2 dg1/myvol1 -snapname toaster:/vol/vol1:snap1 snapdrive snap restore -lvol dg1/myvol1 -fs /fs2 </pre> Example on a Linux host: This example restores a snapshot on a Linux host. <pre> # snapdrive snap restore -dg dg5 -snapname bagel5snapped Starting to restore LUNs in disk group dg5 WARNING: This can take several minutes. DO NOT CONTROL-C! If snap restore is interrupted, the disk group being restored may have inconsistent or corrupted data. For detailed progress information, see the log file /var/log/sd-recovery.log Importing dg5 snap restore: snapshot bagel5snapped contains: disk group dg5 containing host volumes vol1 vol2 vol3 snap restore: restored snapshot surf:/vol/vol1:bagel5snapped # </pre>

Step	Action
	Example of restoring to a non-existent disk group: This example takes you through the steps of creating a snapshot across multiple storage systems and volumes and then restoring it, even though the disk group it contains no longer exists on the host. The example uses disk group dg2, which is on filer2:/vol/vol2, and disk group dg3, which is on filer3:/vol/vol3. The following <code>snap create</code> command was used: <pre>snap create -dg dg2 dg3 -snapname snapdg23</pre> This command created the snapshots <code>filer2:/vol/vol2:snapdg23</code> and <code>filer3:/vol/vol3:snapdg23</code> If dg2 and dg3 still exist on the host, you can use the short name to restore the snapshot: <pre>snap restore -dg dg2 dg3 -snapname snapdg23</pre> Suppose, however, that you had exported both dg2 and dg3. In that case, they no longer exist on the host. You can restore them by specifying either of the snapshots created by the <code>snap create</code> command line, but you must include the long snapshot name on the <code>snap restore</code> command line. You could enter either of the following command lines: <pre>snap restore -dg dg2 dg3 -snapname filer3:/vol/vol3:snapdg23</pre> <pre>snap restore -dg dg2 dg3 -snapname filer2:/vol/vol2:snapdg23</pre> If you only want to restore dg2, enter the one of the following command lines: <pre>snap restore -dg dg2 -snapname filer2:/vol/vol2:snapdg23</pre> <pre>snap restore -dg dg2 -snapname filer3:/vol/vol3:snapdg23</pre> Because dg2 no longer exists on the host, you can specify the snapshot on filer3 even though dg2 was created on filer2. You can only use this inconsistency if the snapshot spans multiple storage systems or volumes and the host entity no longer exists. If dg2 still existed on the host, the second command line would fail.

Step	Action
	Example of warning messages: The following is the standard warning message that appears during a <code>snap restore</code> operation: <pre>Starting to restore LUNs in disk groups: <dg1> WARNING: This can take several minutes. DO NOT CONTROL-C! If snap restore is interrupted, the disk groups being restored may have inconsistent or corrupted data.</pre> SnapDrive for UNIX also displays a warning message during a restore operation if any logical volumes have been added or removed since the snapshot was taken. The message tells you which logical volumes have been added or removed. For example, if you added a volume called <code>/dev/tester2/vol2</code> and removed the volume <code>/dev/tester2/vol3</code> after you took the snapshot <code>snap1</code>, you would see the following message when you executed the <code>snap restore</code> command: <pre>Warning (Warning 0001-178): diskgroup configuration has changed since snapshot lobster:/vol/vol1:snap1 was taken: added hostvol /dev/tester2/vol2 removed hostvol /dev/tester2/vol3</pre>

Restoring a snapshot from a different host

In most cases, you restore a snapshot from the host where you took the snapshot. On occasion, you might need to restore a snapshot using a different, or non-originating, host. To restore a snapshot using a non-originating host, use the same `snapdrive snap restore` command that you would normally use.

Connecting to a snapshot

Connecting to a snapshot

SnapDrive for UNIX lets you connect to a snapshot from a different location on a host. This new location can be on the host where you took the snapshot (the originating host) or on a different host (the non-originating host).

Being able to set up the snapshots in a new location means you can back up a snapshot to another medium, perform maintenance on a disk group, or simply access the snapshot data without disrupting the original copy of the data.

Note

SnapDrive for UNIX does not support using the `snap connect` command to connect to the originating host on Linux and Solaris hosts.

How `snap connect` works

When you use the `snapdrive snap connect` command, it clones the storage (for example, the LUNs) in the disk group you specify and imports it to the host. As a result, the disk group resides on the host with a different name, which is why you can connect to it.

This command always connects the entire target disk group. It then re-activates all of the logical volumes in the target disk group and mounts all the file systems in those logical volumes.

Note

SnapDrive for UNIX always connects whole disk groups. If you specify part of a disk group, it automatically connects the entire disk group. An error occurs if you enter only a subset of the host volumes and/or file systems in each disk group on the command line. You can include the `-force` option to override this error; however, SnapDrive for UNIX then connects the entire disk group.

SnapDrive for UNIX also allows you connect to snapshots that span multiple storage system volumes or multiple storage systems. These volumes can reside on the same storage system or different storage systems.

snap connect and snap restore

Unlike the `snap restore` operation, the `snap connect` operation does not replace the existing contents of the LUNs that make up the host entity with the snapshot contents. It clones the information. Once the connection is made, though, both `snap connect` and `snap restore` perform the same activities in terms of activating logical volumes and mounting file systems.

Limitations on connecting snapshots

The following limitations apply to connecting snapshots:

- ◆ The `snapdrive snap connect` command only works with snapshots created with version 2.x of SnapDrive for UNIX. It does not work with snapshots created using a version 1.x of SnapDrive for UNIX.
- ◆ There is a problem performing `snap connect` operations on an AIX host if the following is true:
 - ❖ The host running AIX 5.2.
 - ❖ The SANpath version is newer than 4.0.5k and has the SANpath workaround enabled (the `failback` attribute for `spdctrl` is set to 3)

In this case, the `snap connect` operation fails. The workaround is to disable SANpath's PVID feature. To do this, enter the following command:

```
setsp -W0
```

**Connecting
snapshots to new
locations**

To connect snapshots to new locations, complete the following step.

Step	Action
1	Execute the <code>snapdrive snap connect</code> command: <pre>snapdrive snap connect <i>fspec_set</i> [<i>fspec_set...</i>] -snapname <i>long_snap_name</i> [-igroup <i>ig_name</i> [<i>ig_name ...</i>]] [-autoexpand] [-autorename]</pre> Note_____ If you have a Linux or Solaris host, do not connect the snapshot to the originating host. A problem can occur if the LVM detects that you are attempting to connect a disk group that has the same ID as an existing disk group. SnapDrive only supports this type of connection if you have an AIX or HP-UX host. _____ Note_____ For details about using the options and arguments in these command line, see the “SnapDrive for UNIX keywords and arguments” on page 247. _____

Step	Action
	Keep the following in mind when you execute this command: ◆ Each <i>fspec_set</i> argument identifies a host LVM entity that was captured in the snapshot specified by <i>long_snap_name</i>. The argument also allows you to specify a set of destination names that SnapDrive for UNIX uses to make the entity available in the newly-connected copy of the snapshot. The argument has the following format: <pre>{-dg -fs -hostvol } src_fspec [dest_fspec] [{-destdg -destvg } dg_name] [{-destlv -desthv } lv_name]</pre> Note _____ The <code>-destdg</code> and <code>-destvg</code> options and the <code>-dg</code> and <code>-vg</code> options are synonyms that reflect the fact that some operating systems refer to disk groups and others refer to volume groups. The <code>-destlv</code> and <code>-desthv</code> options and the <code>-hostvol</code> and <code>-lvvol</code> options are also synonyms. This guide uses <code>-destdg</code> to refer to both disk groups and volume groups and <code>-desthv</code> to refer to both logical volumes and host volumes. For details about using all the options and arguments available with this command, see “SnapDrive for UNIX keywords and arguments” on page 247. ◆ This command fails if you specify a destination name that is already in use. It also fails if you specify a file system name that is being used as a mount point. ◆ The role of <i>src_fspec</i> argument is to specify the connection in the snapshot. It identifies an LVM entity within the snapshot specified by the <i>long_snap_name</i>. SnapDrive for UNIX uses this name to locate the disk group within the snapshot that needs to be connected in order to make the entity available. ◆ The role of <i>dest_fspec</i> argument, if given, is to identify where you’re putting the snapshot information. It specifies the name by which the target entity will be accessible after its disk group is connected. If you do not supply a name for connecting a destination LVM entity, the <code>snap connect</code> command tries to use the name from the source disk group. This means that if you omit the <i>dest_fspec</i> value, <code>snap connect</code> uses the value in <i>src_fspec</i>. ◆ You can use the <code>-destxx</code> options to specify names for destination disk groups or logical volumes if this information is not part of the <i>dest_fspec/src_fspec</i> pair. For example, the <code>-fs</code> option names only a destination mount point so you can use the <code>-destdg</code> option to specify the destination disk group. If you do not specify the name needed to connect an entity in the destination disk group, the <code>snap connect</code> command takes the name from the source disk group.

Step	Action
	◆ The <code>-autoexpand</code> and <code>-autorename</code> options apply to each disk group specified on the <code>snap connect</code> command line and any host LVM entity within that group (i.e., the options apply across every <i>fspec_set</i> argument given on the command line). When you supply either of these options on the command line, you must also include either the name of the disk group or some of the LVM entities in it. The <code>-autoexpand</code> option simplifies the process of identifying the items in the disk group. You can supply a subset of the logical volumes or file systems in the disk group and SnapDrive for UNIX then expands the connection to the rest of the logical volumes or file systems in the disk group. Without the <code>-autoexpand</code> option default, you must specify <i>all</i> effected host volumes and file systems contained in that disk group in order to connect the entire disk group. Note _____ If the value you enter is a disk group, you don't need to enter all the host volumes or file systems because SnapDrive for UNIX knows what the disk group is connecting to _____ ◆ The <code>-autorename</code> option tells SnapDrive for UNIX to generate a new, unused name for the destination entity if the default name is in use. ◆ When you use the <code>-autoexpand</code> option without the <code>-autorename</code> option, the <code>snap connect</code> command fails if the default name for the destination copy of an LVM entity is in use. If you include the <code>-autorename</code> option, SnapDrive for UNIX renames the entity when the default name is in use. This means that with the <code>-autorename</code> option on the command line, the <code>snap connect</code> continues regardless of whether all the necessary names are available. ◆ NetApp strongly recommends that you use the default igroup to contain the LUNs being connected instead of specifying an igroup on the target storage system. Result: SnapDrive for UNIX clones the storage in the disk group you specify and imports that disk group to a new location. You can then use the snapshot in that location. Example 1: The following are command lines you can use with <code>snap connect</code>. <pre># snapdrive snap connect -hostvol srcdg/myvol destdg/myvol -snapname toaster:/vol/vol1:mynsnapshot # snapdrive snap connect -vg myvg -snapname toaster:/vol/vol1:snap1 # snapdrive snap connect -dg src_dgl dest_dgl -snapname toaster:/vol/vol1:mynsnapshot</pre>

Step	Action
	Example 2: The following command line connects a disk group and uses the default names as the destination names (that is, it creates them from the source names). <pre># snapdrive snap connect -vg vg1 -snapname filer1:/vol/vol1:vg1snapshot</pre> <pre>connecting vg1: LUN copy vg1_lun1_0 ... created (original: filer1:/vol/vol1/vg1_lun1) mapping new lun(s) ... done discovering new lun(s) ... done Importing vg1</pre> Example 3: The following command line connects a disk group with a single host volume. It also specifies a name for the destination host volume and disk group. <pre># snapdrive snap connect -lvol vg1/vol1 vg1copy/vol1copy -snapname filer1:/vol/vol1:vg1snapshot</pre> <pre>connecting vg1: LUN copy vg1_lun1_0 ... created (original: filer1:/vol/vol1/vg1_lun1) mapping new lun(s) ... done discovering new lun(s) ... done Importing vg1copy</pre> Example 4: The following command line connects a disk group with two LUNs and two file systems. It specifies a destination name for each of the file systems, the host volume for one of the file systems, and the disk groups for both file systems. <pre># snapdrive snap connect -fs mnt/fs1 /mnt/fs1copy -destvg vg1copy \ -fs /mnt/fs2 /mnt/fs2copy -destlv vg1copy/vol2copy -destvg vg1copy \ -snapname filer1:/vol/vol1:vg1snapshot</pre> <pre>connecting vg1: LUN copy vg1_lun1_0 ... created (original: filer1:/vol/vol1/vg1_lun1) LUN copy vg1_lun2_0 ... created (original: filer1:/vol/vol1/vg1_lun2) mapping new lun(s) ... done discovering new lun(s) ... done Importing vg1copy</pre>

Step	Action
	Example 5: The following command line includes the <code>autoexpand</code> option as it connects a disk group with two file systems. It uses the default names as the destination names (that is, it creates them from the source names). <pre># snapdrive snap connect -lvol mnt/fs1 -snapname filer1:/vol/vol1:vg1snapshot \ -autoexpand</pre> <pre>connecting vg1: LUN copy vg1_lun1_0 ... created (original: filer1:/vol/vol1/vg1_lun1) LUN copy vg1_lun2_0 ... created (original: filer1:/vol/vol1/vg1_lun2) mapping new lun(s) ... done discovering new lun(s) ... done Importing vg1</pre> Example 6: The following command line includes the <code>autorename</code> option as it connects a disk group with two file systems and two LUNs. <pre># snapdrive snap connect -fs mnt/fs1 -snapname filer1:/vol/vol1:vg1snapshot \ -autorename</pre> <pre>connecting vg1: LUN copy vg1_lun1_0 ... created (original: filer1:/vol/vol1/vg1_lun1) LUN copy vg1_lun2_0 ... created (original: filer1:/vol/vol1/vg1_lun2) mapping new lun(s) ... done discovering new lun(s) ... done Importing vg1_0</pre>

Disconnecting a snapshot

Using the disconnect operation

You use `snap disconnect` to remove the mappings between one or more host entities and the LUNs in a snapshot that make up the underlying storage for those entities.

You can use this command to disconnect snapshots that span multiple storage system volumes or multiple storage systems. These volumes can reside on the same storage system or different storage systems.

The disconnect operation does not modify the snapshot except to delete the temporary LUNs that were created during the `snap connect` operation.

Note

In terms of function, this command is equivalent to `snapdrive storage delete`; however, it deletes only the temporary LUNs.

Disconnecting snapshots

To disconnect snapshots, complete the following step.

Step	Action
1	Execute the <code>snapdrive snap disconnect</code> command: <pre data-bbox="206 309 1174 388">snapdrive snap disconnect {-dg -fs -hostvol } file_spec [file_spec ...] [{-dg -fs -hostvol } file_spec [file_spec ...]] [-full]</pre> Note_____ The <code>-dg</code> and <code>-vg</code> options are synonyms. They reflect the fact that some operating systems refer to disk groups and others refer to volume groups. The <code>-lvol</code> and <code>-hostvol</code> options are also synonyms. This guide uses <code>-dg</code> to refer to both disk groups and volume groups and <code>-hostvol</code> to refer to both logical volumes and host volumes. For details about all the options and arguments in this command line, see the “SnapDrive for UNIX keywords and arguments” on page 247. Keep the following in mind when you execute this command: ◆ Use the <code>-full</code> option to tell SnapDrive for UNIX to disconnect the snapshot even if a host-side entity is named that contains other entities; for example, a disk group is named that contains one or more host volumes. If you do not include the <code>-full</code> option, you can only specify empty host-side entities. For example, you would need to specify a disk group that does not have any host volumes in it, a host volume that does not contain any file systems, or a file system. ◆ Make sure the <code>file_spec</code> parameter specifies the name of the host entity you designate. You can specify a disk group (<code>-dg</code>), host volume (<code>-hostvol</code>), or file system (<code>-fs</code>). ◆ An error also occurs if you specify a subset of the host volumes and/or file systems contained in each target disk group. Result: SnapDrive for UNIX removes the mappings for the snapshot and the LUNs it contains from the host entity. Example 1: This command line removes the mappings to all the LUNs underlying the host volume <code>dg5/myvolume</code>. It removes any LUNs that were created with a <code>snap connect</code> operation. <pre data-bbox="206 1249 841 1274"># snapdrive snap disconnect -hostvol dg5/myvolume</pre> Example 2: This command line removes the mappings to all the LUNs underlying the disk groups <code>dg1</code>, <code>dg2</code>, and <code>dg3</code>. It removes any temporary LUNs that might have been created with a <code>snap connect</code> operation. <pre data-bbox="206 1409 763 1433"># snapdrive snap disconnect -dg dg1 dg2 dg3</pre>

Deleting snapshots

Command to use to delete snapshots

The `snapdrive snap delete` command removes the snapshots you specify from a storage system. This command does not perform any operations on the host; it does not remove the mappings and or destroy the disk group. It only removes the snapshot from a storage system, if you have permission to do so. (If you want to keep the LUNs and mappings, see the information on the `snapdrive storage delete` command on [“Deleting storage from the host and storage system”](#) on page 179.)

Reasons to delete snapshots

You might delete older snapshots for the following reasons:

- ◆ To keep fewer stored snapshots than the hard limit of 255 on a storage system volume
After this limit is reached, attempts to create new snapshots fail.
- ◆ To free space on the storage system volume
Even before the snapshot limit is reached, a snapshot fails if the disk does not have enough reserved space for it.

Restriction on deleting snapshots

The following limitations apply to the `snapdrive snap delete` command:

- ◆ This operation fails if any of the snapshots you want to delete are in use were not created by SnapDrive for UNIX. You can override this behavior by including the `-force` option on the command line.
- ◆ If you have a snapshot that spans multiple storage system volumes, you must manually delete the snapshot on each volume.

Deleting a snapshot To delete a snapshot, complete the following step.

Step	Action
1	Execute the <code>snapdrive snap delete</code> command. This command has the following format: <pre>snapdrive snap delete [-snapname] long_snap_name [snap_name ...] [-verbose] [-force [-noprompt]]</pre> <i>long_snap_name</i> is the long form for the name of the snapshot and includes the storage system name, volume, and snapshot name. <i>snap_name</i> can be either the long form of the name of the snapshot or the short form (just the snapshot name without path information). When you supply the short form of the name, the <code>delete</code> operation uses the storage system (and the storage system volume, if not given) last seen in the argument sequence. You can use the <code>-verbose</code> option to display a list of what is being deleted. You can also use the wildcard (*) character in snapshot names. The <code>snap show</code> operation enables you to use the wildcard character to show all snapshot names that match a certain pattern. The following rules apply to using wildcards in snapshot names: ◆ You can use a wildcard at the end of the name only. You cannot use the wildcard at the beginning or the middle of a snapshot name. ◆ You cannot use the wildcard in the storage system or storage system volume fields of a snapshot name. See the examples that follow for a sample command line that includes a wildcard. Note_____ The section “SnapDrive for UNIX keywords and arguments” on page 247 provides details about using the options and arguments in this command line. Result: SnapDrive for UNIX deletes the existing contents of the LUNs you specify in the <code>snap delete</code> command line and replaces them with the contents of the LUNs in the snapshot you specify.

Step	Action
	Example 1: This example displays a list of what is being deleted: <pre data-bbox="494 302 1206 472"># snapdrive snap delete -v filer1:/vol/vol0/snap1 snap2 snap3 snapdrive: deleting filer1:/vol/vol0/snap1 filer1:/vol/vol0/snap2 filer1:/vol/vol0/snap3</pre> Example 2: You can also use the wildcard character to specify the snapshot name: <pre data-bbox="494 600 997 649"># snap delete myfiler:/vol/vol1:mynsnap* /vol/vol2:hissnap* yoursnap* hersnap</pre>

About this chapter

This chapter provides details about working with storage provisioning. It contains information on using SnapDrive for UNIX to create and connect to storage.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Overview of storage provisioning”](#) on page 150
- ◆ [“Creating storage”](#) on page 152
- ◆ [“Displaying information about storage”](#) on page 160
- ◆ [“Increasing the size of the storage”](#) on page 166
- ◆ [“Connecting existing LUNs to the host”](#) on page 170
- ◆ [“Disconnecting LUN mappings from the host”](#) on page 172
- ◆ [“Connecting only the host-side of storage”](#) on page 176
- ◆ [“Disconnecting only the host-side of storage”](#) on page 177
- ◆ [“Deleting storage from the host and storage system”](#) on page 179

Overview of storage provisioning

Using storage provisioning with SnapDrive for UNIX

SnapDrive for UNIX provides end-to-end storage management. With it, you can provision storage from a host to a NetApp storage system and manage that storage using the volume managers on the hosts.

SnapDrive for UNIX lets you perform the following tasks:

- ◆ Create storage by creating LUNs, file systems, logical volumes, and disk groups
- ◆ Display information about storage
- ◆ Connect to storage
- ◆ Resize storage
- ◆ Disconnect from the storage
- ◆ Delete storage

When you use SnapDrive for UNIX to create storage using the `snapdrive storage create` command, it automatically performs all the tasks needed to set up LUNs, including preparing the host, performing discovery mapping, and connecting to each LUN you create.

You can use the `snapdrive storage show` command to display information about the NetApp LUNs underlying one or more disk groups, host volumes, or file systems. You can then determine what will be in a snapshot of a disk group, host volume, or file system.

SnapDrive for UNIX also lets use the `snapdrive storage connect` command to map storage to a new location.

SnapDrive for UNIX also allows you to access existing storage from a different host than the one used to create it. The `snapdrive storage connect` command lets you make one or more existing LUNs accessible on a new host. This operation can be useful if you want to back up the LUNs from the new host to another medium.

The `snapdrive storage resize` command lets you increase the size of your storage in the following ways:

- ◆ Specifying a target size that you want the host entity to reach.
- ◆ Entering a set number of bytes by which you want to increase the storage.

If at some point, you decide you no longer want your storage mapped to its current location, you can use the `snapdrive storage disconnect` command. This command removes the mappings from one or more host locations to the LUNs making up the storage for that location.

You can also delete the storage. If you execute the `snapdrive storage delete` command, SnapDrive for UNIX removes all the host-side entities you specify as well as all their underlying entities and the LUNs associated with them.

Using storage operations across multiple storage system volumes

SnapDrive for UNIX lets you perform many of the storage operations across multiple storage system volume as long as they don't manipulate the LVM. This enables you to work with lists of LUNS that exist across multiple storage system volumes.

Creating storage

Creating storage with SnapDrive for UNIX

You can use SnapDrive for UNIX to create the following:

- ◆ LUNs
- ◆ File systems on the hosts
- ◆ Logical volumes on the hosts
- ◆ Disk groups on the hosts

SnapDrive for UNIX automatically handles all the tasks needed to set up LUNs associated with these entities, including preparing the host, performing discovery mapping, and connecting to each LUN you create.

