



**IBM System Storage N series
SnapDrive 4.2 for UNIX Quick Start Guide
(AIX, HP-UX, Linux, Solaris)**

Contents

Copyright and trademark information	1
Copyright information	1
Trademark information	2
Notices	5
Preface	7
Getting information, help, and services	7
Before you call	7
Using the documentation	7
Web sites	7
Accessing online technical support	8
Supported servers and operating systems	8
Special messages	8
How to send your comments	9
Relevant documentation	9
A quick look at SnapDrive for UNIX	11
What SnapDrive for UNIX is	11
Supported host platforms	11
Terms used in this guide	12
Installing SnapDrive for UNIX	13
Some configuration commands	15
Snapshot copies and storage in SnapDrive for UNIX	19
SnapDrive for UNIX to create Snapshot copies	19
Performing Snapshot operations	19
Creating a Snapshot copy	19
Displaying information about a Snapshot copy	20
Restoring a Snapshot copy	21
Connecting a Snapshot copy to a different location	22
Disconnecting a Snapshot copy from a location on a host	23
Deleting a Snapshot copy	24
SnapDrive for UNIX for storage provisioning	24
Creating a storage entity	25
Displaying information about a storage entity	25
Disconnecting a storage entity that is mapped to a host	26
Connecting a storage entity to a host	26
Disconnecting a storage entity from the host side	27
Connecting a storage entity from the host side	27
Increasing the storage entity size	27
Deleting a storage entity	28
Snapshot copies and storage	28
Setting up a file system on a host entity	29
Displaying information about the storage	30

Increasing the size of storage	31
Making a Snapshot copy	32
Restoring a Snapshot copy	33
Disconnecting the LUN mappings from the host.	33
Connecting an existing LUN to the host	34
Deleting storage from the host.	34
Splitting the volume clone or LUN clone	35
Some SnapDrive for UNIX features in brief	37
Web services and daemon	37
SnapDrive for UNIX configuration checker	37
Automatic detection of host entities	37
FlexClone volumes in SnapDrive for UNIX	38
Volume-based SnapRestore	39
Index	41

Copyright and trademark information

This section includes copyright and trademark information, and important notices.

Copyright information

Copyright © 1994–2010 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2010 IBM Corporation. All rights reserved. Printed in the U.S.A.

U.S. Government Users Restricted Rights—Use, duplication, or disclosure restricted by GSA ADP Schedule Contract with IBM Corporation.

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.

References in this documentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's or NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM and NetApp, are the user's responsibility.

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 NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN

NO EVENT SHALL NETAPP 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.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. 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 NetApp.

The product described in this manual may be protected by one or more U.S.A. 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

IBM, the IBM logo, and [ibm.com](http://www.ibm.com/legal/copytrade.shtml) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

NetApp; the NetApp logo; the Network Appliance logo; Bycast; Cryptainer; Cryptoshred; DataFabric; Data ONTAP; Decru; Decru DataFort; FAServer; FilerView; FlexCache; FlexClone; FlexShare; FlexVol; FPolicy; gFiler; Go further, faster; Manage ONTAP; MultiStore; NearStore; NetCache; NOW

(NetApp on the Web); ONTAPI; RAID-DP; SANscreen; SecureShare; Simulate ONTAP; SnapCopy; SnapDrive; SnapLock; SnapManager; SnapMirror; SnapMover; SnapRestore; SnapValidator; SnapVault; Spinnaker Networks; Spinnaker Networks logo; SpinAccess; SpinCluster; SpinFlex; SpinFS; SpinHA; SpinMove; SpinServer; SpinStor; StorageGRID; StoreVault; SyncMirror; Topio; vFiler; VFM; and WAFL are registered trademarks of NetApp, Inc. in the U.S.A. and/or other countries. Network Appliance, Snapshot, and The evolution of storage are trademarks of NetApp, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. The StoreVault logo, ApplianceWatch, ApplianceWatch PRO, ASUP, AutoSupport, ComplianceClock, DataFort, Data Motion, FlexScale, FlexSuite, Lifetime Key Management, LockVault, NOW, MetroCluster, OpenKey, ReplicatorX, SecureAdmin, Shadow Tape, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, Tech OnTap, Virtual File Manager, VPolicy, and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. Get Successful and Select are service marks of NetApp, Inc. in the U.S.A.

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

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, N.Y. 10504-1785
U.S.A.

For additional information, visit the web at:
<http://www.ibm.com/ibm/licensing/contact/>

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make

improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Preface

Getting information, help, and services

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure that you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure that they are connected properly.
- Check the power switches to make sure that the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.

Using the documentation

Information about N series hardware products is available in printed documents and a documentation CD that comes with your system. The same documentation is available as PDF files on the IBM NAS support Web site:

www.ibm.com/storage/support/nas/

Data ONTAP software publications are available as PDF files on the IBM NAS support Web site:

www.ibm.com/storage/support/nas/

Web sites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates.

- For NAS product information, go to the following Web site:
www.ibm.com/storage/nas/
- For NAS support information, go to the following Web site:
www.ibm.com/storage/support/nas/

- For AutoSupport information, go to the following Web site:
www.ibm.com/storage/support/nas/
- For the latest version of publications, go to the following Web site:
www.ibm.com/storage/support/nas/

Accessing online technical support

For online Technical Support for your IBM N series product, visit the following Web site:

www.ibm.com/storage/support/nas/

Supported servers and operating systems

IBM N series products attach to many servers and many operating systems. To determine the latest supported attachments, follow the link to the Interoperability Matrices from the following Web site:

www.ibm.com/storage/support/nas/

Special messages

This document might contain the following types of messages to alert you to conditions you need to be aware of. Danger notices and caution notices only appear in hardware documentation, where applicable.

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

Attention: An