Note

You do not need to create the LUNs and create a host entity at the same time. If you create the LUNs separately, you can later make them part of a host entity by using the host LVM to add the LUNs to the LVM entities.

If you specify that SnapDrive for UNIX create a logical volume or file system, it also creates the required disk group. It creates the file system based on the types supported with the host logical volume manager (LVM). These types include the following:

- ◆ AIX: JFS and JFS2
- ◆ HP-UX: VxFS
- ◆ Linux: Ext3
- ◆ Solaris: VxFS

When you use SnapDrive for UNIX to create an entity on the host (such as a file system, logical volume, or disk group), you can also specify which LUNs it creates to provide storage for the entity.

Methods for creating storage

To make it easier to create the storage you want, SnapDrive for UNIX provides some basic formats for the `snapdrive storage create` command. This is because the `storage create` operations fall into the following general categories:

- ◆ Creating LUNs. This command automatically creates the LUNs on the storage system, but does not create any host entities. SnapDrive for UNIX performs all of the tasks associated with host preparation and discovery for

each LUN, as well as mapping and connecting to it. See [“Creating LUNS without host entities”](#) on page 153 for the steps required to do this.

- ◆ Creating an entity on the host. This command line lets you create a file system, a logical volume, or a disk group on the host. SnapDrive for UNIX performs all the actions needed to set up the entity, including automatically creating the required disk group and LUNs. You do not need to specify any LUNs for it to create. See [“Creating an entity on the host and setting up LUNs automatically”](#) on page 154 for the steps required to do this.
- ◆ Creating an entity on the host and specifying the LUNs you want associated with it. In this case, you use the command line to specify both the entity you want to set up (file system, logical volume, or disk group) and the LUNs you want associated with that entity. See [“Creating an entity on the host and specifying LUNs”](#) on page 157 for the steps required to do this.

Restrictions for using the storage create command

The `snapdrive storage create` command has the following restriction in SnapDrive for UNIX:

- ◆ You cannot create a disk group, host volume, or file system using LUNs from more than one storage system volume. If you list LUNs from different storage system volumes with the `-lun` option, you cannot include the `-dg`, `-hv`, or `-fs` option on the command line.
- ◆ Some operating systems, such as Linux and Solaris, have limits on how many LUNs you can create. If your host is running one of these operating systems, you may want to run the `storage config` commands. For more information, see [“Preparing hosts for adding LUNs”](#) on page 85.

Creating LUNS without host entities

To provision storage by creating LUNs on the storage system, complete the following step.

Step	Action
1	Execute the following <code>snapdrive storage create</code> command: <code>snapdrive storage create -lun long_lun_name [lun_name ...] -lunsize size [-igroup ig_name [ig_name ...]]</code>

Step	Action
	Keep the following in mind when you execute this command: ◆ Use the <code>-lun</code> argument to specify the storage system and volume where you want to create the LUNs. The first time you specify a LUN, use the long form of the name (i.e., the full pathname). ◆ If you want to specify additional LUNs, you can use the short form of the name if they are on the same storage system and volume as the first LUN. Otherwise, use the long form of the name again. ◆ Use the <code>-lunsize</code> argument to specify the size of each LUN. The default measure for specifying the size is bytes. ◆ By default, SnapDrive for UNIX automatically performs all of the tasks associated with host preparation and discovery for each LUN, as well as mapping and connecting to it. ◆ NetApp strongly recommends that you use the default igroup that SnapDrive for UNIX creates to contain the LUNs being connected instead of specifying an igroup on the target storage system. Result: SnapDrive for UNIX creates the LUNs you specify. Example: This example creates three LUNs on the storage system acctfiler. Each LUN is 10gb. <pre>snapdrive storage create -lun acctfiler:/vol/vol1/lunA lunB lunC -lunsize 10g</pre>

Creating an entity on the host and setting up LUNs automatically

To create an entity on the host such as a file system, logical volume or disk group and have SnapDrive for UNIX automatically create the associated LUNs, complete the following step.

Step	Action
1	Execute the following <code>snapdrive storage create</code> command: <pre>snapdrive storage create <i>host_lvm_fspect</i> -filervol <i>long_filer_path</i> -dgsiz <i>size</i> [-igroup <i>ig_name</i> [<i>ig_name...</i>]]</pre>

Step	Action
	Keep the following in mind when you execute this command: ◆ The <code>host_lvm_fspect</code> argument lets you specify whether you want to create a file system, logical volume, or disk group. This argument has three general formats that you can use individually or combine to create multiple entities. The format you use depends on the entity you want to create. To create a file system, use this format: <pre>-fs file_spec [-fstype type] [-fsopts options]</pre> To create a logical or host volume, use this format: <pre>[-hostvol file_spec] [-dg dg_name]</pre> To create a disk or volume group, use this format: <pre>-dg dg_name</pre> Note The <code>-dg</code> and <code>-vg</code> options are synonyms that reflect the fact that some operating systems refer to disk groups and others refer to volume groups. In addition, <code>-lvol</code> and <code>-hostvol</code> are also synonyms. This guide uses <code>-dg</code> to refer to both disk groups and volume groups and <code>-hostvol</code> to refer to both logical volumes and host volumes. For details about using all the options and arguments available with this command, see “SnapDrive for UNIX keywords and arguments” on page 247. If you create a file system, you may also include the host volume specifications, the disk group specifications, or both specifications to indicate the host volume and/or disk group on which the file system will be based. If you do not include these specifications, SnapDrive for UNIX automatically generates the names for the host volume and/or disk group. ◆ When you specify a host volume, SnapDrive for UNIX creates a concatenated host volume. While this is the only format SnapDrive for UNIX supports when creating host volumes, it does allow you to manipulate existing striped host volumes. ◆ When you create a file system, you can specify any type (<code>-fstype</code>) that SnapDrive for UNIX supports on the host LVM for that operating system. Current values that you can include with this argument are <code>vxfs</code> (Solaris and HP-UX), <code>jfs2</code> (AIX), or <code>ext3</code> (Linux). ◆ You can use the <code>-fsopts</code> option to specify options that you want passed to the host operation that creates the new file system (frequently, this operation is the host command <code>mkfs</code>). The argument you supply with this option usually needs to be specified as a quoted string and must contain exactly the text to be passed to the host command.

Step	Action
	◆ If you are creating an entity such as a file system, you do not need to supply a value for a disk or volume group. SnapDrive for UNIX automatically creates one. ◆ Use the <code>-filervol</code> option to specify the storage system and volume where you want the LUNs created. ◆ Do not specify any LUNs. SnapDrive for UNIX creates them automatically when you use this form of the <code>snapdrive storage create</code> command. It uses system defaults to determine the number of LUNs, the LUN IDs, and the size of each LUN. It bases the names of the associated disk/volume groups on the name of the host volume or file system. ◆ To control the size of the host entity, use the <code>-dgsiz</code> option to specify the size in bytes of the underlying disk group. ◆ NetApp strongly recommends that you use the default igroup that SnapDrive for UNIX creates to contain the LUNs being connected instead of specifying an igroup on the target storage system. Result: SnapDrive for UNIX creates the host entity you specify and creates LUNs for it on the storage system you specify. It performs all of the tasks associated with host preparation and discovery for each of the LUNs, as well as mapping and connecting the LUNs to the host entity. Example 1: This example creates the file system <code>acctfs</code> with a Solaris file type of <code>VxFS</code>. It sets up LUNs on the storage system <code>acctfiler</code> and creates a disk group that is 1gb. <pre>snapdrive storage create -fs /mnt/acctfs -fstype vxfs -filervol acctfiler:/vol/acct -dgsiz 1g</pre> Example 2: This example on an HP-UX host creates the file system <code>acctfs</code> and specifies only the storage system volume. SnapDrive for UNIX internally creates and names the entities under the file system. <pre>snapdrive storage create -fs /mnt/acctfs -filervol acctfiler:/vol/vol1 -vgsiz 10g</pre> <pre> LUN acctfs_SdLun ... created mapping new lun(s) ... done discovering new lun(s) ... done LUN to device file mappings: - acctfiler:/vol/vol1/acctfs_SdLun => /dev/rdisk/c9t0d6 disk group acctfs_SdDg created host volume acctfs_SdHv created file system /mnt/acctfs created </pre>

Creating an entity on the host and specifying LUNs

If you want to both create a host entity such as a file system, logical volume, or disk group and specify the LUNs that are created as part of it, complete the following step.

Step	Action
1	Execute the following <code>snapdrive storage create</code> command: <pre>snapdrive storage create host_lvm_fspect -lun long_lun_name [lun_name ...] -lunsize size [-igroup ig_name [ig_name ...]]</pre> Keep the following in mind when you execute this command: ◆ The <code>host_lvm_fspect</code> argument lets you specify whether you want to create a file system, logical volume, or disk group. This argument has the three formats. The format you use depends on the entity you want to create. To create a file system, use this format: <pre>-fs file_spec [-fstype type] [-fsopts options]</pre> To create a host volume, use this format: <pre>[-hostvol file_spec] [-dg dg_name]</pre> To create a disk/volume group, use this format: <pre>-dg dg_name</pre> Note The <code>-dg</code> and <code>-vg</code> options are synonyms that reflect the fact that some operating systems refer to disk groups and others refer to volume groups. In addition, <code>-lvol</code> and <code>-hostvol</code> are also synonyms. This guide uses <code>-dg</code> to refer to both disk groups and volume groups and <code>-hostvol</code> to refer to both logical volumes and host volumes. For details about using all the options and arguments available with this command, see “SnapDrive for UNIX keywords and arguments” on page 247. ◆ When you create a file system, you can specify any type (<code>-fstype</code>) that SnapDrive for UNIX supports on the host LVM for that operating system. Current values that you can include with this argument are <code>vxfs</code> (Solaris and HP-UX), <code>jfs2</code> (AIX), or <code>ext3</code> (Linux). ◆ You can use the <code>-fsopts</code> option to specify options that you want passed to the host operation that creates the new file system (frequently, this operation is the host command <code>mkfs</code>). The argument you supply with this option usually needs to be specified as a quoted string and must contain exactly the text to be passed to the host command. ◆ SnapDrive for UNIX creates a disk/volume group to hold the LUNs based on the value you enter with the <code>-dg</code> option. The name you supply for the group must not exist. ◆ Use the <code>-lun</code> argument to specify the storage system and volume where you want to create the LUNs.

Step	Action
	◆ Use the <code>-lunsize</code> argument to specify the size in bytes of each LUN. The size of the LVM entity depends on the aggregated size of the LUNs you request. Result: SnapDrive for UNIX creates the host entity and the LUNs you specify. Example 1: This example creates the file system <code>acctfs</code> with an AIX file system type of <code>jfs2</code>. It sets up three LUNs on the storage system <code>acctfiler</code>. Each LUN is 10gb. <pre># snapdrive storage create -fs /mnt/acctfs -fstype jfs2 -lun acctfiler:/vol/vol1/lunA lunB lunC -lunsize 10g LUN acctfiler:/vol/vol1/lunA ... created LUN acctfiler:/vol/vol1/lunB ... created LUN acctfiler:/vol/vol1/lunC ... created mapping new lun(s) ... done discovering new lun(s) ... done LUN to device file mappings: - acctfiler:/vol/vol1/lunA => hdisk2 - acctfiler:/vol/vol1/lunB => hdisk3 - acctfiler:/vol/vol1/lunC => hdisk4 disk group acctfs_SdDg created host volume acctfs_SdHv created file system /mnt/acctfs created</pre>

Step	Action
	Example 2: This example on a Solaris host creates the file system acctfs on three LUNs and explicitly names the volume group and host volume underneath it. Each LUN is 10gb. <pre># snapdrive storage create -fs /mnt/acctfs -hostvol acctfsdg/acctfshv -lun acctfiler:/vol/vol1/lunA lunB lunC -lunsize 10g LUN acctfiler:/vol/vol1/lunA ... created LUN acctfiler:/vol/vol1/lunB ... created LUN acctfiler:/vol/vol1/lunC ... created mapping new lun(s) ... done discovering new lun(s) ... done LUN to device file mappings: - acctfiler:/vol/vol1/lunA => /dev/vx/rdmp/c4t0d3s2 - acctfiler:/vol/vol1/lunB => /dev/vx/rdmp/c4t0d7s2 - acctfiler:/vol/vol1/lunC => /dev/vx/rdmp/c4t0d8s2 disk group acctfsvg created host volume acctfshv created file system /mnt/acctfs created</pre>

Displaying information about storage

Command to use to display available storage

The `snapdrive storage show` (or `list`) command shows the NetApp LUNs underlying one or more disk groups, host volumes, or file systems. It is a good way to find out what will be in a snapshot of a disk group, host volume, or file system.

With this command, you can display the following:

- ◆ LUNs available for specific storage systems or storage system volumes
- ◆ LUNs associated with the file systems, host volumes or disk or volume groups you specify.

Methods for displaying storage information

To make it easier to display information about storage, SnapDrive for UNIX provides several formats for the `snapdrive storage show` command. This is because the storage show operations fall into the following general categories:

- ◆ Displaying information about a specific LUN. See [“Displaying information about a specific LUN”](#) on page 161 for the steps to do this.
- ◆ Listing information about LUNs available for specific storage systems or storage system volumes. See [“Displaying information based on the storage system”](#) on page 162 for the steps required to do this.
- ◆ Displaying information about LUNs associated with the arguments you specify. These arguments can include file systems, host volumes, or disk or volume groups. If you use the `-verbose` option on the command line, SnapDrive for UNIX provides detailed output, such as showing the storage hierarchy including the backing LUNs. See [“Displaying LUN information based on a host entity”](#) on page 163 for the steps required to do this.

Restrictions for using the storage show command

The `snapdrive storage show` command has the following restrictions in SnapDrive for UNIX:

- ◆ The adapter name is not displayed.
- ◆ The status of non-primary paths is displayed as a hyphen (-) instead of secondary (S).
- ◆ If there is no volume manager or host volume underlying a file system, SnapDrive for UNIX displays the following:

```
# snapdrive storage show -fs /pub
fs: /dev/rdsk/clt0d1s2          mountpoint: /pub
```

```

device filename adapter path size    state    lun path
-----
/dev/rdsdsk/c1t0d1s2 -    P    10g    online  myfiler:/vol0/lun-1
/dev/rdsdsk/c2t0d1s2 -          10g    online  myfiler:/vol0/lun-1

```

- ◆ The storage show operation displays the same basic output for storage systems, storage system volumes, LUNs, or raw disks. It limits the output displayed to only the storage systems, storage system volumes, LUNs, or raw disks. If a disk group has LUNs from more than just the specified storage system, storage system filer volume, LUN or raw disk, the operation displays output for all.

Displaying information about a specific LUN

To display information about specific LUNs, complete the following step.

Step	Action
1	Execute the <code>snapdrive storage show list</code> command. To display information about a LUN, use the following format: <pre> snapdrive storage { show list } -lun long_lun_name [lun_name ...] [-verbose] </pre> Note _____ You can use either <code>snapdrive storage show</code> or <code>snapdrive storage list</code> in the command line. These commands are synonyms. For details about using all the options and arguments available with this command, see “SnapDrive for UNIX keywords and arguments” on page 247. Keep the following in mind when you execute this command: ◆ It might take a few minutes for SnapDrive for UNIX to display information. This is because it must gather information. Result: Depending on the command options you select, SnapDrive for UNIX either displays basic information about the LUNs or lists the LUNs associated with the host entity. Examples: The following are examples of possible command lines. <pre> # snapdrive storage show -verbose -lun BigFiler3:/vol/vol1:lunA lunB lunC </pre>

Displaying information based on the storage system

To display information about the available LUNs on a certain storage system or to get a list of the LUNs, complete the following step.

Step	Action
1	Execute the <code>snapdrive storage show</code> command. This command has the following formats: <pre> snapdrive storage { show list } -filer <i>filer_name</i> [<i>filer_name...</i>] [-verbose] snapdrive storage { show list } -filervol <i>long_filer_path</i> [<i>filer_path ...</i>] [-verbose]</pre> Note _____ You can use either <code>snapdrive storage show</code> or <code>snapdrive storage list</code> in the command line. These commands are synonyms. For details about using all the options and arguments available with this command, see “SnapDrive for UNIX keywords and arguments” on page 247. _____ Keep the following in mind when you execute this command: ◆ If you want to find out about LUNs on the storage system, use the <code>-filer</code> option. Use the <code>-filervol</code> option if you want to learn about the LUNs available on the storage system volume. ◆ It might take a few minutes for SnapDrive for UNIX to display information. This is because it must gather information. Result: Depending on the command options you select, SnapDrive for UNIX either displays basic information about the LUNs or lists the LUNs associated with the host entity. Examples: The following examples of possible command lines. <pre> # snapdrive storage show -verbose -filer BigFiler3 BigFiler4 # snapdrive storage show -verbose -filervol BigFiler3:/vol/vol1</pre>

Displaying LUN information based on a host entity

To display information about the LUNs associated with a host entity such as a file system, disk group, volume group, or host volume, complete the following step.

Step	Action
1	Execute the <code>snapdrive storage show</code> command: <pre>snapdrive storage {show list} {-dg -fs -hostvol } file_spec [file_spec ...] [{-dg -fs -hostvol } file_spec [file_spec ...]] [-verbose]</pre> Note _____ You can use either <code>snapdrive storage show</code> or <code>snapdrive storage list</code> in the command line. These commands are synonyms. For details about using all the options and arguments available with this command, see “SnapDrive for UNIX keywords and arguments” on page 247. Keep the following in mind when you execute this command: ◆ This form of the command provides information about the specified disk groups, volume groups, file systems, or host volumes, so you must supply a value for either <code>-dg</code>, <code>-fs</code>, or <code>-hostvol</code>. SnapDrive for UNIX looks for this entity on the current host. Note _____ The <code>-dg</code> and <code>-vg</code> options are synonyms that reflect the fact that some operating systems refer to disk groups and others refer to volume groups. In addition, <code>-lvol</code> and <code>-hostvol</code> are also synonyms. This guide uses <code>-dg</code> to refer to both disk groups and volume groups and <code>-hostvol</code> to refer to both logical volumes and host volumes. For details about using all the options and arguments available with this command, see “SnapDrive for UNIX keywords and arguments” on page 247. ◆ It might take a few minutes for SnapDrive for UNIX to display information. This is because it must gather information. Result: SnapDrive displays basic information about the host. This information is the same regardless of whether it is for a disk group, a host volume, or a file system. For example, SnapDrive displays the following output for each disk group: ◆ The name of the disk group. ◆ The name of each host volume in the disk group and what its state is (for example, available). Each host volume is displayed on a separate line.

Step	Action																		
	◆ Each file system associated with the host volume (or volumes) and its mount point. Each file system is displayed on a separate line. If the file system is not hosted on a host volume (and therefore not part of a disk group), no information on the file system is displayed. ◆ Information about the LUNs: ❖ The device file name for the LUN ❖ Whether the path is primary (P) or secondary (-) ❖ The LUN state, such as online ❖ The path to the LUN Example 1: These are examples of command lines that use the <code>storage show</code> command. <pre># snapdrive storage list -dg dg1 dg2 # snapdrive storage show -quiet -dg QA_dg_test # snapdrive storage show -fs /mnt/fs21 /mnt/fs41 # snapdrive storage show -hostvol dg2/vol1 dg4/myvol3 # snapdrive storage list -dg dg2 -fs /mnt/fs42 # snapdrive storage show -hostvol dg2/vol3 -fs /mnt/fs42</pre> Example 2: This example shows output that appears when you display information about the disk group <code>dg1</code> on a Solaris host (the output varies slightly depending on which host you use): <pre>snapdrive storage show -dg dg1 dg: dg1 hostvol: /dev/vx/dsk/dg1/vol1 state: AVAIL hostvol: /dev/vx/dsk/dg1/vol3 state: AVAIL hostvol: /dev/vx/dsk/dg1/vol4 state: AVAIL fs: /dev/vx/dsk/dg1/vol1 mount point: /pub fs: /dev/vx/dsk/dg1/vol3 mount point: NOT MOUNTED fs: /dev/vx/dsk/dg1/vol4 mount point: NOT MOUNTED</pre> <table><thead><tr><th>device filename</th><th>adapter</th><th>path</th><th>size</th><th>state</th><th>lun path</th></tr></thead><tbody><tr><td>/dev/vx/rdmp/c1t0d2s2</td><td>-</td><td>P</td><td>10g</td><td>online</td><td>myfiler:/vol/vol1/lun2</td></tr><tr><td>/dev/vx/rdmp/c1t0d7s2</td><td>-</td><td>P</td><td>10g</td><td>online</td><td>myfiler:/vol/vol1/lun7</td></tr></tbody></table>	device filename	adapter	path	size	state	lun path	/dev/vx/rdmp/c1t0d2s2	-	P	10g	online	myfiler:/vol/vol1/lun2	/dev/vx/rdmp/c1t0d7s2	-	P	10g	online	myfiler:/vol/vol1/lun7
device filename	adapter	path	size	state	lun path														
/dev/vx/rdmp/c1t0d2s2	-	P	10g	online	myfiler:/vol/vol1/lun2														
/dev/vx/rdmp/c1t0d7s2	-	P	10g	online	myfiler:/vol/vol1/lun7														

Step	Action																																																
	Example 3: This example uses an HP-UX host. In this example, SnapDrive for UNIX displays information about the volume group vg01. <pre># snapdrive storage show -vg vg01 dg: vg01 hostvol: /dev/vg01/lvol1 state: AVAIL fs: /dev/vg01/lvol1 mount point: NOT MOUNTED</pre> <table><tr><th>device filename</th><th>adapter</th><th>path</th><th>size</th><th>state</th><th>lun path</th></tr><tr><td>/dev/rdsk/c22t0d0</td><td>-</td><td>P</td><td>10g</td><td>online</td><td>toaster:/vol/vol1/myhp7-1</td></tr><tr><td>/dev/rdsk/c20t0d2</td><td>-</td><td>P</td><td>5g</td><td>online</td><td>toaster:/vol/vol1/myhp7-3</td></tr></table> Example 4: This example uses a Solaris host. In this example, SnapDrive for UNIX displays information about the disk group dg1. <pre># snapdrive storage show -vg dg1 saving name dg1 dg: dg1 hostvol: /dev/vx/dsk/dg1/vol1 state: AVAIL fs: /dev/vx/dsk/dg1/vol1 mount point: NOT MOUNTED</pre> <table><tr><th>device filename</th><th>adapter</th><th>path</th><th>size</th><th>state</th><th>lun path</th></tr><tr><td>/dev/vx/rdmp/c1t0d0s2</td><td>-</td><td>P</td><td>500m</td><td>online</td><td>sky:/vol/vol1/matilda_lun0</td></tr></table> Example 5: This example uses an AIX host. In this example, SnapDrive for UNIX displays information about the disk group dg1. <pre># snapdrive storage show -dg dg1 dg: dg1 hostvol: /dev/dg1_log state: AVAIL hostvol: /dev/dg1_voll state: AVAIL fs: /dev/dg1_log mount point: MOUNTED fs: /dev/dg1_voll mount point: MOUNTED</pre> <table><tr><th>device filename</th><th>adapter</th><th>path</th><th>size</th><th>state</th><th>lun path</th></tr><tr><td>/dev/rhdisk2</td><td>-</td><td>P</td><td>2g</td><td>online</td><td>flip:/vol/vol1/nate_nane-aix5_lun0</td></tr><tr><td>/dev/rhdisk3</td><td>-</td><td>P</td><td>2g</td><td>online</td><td>flip:/vol/vol1/nate_nane-aix5_lun1</td></tr></table>	device filename	adapter	path	size	state	lun path	/dev/rdsk/c22t0d0	-	P	10g	online	toaster:/vol/vol1/myhp7-1	/dev/rdsk/c20t0d2	-	P	5g	online	toaster:/vol/vol1/myhp7-3	device filename	adapter	path	size	state	lun path	/dev/vx/rdmp/c1t0d0s2	-	P	500m	online	sky:/vol/vol1/matilda_lun0	device filename	adapter	path	size	state	lun path	/dev/rhdisk2	-	P	2g	online	flip:/vol/vol1/nate_nane-aix5_lun0	/dev/rhdisk3	-	P	2g	online	flip:/vol/vol1/nate_nane-aix5_lun1
device filename	adapter	path	size	state	lun path																																												
/dev/rdsk/c22t0d0	-	P	10g	online	toaster:/vol/vol1/myhp7-1																																												
/dev/rdsk/c20t0d2	-	P	5g	online	toaster:/vol/vol1/myhp7-3																																												
device filename	adapter	path	size	state	lun path																																												
/dev/vx/rdmp/c1t0d0s2	-	P	500m	online	sky:/vol/vol1/matilda_lun0																																												
device filename	adapter	path	size	state	lun path																																												
/dev/rhdisk2	-	P	2g	online	flip:/vol/vol1/nate_nane-aix5_lun0																																												
/dev/rhdisk3	-	P	2g	online	flip:/vol/vol1/nate_nane-aix5_lun1																																												

Increasing the size of the storage

Resizing the storage to make it larger

SnapDrive for UNIX lets you increase the size of the storage system volume group or disk group. You use the `snapdrive storage resize` command to do this.

Note

This command does not let you resize host volumes or file systems. You need to use the LVM commands to perform those tasks after you have resized the underlying disk group. For information on performing those tasks, see [“Resizing host volumes and file systems”](#) on page 169.

You can put the storage resize operations into the following general categories:

- ◆ Setting a target size in bytes to which you want to increase the storage.
- ◆ Specifying a number of bytes by which you want to increase the storage.

SnapDrive for UNIX adds a system-generated LUN. If you specify an amount by which you want to increase the storage, such as 50MB, it makes the LUN 50MB. If you specify a target size for the storage, it calculates the difference between the current size and the target size. The difference becomes the size of the LUN it then creates.

Limitations on using storage resize

SnapDrive for UNIX has the following limitations on using the storage resize command:

- ◆ The storage resize operation can only increase the size of storage. You cannot use it to decrease the size of an entity.
- ◆ All LUNs must reside in the same storage system volume.
- ◆ The resize operation is not supported directly on logical host volumes or file systems. In those cases, you must use the LVM commands to resize the storage.
- ◆ You cannot resize a LUN; you must use the `-addlun` option to add a new LUN.

Increasing the size of the storage To increase the size of your storage, complete the following step.

Step	Action
1	Execute the <code>snapdrive storage resize</code> command: <pre>snapdrive storage resize -dg file_spec { -growby -growto } size [-addlun [-igroup ig_name [ig_name ...]]]</pre> Note_____ The <code>-dg</code> and <code>-vg</code> options are synonyms that reflect the fact that some operating systems refer to disk groups and others refer to volume groups. This guide uses <code>-dg</code> to refer to both disk groups and volume groups. For details about using all the options and arguments available with this command, see “SnapDrive for UNIX keywords and arguments” on page 247. _____ Keep the following in mind as you use this command: ◆ You must use the <code>-addlun</code> option to increase the entity size. This option adds a new, internally-generated LUN to the underlying disk group. ◆ Use the <code>-growby</code> option to add <i>size</i> bytes to the entity. ◆ Use the <code>-growto</code> option to increase the size of the entity by the value specified in the <i>size</i> argument. You can enter this value as bytes, kilobytes, megabytes, or terabytes. ◆ NetApp strongly recommends that you use the default igroup that SnapDrive for UNIX creates to contain the LUNs being connected instead of specifying an igroup on the target storage system. Result: This command increases the size of the host-side <i>file_spec</i> entity (logical volume or disk group) either by the number of bytes you specify or to the number of bytes you specify. ◆ Adding bytes to storage (<code>-growby</code>) ◆ Increasing it to the byte size you specify (<code>-growto</code>). Example 1: This command line increases the size of the disk group <code>my_dg</code> by having it increase to 155mb. In this case, SnapDrive for UNIX automatically increases the size of the last LUN in the disk group until the accumulated disk group size is 155mb. <pre># snapdrive storage resize -dg my_dg -growto 155m discovering filer LUNs in disk group my_dg...done resizing LUN x-parrot:/vol/vol0/my_dg_1_SdLun...done adapting disk group my_dg to lun resize ...done Disk group my_dg has been resized #</pre>

Step	Action
	Example 2: This command line is the same one as in example 1, except that it uses the <code>-addlun</code> option to increase the size of the disk group <code>my_dg</code> to 155mb. When you supply the <code>-addlun</code> option, SnapDrive for UNIX automatically creates a new LUN in the disk group. It makes the LUN large enough so that the accumulated disk group size reaches 155mb. You do not need to specify any values for the <code>-addlun</code> option; SnapDrive for UNIX sets it up automatically. <pre># snapdrive storage resize -dg my_dg -addlun -growto 155m discovering filer LUNs in disk group my_dg...done LUN x-parrot:/vol/vol0/my_dg_2_SdLun ... created mapping new lun(s) ... done discovering new lun(s) ... done. mapping LUN(s)...done. initializing LUN(s) and adding to disk group my_dg...done Disk group my_dg has been resized #</pre> Example 3: This command line uses the <code>-growby</code> option and increases the size of the disk group <code>testdg</code> to 100m. <pre># snapdrive storage resize -dg testdg -growby 100m discovering filer LUNs in disk group testdg...done resizing LUN toaster:/vol/vol0/testdg_SdLun...done adapting disk group testdg to lun resize ...done Disk group testdg has been resized #</pre> Example 4: This command line uses the <code>-growby</code> option with the <code>-addlun</code> option. SnapDrive for UNIX increases the size of the storage by adding a LUN that is 100MB. <pre># snapdrive storage resize -dg testdg -growby 100m -addlun discovering filer LUNs in disk group testdg...done LUN toaster:/vol/vol0/testdg_1_SdLun ... created mapping new lun(s) ... done discovering new lun(s) ... done. mapping LUN(s)...done. initializing LUN(s) and adding to disk group testdg...done Disk group testdg has been resized #</pre>

Resizing host volumes and file systems

The `snappedrive storage resize` command applies only to storage system disk groups and volume groups. If you want to increase the size of your host volume or file system, you must use LVM commands. The following table summarizes the LVM commands you can use on the different platforms. For more information on these commands, see their man pages.

Host	Host volume	File systems
AIX	<code>extendlv</code>	<code>chfs</code>
HP-UX	<code>lvextend</code>	<code>extendfs/fsadm</code>
Linux	<code>lvextend</code>	<code>resize2fs</code>
Solaris	<code>vxassist</code>	<code>fsadm</code>

Connecting existing LUNs to the host

Using the storage connect command

The `storage connect` operation maps existing LUNs to the host. When you execute the `snapdrive storage connect` command, SnapDrive for UNIX performs the necessary discovery on the host to connect the host to the LUNs. It does not modify LUN contents.

Note SnapDrive for UNIX does not reconstruct LVM entities such as disk groups and their contents that are present on the LUNs. To reconstruct LVM entities, you must use host-side LVM commands.

Connecting storage to the host

To use the `storage connect` command to map LUNs to the host, complete the following step.

Step	Action
1	Execute the <code>snapdrive storage connect</code> command. This command has the following format: <pre>snapdrive storage connect -lun long_lun_name [lun_name ...] [-igroup ig_name [ig_name ...]]</pre> Note For details about using the options and arguments in these command line, see the “SnapDrive for UNIX keywords and arguments” on page 247. ◆ The first value you supply with the <code>-lun</code> option must include the storage system name, volume, and LUN name. To connect multiple LUNs on the same volume, you can use relative pathnames for the <code>-lun</code> option after you supply the complete information in the first pathname. When SnapDrive for UNIX encounters a relative pathname, it looks for the LUN on the same volume as the first LUN. To connect additional LUNs that are not on the same volume, enter the full pathname to each LUN. ◆ NetApp strongly recommends that you use the default igroup that SnapDrive for UNIX creates to contain the LUNs being connected instead of specifying an igroup on the target storage system.

Step	Action
	Example 1: The following command line connects a single LUN (toaster_lun2) on the storage system filer1 to the host <pre># snapdrive storage connect -lun filer1:/vol/vol1/toaster_lun2</pre> Example 2: This command line connects multiple LUNs from the same storage system to the host. Because toaster_lun3 and toaster_lun4 don't specify full pathnames, SnapDrive for UNIX looks for them on the same volume of storage system filer1 as toaster_lun2. <pre># snapdrive storage connect -lun filer1:/vol/vol1/toaster_lun2 toaster_lun3 toaster_lun4</pre> Example 3: This command line connects LUNs from different storage systems to the host. <pre># snapdrive storage connect -lun filerA:/vol/vol1/lunA filerB:/vol/vol1/lunB</pre>

Disconnecting LUN mappings from the host

Using storage disconnect

The `storage disconnect` operation removes from the host entities the LUN mappings that were established during a `storage create` or `storage connect` operation.

To use this command you can either specify the LUNs themselves or the disk group, file system or host volume to which they belong.

When SnapDrive for UNIX removes the LUN mappings, it exports the disk groups. This action marks the disk groups as having been exported. That is the only change that disconnecting the mappings has on the contents of the LUNs. It does not change the contents in any other way.

Methods for disconnecting storage

To make it easier to disconnect the storage, SnapDrive for UNIX provides several formats for the `snapdrive storage disconnect` command. This is because the disconnect operations fall into the following general categories:

- ◆ Specifying the LUNs that you want to disconnect from the host. See [“Disconnecting specific LUNs from the host”](#) on page 173 for the steps required to do this.
- ◆ Specifying a disk group or other host entity. SnapDrive for UNIX then disconnects all the LUNs associated with that host. See [“Disconnecting LUNs associated with host entities”](#) on page 174 for the steps required to do this.

Limitations for using the storage disconnect command

The `snapdrive storage disconnect` command has the following restrictions in SnapDrive for UNIX:

- ◆ If you use the `-lun` option to specify the name of a LUN that is a member of a host disk group, the `storage disconnect` command fails.
- ◆ If you use `-lun` option to specify the name of the LUN that is not discovered by multipathing software on the host, the `storage disconnect` command fails. For example, on Solaris hosts, the LUN has to be under DMP control. In other words, the LUN has to have a corresponding `/dev/vx/dmp` device.

**Tips for using
storage disconnect**

When you use the `storage disconnect` command on some operating systems, you lose information such as the host volume names, the file system mount point, the storage system volume names, and the names of the LUNs. Without this information, reconnecting the storage at a later point in time is difficult.

To avoid this potential problems, NetApp recommends that you first create a snapshot of the storage using the `snapdrive snap create` command **before** you execute the `snapdrive storage disconnect` command.

That way, if you want to reconnect the storage later, you can use the following workaround:

- 1. Execute the command the following command:

```
snapdrive snap restore filespec -snapname long_snap_name
```

Include the full path to the snapshot in this command
- 2. Now you can remove the snapshot, if you choose, by executing the `snapdrive snap delete` command.

**Disconnecting
specific LUNs from
the host**

To use the `storage disconnect` command to remove the mappings for the LUNS you specify, complete the following step.

Step	Action
1	Execute the <code>snapdrive storage disconnect</code> command. <pre>snapdrive storage disconnect -lun long_lun_name [lun_name...]</pre> Note _____ For details about using the options and arguments in these command line, see the “SnapDrive for UNIX keywords and arguments” on page 247. _____ Keep the following in mind as you use this command: ◆ The first value you supply with the <code>-lun</code> option must include the storage system name, volume, and LUN name. To disconnect multiple LUNs on the same volume, you can use relative pathnames for the <code>-lun</code> option after you supply the complete information in the first pathname. When SnapDrive for UNIX encounters a relative pathname, it looks for the LUN on the same volume as the first LUN. To disconnect additional LUNs that are not on the same volume, enter the full pathname to each LUN.

Step	Action
	Example 1: The following command line disconnects a single LUN (toaster_lun2) on the storage system filer1. <pre># snapdrive storage disconnect -lun filer1:/vol/vol1/toaster_lun2</pre> Example 2: This command line disconnects LUNs from different storage systems to the host. <pre># snapdrive storage disconnect -lun filerA:/vol/vol1/lunA filerB:/vol/vol1/lunB</pre>

Disconnecting LUNs associated with host entities

To use the storage disconnect command to remove the mappings for the LUNS associated with specific host entities such as disk groups, complete the following step.

Step	Action
1	Execute the snapdrive storage disconnect command: <pre>snapdrive storage disconnect -dg <i>file_spec</i> [<i>file_spec</i> ...] [{-dg -fs -hostvol} <i>file_spec</i> [<i>file_spec</i> ...] ...] [-full]</pre> Note— The -dg and -vg options are synonyms that reflect the fact that some operating systems refer to disk groups and others refer to volume groups. In addition, -lvol and -hostvol are also synonyms. This guide uses -dg to refer to both disk groups and volume groups and -hostvol to refer to both logical volumes and host volumes. For details about using all the options and arguments available with this command, see “SnapDrive for UNIX keywords and arguments” on page 247. Keep the following in mind as you use this command: ◆ This form of the command disables the volume group or disk group you specify and disconnects the LUNs associated with it. ◆ If you include the -full option, SnapDrive for UNIX disconnects host-side entities that contain other entities, such as a disk group containing one or more host volumes. Without the -full option, this command fails if any of the specified entities contain other entities. For example, the command fails if there are disk groups that contain host volumes or host volumes that contain file systems.

Step	Action
	Example 1: The following command line disconnects a single LUN (toaster_lun2) on the storage system filer1. <pre># snapdrive storage disconnect -lun filer1:/vol/vol1/toaster_lun2</pre> Example 2: This command line disconnects LUNs from different storage systems to the host. <pre># snapdrive storage disconnect -lun filerA:/vol/vol1/lunA filerB:/vol/vol1/lunB</pre>

Connecting only the host-side of storage

Connecting host entities to storage

The `snapdrive host connect` command connects the LUNs you specify to the host without changing their state on their state on the storage system. This command performs the host portion of the actions taken by the `storage connect` command. As a result, it assumes that the LUNs have previously been mapped on the storage system to the current host.

You can use this command to reconnect LUNs that were disconnected by executing a `host disconnect` command. You can also use this command to connect to the host LUNs that have been fully-provisioned and mapped to the host on the storage system outside SnapDrive for UNIX.

Using host connect

To use the `host connect` command to map LUNs to the host, complete the following step.

Step	Action
1	Execute the <code>snapdrive host disconnect</code> command: <pre>snapdrive host connect -lun long_lun_name [lun_name ...]</pre> Note For details about using the options and arguments in these command line, see the “SnapDrive for UNIX keywords and arguments” on page 247. Keep the following in mind as you use this command: Specify the LUNs that you want to connect to the host. The first value you supply with the <code>-lun</code> option must include the storage system name, volume, and LUN name. To connect multiple LUNs on the same volume, you can use relative pathnames for the <code>-lun</code> option after you supply the complete information in the first pathname. When SnapDrive for UNIX encounters a relative pathname, it looks for the LUN on the same volume as the first LUN. To connect additional LUNs that are not on the same volume, enter the full pathname to each LUN. Example: The following command line connects the LUN lunA. <pre># snapdrive host connect -lun acctfiler:/vol/vol1/lunA</pre>

Disconnecting only the host-side of storage

Disconnecting host entities from storage

The `snapdrive host disconnect` command disconnects the host-side entities you specify without changing the state on the storage system where the underlying LUNs reside.

This command performs the host portion of the `storage disconnect` actions. Because it does not perform the storage system actions, you do not risk disconnecting other hosts from the LUNS if all the hosts share an igroup on the storage system. It only affects the host.

Tips for using host disconnect

When you use the `snapdrive host disconnect` command on some operating systems, you lose information such as the host volume names, the file system mount point, the storage system volume names, and the names of the LUNs. Without this information, reconnecting the storage at a later point in time is difficult.

To avoid this potential problems, NetApp strongly recommends that you first create a snapshot of the storage using the `snapdrive snap create` command **before** you execute the `snapdrive host disconnect` command.

That way, if you want to reconnect the storage later, you can use the following workaround:

1. First execute the following command:

```
snapdrive host connect -lun long_lun_name
```

2. Now execute the command:

```
snapdrive snap restore filespec -snapname long_snap_name
```

Include the full path to the snapshot in this command

3. Now you can remove the snapshot, if you choose, by executing the `snapdrive snap delete` command.

Using host
disconnect

To use the `host disconnect` command, complete the following step.

Step	Action
1	Execute the <code>snapdrive host disconnect</code> command: <pre>snapdrive host disconnect {-dg -fs -hostvol} file_spec [file_spec ...] [{-dg -fs -hostvol } file_spec [file_spec ...] ...] [-full]</pre> Note _____ For details about using the options and arguments in these command line, see the “SnapDrive for UNIX keywords and arguments” on page 247. Keep the following in mind as you use this command: ◆ You can disconnect a disk group, file system, or host volume. Enter the name of the entity as the value for the <i>file_spec</i> argument. The command disables the entity you specify and disconnects the LUNs associated with it. Note _____ The <code>-dg</code> and <code>-vg</code> options are synonyms that reflect the fact that some operating systems refer to disk groups and others refer to volume groups. In addition, <code>-lvol</code> and <code>-hostvol</code> are also synonyms. This guide uses <code>-dg</code> to refer to both disk groups and volume groups and <code>-hostvol</code> to refer to both logical volumes and host volumes. For details about using all the options and arguments available with this command, see “SnapDrive for UNIX keywords and arguments” on page 247. ◆ If you include the <code>-full</code> option, SnapDrive for UNIX disconnects host-side entities that contain other entities, such as a disk group containing one or more host volumes. Without the <code>-full</code> option, this command fails if any of the specified entities contain other entities. For example, the command fails if there are disk groups that contain host volumes or host volumes that contain file systems. Example 1: The following command line disconnects the disk groups <code>dg1</code>, <code>dg2</code>, and <code>dg3</code>. <pre># snapdrive host disconnect -dg dg1 dg2 dg3</pre> Example 2: This command line disconnects the logical volume <code>dg1/mylvoll1</code>. <pre># snapdrive host disconnect -lvol dg1/mylvoll1</pre>

Deleting storage from the host and storage system

Using storage delete

The `snapdrive storage delete` command removes the storage entities on the host as well as all underlying host-side entities and storage system LUNs backing them.

Note

This command deletes data. Use caution in running it.

Methods for deleting storage

To make it easier to delete the storage, SnapDrive for UNIX provides several formats for the `snapdrive storage delete` command. This is because the delete operations fall into the following general categories:

- ◆ Specifying the LUNs that you want to delete from the host. See “[Deleting specific LUNS](#)” on page 180 for the steps required to do this.
- ◆ Specifying the host-side entity for which you want to delete the storage, including all associated host-side objects and all storage system LUNs. See “[Deleting specific host entities](#)” on page 181 for the steps required to do this.

Limitations for using the storage delete command

The `snapdrive storage delete` command has the following restrictions in SnapDrive for UNIX:

- ◆ If you use the `-lun` option to specify the name of a LUN that is a member of a host disk group, the `storage delete` command fails.
- ◆ If you use `-lun` option to specify the name of the LUN that is not discovered by multipathing software on the host, the `storage delete` command fails. For example, on Solaris hosts, the LUN has to be under DMP control. In other words, the LUN has to have a corresponding `/dev/vx/dmp` device.

Deleting specific LUNS

To delete specific LUNs, complete the following step.

Step	Action
1	Execute the snapdrive storage delete command. <pre>snapdrive storage delete -lun long_lun_name [lun_name ...]</pre> Note For details about using the options and arguments in these command line, see the “SnapDrive for UNIX keywords and arguments” on page 247. Keep the following in mind as you use this command. ◆ This form of the command deletes the LUNs you specify. You can use the command to delete LUNs on a specific storage system and volume or on multiple storage systems. Note This command deletes data. Use caution in running it. ◆ The first value you supply with the -lun option must include the storage system name, volume, and LUN name. To delete multiple LUNs on the same volume, you can use relative pathnames for the -lun option after you supply the complete information in the first pathname. When SnapDrive for UNIX encounters a relative pathname, it looks for the LUN on the same volume as the first LUN. To delete additional LUNs that are not on the same volume, enter the full pathname to each LUN. Example 1: The following command line tells SnapDrive for UNIX to delete LUNs from storage systems. <pre># snapdrive storage delete -lun filer1:/vol/vol1/lun1 lun2</pre>

Deleting specific host entities

To delete a host entity such as a disk group and its associated LUNs, complete the following step.

Step	Action
1	Execute the <code>snapdrive storage delete</code> command. <pre>snapdrive storage delete {-dg -fs -hostvol } file_spec [file_spec ...] [{-dg -fs -hostvol } file_spec [file_spec ...]] [-full]</pre> Note_____ For details about using the options and arguments in these command line, see the “SnapDrive for UNIX keywords and arguments” on page 247. _____ Keep the following in mind as you use this command: ◆ This form of the command deletes all the host and storage system components for the host-side entities that you specify. The host-side entity can be a logical volume, file system, or disk group. ◆ If you include the <code>-full</code> option, SnapDrive for UNIX deletes host-side entities that contain other entities, such as a disk group containing one or more host volumes. Without the <code>-full</code> option, this command fails if any of the specified entities contain other entities. For example, the command fails if there are disk groups that contain host volumes or host volumes that contain file systems. Example: The following command lines tell SnapDrive for UNIX to delete the objects from specific host entities, such as disk groups, logical volumes and file systems. <pre># snapdrive storage delete -full -dg dg1 dg2 dg3 # snapdrive storage delete -lvol mydg/vol3 mydg/vol5 # snapdrive storage delete -fs /mnt/acct2</pre>

About this chapter This chapter provides information about the troubleshooting tools available. These tools are useful for gathering information as well as solving problems.

Topics in this chapter This chapter discusses the following topics:

- ◆ “[Data collection utility](#)” on page 184
- ◆ “[General troubleshooting tips for SnapDrive for UNIX](#)” on page 187
- ◆ “[Understanding error messages](#)” on page 190
- ◆ “[Common error conditions](#)” on page 192
- ◆ “[Standard exit status values](#)” on page 194

Data collection utility

About the data collection utility

SnapDrive for UNIX provides a data collection utility (snapdrive.dc) that collects diagnostic information about SnapDrive for UNIX and your system setup. It does this by running NetApp diagnostic utilities and copying SnapDrive for UNIX log files to a special directory. Then it creates a compressed file containing this information that you can send to technical support for analysis.

Note

This utility only gathers basic information about the system and the configuration of SnapDrive for UNIX. It does not copy the file containing login information for the storage systems. It also does not make any configuration changes.

Tasks performed by snapdrive.dc

The snapdrive.dc utility performs the following tasks:

- ◆ Runs the `<host>_info` utility, which comes with your Network Appliance host attach or support kit. For example, the Solaris kit includes the `solaris_info` utility, while the HP-UX kit includes the `hpux_info` utility. On AIX, it is the `aix_info` utility and on Linux, it is the `linux_info` utility. The `<host>_info` utility provides information about the host and saves the information to a compressed file.
- ◆ Creates a directory called `<directory>/ntap_snapdrive_<name>` (the directory pathname can vary depending your host; see the kit documentation for more information on this pathname). It places copies of the following files in it:
 - ❖ SnapDrive version, as indicated by running the `snapdrive version` command
 - ❖ The `snapdrive.conf` file
 - ❖ The audit log file
 - ❖ The trace log file
 - ❖ The recovery log file
 - ❖ The `snapdrive.conf` file
 - ❖ The file created by the `<host>_info` utility
- ◆ Creates a compressed file of the directory contents and displays a message suggesting you send this file to technical support.

Executing the data collection utility

To execute the data collection utility, complete the following steps.

Step	Action
1	Log in as root.
2	Change to the SnapDrive for UNIX diagnostic directory. The path is: <pre><install_directory>/diag</pre> <i>install_directory</i> is the SnapDrive for UNIX installation directory for your host operating system. This directory can vary depending on your host operating system. See the installation steps to determine where this directory is on your host.
3	At the command prompt, enter the following command: <pre>snapdrive.dc [-d <directory>] [-n <file_name>] [-f]</pre> -d <directory> specifies the location for the compressed file that this utility creates. The default location is /tmp/netapp. -n <file_name> specifies a string to be included in the name for the directory and compressed output file. If you supply a value for this argument, the snapdrive.dc utility creates a directory called ntap_snapdrive_<name> and a file name called ntap_snapdrive_<name>.tar.Z. The default path name is /tmp/netapp/ntap_snapdrive_info.tar.Z. -f forces SnapDrive for UNIX to overwrite the files if they currently exist. Without the -f option, the snapdrive.dc utility returns an error message if the output directory or the file name exists. It does not overwrite any directories or file names. Result: This utility runs the <host>_info script for your host operating system, as well as copying log files and the configuration file for SnapDrive for UNIX. It places this information in a directory, and then creates a compressed file containing the contents of the directory.
4	Send the <directory>/ntap_snapdrive_<name>.tar.Z file to technical support for analysis.

Examples of using snapdrive.dc

The following are examples of using the snapdrive.dc utility.

Example 1: This example runs the snapdrive.dc utility without supplying any command-line arguments:

```
# snapdrive.dc
SnapDrive configuration info and logs are in directory
/tmp/netapp/ntap_snapdrive_info.
Compressed file is /tmp/netapp/ntap_snapdrive_info.tar.Z.
Please send this file to technical support for analysis.
```

Example 2: This example uses the command-line options to specify a directory and a name for the resulting file:

```
# snapdrive.dc -d . -n mysystem
SnapDrive configuration info and logs are in directory
./ntap_snapdrive_mysystem.
Compressed file is ./ntap_snapdrive_mysystem.tar.Z.
Please send this file to technical support for analysis.
```

General troubleshooting tips for SnapDrive for UNIX

Attempts to create a snapshot fail

The following conditions can affect the `snapdrive snap create` command's ability to create a snapshot:

- ◆ The `snapdrive snap create` command must be able to perform the following tasks:
 - ❖ Query all disk groups for a list of LUNs
 - ❖ Query all LUNs via SCSI commands
- ◆ The LUNs in the disk group must be online and mapped to the disk group before you take the snapshot. If any of the LUNs are offline or unmapped, the create operation fails.

Make sure that all the LUNs are online and mapped to the host before you attempt to take a snapshot.
- ◆ The access permissions do not allow the host to create a snapshot of information on that storage system. For more information on access permissions, see [“Setting up access control”](#) on page 101.
- ◆ Make sure you are logged in as root. SnapDrive for UNIX requires this for its commands. In addition, because it works with LVM entries, you must be logged on to the host with sufficient permission for it to manipulate the LVM entries.

Attempts to restore a snapshot fail

The following can affect the `snapdrive snap restore` command's ability to restore a snapshot:

- ◆ The `snapdrive snap restore` command must be able to perform the following tasks:
 - ❖ Query all disk groups for a list of LUNs
 - ❖ Query all LUNs via SCSI commands

The host volumes and file systems do not need to be available and mounted for the `snapdrive snap restore` command to succeed.
- ◆ The access permissions do not allow the host to create and restore a snapshot of information on that storage system. For more information on access permissions, see [“Setting up access control”](#) on page 101.
- ◆ If you add or remove a LUN from a disk group after taking a snapshot and then try to restore it, the `snapdrive snap restore` operation fails. SnapDrive for UNIX cannot guarantee data consistency across a restore in this case.

Caution

You should always take another snapshot after you add a LUN to or remove a LUN from a disk group.

- ◆ If you add a host volume or file systems to a disk group after taking a snapshot and then try to restore the snapshot, you receive an error message. You must include the force option (`-f`) on the command line for the restore operation to succeed in this situation.

Caution

Host volumes and file systems that you add after taking a snapshot are no longer accessible when you restore the snapshot.

- ◆ If you rename a host volume or a file system or change the mount point after taking a snapshot and then try to restore it, the restore operation fails. You must include the force option (`-f`) on the command line for the restore operation to succeed in this situation.

Caution

Use the `-f` option with great care to make sure you don't accidentally overwrite something that you didn't intend to overwrite.

- ◆ If you rename a LUN after taking a snapshot and then try to restore that snapshot, the restore operation fails. After you rename a LUN, you should take a new snapshot.

WARNING

After you start a `snapdrive snap restore` operation, do not halt it. Halting this operation might leave the system in an inconsistent state. You then might need to perform a manual recovery.

Some commands delayed when storage cluster is in failover mode

SnapDrive for UNIX commands that create, connect, delete, or disconnect a LUN take several minutes to complete on Solaris host while one of the clustered storage systems is down.

Such commands are `storage create`, `storage connect`, `storage delete`, `storage disconnect`, `snap connect`, `snap disconnect` and, in some cases, `snap restore`.

This problem affects `snap restore` only when SnapDrive for UNIX needs to restore LUNs that were deleted or disconnected.

Data distribution on disks changes after a restore operation

Symptom: The data distribution on a disk changes after you restore a snapshot.

Explanation: The restore operation replaces the data on a disk based on the way it was when you took the snapshot. If you take a snapshot and then change the way data is distributed among disks in the disk group, the restored snapshot returns the data to the previous disk configuration. This means that changes to striping, mirroring, and so on can take place as a result of a restore operation.

What to do: If you change the data distribution on the disks in the disk group, take a new snapshot.

Understanding error messages

Error message locations

SnapDrive for UNIX provides information about error messages in the following places:

- ◆ The command line. It displays all messages to the standard error output of the SnapDrive for UNIX command.
- ◆ The system log. SnapDrive for UNIX logs all errors that have a severity level of Fatal and Admin error to the system log using the syslog(3) mechanism.

In addition, SnapDrive for UNIX records information about the commands and their execution in the log files, specifically:

- ◆ The audit log file. The audit log records the following information for each SnapDrive for UNIX command:
 - ❖ Who issued it.
 - ❖ When it was issued.
 - ❖ What its exit status was. This is very useful in determining what actually happened on a system.
- ◆ The trace log file. The trace log records more detailed information about any errors that occur. Technical support uses this log when diagnosing problems.

Error message format

SnapDrive for UNIX error messages conform to the following format:

`severity (<error type> <message-ID>): message text`

where:

severity - specifies the severity of the error. Values are Error, Warning and Information.

error type - specifies the type of error that SnapDrive encountered. Return values include:

Warning - Snapdrive for UNIX executed the command but issued a warning about conditions that might require your attention.

Command - Snapdrive for UNIX failed to execute the command due to an error in the command line. Check the command line format and variables to ensure they are correct.

Admin - Snapdrive for UNIX failed to execute the command due to incompatibilities in the system configuration. Contact your System Administrator to review your configuration parameters.

Fatal - Snapdrive for UNIX failed to execute the command due to an unexpected condition. Fatal errors are rare. If a fatal error occurs and you have problems resolving it, contact technical support for assistance in determining the steps you need to take to recover correctly and fix any error condition.

message-ID - a unique identifier used by technical support to locate the specific code that produced the error. If you need to call technical support, NetApp recommends that you record the message-ID that accompanied the error message.

message text - text that explains the error. This text might include information from another component to provide more detail about the error. For example, if a command-line argument such as a disk group is missing, the error message tells you what is missing. Or the ManageONTAP APIs that SnapDrive for UNIX uses to control the storage system might supply additional text to help explain the error. In this case, the text follows the basic SnapDrive for UNIX error message.

Sample error message

The following command level message indicates a problem on the command line. The message-ID is 0001-426.

```
Error (Command error 0001-426): File system mount point
/u/rsmith/.bashrc is not available for use. Please choose another
mount point
```

Common error conditions

Introduction

This section provides information about common error messages that you may encounter. This chapter does not include information for every component or error condition.

Operating system limits on open files

SnapDrive for UNIX checks for operating system limitations on the number of files opened by a process.

Note

The default limit for the number of file handles opened simultaneously by one process varies based on your operating system. Check your operating system documentation to determine the limit.

If the number of open LUNs for one operation exceeds the operating system limit on the number of file handles opened simultaneously by one process, SnapDrive exits with the following error message:

```
Error (Admin error 0001-001): Unable to open device <path-to-device>
```

Example: You see an error message similar to the following one if this limit is exceeded on a Solaris host:

```
Error (Admin error 0001-001): Unable to open device  
/dev/rdsk/c1t1d26s2
```

SnapDrive for UNIX reports an error if iSCSI is not running

Symptom: SnapDrive for UNIX running on a Linux host returns an error message stating that the device is not a NetApp LUN. This happens after one of the following situations occurs:

- ◆ An iSCSI stop was issued.
- ◆ A takeover has happened and the iSCSI service hasn't started yet.
- ◆ A giveback is happening and the iSCSI service hasn't started yet.

Explanation: For SnapDrive for UNIX to execute, the iSCSI service must be running. If the service is halted, SnapDrive for UNIX returns an error message. For example, if you executed a `snapdrive storage show` command when the iSCSI service is not running, you might see output similar to the following:

```
# snapdrive storage show -dg toaster1
```

```
dg: toaster1
hostvol: /dev/toaster1/lvol1      state: AVAIL
hostvol: /dev/taoster1/lvol2      state: AVAIL
fs: /dev/toaster1/lvol1  mount point: /mnt/uml
fs: /dev/toaster1/lvol2  mount point: NOT MOUNTED
```

device	filename	adapter	path	size	state	lun	path
/dev/sdb	-	P	34:device	/dev/sdb	is not	NetApp	LUN ()
/dev/sdc	-	P	34:device	/dev/sdc	is not	NetApp	LUN ()

```
Error (Command error 0001-185): storage show failed: no NetApp devices to show
```

What to do: Make sure the iSCSI service is running. If the filer is in the process of a takeover or giveback, wait until it starts the iSCSI service. If the filer is in normal running mode, restart the iSCSI service. Once the service is running, reissue the `snapdrive storage show` command. This time you should see output similar to the following:

```
# snapdrive storage show -dg toaster1
```

```
dg: toaster1
hostvol: /dev/toaster1/lvol1      state: AVAIL
hostvol: /dev/toaster1/lvol2      state: AVAIL
fs: /dev/toaster1/vol1  mount point: /mnt/uml
fs: /dev/toaster1/lvol2  mount point: NOT MOUNTED
```

device	filename	adapter	path	size	state	lun	path
/dev/sdb	-	P	2g	online	eccentric:/vol/vol1/lun1		
/dev/sdc	-	P	2g	online	eccentric:/vol/vol1/lun2		

```
#
```

Standard exit status values

Understanding exit status values of error messages

Each SnapDrive error message ID is linked to an exit status value. Exit status values contain the following information:

Exit status value - indicates the basic cause of the error condition.

Severity - indicates how serious this error is. The level of seriousness depends on the message, not the value. Possible values are:

Warning - Snapdrive for UNIX executed the command but issued a warning about conditions that might require your attention.

Command - Snapdrive for UNIX failed to execute the command due to an error in the command line. Check the command line format to ensure they are correct.

Admin - Snapdrive for UNIX failed to execute the command due to incompatibilities in the system configuration. Contact your System Administrator to review your configuration parameters.

Fatal - Snapdrive for UNIX failed to execute the command due to an unexpected condition. Fatal errors are rare. If a fatal error occurs and you have problems resolving it, contact technical support for assistance in determining the steps you need to take to recover correctly and fix any error condition.

Using exit status values

Exit status values are used in scripts to determine the success or failure of an SnapDrive command.

- ◆ A value of zero indicates that the command completed successfully.
- ◆ A non-zero value indicates that the command did not complete, and provides information about the cause and severity of the error condition.

Script example

The following script uses SnapDrive for UNIX exit status values:

```
#!/bin/sh
# This script demonstrates a SnapDrive
# script that uses exit codes.

RET=0;
#The above statement initializes RET and sets it to 0
```

```

snapdrive snap create -dg vg22 -snapname vg22_snap1;
# The above statement executes the snapdrive command

RET=$?;
#The above statement captures the return code.
#If the operation worked, print
#success message. If the operation failed, print
#failure message and exit.

if [ $RET -eq 0 ]; then
echo "snapshot created successfully"
else
echo "snapshot creation failed, snapdrive exit code was $RET"
exit 1
fi

exit 0;

```

If RET=0, the command executed successfully and the script outputs the following:

```

# ./tst_script
snap create: snapshot vg22_snap1 contains:
disk group vg22 containing host volumes
lvoll
snap create: created snapshot betty:/vol/vol2:vg22_snap1
snapshot created successfully

```

If RET= a non-zero value, the command did not execute successfully. The following example shows typical output.

```

# ./tst_script
Error (Command error 0001-155): Snapshot betty:/vol/vol2:vg22_snap1
already exists on betty:/vol/vol2.
Please use -f (force) flag to overwrite existing snapshot
snapshot creation failed, snapdrive exit code was 4

```

Exit status values

The following table contains information about exit status values. The exit status values are numbered sequentially. If SnapDrive for UNIX does not currently implement an error, that exit status value is not included in the table. As result, there may be some gaps in the numbers.

Exit value	Error name	Severity	Description
1	Not supported	Command error	A function was invoked that is not supported in this version of SnapDrive for UNIX.
2	No memory	Fatal	The system has run out of memory. SnapDrive for UNIX cannot proceed until you free enough memory for it to work. Check other applications running to verify that they are not consuming excessive memory.
3	Invalid command	Command error	You issued an invalid command; this is likely to be a syntax error in the text of the command you entered.
4	Already exists	Command error	You requested that something be created that already exists. Usually, this error refers to a snapshot name, which must not exist on the storage system volume where you are taking the snapshot.
5	Create thread failed	Admin error	SnapDrive for UNIX could not create a process thread. Check the other processes running on the system to make sure that enough thread resources are available.
6	Not found	Command error	You included a file, data group, host volume, file system, or other argument on the SnapDrive for UNIX command line that does not exist.
7	Not a mounted file system	Command error	The file system you want to access either is not a valid file system or is not mounted.
9	Volume manager error	Command error	An error was returned when accessing the volume manager. See the specific error message to get details of which error, and why.

Exit value	Error name	Severity	Description
10	Invalid name	Command error	You supplied a name on the command line that was not correctly formatted. For example, a storage system volume was not specified as “<filer>:/vol/<volname>”. This message also occurs when an invalid character is given in either a storage system or a volume manager-based name.
11	Device not found	Admin error	SnapDrive for UNIX cannot access a LUN in the disk group that you want to take a snapshot of. Check the status of all LUNs, both on the host and on the storage system. Also check that the storage system volume is online, and that the storage system is up and connected to the host.
12	Busy	Command error	The LUN device, file, directory, disk group, host volume, or other entity was busy. This is generally a nonfatal error that goes away when you retry the command. It sometimes indicates that a resource or process is hung, causing the object to be busy and unavailable for SnapDrive for UNIX to use. It could also indicate you are trying to take a snapshot during a period when there is too much heavy I/O traffic for the snapshot to succeed.
13	Unable to initialize	Fatal	SnapDrive for UNIX could not initialize third-party material that it needs. This can refer to file systems, volume managers, cluster software, multipathing software, and so on.

Exit value	Error name	Severity	Description
14	SnapDrive busy	Command error	Another user or process is performing an operation on the same hosts or storage systems at the same time you asked SnapDrive for UNIX to perform an operation. Retry your operation. Occasionally this message means that the other process is hung and you must kill it. Note _____ The <code>snap restore</code> operation can take a long time under some circumstances. Be sure that the process you think is hung is not just waiting for a <code>snap restore</code> operation to be completed. _____
15	Config file error	Fatal	The <code>snapdrive.conf</code> file has invalid, inadequate, or inconsistent entries. See the specific error message for details. You must correct this file before SnapDrive for UNIX can continue. See “Setting values in snapdrive.conf” on page 82 for information about modifying the <code>snapdrive.conf</code> file.
17	Bad permissions	Command error	You do not have permission to execute this command. You must be logged in as root to run SnapDrive for UNIX.
18	No filer	Admin error	SnapDrive for UNIX cannot contact the storage system needed for this command. Check the connectivity to the storage system indicated in the error message.
19	Bad filer login	Admin error	SnapDrive for UNIX cannot log in to the storage system using the login information you supplied.
20	Bad license	Admin error	A service SnapDrive for UNIX requires is not licensed to run on this storage system.

Exit value	Error name	Severity	Description
22	Cannot freeze fs	Admin error	A <code>snap create</code> operation failed because SnapDrive for UNIX could not freeze the file systems specified in order to make the snapshot. Confirm that the system I/O traffic is light enough to freeze the file system and then retry the command.
26	Resource busy on restore	Admin error	The <code>snap restore</code> operation failed because a resource is busy that must be disconnected or disabled before the operation can begin. Examples of potentially busy resources include a file system or logical volume. Either repeat the command and include the <code>-f</code> (force) option on the command line or terminate the processes using the resource. Note— On some systems, the <code>-f</code> option does not disconnect the resource.
27	Inconsistent snapshot	Admin error	The <code>snap restore</code> operation failed because you requested a restore from a snapshot with inconsistent images of the disk group. Inconsistent images can occur in the following cases: ◆ You did not take the snapshot using SnapDrive for UNIX. ◆ The <code>snap create</code> operation was interrupted before it set consistent bits and, thus, could not clean up (as in the case of a catastrophic system failure). ◆ Some type of data problem occurred with the snapshot after it was taken.
28	HBA failure	Admin error	SnapDrive for UNIX encountered an error while trying to retrieve information from the HBA.

Exit value	Error name	Severity	Description
29	Bad metadata	Admin error	SnapDrive for UNIX encountered an error in the snapshot metadata that it wrote when it created the snapshot.
30	No snapshot metadata	Admin error	SnapDrive for UNIX can't perform a restore operation because the metadata does not contain all requested disk groups.
31	Bad password file	Admin error	The password file has a bad entry. Use the <code>snapdrive config delete</code> command to delete the login entry for this storage system. Then reenter the login information using the <code>snapdrive config set user_name</code> command. For information on specifying user logins to storage systems, see “Specifying the current login information for storage systems” on page 107.
32	No password file	Admin error	Either there is no password file containing the storage system login information or the path to the password file is incorrect. Check the path given for the password file in the <code>snapdrive.conf</code> file and make sure that the password file has been initialized with logins for your filers. You can use the <code>snapdrive config show</code> command to display the configuration file. For information on updating this file, see “Setting values in snapdrive.conf” on page 82.
33	No password file entry	Admin error	The password file has no entry for this storage system. Run <code>snapdrive config set user_name <filer_name></code> for every storage system on which you need to run SnapDrive for UNIX. Then try this operation again.
34	Not a NetApp LUN	Admin error	A SnapDrive for UNIX command encountered a LUN that is not on a NetApp storage system.

Exit value	Error name	Severity	Description
35	User aborted	Admin error	The system displayed a prompt asking you to confirm an operation and you indicated that you did not want the operation performed.
36	I/O stream error	Admin error	The system input or system output routines returned an error that SnapDrive did not understand. Run snapdrive.dc and send that information to technical support so that they can help you determine which steps to perform to complete the recovery.
37	File system full	Admin error	An attempt to write a file failed because there was insufficient space on the file system. SnapDrive for UNIX can proceed when you free enough space on the appropriate file system.
38	File error	Admin error	An I/O error occurred when SnapDrive for UNIX was reading or writing a system configuration file or a temporary file.
39	Duplicate diskgroup	Command error	SnapDrive for UNIX got a duplicate minor node number when trying to activate a disk group.
40	File system thaw failed.	Admin error	A snap create command failed due to system activity on the file system. This usually occurs when the SnapDrive file system freeze, required for the snapshot, times out before the snapshot is complete.

Exit value	Error name	Severity	Description
41	System configuration query error	Admin error	SnapDrive for UNIX encountered a problem accessing the Object Data Manager (AIX systems only). This can be caused by: ◆ An incorrect ODMPATH environment variable. ◆ An ODM lock, open and held by a crashed or hung process ◆ Corrupt ODM files
43	Name already in use	Command error	SnapDrive for UNIX attempted to create a disk group, host volume, file system or LUN but the name was already in use. To correct, select a name that is not in use, and re-enter the SnapDrive command.
44	File system manager error	Fatal	SnapDrive for UNIX encountered an unexpected error from the file system when: ◆ attempting to create the file system ◆ making an entry in the file system mount table to automatically mount the file system at boot. The text of the error message displayed with this code describes the error that the file system encountered. Record the message, and send it to technical support so that they can help you determine which steps to perform to complete the recovery.
45	Mount point error	Admin error	The file system mount point appeared in the system mount table file. To correct, select a mount point that is not in use or listed in the mount table, and re-enter the SnapDrive for UNIX command.
46	LUN not found	Command error	A SnapDrive for UNIX command attempted to access a LUN that did not exist on the filer. To correct, check that the LUN exists and that the name of the LUN is entered correctly.

Exit value	Error name	Severity	Description
47	Initiator group not found	Admin error	A storage system initiator group could not be accessed as expected. As a result, SnapDrive for UNIX cannot complete the current operation. The specific error message describes the problem and the steps you need to perform to resolve it. Fix the problem and then repeat the command.
48	Object offline	Admin error	Snapdrive for UNIX attempted to access an object (such as a volume) but failed because the object was offline.
49	Conflicting entity	Command error	SnapDrive for UNIX attempted to create an igroup, but encountered an igroup of the same name.
50	Cleanup error	Fatal	Snapdrive for UNIX encountered an item that should be removed but is still there.
51	Disk group ID conflict	Command error	A SnapDrive for UNIX or storage connect command requested a disk group ID that conflicts with an existing disk group. This usually means that a <code>snap connect</code> originating host operation is being attempted on a system that does not support it. To fix this problem, attempt the operation from a different host.
52	LUN not mapped to any host	Admin error	A LUN is not mapped to any host. In other words, it does not belong to a storage system initiator group. To be accessible, the LUN must be mapped to the current host outside SnapDrive for UNIX.

Exit value	Error name	Severity	Description
53	LUN not mapped to local host	Admin error	A LUN is not mapped to the current host. In other words, it does not belong to a storage system initiator group that includes initiators from the current host. To be accessible, the LUN must be mapped to the current host outside SnapDrive for UNIX.
54	LUN is mapped using foreign igroup	Admin error	A LUN is mapped using a foreign storage system initiator group. In other words, it belongs to a storage system igroup containing only initiators not found on the local host. As a result, SnapDrive for UNIX cannot delete the LUN. To use SnapDrive for UNIX to delete a LUN, the LUN must belong only to local igroups; that is, igroups containing only initiators found on the local host.
55	LUN is mapped using mixed igroup	Admin error	A LUN is mapped using a mixed storage system initiator group. In other words, it belongs to a storage system igroup containing both initiators found on the local host and initiators not found there. As a result, SnapDrive for UNIX cannot disconnect the LUN. To use SnapDrive for UNIX to disconnect a LUN, the LUN must belong only to local igroups or foreign igroups; not mixed igroups. (Local igroups contain only initiators found on the local host; foreign igroups contain initiators not found on the local host.)

Exit value	Error name	Severity	Description
56	Snapshot restore failed	Admin error	Snapdrive for UNIX attempted a <code>snap restore</code> operation, but it failed without restoring any LUNs in the snapshot. The specific error message describes the problem and the steps you need to perform to resolve it. Fix the problem and then repeat the command.
58	Host reboot needed	Admin error	The host operating system requires a reboot in order to update internal data. SnapDrive for UNIX has prepared the host for this update, but cannot complete the current operation. Reboot the host and then re-enter the SnapDrive for UNIX command line that caused this message to appear. After the reboot, the operation will be able to complete.

Exit value	Error name	Severity	Description
59	Host, LUN preparation needed	Admin error	The host operating system requires an update to internal data in order to complete the current operation. This update is required to allow a new LUN to be created. SnapDrive for UNIX cannot perform the update, because automatic host preparation for provisioning has been disabled because the <code>snapdrive.conf</code> variable <code>enable-implicit-host-preparation</code> is set to off. (For more information, see “Setting configuration information” on page 62.) With automatic host preparation disabled, you should use either the <code>snapdrive config prepare luns</code> command to prepare the host to provision LUNs or perform the preparation steps manually. To avoid this error message, set the <code>enable-implicit-host-preparation</code> value to On in the <code>snapdrive.conf</code> file. Note _____ Currently, only Linux and Solaris platforms require host preparation. _____
126	Unknown error	Admin error	An unknown error occurred that might be serious. Run the <code>snapdrive.dc</code> utility and send its results to technical support for analysis.
127	Internal error	Fatal	A SnapDrive for UNIX internal error occurred. Run <code>snapdrive.dc</code> and send its results to technical support for analysis.

About this chapter

This chapter provides details about creating and using SnapDrive for UNIX snapshots.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Collecting information needed by SnapDrive for UNIX Commands”](#) on page 208
- ◆ [“Checklist for using the snapshot create command”](#) on page 209
- ◆ [“Checklist for using the snapshot rename command”](#) on page 211
- ◆ [“Checklist for using the snapshot show command”](#) on page 213
- ◆ [“Checklist for using the snapshot restore command”](#) on page 215
- ◆ [“Checklist for using the snapshot delete command”](#) on page 217
- ◆ [“Checklist for using the snapshot connect command”](#) on page 219
- ◆ [“Checklist for using the snapshot disconnect command”](#) on page 223
- ◆ [“Checklist for using the storage create command”](#) on page 225
- ◆ [“Checklist for using the storage resize command”](#) on page 229
- ◆ [“Checklist for using the storage show command”](#) on page 231
- ◆ [“Checklist for using the storage connect command”](#) on page 234
- ◆ [“Checklist for using the storage disconnect command”](#) on page 235
- ◆ [“Checklist for using the host disconnect command”](#) on page 238
- ◆ [“Checklist for using the storage delete command”](#) on page 240

Collecting information needed by SnapDrive for UNIX Commands

Collecting information needed by commands

This chapter provides checklists you can use as you execute Snapdrive for UNIX. For each command, it supplies the following:

- ◆ Recommended usage formats
- ◆ Information on the keywords, options, and arguments available with the commands and the values you should supply
- ◆ Examples of the commands

These checklists provide a quick overview of using these commands. For more information on the commands, see the chapter “[Creating and Using Snapshots](#)” on page 113 and the chapter “[Provisioning and Managing Storage](#)” on page 149.

General notes about the commands

The following are general notes about the commands:

- ◆ The `-dg` and `-vg` options in the command lines are synonyms that reflect the fact that some operating systems refer to disk groups and others refer to volume groups. This guide uses `-dg` to refer to both disk groups and volume groups
- ◆ The `-lvol` and `-hostvol` options in the command lines are synonyms that reflect the fact that some operating systems refer to logical volumes and others refer to host volumes. This guide uses `-hostvol` to refer to both logical volumes and host volumes.
- ◆ NetApp strongly recommends that you use the default `igroup` and not specify an `igroup` explicitly by including the `-igroup` option on your command line. For information on how SnapDrive for UNIX handles `igroup` names, see “[Command-line keywords](#)” on page 250 in the “[Command Reference](#)” chapter. If you need to specify an `igroup`, see the SnapDrive for UNIX man page for information on doing that.

Checklist for using the snapshot create command

Checklist for creating snapshots

This checklist is a planning aid for creating snapshots of one or more storage system volumes.

Usage format for snapdrive snap create:

```
snapdrive snap create
  {-dg | -hostvol | -fs} file_spec [file_spec ...]
  [{-dg | -hostvol | -fs} file_spec [file_spec ...] ...]
  -snapname snap_name
  [-force [-noprompt]]
```

The following is a summary of the information you need to supply when you use the snapdrive snap create command.

Requirement	Argument	Value you want to supply
Determine which type of host entity that you want to capture in the snapshot. Supply that entity’s name with the appropriate argument. This is the value for the <i>file_spec</i> argument. This argument translates into a set of disk groups on the storage system. SnapDrive for UNIX freezes any file systems in host volumes in those disk groups and takes a snapshot of all storage system volumes containing LUNs in those disk groups.		
◆ Disk group (-dg <i>file_spec</i>) or volume group (-vg <i>file_spec</i>)	<i>name of the disk or volume group</i>	
◆ File system (-fs <i>file_spec</i>)	<i>filesystem_name</i>	
◆ Host volume (-hostvol <i>file_spec</i>) or logical volume (-lvol <i>file_spec</i>)	<i>name of the host or logical volume</i> Note _____ You must supply both the requested volume and the disk group containing it; for example, -hostvol dg3/acct_volume.	
Specify the name for the snapshot. This can be either the long version of the name that includes the storage system and volume with the snapshot name or the short version that is just the snapshot name.		

Requirement	Argument	Value you want to supply
◆ Snapshot name (-snapname <i>snap_name</i>)	<i>snapshot_name</i>	
(Optional) Decide if you want to overwrite an existing snapshot. Without this option, this operation halts if you supply the name of an existing snapshot. When you supply this option and specify the name of an existing snapshot, the command prompts you to confirm that you want to overwrite the snapshot. To prevent SnapDrive for UNIX from displaying the prompt, include the -noprompt option also. (You must always include the -force option if you want to use the -noprompt option.)		
◆ -force	~	
◆ -noprompt	~	

Examples of snap create command lines

The following are some sample command lines that use the `snapdrive snap create` command.

```
# snapdrive snap create -dg -dg_name1 dg_name2 -snapname snap_name
# snapdrive snap create -fs /fs_mount_dir -lvol dg1/logical_volname
  -snapname snap1
```

Checklist for using the snapshot rename command

Checklist for renaming snapshots

This checklist is a planning aid for renaming snapshots.

Usage format for snapdrive snap rename:

```
snapdrive snap rename [-snapname] old_long_snap_name new_snap_name
[-force [-noprompt]]
```

Note _____
An error occurs if you try to rename a snapshot to a different storage system volume.

An error occurs if the new name for the snapshot already exists. You can use the `-force` option to force SnapDrive for UNIX to change the name without generating an error.

The following is a summary of the information you need to supply when you use the `snapdrive snap rename` command.

Requirement	Argument	Value you want to supply
Enter the <code>-snapname</code> option followed by the name of the existing snapshot and then the name to which you want to change it. You must use the long form of the snapshot name for the <i>old_long_snap_name</i> , where you enter the storage system name, volume, and snapshot name. You can use the short form of the name for the <i>new_snap_name</i> because it must exist on the same storage system volume as the original snapshot.		
◆ <code>-snapname</code>		
◆ Current name of the snapshot; use the long form of this name	<i>old_long_snap_name</i>	
◆ New name of the snapshot	<i>new_snap_name</i>	
(Optional) Decide if you want to overwrite an existing snapshot. Without this option, this operation halts if you supply the name of an existing snapshot. When you supply this option and specify the name of an existing snapshot, it prompts you to confirm that you want to overwrite the snapshot. To prevent SnapDrive for UNIX from displaying the prompt, include the <code>-noprompt</code> option also. (You must always include the <code>-force</code> option if you want to use the <code>-noprompt</code> option.)		
◆ <code>-force</code>	~	

Requirement	Argument	Value you want to supply
◆ -noprompt	~	

Examples of snap rename command lines

The following are some sample command lines that use the `snapdrive snap rename` command.

```
# snapdrive snap rename -force -snapname filer1:/vol/vol1:oldsnap
newsnapshot
```

```
# snapdrive snap rename -snapname filer1:/vol/vol1:FridaySnap
filer1:/vol/vol1:Snap040130
```

Checklist for using the snapshot show command

Checklist for displaying snapshots

This checklist is a planning aid for displaying information about snapshots.

Note

You can use either `snapdrive snap show` or `snapdrive snap list` in these command lines. These commands are synonyms. The usage formats use `snapdrive snap show` to make this guide easier to read.

Usage formats for `snapdrive snap show`:

Display information about snapshots residing on a storage system:

```
snapdrive snap show -filer filename [filename...]  
[-verbose]
```

Display information about snapshots of a storage system volume:

```
snapdrive snap show -filervol filervol [filervol...]  
[-verbose]
```

Display information about snapshots of a disk or volume group, file system, or host or logical volume:

```
snapdrive snap show {-dg | -fs | -hostvol } file_spec  
[file_spec ...]  
[-verbose]
```

Display information about a snapshot:

```
snapdrive snap show [-snapname] long_snap_name [snap_name ...]  
[-verbose]
```

Note

The `show` and `list` forms of this command are synonymous.

For SnapDrive for UNIX 2.0, you must use the long form of the snapshot name when you display information about snapshots.

The following is a summary of the information you need to supply when you use the `snapdrive show | list` command.

Requirement	Argument	Value you want to supply
Based on the command you enter, you can display information about any of the following: ◆ Storage systems ◆ Storage system volumes ◆ Disk or volume groups) ◆ File systems ◆ Host or logical volumes ◆ Snapshots The value you enter for the <code>file_spec</code> argument must identify the host entity about which you want to display information. The command assumes the entities are on the current host.		
◆ Storage system (<code>-filer</code>)	<i>filename</i>	
◆ A volume on the storage system (<code>-filervol</code>)	<i>filervol</i>	
◆ Disk group (<code>-dg file_spec</code>) or volume group (<code>-vg file_spec</code>)	<i>name of the disk or volume group</i>	
◆ File system (<code>-fs file_spec</code>)	<i>filesystem_name</i>	
◆ Host volume (<code>-hostvol file_spec</code>) or logical volume (<code>-lvol file_spec</code>)	<i>name of the host or logical volume</i>	
If you want to display information about a snapshot, specify the name for the snapshot. For the first snapshot, <i>long_snap_name</i>, enter the long version of the name, which includes the storage system name, volume, and snapshot name. You can use the short version of the snapshot name if it is on the same storage system and volume.		
◆ Snapshot name (<code>-snapname long_snap_name</code>)	<i>long_snap_name</i>	
◆ Additional snapshot names	<i>snap_name</i> (long or short version)	
To display additional information, include the <code>-verbose</code> option.		
<code>-verbose</code>	~	

Checklist for using the snapshot restore command

Checklist for restoring snapshots

This checklist is a planning aid for restoring snapshots.

Usage format for snapdrive snap restore:

```
snapdrive snap restore
{-dg | -hostvol | -fs} file_spec [file_spec ...]
[{-dg | -hostvol | -fs} file_spec [file_spec ...] ...]
-snapshotname long_snap_name
[-force [-noprompt]]
```

The following is a summary of the information you need to supply when you use the snapdrive snap restore command.

Requirement	Argument	Value you want to supply
Decide the type of host entity that you wish to restore and enter that entity’s name with the appropriate argument. SnapDrive for UNIX restores the entire disk group containing the entity, even if the entity is a host volume or file system. This is the value for the <i>file_spec</i> argument. This argument translates into a set of disk groups on the storage system. SnapDrive for UNIX freezes any file systems hosted on volumes in those disk groups and takes a snapshot of all storage system volumes containing LUNs in those disk groups.		
◆ Disk group (-dg <i>file_spec</i>) or volume group (-vg <i>file_spec</i>)	<i>name of the disk or volume group</i>	
◆ File system (-fs <i>file_spec</i>)	<i>filesystem_name</i>	
◆ Host volume (-hostvol <i>file_spec</i>) or logical volume (-lvol <i>file_spec</i>)	<i>name of the host or logical volume</i> Note _____ You must supply both the requested volume and the disk group containing it; for example, -hostvol dg3/acct_volume. _____	

Requirement	Argument	Value you want to supply
Specify the name for the snapshot. If any of the <i>file_spec</i> arguments you supply on the command line currently exist on the local host, you can use a short form of the snapshot name. If none of the <i>file_spec</i> arguments exist on the host, you must use a long form of the snapshot name where you enter the storage system name, volume, and snapshot name. The following is an example of a long snapshot name: <pre>big_filer:/vol/account_vol:snap_20031115</pre> The short form of the same snapshot name would omit the storage system and storage system volume name, so it would appear as: <pre>snap_20031115</pre>		
◆ Snapshot name (-snapname)	<i>snapshot_name</i>	
(Optional) Decide if you want to overwrite an existing snapshot. Without this option, this operation halts if you supply the name of an existing snapshot. When you supply this option and specify the name of an existing snapshot, it prompts you to confirm that you want to overwrite the snapshot. To prevent SnapDrive for UNIX from displaying the prompt, include the <code>-noprompt</code> option also. (You must always include the <code>-force</code> option if you want to use the <code>-noprompt</code> option.)		
◆ -force	~	
◆ -noprompt	~	

Examples of snap restore command lines

The following are some sample command lines that use the `snapdrive snap restore` command.

```
# snapdrive snap restore -vg volgrp1 volgrp2 -snapname
filer:/vol/volX:snap_vg1_vg2

# snapdrive snap restore -hostvol dg1/hostvol1 -fs /fs_mnt_dir
-snapname filer:/vol/volX:snap_name
```

Checklist for using the snapshot delete command

Checklist for deleting snapshots

This checklist is a planning aid for deleting snapshots.

Usage format for snapdrive snap delete:

```
snapdrive snap delete
[-snapname} long_snap_name [snap_name...]
[-verbose]
```

Note

If the snapshot you specify is in use, this operation fails. SnapDrive for UNIX only reports that this operation completed successfully if all the snapshots are removed.

The following is a summary of the information you need to supply when you use the `snapdrive snap delete` command.

Requirement	Argument	Value you want to supply
Specify the name for the snapshot. Use the long form of the snapshot name where you enter the storage system name, volume, and snapshot name. The following is an example of a long snapshot name: big_filer:/vol/account_vol:snap_20031115 If you want to specify additional snapshots, you can use the short form of the name if they are on the same storage system and volume as the first snapshot. Otherwise, use the long form of the name again.		
◆ Snapshot name (-snapname)	long_snapshot_name	
◆ Additional snapshots	snapshot_name (either long or short form)	
To display a list of the snapshots being deleted, include the -verbose option. This option fills in the missing storage system and volume information in cases where you used the short form of the snapshot name.		
◆ -verbose	~	

Example of snap delete command lines

The following command line uses the `snapdrive snap delete` command.

```
# snapdrive snap delete -v filer1:/vol/vol0/snap1 snap2 snap3
snapdrive: deleting
filer1:/vol/vol0/snap1
filer1:/vol/vol0/snap2
filer1:/vol/vol0/snap3
```

Checklist for using the snapshot connect command

Checklist for connecting snapshots

This checklist is a planning aid for connecting snapshots. It summarizes the information you need to provide.

Usage format for snapdrive snap connect:

```
snapdrive snap connect fspec_set [fspec_set...]
    -snapname long_snap_name
    [-igroup ig_name [ig_name ...]]
    [-autoexpand] [-autorename]
```

where *fspec_set* has the format:

```
{-dg | -fs | -hostvol} src_file_spec    [dest_file_spec]
    [{-destdg | -destvg} dgname]
    [{-destlv | -desthv} lvname]
```

Note

This command must always start with the name of a disk group or a host logical volume manager entity (for example, -dg, -hostvol, or -fs).

Any destination names you supply must not currently be in use, or the command fails.

If you included the -autoexpand option on the command line and a name it needs in order to connect the destination copy of an LVM entity is in use, the snap connect command fails unless you also included the -autorename option on the command line.

If you do not include -autoexpand and you do not specify all the host volumes in all the disk groups referred to on the command line, either by specifying the host volume itself or the file system), the command fails.

The following is a summary of the information you need to supply when you use the snapdrive snap connect command.

Requirement	Argument	Value you want to supply
Decide the type of host entity that you want to use to attach the snapshot and supply that entity's name with the appropriate argument. This is the value for the <i>src_fspect</i> argument. It must be the name of the LVM entity within the snapshot that you want to connect. SnapDrive for UNIX uses the name to locate the required disk group within the snapshot.		
◆ Disk group (-dg <i>file_spec</i>) or volume group (-vg <i>file_spec</i>)	<i>name of the disk or volume group</i>	
◆ File system (-fs <i>file_spec</i>)	<i>filesystem_name</i>	
◆ Host volume (-hostvol <i>file_spec</i>) or logical volume (-lvol <i>file_spec</i>)	<i>name of the host or logical volume</i>	
(Optional) Supply a name by which the target entity will be accessible after the disk group is connected. SnapDrive for UNIX uses this name to connect the destination LVM entity. This is the <i>dest_file_spec</i> argument. If you omit this name, the <code>snap connect</code> command uses the value you supplied for <i>src_fspect</i>.		
◆ Name of target entity	<i>dest_file_spec</i>	
(Optional) Specify the names for the destination disk groups or host volumes. If you included this information as part of the <i>dest_fspect/src_fspect</i> pair, you do not need to enter it here. If you do not specify the name needed to connect an entity in the destination disk group, the <code>snap connect</code> command takes the name from the source disk group. If it cannot use that name, the operation fails, unless you included -autorename on the command line.		
◆ Destination disk group (-destdg) or destination volume group (-destvg)	<i>dgname</i>	
◆ Destination logical volume (-destlv) or destination host volume (-desthv)	<i>lvname</i>	

Requirement	Argument	Value you want to supply
Specify the name for the snapshot. Use the long form of the name where you enter the storage system name, volume, and snapshot name.		
◆ Snapshot name (-snapname)	<i>long_snap_name</i>	
(Optional: NetApp recommends that you use the default igroup for your host instead of supplying an igroup name.)		
◆ Igroup name (-igroup)	<i>ig_name</i>	
To shorten the amount of information you must supply when connecting to a volume group, include the -autoexpand option on the command line. This option lets you name only a subset of the logical volumes or file systems in the volume group. It then expands the connection to the rest of the logical volumes or file systems in the disk group. This way you do not need to specify each logical volume or file system. SnapDrive for UNIX uses this information to generate the name of the destination entity. This option applies to each disk group specified on the command line and all host LVM entities within the group. NetApp recommends that, if you include this option, you should also include the -autorename option. If the -autoexpand option needs to connect the destination copy of an LVM entity, but the name is already in use, the snap connect command fails unless the -autorename option is on the command line.		
-autoexpand	~	
The -autorename option works with the -autoexpand option. This option enables the snap connect command to rename a destination LVM entity if the default name set by the -autoexpand option is already in use. This option applies to all host LVM entities specified on the command line. If you include the -autorename option on the command line, it implies the -autoexpand option, even if you don't include that option.		
-autorename	~	

Examples of snap connect command lines

The following are some sample command lines that use the `snapdrive snap connect` command.

```
# snapdrive snap connect -hostvol srcdg/myvol destdg/myvol
# snapdrive snap connect -vg myvg -snapname snap1
```

```
# snapdrive snap connect -dg src_dg1 dest_dg1  
-snapname filer1:/vol/vol1:mynsnapshot
```

Checklist for using the snapshot disconnect command

Checklist for disconnecting snapshots

This checklist is a planning aid for disconnecting snapshots. It summarizes the information you need to provide.

Usage format for snapdrive snap disconnect:

```
snapdrive snap disconnect
  {-dg | -fs | -hostvol} file_spec [file_spec ...]
  [{-dg | -fs | -hostvol } file_spec [file_spec ...]]
  [-full]
```

Note

This command must always start with the name of a disk group or a host logical volume manager entity (for example, -dg, -hostvol, or -fs).

An error occurs if the host entity is using LUNs that are not part of the snapshot.

An error also occurs if you specify a subset of the host volumes and/or file systems contained in each target disk group.

The following is a summary of the information you need to supply when you use the snapdrive snap disconnect command.

Requirement	Argument	Value you want to supply
Specify the type of host entity that you want to use to disconnect the snapshot and supply that entity's name with the appropriate argument. This is the value for the <i>file_spec</i> argument.		
◆ Disk group (-dg <i>file_spec</i>) or volume group (-vg <i>file_spec</i>)	<i>name of the disk or volume group</i>	
◆ File system (-fs <i>file_spec</i>)	<i>filesystem_name</i>	
◆ Host volume (-hostvol <i>file_spec</i>) or logical volume (-lvol <i>file_spec</i>)	<i>name of the host or logical volume</i>	

Requirement	Argument	Value you want to supply
Include the <code>-full</code> option on the command line if you want SnapDrive for UNIX to disconnect the objects from the snapshot even if a host-side entity on the command line contains other entities (such as a disk group that contains one or more host volumes). If you do not include this option, you must specify only empty host-side entities.		
<code>-full</code>	<code>~</code>	

Examples of snap disconnect command lines

The following are some sample command lines using the `snapdrive snap disconnect` command.

```
# snapdrive snap disconnect -hostvol dg5/myvolume
```

```
# snapdrive snap disconnect -vg vg1 vg10
```

```
# snapdrive snap disconnect -quiet -dg dg1 dg2 dg3
```

Checklist for using the storage create command

Checklist for creating storage

This checklist is a planning aid for creating storage. It summarizes the information you need to provide.

Usage format for snapdrive storage create:

Creating LUNS as storage:

```
snapdrive storage create -lun long_lun_name [lun_name...]
-lunsize size
[-dg dg_name]
[-igroup ig_name [ig_name ...]]
```

Creating host or logical volumes or disk or volume groups as storage and having SnapDrive for UNIX automatically set up the underlying LUNs:

```
snapdrive storage create -hostvol file_spec
[-dg dg_name]
-dgsize size
-filervol long_filer_path
[-igroup ig_name [ig_name...]]
```

Creating file systems as storage and having SnapDrive for UNIX automatically set up the underlying LUNs:

```
snapdrive storage create -fs file_spec -dgsize size
[-hostvol file_spec]
[-dg dg_name]
[-fstype type]
[-fsopts options]
-filervol long_filer_path
[-igroup ig_name [ig_name...]]
```

Creating LVM entities and naming the LUNs that will be under them:

```
snapdrive storage create -lun long_lun_name [lun_name...]
-lunsize size
-hostvol file_spec
[-dg dg_name]
[-igroup ig_name [ig_name...]]
```

```
snapdrive storage create -lun long_lun_name [lun_name...]
-lunsize size
-fs file_spec
```

```
[-hostvol file_spec]
[-dg dg_name]
[-fstype type]
[-fsopts options]
[-igroup ig_name [ig_name...]]
```

Note

It is not possible to construct a disk group (or, therefore, a logical volume or file system) using LUNs from more than one filer volume. If you supply multiple volumes in the list of LUNs given to `-lun`, do not include the `-dg` option.

The following is a summary of the information you need to supply when you use the `snapdrive storage create` command.

Requirement	Argument	Value you want to supply
Decide the type of storage you want to provision. Based on the command you enter, you can create any of the following: ◆ LUNs If you create one or more LUNs, the first argument must use the long form of the LUN name, which specifies the storage system name, the volume name, and the name of the LUN within the volume. To specify additional LUNs, you can use the LUN name alone if the new LUN is on the same storage system and volume as the previous LUN. Otherwise, you can specify a new storage system name and volume name (or just a volume name) to replace the previous values. ◆ Disk or volume groups ◆ File systems ◆ Host or logical volumes When you specify a disk or volume group, file system, or host or logical volume, SnapDrive for UNIX performs all the actions necessary to create the entity you specify. You can either explicitly specify the LUNs it contains or just supply the storage system and volume information and let SnapDrive for UNIX automatically create the LUNs.		
◆ A LUN (<code>-lun</code>)	<i>long_lun_name</i>	
◆ Additional LUNs	<i>lun_name</i> (long or short form)	
◆ Disk group (<code>-dg dgname</code>) or volume group (<code>-vg vgname</code>)	<i>disk or volume group name</i>	

Requirement	Argument	Value you want to supply
◆ Host volume (-hostvol <i>file_spec</i>) or logical volume (-lvol <i>file_spec</i>)	<i>Host or logical volume name</i>	
◆ File system (-fs <i>file_spec</i>)	<i>filesystem_name</i>	
Specify the size in bytes of each LUN being created.		
◆ Lun size (-lunsize)	<i>size</i>	
Specify the size in bytes of the volume group or disk group being created.		
◆ Disk group size (-dgsiz) or volume group size (-vgsiz)	<i>size</i>	
Specify the storage system and its volume where you want SnapDrive for UNIX to create the LUNs automatically. Use the -filervol option when you use the command line format that has SnapDrive for UNIX automatically create the LUNs.		
◆ Path to storage system volume (-filervol)	<i>long_filer_path</i>	
If you are creating a file system, supply the string representing the file system type. The following are the values SnapDrive for UNIX accepts for each host platform: ◆ Solaris: VxFS ◆ HP-UX: VxFS ◆ AIX: either JFS or JFS2 ◆ Linux: ext3 Note_____ By default, SnapDrive for UNIX supplies this value if there is only one file system type for your host platform. In that case, you do not need to enter it. _____		
◆ File system type (-fstype)	<i>type</i>	
If you are creating a file system, you can use -fsopts to specify options you want to pass to the host command used to create the file systems. For example, you might supply options that the <code>mkfs</code> command would use. The value you supply usually needs to be specified as a quoted string and must contain exactly the text to be passed to the command.		

Requirement	Argument	Value you want to supply
◆ Options for creating the host file system (-fsopts <i>options</i>)	<i>option name and value</i>	
(Optional: NetApp recommends that you use the default igroup for your host instead of supplying an igroup name.)		
◆ Igroup name (-igroup)	<i>ig_name</i>	

Examples of storage create command lines

The following are some sample command lines that use the snapdrive storage create command.

```
# snapdrive storage create -lun kipper:/vol/vol1/lunA lunB lunC
-dg data_dg -lunsize 10g
```

```
# snapdrive storage create -lvol dg5/myvol -filervol
filer2:/vol/acct -dgsiz 500m
```

```
# snapdrive storage create -fs /mnt/acctfs -filervol
filer2:/vol/acct -dgsiz 100m
```

```
# snapdrive storage create -fs /mnt/acctfs -fstype vxfs
-filervol acctfiler:/vol/acct -dgsiz 1g
```

Checklist for using the storage resize command

Checklist for resizing storage

This checklist is a planning aid for resizing storage.

Usage format for snapdrive storage resize:

```
snapdrive storage resize
  -dg file_spec [file_spec ...]
  {-growby | -growto} size
  [-addlun [-igroup ig_name [ig_name ...]]]
```

Note

You cannot use the `snapdrive storage resize` command to reduce the size of an entity. You can only increase the size.

The `snapdrive storage resize` command is not supported directly on logical volumes or file systems.

The following is a summary of the information you need to supply when you use the `snapdrive storage resize` command.

Requirement	Argument	Value you want to supply
Decide whether you want to increase the size of a disk or volume group and enter that entity’s name with the appropriate argument.		
◆ Disk group (-dg file_spec) or volume group (-vg file_spec)	name of the disk or volume group	
Decide how you want to increase the storage size. Keep the following in mind as you use this command:		
◆ Use the -growby option to increase the entity size by the bytes specified in the size argument.		
◆ Use the -growto option to increase the entity size so that the new total size is the number of bytes specified in the size argument.		
◆ Use the -addlun option to increase the entity size by adding a new, internally-generated LUN to the underlying disk group. If you do not use this argument, SnapDrive for UNIX increases the size of the last LUN in the disk group to meet the byte size specified in either the -growby option or the -growto option.		
◆ Specify the number of bytes by which you want to increase the storage (-growby size)	number_of_bytes	

Requirement	Argument	Value you want to supply
◆ Specify the size in bytes that you want the storage to reach (-growto <i>size</i>)	<i>number_of_bytes</i>	
◆ Tell SnapDrive for UNIX to increase the size by adding a new LUN to the disk group (-addlun)	~	
(Optional: NetApp recommends that you use the default igroup for your host instead of supplying an igroup name.)		
◆ Igroup name (-igroup)	<i>ig_name</i>	

Examples of storage resize command lines

The following are some sample command lines that use the `snapdrive storage resize` command.

```
# snapdrive storage resize -dg my_dg -addlun -growby 155m
# snapdrive storage resize -dg my_dg -growto 155m
```

Checklist for using the storage show command

Checklist for displaying storage information

This checklist is a planning aid for displaying information about storage.

Usage formats for snapdrive storage show:

Display information about LUNs on a storage system:

```
snapdrive storage { show | list } -filer filer_name  
[filer_name ...]  
[-verbose]  
[-quiet]
```

Display information about LUNs on a storage system volume:

```
snapdrive storage { show | list } -filervol long_filer_path  
[filer_path...]  
[-verbose]  
[-quiet]
```

Display information about a LUN:

```
snapdrive storage { show | list } -lun long_lun_name  
[lun_name ...]  
[-verbose]  
[-quiet]
```

Display information about a disk or volume group, file system, or host or logical volume:

```
snapdrive storage { show | list } {-dg | -fs | -hostvol } file_spec  
[file_spec ...]  
[{-dg | -fs | -hostvol } file_spec [file_spec ...]]  
[-verbose]  
[-quiet]
```

Note

The `show` and `list` forms of this command are synonymous.

For SnapDrive for UNIX 2.0, you must use the long form of the snapshot name when you display information about snapshots.

The following is a summary of the information you need to supply when you use the `snapdrive show` command.

Examples of

Requirement	Argument	Value you want to supply
Based on the command you enter, you can display information about any of the following: ◆ LUNs on a storage system accessible to the current host ◆ LUNs on storage system volumes accessible to the current host ◆ Specific LUNs ◆ Disk or volume groups ◆ File systems ◆ Host or logical volumes The value you enter for the <i>file_spec</i> argument must identify the host entity about which you want to display information. The command assumes the entities are on the current host. When you display information about a LUN on a storage system volume, you must supply the path to the LUN, including the name of the storage system and the volume, as the value for <i>long_filer_path</i> and <i>long_LUN_name</i>.		
◆ Storage system (-filer)	<i>filer_name</i>	
◆ LUNs on a storage system volume on the storage system (-filervol)	<i>long_filer_path</i>	
◆ Specific LUNs (-lun)	<i>long_LUN_name</i>	
◆ Disk group (-dg <i>file_spec</i>) or volume group (-vg <i>file_spec</i>)	<i>name of the disk or volume group</i>	
◆ File system (-fs <i>file_spec</i>)	<i>filesystem_name</i>	
◆ Host volume (-hostvol <i>file_spec</i>) or logical volume (-lvol <i>file_spec</i>)	<i>name of the host or logical volume</i>	
To display additional information, include the -verbose option.		
-verbose	~	
To prevent all output, both diagnostic and normal output, include the -quiet option. This option overrides the -verbose, which can be useful for existence checking within a script.		

Requirement	Argument	Value you want to supply
-quiet	~	

**storage show
command lines**

The following are sample snapdrive storage show | list command lines.

```
# snapdrive storage show -verbose -filer BigFiler3 BigFiler4

# snapdrive storage show -dg dg1 dg2

# snapdrive storage show -quiet -dg QA_dg_test

# snapdrive storage show -verbose -filervol BigFiler3:/vol/vol1

# snapdrive storage show -lun BigFiler3:/vol/vol1:lunA lunB lunC

# snapdrive storage show -filer BigFiler3 BigFiler4

# snapdrive storage show -fs /mnt/fs21 /mnt/fs41

# snapdrive storage show -hostvol dg2/vol1 dg4/myvol3

# snapdrive storage show -dg dg2 dg4

# snapdrive storage show -dg dg2 -fs /mnt/fs42

# snapdrive storage show -hostvol dg2/vol3 -fs /mnt/fs42
```

Checklist for using the storage connect command

Checklist for connecting storage

This checklist is a planning aid for connecting storage. It summarizes the information you need to provide.

Usage format for snapdrive storage connect:

```
snapdrive storage connect -lun long_lun_name [lun_name...]  
[-igroup ig_name [ig_name...]]
```

The following is a summary of the information you need to supply when you use the snapdrive storage connect command.

Requirement	Argument	Value you want to supply
Specify one or more LUNs that you want to connect to the host. If you connect one or more LUNs, the first argument must use the long form of the LUN name, which specifies the storage system name, the volume name, and the name of the LUN within the volume. To specify additional LUNs, you can use the LUN name alone if the new LUN is on the same storage system and volume as the previous LUN. Otherwise, you can specify a new storage system name and volume name (or just a volume name) to replace the previous values.		
◆ A LUN (-lun)	long_lun_name	
◆ Additional LUNs	lun_name (long or short form)	
(Optional: NetApp recommends that you use the default igroup for your host instead of supplying an igroup name.)		
◆ Igroup name (-igroup)	ig_name	

Examples of storage connect command lines

The following are some sample command lines that use the snapdrive storage connect command.

```
# snapdrive storage connect -lun filer1:/vol/vol1/toaster_lun2  
  
# snapdrive storage connect -lun filer1:/vol/vol1/toaster_lun2  
toaster_lun3 toaster_lun4
```

Checklist for using the storage disconnect command

Checklist for disconnecting storage

This checklist is a planning aid for disconnecting storage. It summarizes the information you need to provide.

Usage format for snapdrive storage disconnect:

```
snapdrive storage disconnect -lun long_lun_name [lun_name...]

snapdrive storage disconnect
  {-dg | -fs | hostvol} file_spec [file_spec ...]
  [{-dg | -fs | -hostvol} file_spec
  [file_spec ...] ...]
  [-full]
```

The following is a summary of the information you need to supply when you use the snapdrive storage disconnect command.

Requirement	Argument	Value you want to supply
Based on the command you enter, you can remove mappings from any of the following: ◆ LUNs If you create one or more LUNs, the first argument must use the long form of the LUN name, which specifies the storage system name, the volume name, and the name of the LUN within the volume. To specify additional LUNs, you can use the LUN name alone if the new LUN is on the same storage system and volume as the previous LUN. Otherwise, you can specify a new storage system name and volume name (or just a volume name) to replace the previous values. ◆ Disk or volume groups ◆ File systems ◆ Host or logical volumes The value you enter for the <i>file_spec</i> argument must identify the host entity you’re disconnecting.		
◆ A LUN (-lun)	long_lun_name	
◆ Additional LUNs	lun_name (long or short form)	
◆ Disk group (-dg file_spec) or volume group (-vg file_spec)	name of the disk or volume group	
◆ File system (-fs file_spec)	filesystem_name	

Requirement	Argument	Value you want to supply
◆ Host volume (-hostvol <i>file_spec</i>) or logical volume (-lvol <i>file_spec</i>)	<i>name of the host or logical volume</i>	
If you want SnapDrive for UNIX to disconnect the storage you specify even if you include on the command line a host-side entity that contains other entities (such as a disk group that contains one or more host volumes), include the -full option on the command line. If you do not include this option, you must specify only empty host-side entities.		
-full	~	

Checklist for using the host connect command

Checklist for connecting storage to the host

This checklist is a planning aid for connecting the host to the existing storage. This command applies only to the host side; it does not affect the storage system. Here is a summary of the information you need to provide.

Usage format for snapdrive host disconnect:

```
snapdrive host connect -lun long_lun_name [lun_name ...]
```

The following is a summary of the information you need to supply when you use the snapdrive host connect command.

Requirement	Argument	Value you want to supply
Specify one or more LUNs that you want to connect to the host. If you connect one or more LUNs, the first argument must use the long form of the LUN name, which specifies the storage system name, the volume name, and the name of the LUN within the volume. To specify additional LUNs, you can use the LUN name alone if the new LUN is on the same storage system and volume as the previous LUN. Otherwise, you can specify a new storage system name and volume name (or just a volume name) to replace the previous values.		
◆ A LUN (-lun)	long_lun_name	
◆ Additional LUNs	lun_name (long or short form)	

Checklist for using the host disconnect command

Checklist for disconnecting storage from the host

This checklist is a planning aid for disconnecting the host from the existing storage. This command applies only to the host side; it does not affect the storage system. Here is a summary of the information you need to provide.

Usage format for snapdrive host disconnect:

```
snapdrive host disconnect
  {-dg | -fs | -hostvol} file_spec [file_spec ...]
  [{-dg | -fs | -hostvol} file_spec
  [file_spec ...] ...]
  [-full]
```

The following is a summary of the information you need to supply when you use the snapdrive host disconnect command.

Requirement	Argument	Value you want to supply
Based on the command you enter, you can disconnect any of the following host-side entities without changing the state of the storage system where the underlying LUNs reside. ◆ Disk or volume groups ◆ File systems ◆ Host or logical volumes The value you enter for the <i>file_spec</i> argument must identify the host entity you're disconnecting.		
◆ Disk group (-dg <i>file_spec</i>) or volume group (-vg <i>file_spec</i>)	<i>name of the disk or volume group</i>	
◆ File system (-fs <i>file_spec</i>)	<i>filesystem_name</i>	
◆ Host volume (-hostvol <i>file_spec</i>) or logical volume (-lvol <i>file_spec</i>)	<i>name of the host or logical volume</i>	
If you want SnapDrive for UNIX to disconnect the storage you specify even if you include on the command line a host-side entity that contains other entities (such as a disk group that contains one or more host volumes), include the -full option on the command line. If you do not include this option, you must specify only empty host-side entities.		

Requirement	Argument	Value you want to supply
-full	~	

Checklist for using the storage delete command

Checklist for deleting storage

This checklist is a planning aid for deleting storage.

Usage format for snapdrive storage delete:

```
snapdrive storage delete -lun long_lun_name [lun_name...]

snapdrive storage delete
  {-dg | -fs | -hostvol} file_spec [file_spec ...]
[{-dg | -fs | -hostvol} file_spec [file_spec ...] ...]
[-full]
```

Note

If the storage you specify is in use, this operation fails.

The following is a summary of the information you need to supply when you use the `snapdrive storage delete` command.

Requirement	Argument	Value you want to supply
Based on the command you enter, you can delete any of the following types of storage: ◆ LUNs Specify one or more LUNs that you want to delete from the storage system. If you create one or more LUNs, the first argument must use the long form of the LUN name, which specifies the storage system name, the volume name, and the name of the LUN within the volume. To specify additional LUNs, you can use the LUN name alone if the new LUN is on the same storage system and volume as the previous LUN. Otherwise, you can specify a new storage system name and volume name (or just a volume name) to replace the previous values. ◆ Disk or volume groups ◆ File systems ◆ Host or logical volumes The value you enter for the <i>file_spec</i> argument must identify the host entity you’re deleting.		
◆ A LUN (-lun)	<i>long_lun_name</i>	
◆ Additional LUNs	<i>lun_name</i> (long or short form)	
◆ Disk group (-dg <i>file_spec</i>) or volume group (-vg <i>file_spec</i>)	<i>name of the disk group or volume group</i>	

Requirement	Argument	Value you want to supply
◆ File system (-fs <i>file_spec</i>)	<i>filesystem_name</i>	
◆ Host volume (-hostvol <i>file_spec</i>) or logical volume (-lvol <i>file_spec</i>)	<i>name of the host volume or logical volume</i> Note _____ You must supply both the requested volume and the disk group containing it; for example, -hostvol dg3/acct_volume. _____	
If you want SnapDrive for UNIX to delete the storage you specify even if you include on the command line a host-side entity that contains other entities (such as a disk group that contains one or more host volumes), include the -full option on the command line. If you do not include this option, you must specify only empty host-side entities.		
-full	~	

About this chapter

This chapter contains a list of the commands supported in SnapDrive for UNIX and information about the keywords and arguments that work with them.

Topics in this chapter

This chapter discusses the following topics:

- ◆ [“Summary of the SnapDrive for UNIX commands”](#) on page 244
- ◆ [“SnapDrive for UNIX keywords and arguments”](#) on page 247

Summary of the SnapDrive for UNIX commands

Command summary SnapDrive for UNIX supports the following command lines:

```
snapdrive config access {show | list} filer-name
snapdrive config check luns
snapdrive config delete filename [filename ...]
snapdrive config list
snapdrive config prepare luns -count count
snapdrive config set user_name filename [filename ...]
snapdrive config show [host_file_name]
snapdrive host connect -lun long_lun_name [lun_name ...]
snapdrive host disconnect {-vg | -dg | -fs | -lvol | -hostvol}
file_spec [file_spec ...] [{-vg | -dg | -fs | -lvol | -hostvol}
file_spec [file_spec ...] ...] [-full]
snapdrive snap connect fspec_set [fspec_set...] -snapname
long_snap_name [-igroup ig_name [ig_name ...]] [-autoexpand]
[-autorename]
```

Note

The *fspec_set* argument has the following format:

```
{-vg | -dg | -fs | -lvol | -hostvol} src_fspec [dest_fspec]
{-destdg | -destvg } dg_name] [{-destlv | -desthv } lv_name]
```

```
snapdrive snap create { -vg | -dg | -fs | -lvol |
-hostvol } file_spec [file_spec ...] [ { -vg | -dg | -fs | -lvol |
-hostvol } file_spec [file_spec...] ] -snapname snap_name
[-force [-noprompt]]

snapdrive snap delete [-snapname] long_snap_name [snap_name...]
[-verbose] [-force [-noprompt]]

snapdrive snap disconnect {-vg | -dg | -fs | -lvol | -hostvol}
file_spec [file_spec ...] [{-vg | -dg | -fs | -lvol | -hostvol}
file_spec [file_spec ...]] [-full]

snapdrive snap rename [-snapname] old_long_snap_name new_snap_name
[-force [-noprompt]]
```

```

snapdrive snap restore -snapname long_snap_name {-dg | -vg | -
hostvol | -lvol | -fs} file_spec [file_spec ...]
[{-dg | -vg | -hostvol | -lvol | -fs} file_spec
[file_spec ...] ...] [-force [-noprompt]]

snapdrive snap { show | list } -filer filename [filename...]
[-verbose]

snapdrive snap { show | list } -filervol filervol [filervol...]
[-verbose]

snapdrive snap { show | list } { -vg | -dg | -fs | -lvol |
-hostvol } file_spec [file_spec ...] [-verbose]

snapdrive snap { show | list } [-snapname] long_snap_name
[snap_name ...] [-verbose]

snapdrive storage connect -lun long_lun_name [lun_name ...]
[-igroup ig_name [ig_name ...]]

snapdrive storage create -lun long_lun_name [lun_name ...] -lunsize
size [{ -dg | -vg } dg_name] [-igroup ig_name [ig_name ...]]

snapdrive storage create host_lvm_fs_spec -filervol long_filer_path -
dgsize size [-igroup ig_name [ig_name ...]]

snapdrive storage create host_lvm_fs_spec -lun long_lun_name
[lun_name ...] -lunsize size [-igroup ig_name [ig_name ...]]

```

Note

You can use one of the following three formats for the *host_lvm_fs_spec* argument, depending on the type of storage you want to create. (Remember that -dg is a synonym for -vg and -hostvol is a synonym for -lvol.)

To create a file system:

```
-fs file_spec [-fstype type] [-fsopts options] [-hostvol
file_spec] [-dg dg_name]
```

To create a logical or host volume, use this format:

```
[-hostvol file_spec] [-dg dg_name]
```

To create a disk group, use this format:

```
-dg dg_name
```

```

snapdrive storage delete -lun long_lun_name [lun_name ...]

snapdrive storage delete { -vg | -dg | -fs | -lvol | -hostvol }
file_spec [file_spec ...] [{ -vg | -dg | -fs | -lvol | -hostvol }
file_spec [file_spec ...]] [-full]

snapdrive storage disconnect -lun long_lun_name [lun_name...]

```

```

snapdrive storage disconnect {-vg | -dg | -fs | -lvol | -hostvol}
file_spec [file_spec ...] [{-vg | -dg | -fs | -lvol | -hostvol}
file_spec file_spec ...] ...] [-full]

snapdrive storage resize {-dg | -vg} file_spec [file_spec ...]
{-growby | -growto} size [-addlun [-igroup ig_name [ig_name ...]]]

snapdrive storage { show | list } -filer filer_name
[filer_name ...] [-verbose] [-quiet]

snapdrive storage { show | list } -filervol long_filer_path
[filer_path...] [-verbose] [-quiet]

snapdrive storage { show | list } -lun long_lun_name
[lun_name ...] [-verbose] [-quiet]

snapdrive storage { show | list } { -vg | -dg | -fs | -lvol |
-hostvol } file_spec [file_spec ...] [{ -vg | -dg | -fs | -lvol | -
hostvol } file_spec [file_spec ...]] [-verbose] [-quiet]

snapdrive version

```

SnapDrive for UNIX keywords and arguments

Command-line options

SnapDrive for UNIX enables you to include the following options as appropriate with its commands. In certain cases, you can abbreviate these options. For example, you can use `-h` instead of `-help`.

Option	Description
<code>-addlun</code>	Tells SnapDrive for UNIX to add a new, internally-generated LUN to a storage entity in order to increase its size.
<code>-autoexpand</code>	Used with <code>snap connect</code> to enable you to request that a disk group be connected when you supply a subset of the logical volumes or file systems in the disk group.
<code>-autorename</code>	Used with <code>snap connect</code> to enable the command to rename any newly-connected LVM entities for which the default name is already in use.
<code>-dgsiz</code> (synonymous with <code>-vgsiz</code>)	Used with the <code>storage create</code> command to specify the size in bytes of the disk group you want to create.
<code>-force</code> (or <code>-f</code>)	Causes operations to be attempted that SnapDrive for UNIX would not undertake ordinarily. SnapDrive for UNIX prompts you to ask for confirmation before it executes the operation.

Option	Description
<code>-fsopts</code>	The options you want passed to the host operation that creates the new file system. Depending on your host operating system, this host operation might be a command such as the <code>mkfs</code> command. The argument you supply with this option usually needs to be specified as a quoted string and must contain exactly the text to be passed to the command. For example, you might enter <code>-o largefiles</code> as the option you want passed to the host operation.
<code>-fstype</code>	The type of file system you are creating. The file system must be a type that SnapDrive for UNIX supports for your operating system. It currently supports the following types: ◆ AIX: JFS and JFS2 ◆ HP-UX: VxFS ◆ Linux: Ext3 ◆ Solaris: VxFS Current values that you can include with this argument are <code>vxfs</code> (Solaris and HP-UX), <code>jfs2</code> (AIX), or <code>ext3</code> (Linux)
<code>-full</code>	Allows operations on a specified host-side entity to be performed even if the entity is not empty (for example, the entity might be a volume group containing one or more logical volumes).
<code>-growby</code>	The number of bytes you want to add to a LUN or disk group in order to increase its size.

Option	Description
<code>-growto</code>	The target size in bytes for a LUN, disk group, or volume group. SnapDrive for UNIX automatically calculates the number of bytes necessary to reach the target size and increases the size of the object by that number of bytes.
<code>-help</code> (or <code>-h</code>)	Prints out the usage message for this command and operation. Enter this option by itself without other options.
<code>-lunsize</code>	The size of the LUN in bytes to be created by a given command.
<code>-noprompt</code>	Suppresses prompting during command execution. By default, any operation that might have dangerous or non-intuitive side effects prompts you to confirm that the SnapDrive for UNIX should be attempted. This option overrides that prompt; when combined with the <code>-force</code> option, SnapDrive for UNIX performs the operation without asking for confirmation.
<code>-quiet</code> (or <code>-q</code>)	Suppresses the reporting of errors and warnings, regardless of whether they are normal or diagnostic. It returns zero (success) or non-zero status. The <code>-quiet</code> option overrides the <code>-verbose</code> option.
<code>-verbose</code> (or <code>-v</code>)	Displays detailed output, where appropriate. All commands and operations accept this option, although some might ignore it.
<code>-vgsize</code> (synonymous with <code>-dgsiz</code>)	Used with the <code>storage create</code> command to specify the size in bytes of the volume group you want to create.

Rules for keywords

SnapDrive for UNIX uses keywords to specify sequences of strings corresponding to the host and storage system objects with which you are working. The following rules apply to SnapDrive for UNIX keywords:

- ◆ Precede each keyword with a hyphen (-).
- ◆ Do not concatenate keywords.
- ◆ Enter the entire keyword and hyphen, not an abbreviation.

Command-line keywords

Here are the keywords you can use with the SnapDrive for UNIX commands. You use them to specify the targets of the SnapDrive for UNIX operations. These keywords can take one or more arguments.

Note—
Some LVMs refer to *disk groups* and some refer to *volume groups*. In SnapDrive for UNIX, these terms are treated as synonyms. Moreover, some LVMs refer to *logical volumes* and some refer to *volumes*. SnapDrive for UNIX treats the term *host volume* (which was created to avoid confusing *host logical volumes* with *storage system volumes*) and the term *logical volume* as synonymous.

Keyword	Argument used with this keyword
-dg (synonymous with -vg)	The name of the host disk group. You can enter the name of either a disk group or a volume group with this option.
-destdg -desthv -destlv -destvg	The destination group or volume.
-filer	The name of a storage system.
-filervol	The name of the storage system and a volume on it.
-fs	The name of a file system on the host. The name used is the directory where the file system is currently mounted or is to be mounted (i.e., the mount point).

Keyword	Argument used with this keyword
-hostvol (synonymous with -lvol)	The host volume name, including the disk group that contains it. For example, you might enter large_vg/accounting_lvol.
-igroup	The name of an initiator group (igroup). NetApp strongly recommends that you use the default igroup that SnapDrive for UNIX creates instead of specifying an igroup on the target storage system. The default igroup is <code>hostname_protocol_SbIg</code> where: ◆ <i>hostname</i> is the local (non-domain qualified) name of the current host. ◆ <i>protocol</i> is either <code>fc</code> or <code>iscsi</code>, depending which protocol the host is using. If the igroup <code>hostname_protocol_SbIg</code> does not exist, SnapDrive for UNIX creates it and places all the initiators for the host in it. If it exists and has the correct initiators, SnapDrive for UNIX uses the existing igroup. If the igroup exists, but does not contain the initiators for this host, SnapDrive for UNIX creates a new igroup with a different name and uses that igroup in the current operation. To avoid using the same name, SnapDrive for UNIX includes a unique number when it creates the new name. In this case, the name format is <code>hostname-number_protocol_SbIg</code>. If you supply your own igroup name, SnapDrive for UNIX does not validate the contents of the igroup. This is because it cannot always determine which igroups are on the hosts. For details on specifying igroups, see the SnapDrive for UNIX man page.

Keyword	Argument used with this keyword
-lun	The name of a LUN on a storage system. For the first LUN name you supply with this keyword, you must supply the full pathname (storage system name, volume name, and LUN name). For additional LUN names, you can specify either only the names within their volume (if the volume stays unchanged) or a path to indicate a new storage system name or a new volume name (if you just want to switch volumes).
-lvol (synonymous with -hostvol)	The logical volume name, including the volume group that contains it. For example, you might enter large_vg/accounting_lvol as the logical volume name.
-raw	The name of a raw device file on the host.
-snapname	The name of a snapshot.
-vg (synonymous with -dg)	The name of the host volume group. You can enter the name of either a disk group or a volume group with this option.

Command-line arguments

The following table describes the arguments you can specify with the keywords.

Use the format `snapdrive type_name operation_name [<keyword/option> <arguments>]`; for example, if you wanted to create a snapshot called `snap_hr` from the host file system `/mnt/dir`, you would enter the following command line:

```
snapdrive snap create -fs /mnt/dir -snapname snap_hr
```

Argument:	Description:
<i>dest_fspect</i>	The name by which the target entity will be accessible after its disk group is connected.
<i>dgname</i>	The name of a disk group or volume group.
<i>filename</i>	The name of a storage system.

Argument:	Description:
<i>filer_path</i>	A pathname to a filer object. This name can contain the filer name and volume, but it doesn't have to if SnapDrive for UNIX can use default values for the missing components based on values supplied in the previous arguments. The following are examples of pathnames: <pre>test_filer:/vol/vol3/qtrees_2</pre> <pre>/vol/vol3/qtrees_2</pre> <pre>qtrees_2</pre>
<i>file_spec</i>	The name of a host object such as a host volume, a disk or volume group, or a file system. In general, you use the <i>file_spec</i> argument as one of the following: ◆ An object you want SnapDrive for UNIX to take a snapshot of or to restore from a snapshot. ◆ An object that you want to either create or use when provisioning storage. The objects do not have to be all the same type. If you supply multiple host volumes, they must all belong to the same volume manager. If you supply values for this argument that resolve to redundant disk groups or host volumes, the command fails. Example of incorrect usage: This example assumes dg1 has host volumes hv1 and hv2, with file systems fs1 and fs2. As a result, the following arguments would fail because they involve redundant disk groups or host volumes. <pre>-dg dg1 -hostvol dg1/hv1</pre> <pre>-dg dg1 -fs /fs1</pre> <pre>-hostvol dg1/hv1 -fs /fs1</pre> Example of correct usage: This example shows the correct usage for this argument. <pre>-hostvol dg1/hv1 dg1/hv2</pre> <pre>-fs /fs1 /fs2</pre> <pre>-hostvol dg1/hv1 -fs /fs2</pre>

Argument:	Description:
<i>fspec_set</i>	Used with the <code>snap connect</code> command to identify a host LVM entity that was captured in the snapshot specified by the <i>long_snap_name</i> argument. The argument also lets you specify a set of destination names that SnapDrive for UNIX uses when it makes the entity available in the newly connected copy of the snapshot. The format of for <i>fspec_set</i> is: <pre>{ -vg -dg -fs -lvol -hostvol } src_fspec [dest_fspec] [{ -destdg -destvg } dg_name] [{ -destlv -desthv } lv_name]</pre> For more information, see the discussion for “Checklist for using the snapshot connect command” on page 219.
<i>host_lvm_fspec</i>	Lets you specify whether you want to create a file system, logical volume, or disk group when you are executing the <code>storage create</code> command. This argument may have any of the three formats shown below. The format you use depends on the entity you want to create. Note The <code>-dg</code> and <code>-vg</code> options are synonyms that reflect the fact that some operating systems refer to disk groups and others refer to volume groups. In addition, <code>-lvol</code> and <code>-hostvol</code> are also synonyms. This guide uses <code>-dg</code> to refer to both disk groups and volume groups and <code>-hostvol</code> to refer to both logical volumes and host volumes. To create a file system, use this format: <pre>-fs file_spec [-fstype type] [-fsopts options] [-hostvol file_spec] [-dg dg_name]</pre> To create a logical or host volume, use this format: <pre>[-hostvol file_spec] [-dg dg_name] -hostvol</pre> To create a disk or volume group, use this format: <pre>file_spec [-dg dg_name] -dg dg_name</pre>

Argument:	Description:
	You must name the top level entity that you are creating. you may not need to supply names for any underlying entities. If you don't supply names for the underlying entities, SnapDrive creates them with internally generated names. If you specify that SnapDrive for UNIX create a file system, you must specify a type that SnapDrive for UNIX supports with the host LVM. These types include the following: ◆ AIX: JFS and JFS2 ◆ HP-UX: VxFS ◆ Linux: Ext3 ◆ Solaris: VxFS Current values that you can include with the <code>-fstype</code> argument are vxfs (Solaris and HP-UX), jfs2 (AIX), or ext3 (Linux). The option <code>-fsopts</code> is used to specify options to be passed to whatever host operation creates the new file system (e.g., <i>mkfs</i>).
<i>ig_name</i>	The name of an initiator group.
<i>long_filer_path</i>	A pathname that includes the filer name, volume name, and possibly other directory and file elements within that volume. The following are examples of long pathnames: <pre>test_filer:/vol/vol3/qtrees_2 10.10.10.1:/vol/vol4/lun_21</pre>
<i>long_lun_name</i>	A name that includes the storage system name, volume, and LUN name. The following is an example of a long LUN name: <pre>test_filer:/vol/vol1/lunA</pre>

Argument:	Description:
<i>long_snap_name</i>	A name that includes the storage system name, volume, and snapshot name. The following is an example of a long snapshot name: <pre>test_filer:/vol/account_vol:snap_20040202</pre> With the <code>snap show</code> and <code>snap delete</code> operations, you can use the asterisk (*) character as a wildcard to match any part of a snapshot name. If you use a wildcard character, you must place it at the end of the snapshot name. SnapDrive for UNIX displays an error message if you use a wildcard at any other point in a name. Example: This example uses wildcards with both the <code>snap show</code> command and the <code>snap delete</code> command: <pre>snap show myfiler:/vol/vol2:mysnap* myfiler:/vol/vol2:/yoursnap* snap show myfiler:/vol/vol1/qtreet1:qtree_snap* snap delete 10.10.10.10:/vol/vol2:mysnap* 10.10.10.11:/vol/vol3:yoursnap* hersnap</pre> Limitation for wildcards: You cannot enter a wildcard in the middle of a snapshot name. For example, the following command line produces an error message because the wildcard is in the middle of the snapshot name: <pre>banana:/vol/vol1:my*snap</pre>
<i>lun_name</i>	The name of a LUN. This name does not include the storage system and volume where the LUN is located. The following is an example of a LUN name: <pre>lunA</pre>
<i>path</i>	Any path name.

Argument:	Description:
<i>short_snap_name</i>	A name that consists of only the snapshot name. You can use the short form of the snapshot name when you do not need to include the storage system and storage system volume as part of the name. For example, when you create a snapshot, the storage system and volume must be the same as those of the original data. You can use wildcards with the <code>snap show</code> and <code>snap delete</code> operations. If you use a wildcard, you must place it at the end of the snapshot name. SnapDrive for UNIX displays an error message if you enter a wildcard in the middle of a snapshot name. Example: This example uses wildcards with the <code>snap show</code> command: <pre>snap show toaster:/vol/vol1:mysnap*</pre> Limitations for wildcards: You cannot enter a wildcard in the middle of a snapshot name. For example, the following snapshot name is incorrect: <code>toaster:/vol/vol1:test*snap</code>
<i>size</i>	The size in bytes of an object such as a LUN or a disk group. You can abbreviate sizes; for example, you might enter use 2g or 500m.
<i>snap_name</i>	A snapshot name in either the long or short form.
<i>src_fspect</i>	Identifies an LVM entity within the snapshots you want to connect to a different location using <code>snap connect</code>. This name is used to locate the disk group within the snapshot that needs to be connected in order to make the entity available.
<i>volname</i>	The name of the filer volume.

Glossary

Attach Kit (HAK)	The attach kit or host attach kit (HAK) is the standard software and documentation NetApp supplies to allow you to connect your host to a NetApp storage system.
Cluster	You can have a storage system cluster or a host cluster. The standard storage system cluster involves having a pair of storage systems attached via fabric to a switch in such a way that one can serve its partner's data if the partner fails. A host cluster refers to a host with multiple nodes. See the Network Appliance attach or support kit documentation for your operating system for information on the number of supported nodes on a host.
Clustered failover (CFO)	This is a method that ensures data availability by transferring the data service of a failed storage system to another storage system in the cluster. Transfer of data service is often transparent to users and applications
Disk group	A unit of storage that can contain one or more volumes and one or more file systems. A set of LUNs managed by a Logical Volume Manager (LVM) on the host. LUNs in a disk group are combined in stripes, mirrors, and other RAID combinations by the host LVM. These combinations can then be divided into host volumes, which can be used to store data directly (as if they were raw disks), or can be used to host a file system. The data on a disk group is generally spread over all LUNs in the disk group, making it difficult to impossible to restore only a single LUN's data. SnapDrive for UNIX uses term <i>disk group</i> interchangeably with the term <i>volume group</i>.
Failover	Failover refers to situations where a system component fails and another component takes over its functions while the system continues to operate. Failover can occur between storage system heads configured in a CFO cluster, among redundant FCP or iSCSI paths from the host to storage system LUNs, or from one clustered host to another.

In a clustered failover, an alternate system takes over and emulates the primary system if the primary system becomes unusable.

FCP initiator

See Initiator.

File system

A file system in the SnapDrive for UNIX context refers to the file organization system supported by NetApp on the host operating system.

Host

A host is the machine running a UNIX operating system that accesses storage on a filer.

Host bus adapter (HBA)

A Host Bus Adapter (HBA) refers to the adapters used to connect hosts and storage systems in a NetApp SAN so that hosts can access LUNs on the storage systems using FCP. *See also* LUN (Logical Unit Number).

Host volume

A host volume exists within a logical volume manager. It is a subset of a disk group (volume group) that acts as a disk to applications accessing it. The data in a host volume can be distributed over all or a subset of the disks in the disk group to which it belongs. Use of a host volume allows the administrator to easily resize the amount of storage assigned to a file system or application, without having to change the *disk* the application is using. The term *host volume* is used interchangeably with *logical volume*. The term *volume*, which is often used in conjunction with host logical volume managers, is avoided in this document, to avoid confusion with storage system volumes.

initiator

The hardware that initiates data exchange.

initiator group (igroup)

Initiators that are grouped as authorized to access particular targets.

There are three categories of igroups that apply to SnapDrive for UNIX

- ◆ Foreign. This igroup contains only initiators not found on the local host.
- ◆ Local. This igroup contains only initiators found on the local host.

- ◆ Mixed. This igroup contains both initiators found on the local host and initiators not found there.

Logical volume See host volume.

LUN (Logical Unit Number) LUN refers to a logical unit of storage identified by a number.

LUN ID This is the numerical identifier for a LUN.

Network Interface Card (NIC) A network interface card (NIC) refers to a Gigabit Ethernet (commonly known as GbE) or a Fast Ethernet card that is compliant with the IEEE 802.3 standards.

These cards can provide the following connectivity functions:

- ◆ Connect hosts and filers to a local area network (LAN)
- ◆ Connect hosts and filers to data-center switching fabrics, specifically, enable hosts to connect to LUNs on filers using iSCSI

Raw Data that is passed to an I/O device without being interpreted or processed. Raw access to a disk device implies that it is being used directly by an application to store data, rather than having a file system on it and letting the application store data in files in the file system.

Rollback This is a snapshot of the data on the storage system that SnapDrive for UNIX makes before it begins a `snap restore` operation. Having a rollback snapshot means that, in the event of a problem with a `snap restore` operation, you can restore the data on the storage system to the state it was in before the operation began. This option is enabled in the `snapdrive.conf` file by default (see “[Setting values in snapdrive.conf](#)” on page 82).

SAN (Storage Area Network) A SAN (Storage Area Network) is a storage setup composed of one or more storage systems and connected to one or more hosts in an FCP or an iSCSI environment.

To a host running SnapDrive for UNIX, a connected SAN is just another target storage device within which SnapDrive for UNIX can create and manage LUNs.

Snapshot

A snapshot refers to the NetApp Snapshot™ technology. This technology enables the recovery after accidental deletion or modification of the data stored on a NetApp storage system by referencing a point-in-time image of that data.

storage system

A storage system (sometimes called a filer) is a NetApp storage system that supports the FCP (Fibre Channel Protocol), iSCSI, and/or GbE (Gigabit Ethernet) protocols.

storage system volume

A storage system (or filer) volume is a functional unit of filer storage. It comprises a collection of physical disks. A volume can be composed of one or more RAID groups to ensure data integrity and availability if multiple disks fail simultaneously within the same volume. For more information about storage system volumes, see the *Data ONTAP Storage Management Guide*.

Support Kit

The support kit is the standard software and documentation NetApp supplies to allow you to connect your host to a NetApp storage system using an iSCSI protocol.

Target port

For NetApp SANs, a target is the HBA port on the storage system used to receive the SCSI I/O commands that the host initiator port sends. *See also* Initiator.

Volume

See either host volume or storage system volume.

Index

Symbols

160, 172, 221

<host>_info utility 184

A

access permissions 4

control levels 101

examples 105

permissions file 101

setting 101

setting configuration option 66

steps for setting 102

steps for viewing 105

-addlun command option 247

AIX host

example of successful installation 29

example of successful uninstallation 36

example of uninstall warning 37

files installed 55

installation steps for SnapDrive for UNIX 23

installing SnapDrive for UNIX 22

requirements 22

uncompressing NOW file 23

uninstalling SnapDrive for UNIX 31

using SMIT 24

verify installation 30

arguments

file_spec 253

filer_path 253

long_filer_path 255

path 256

rules for using 252

short_snap_name 257

snap_name 257

attach kit

definition 259

install 19

audit log file 66, 86

disabling 87

example 90

maximum size 67

number of old files 67

specifying pathname 87

using 89

audit-log-file option 66

audit-log-save option 67

-autoexpand command option 247

-autorename command option 247

autosupport

configuration options 68

example of a message 95

setting up 95

snapdrive.conf option 68

using 94

autosupport-enabled option 68

autosupport-filer 68

available-lun-reserve option 69

B

bytes

adding to storage 167

C

CD-ROM

getting software 21

CFO

definition 259

checklist

general information 208

host connect command 237

host disconnect command 238

snap connect command 219

snap create command 209

snap delete command 217

snap disconnect command 223

snap rename command 211

snap restore command 215

snap show command 213

storage connect command 234

storage create command 225

storage delete command 240

storage disconnect command 235

- storage resize command 229
- storage show command 231
- cluster
 - definition of 259
 - host 259
 - storage system 259
 - support in SnapDrive for UNIX 6
- clustered failover
 - definition 259
- command line mode
 - using 97
- commands
 - addlun option 247
 - arguments 252
 - autoexpand option 247
 - autorename option 247
 - config access 63
 - config check luns 63
 - config delete 63
 - config list 62, 109
 - config prepare luns 63
 - config set 63
 - config show 62
 - dg keyword 250
 - disconnect 144
 - file_spec arguments 253
 - filer_path argument 253
 - force option 131, 247
 - fsopts option 248
 - fstype option 248
 - full option 248
 - growby option 248
 - growto option 249
 - help option 249
 - host connect 176
 - host disconnect 178
 - hostvol keyword 251
 - igroup keyword 251
 - keyword 250
 - keyword rules 250
 - long_filer_path argument 255
 - lunsize option 249
 - noprompt option 249
 - options 247
 - path argument 256
 - quiet option 249
 - running from the command line 97
 - short_snap_name argument 257
 - snap connect 139
 - snap create 117
 - snap delete 147
 - snap list 126
 - snap rename 127
 - snap restore 131
 - snap show 120
 - snap_name argument 257
 - snapdrive config delete 109
 - snapdrive config set 108, 109
 - snapdrive version 84
 - snapname keyword 252
 - storage show 161, 162, 163
 - summary of all commands 244
 - swverify 42, 48
 - verbose option 249
- Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products 11
- config command
 - access 63
 - check luns 63
 - delete 63
 - list 62
 - options 62
 - prepare luns 63
 - set 63
 - show 62
- config delete command
 - SnapDrive for UNIX 109
- config list command 109
- config list operation 62
- config set command 108
 - SnapDrive for UNIX 109
- config show operation 62
- configuration options
 - all-access-if-rbac-unspecified 66
 - audit-log-file 66
 - audit-log-max-size option 67
 - audit-log-save 67
 - autosupport-enabled 68
 - autosupport-filer 68
 - available-lun-reserve 69

- default-noprompt 69
- default-transport 69
- device-retries option 70
- device-retry-sleep-secs 70
- enable-implicit-host-preparation 71
- filer-restore-retries 72
- filer-restore-retry-sleep 72
- filesystem-freeze-timeout 72
- mgmt-retries 72
- mgmt-retry-sleep 72
- mgmt-sleep-long-secs 73
- multipathing-type 73
- password-file 74
- path 74
- prefix-filer-lun 74
- prepare-lun-count 75
- recovery-log-file 75
- recovery-log-save 76
- snapcreate-consistency-retries 77
- snapdelete-delete-rollback-with-snap 77
- snaprestore-delete-rollback-after-restore 77
- snaprestore-make-rollback 78
- snaprestore-must-make-rollback 79
- snaprestore-must-make-snapinfo-on-qtree 76
- snaprestore-snapmirror-check 79
- space-reservations-enabled 79
- trace-enabled 80
- trace-level 80
- trace-log-file 80, 81
- trace-log-save 81
- use-https-to-filer 82
- viewing with config show 62
- configuring
 - storage system volumes 17
- connecting snapshots
 - examples 141
 - restrictions 138
 - steps for connecting 139
- conventions
 - formatting x
 - keyboard xi
- creating snapshots
 - examples 119
 - snap create fails 187
 - steps for creating 117

- using SnapDrive for UNIX 117
- creating storage 152
- Ctrl-C
 - disrupts snapshot restoration 130

D

- data
 - changes after snap restore 189
 - raw 261
- data collection utility
 - examples 186
 - executing 185
 - information collected 184
- Data ONTAP
 - version 14
- Data ONTAP Block Access Management Guide 12
- default-noprompt option 69
- default-transport option 69
- deleting snapshots
 - examples 148
 - reasons 146
 - restrictions 146
 - steps 147
- depot file
 - HP-UX installation 39
- device-retries option 70
- device-retry-sleep-secs option 70
- dg 250
- disconnecting snapshots
 - examples 145
 - steps 144
- disconnecting storage 172, 174
- disk
 - data changes after restore 189
- disk group
 - definition 259
 - increasing size 166
- displaying snapshots
 - examples 122
 - steps for displaying 120
- documentation
 - additional 11
 - Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products 11

- Data ONTAP Block Access Management Guide 12
- FCP host attach kit 11
- iSCSI host support kit 12
- SnapDrive for UNIX Compatibility Matrix 11
- SnapDrive for UNIX man page 11
- SnapDrive for UNIX Quick Start 11
- SnapDrive for UNIX Release Notes 11
- storage system 14
- System Configuration Guide 12

E

- enable-implicit-host-preparation option 71
- error messages 190
 - device not NetApp LUN 192
 - limit on open files 192
 - return codes 194
- examples 148
 - snap connect command 221
 - snap create command 210
 - snap delete command 218
 - snap disconnect command 224
 - snap rename command 212
 - snap restore command 216
 - storage connect command 234
 - storage create command 228
 - storage resize command 230
 - storage show command 233

F

- failover
 - definition 259
 - NetApp cluster 6
- FCP
 - license required 15
- FCP host attach kit documentation 11
- Fibre Channel Protocol *See* FCP
- file system
 - creating 153
 - definition 260
- file systems
 - resizing 169

- file_spec 253
- filer
 - cluster 6
 - defined ix
- filer *See* storage system
- filer volume *See* storage system volume
- filer_path
 - example 253
- filer-restore-retries option 72
- filer-restore-retry-sleep option 72
- FilerView
 - setting snap reserve 17
- files
 - .rmp (Linux) 46
 - <host>_info utility 184
 - access permissions 101
 - audit log 66, 67, 86
 - audit log file 89, 90
 - audit-log-save 67
 - depot (HP-UX) 39
 - installed on hosts 55
 - limit open files 192
 - RBAC permissions 66
 - recovery log 76, 86
 - recovery log file 90, 92
 - sdhost-name.prbac 101
 - snapdrive.conf 62
 - snapdrive.dc 184
 - trace log 80, 81, 86
- filesystem-freeze-timeout option 72
- force
 - command option 247
- fsopts
 - command option 248
- fstype
 - command option 248
- full command option 248

G

- growby
 - command option 248
- growby
 - command option 167

- growto
 - command option 249
- growto
 - command option 167
- guidelines
 - creating storage system volumes 16

H

- HAK (Host Attach Kit)
 - install 19
- help
 - command option 249
- host
 - <host>_info utility 184
 - cluster 259
 - clustered 6
 - displaying storage information 163
 - limit on open files 192
 - non-originating 136
 - volume 260, 262
- Host Attach Kit *See* HAK
- Host Bus Adapter (HBA)
 - definition 260
- host connect command
 - checklist 237
 - steps for using 176
- host disconnect command
 - checklist 238
 - examples 178
 - steps for using 178
- host entities
 - creating
 - storage create command
 - creating host entities 153
 - deleting 181
 - disconnecting LUNs 174
- host volume 3
 - definition 260
- hosts
 - AIX requirements 22
 - comparing platforms 9
 - differences in commands 244
 - differences in platforms 10
 - HP-UX requirements 38

- installation steps on AIX host 23
- installation steps on HP-UX host 39
- installation steps on Linux host 46
- installation steps on Solaris host 51
- installing on AIX 22
- installing on HP-UX 38
- installing on Linux 45
- installing on Solaris 50
- Linux requirements 45
- preparing for LUNs 85
- resizing 169
- Solaris requirements 50
- terminology differences 10

- hostvol 251

- HP-UX host
 - depot file 39
 - example of successful installation 42
 - example of successful uninstall 44
 - files installed 55
 - installation steps for SnapDrive for UNIX 39
 - installing SnapDrive for UNIX 38
 - moving file to host 39
 - requirements 38
 - uninstalling SnapDrive for UNIX 43

- HTTPS 3
 - enabling 111

I

- igroup keyword 251
- increasing storage size 166
- installation
 - verify 42
- installing
 - completing the installation 54
 - on AIX host 22
 - on HP-UX host 38
 - on Linux host 45
 - on Solaris host 50
 - steps for AIX host 23
 - steps for HP-UX host 39
 - steps for Linux host 46
 - steps for Solaris host 51
 - upgrading SnapDrive for UNIX 59
 - verify 42, 48

- IP address
 - partner 15
- IP interface
 - used by SnapDrive for UNIX 3
- iSCSI
 - license required 15
 - problem if halted 192
- iSCSI host support kit documentation 12

K

- keywords 250
 - dg 250
 - hostvol 251
 - igroup 251
 - rules for using 250
 - snapname 252

L

- license
 - FCP 15
 - iSCSI 15
 - storage system 15
- Linux host
 - .rpm file 46
 - adding host entries for LUNs 85
 - error if iSCSI halted 192
 - example of successful installation 47
 - files installed 55
 - installation steps for SnapDrive for UNIX 46
 - installing SnapDrive for UNIX 45
 - moving downloaded file 46
 - preparing for LUNs 85
 - requirements 45
 - uninstalling SnapDrive for UNIX 49
 - verify installation 48
- log files
 - audit 67
 - audit log 66, 86, 89, 90
 - disabling 87
 - enabling 86, 87
 - recovery 76
 - recovery log 86, 90, 92

- specifying pathnames 87
- trace 80, 81
- trace log 86
- Logical Unit Numbers *See* LUN
- logical volume
 - creating
 - disk group
 - creating 153
- Logical volume manager *See* LVM
- logical volumes 3
 - definition 261
- logins
 - storage system 107
- long_filer_path
 - example 255
- LUN ID
 - defined ix
 - definition 261
- LUNs
 - adding host entries for LUNs 85
 - connecting to host 176
 - connecting to storage 170
 - defined ix
 - definition 261
 - deleting 179
 - disconnecting from host entity 174
 - disconnecting from storage 172
 - disconnecting to host 178
 - displaying storage information 161
 - not recognized 192
 - preparing host 85
 - preparing the host 63
- lunsize
 - command option 249
- LVM
 - using to resize 169
 - works with SnapDrive for UNIX 3

M

- man page
 - SnapDrive for UNIX 11

- mgmt-retries option 72
- mgmt-retry-sleep option 72
- mgmt-sleep-long-secs option 73
- multipathing-type option 73

N

- name/value pair
 - snapdrive.conf file 83, 112
- Network Interface Card (NIC), definition 261
- non-originating host 136
- noprompt
 - command option 249
- NOW description page
 - most up-to-date requirements 5
- NOW site
 - downloading software 20

O

- open files
 - limit 192
- operating system
 - limit on open files 192
 - storage system 14
- options
 - addlun 247
 - autoexpand 247
 - autorename 247
 - force 131, 247
 - fsopts 248
 - fstype 248
 - full 248
 - growby 248
 - growto 249
 - help 249
 - lunsize 249
 - noprompt 249
 - quiet 249
 - using on command lines 247
 - verbose 249

P

- partner
 - specify IP address 15

- password-file option 74
- passwords
 - storage system 107
- path 256
- path configuration option 74
- prefix-filer-lun option 74
- prepare-lun-count option 75

Q

- Quick Start Guide 11
- quiet
 - command option 249

R

- raw
 - definition 261
- RBAC *See* access permissions
- recommendations
 - snap reserve 17
 - storage system volumes 17
- recovery log file 76, 86
 - disabling 87
 - enabling 87
 - example 92
 - specifying pathname 87
 - using 90
- recovery-log-fileoption 75
- recovery-log-save option 76
- Red Hat Linux host
 - see Linux host 45
- renaming snapshots
 - examples 128
 - steps for renaming 127
- requirements
 - AIX host 22
 - HP-UX host 38
 - Linux host 45
 - most current 5
 - Solaris host 50
 - storage system 14
- restoring snapshots
 - examples 133
 - force option 131
 - interrupting with Ctrl-C 130

- restoring from a different host 136
- restrictions 129
- snap restore fails 187
- steps for restoring 130
- time required 130
- warning 188
- restrictions
 - on connecting snapshots 138
 - on restoring snapshots 129
 - snapshots 116
- return codes
 - error messages 194
- rollback 77, 78, 79
 - definition 261
- rpm file
 - Linux installation 46
- rsh
 - not used 4

S

- SAN (Storage Area Network), definition 261
- security
 - considerations 3
 - features 100
- short_snap_name 257
- SMIT
 - using on AIX 24
 - using on AIX host 23
- snap connect command 139
 - checklist 219
 - examples 141
- snap create command 117
 - checklist 209
 - examples 119, 210
 - fails 187
- snap delete command 147, 148
 - checklist 217
 - commands
 - snap delete 146
 - example 218
- snap disconnect command 144
 - checklist 223
 - examples 145, 224
- snap rename command 127
 - checklist 211
 - examples 128, 212
- snap reserve 16
 - recommended value 16
 - resetting 17
- snap restore command 131
 - checklist 215
 - data on disk changes 189
 - examples 133, 216
 - fails 187
 - warning 188
- snap show command
 - checklist 213
 - examples 122
- snap_name 257
- snapcreate-consistency-retries option 77
- snapdelete-delete-rollback-with-snap option 77
- SnapDrive for UNIX
 - access permissions 4, 101
 - additional documentation 11
 - AIX requirements 22
 - audit log 89
 - autosupport 94
 - capabilities 2
 - cautions 16
 - checking version 84
 - command arguments 252
 - command keywords 250
 - command options 247
 - comparing platforms 9
 - config access 63
 - config check luns 63
 - config delete 63
 - config delete command 109
 - config list 62
 - config prepare luns 63
 - config set 63
 - config set command 108, 109
 - config show 62
 - configurations 6
 - connecting snapshots 137
 - creating snapshots 117
 - creating storage 153, 154, 157
 - data collection utility 184

- deleting snapshots 146, 147
- deleting user logins 109
- disabling log files 87
- disconnecting snapshots 144
- displaying information on storage 160
- displaying snapshots 120
- downloading software 20
- enabling log files 86
- error messages 190
- files installed on host 55
- getting software 19
- getting software from CD-ROM 21
- how it works 2
- HP-UX requirements 38
- HTTPS 3
- installation steps on AIX 23
- installation steps on HP-UX 39
- installation steps on Linux 46
- installation steps on Solaris 51
- installing on AIX 22
- installing on HP-UX 38
- installing on Linux 45
- installing on Solaris 50
- Linux requirements 45
- log files 86
- logical volumes 3
- multiple platforms 9
- NetApp Snapshot technolog 2
- prerequisites 4
- recommendations 16
- recovery log 90
- renaming snapshots 127
- requires attach kit 19
- restoring snapshots 129
- running from the command line 97
- security 3
- security features 100
- setting snap reserve 17
- setting up autosupport 95
- setting values in snapdrive.conf 82
- snapshot restrictions 116
- Solaris requirements 50
- specifying log file pathnames 87
- specifying logins for storage system 107
- stack 5

- storage 2
- storage system volumes 16, 17
- summary of all commands 244
- supported snapshots 114
- time to restore snapshots 130
- troubleshooting tips 187
- uninstalling from AIX host 31
- uninstalling from HP-UX host 43
- uninstalling from Linux host 49
- uninstalling from Solaris host 52
- upgrading versions 59
- version command 84
- volume manager terminology 10
- works with LVM 3
- SnapDrive for UNIX Compatibility Matrix 11
- SnapDrive for UNIX Quick Start 11
- SnapDrive for UNIX Release Notes 11
- snapdrive.conf
 - all-access-if-rbac-unspecified option 66
 - audit-log-file option 66
 - audit-log-max-size option 67
 - audit-log-save option 67
 - autosupport-enabled option 68
 - autosupport-filer option 68
 - available-lun-reserve option 69
 - changes with version 2.0 59
 - default values 63
 - default-noprompt option 69
 - default-transport option 69
 - device-retries option 70
 - device-retry-sleep-secs option 70
 - enable-implicit-host-preparation option 71
 - filer-restore-retries option 72
 - filer-restore-retry-sleep option 72
 - filesystem-freeze-timeout option 72
 - mgmt-retries option 72
 - mgmt-retry-sleep option 72, 73
 - multipathing-type option 73
 - name/value pair 83, 112
 - password-file option 74
 - path option 74
 - prefix-filer-lun option 74
 - prepare-lun-count option 75
 - recovery-log-file option 75
 - recovery-log-save option 76

- setting values 82
- snapcreate-consistency-retries option 77
- snapdelete-delete-rollback-with-snap option 77
- snaprestore-delete-rollback-after-restore option 77
- snaprestore-make-rollback option 78
- snaprestore-must-make-rollback option 79
- snaprestore-must-make-snapinfo-on-qtree option 76
- snaprestore-snapmirror-check option 79
- space-reservations-enabled option 79
- trace-enabled option 80
- trace-level option 80
- trace-log-file option 80, 81
- trace-logs-ave option 81
- use-https-to-fileroption 82
- viewing 62
- snapdrive.dc 184
 - examples 186
 - executing 185
 - tasks performed 184
- SnapDriver for UNIX
 - snapshots 2
- snapname 252
- SnapRestore
 - license required 15
- snaprestore-delete-rollback-after-restore option 77
- snaprestore-make-rollback option 78
- snaprestore-must-make-rollback 79
- snaprestore-must-make-snapinfo-on-qtree option 76
- snaprestore-snapmirror-check option 79
- snapshot
 - about 114
 - created with SnapDrive for UNIX 2
 - creating 117
 - definition 262
 - deleting 146, 147
 - disconnecting 144
 - displaying information 120
 - examples of snap connect command 141
 - examples of snap create command 119
 - examples of snap delete command 148
 - examples of snap disconnect command 145
 - examples of snap rename command 128
 - examples of snap restore command 133
 - examples of snap show command 122
 - interrupting restoration 130
 - listing on storage system 126
 - listing with wildcards 126
 - reasons to delete 146
 - renaming a snapshot 127
 - restoring 129
 - restoring from a different host 136
 - restrictions 116
 - restrictions on connecting 138
 - restrictions on deleting 146
 - restrictions on restoring 129
 - rollback 77, 78, 261
 - rollback option 79
 - spanning storage systems 115
 - spanning volumes 115
 - steps for connecting 139
 - steps for displaying 120
 - steps for renaming 127
 - steps for restoring 130
 - time to restore 130
 - using wildcards 257
- snapshots
 - connecting from different locations 137
- software
 - downloading 20
 - getting 19
 - getting from CD-ROM 21
 - installing 22, 38, 45
 - installing on Solaris 50
 - uncompressing 23
 - uncompressing on Solaris 51
- Solaris host
 - adding host entries for LUNs 85
 - files installed 55
 - installation steps for SnapDrive for UNIX 51
 - installing SnapDrive for UNIX 50
 - preparing for LUNs 85
 - requirements 50
 - uncompressing SnapDrive for UNIX software 51
 - uninstalling SnapDrive for UNIX 52

- space reservations 16
- space-reservations-enabled option 79
- special messages xi
- stack
 - required 5
- storage
 - connecting 170
 - connecting LUNs to host 176
 - created with SnapDrive for UNIX 2
 - creating 152
 - creating specific LUNs 153, 157
 - deleting 179, 180, 181
 - disconnecting 172, 174
 - disconnecting LUNs to host 178
 - disconnecting specific LUNs 173
 - displaying information 160
 - examples of storage show command 164
 - growing 167
 - resizing 166
 - setting up LUNs automatically 154
 - steps for deleting host entities 181
 - steps for deleting LUNs 180
 - steps for displaying information 161
 - steps for displaying information on host entities 163
 - steps for displaying information on storage systems 162
- Storage Area Network *See* SAN
- storage commands
 - using across storage systems 151
- storage connect command
 - checklist 234
 - examples 171, 234
 - steps for using 170
- storage create command 153, 154, 157
 - checklist 225
 - examples 158, 228
 - 156
- storage delete command 180, 181
 - checklist 240
 - limitations 179
- storage disconnect command 174
 - checklist 235
 - examples 174, 175
 - limitations 172
 - steps for using 173
 - tips 173
- storage provisioning 150
- storage resize command 167
 - checklist 229
 - examples 167, 230
 - limitations 166
 - steps for using 167
- storage show command 161, 162, 163
 - checklist 231
 - examples 164, 233
 - restrictions 160
- storage system
 - cluster 259
 - clustered 6
 - configuring volumes 17
 - definition 262
 - deleting SnapDrive for UNIX logins 109
 - displaying storage information 162
 - documentation 14
 - failover 6
 - guides for creating volumes 16
 - login required 3
 - logins 107
 - operating system 14
 - preparing 14
 - required license 15
 - requirements 14
 - resetting snap reserve 17
 - setup program 15
 - specify login information 107
 - specify partner IP address 15
 - target 262
 - volume 262
 - volume preparation 16
- storage system volume
 - definition 262
- storage systems
 - spanning 115, 151
 - verifying user names 109
- support kit
 - definition 262
- swverify command
 - verify installation 42, 48

T

target

- definition 262

Terminology ix

trace log file 80, 81, 86

- disabling 87

- enabling 87

- specifying pathname 87

trace-enabled option 80

trace-level option 80

trace-log-file option 80, 81

trace-log-save option 81

troubleshooting

- data collection utility 184

- data on disk changes 189

- device not recognized 192

- error message return codes 194

- error messages 190

- iSCSI halted 192

- limits on open files 192

- snap create fails 187

- snap restore fails 187

- tips 187

U

uninstalling

- SnapDrive for UNIX from AIX host 31

- SnapDrive for UNIX from HP-UX host 43

- SnapDrive for UNIX from Linux host 49

- SnapDrive for UNIX from Solaris host 52

upgrading

- SnapDrive for UNIX 59

use-https-to-filer option 82

user logins

- deleting from SnapDrive for UNIX 109

- displaying names with config list 109

utilities

- <host>_info 184

- data collection 184

V

-verbose

- command option 249

version command

- SnapDrive for UNIX 84

volume

- logical 261

volume group

- increasing size 166

volumes

- definition 262

- host 260

- preparing for storage system 16

- recommendations for storage system 17

- spanning 115

- storage system 16, 17, 262

W

wildcards

- example 256, 257

- using with snapshots 257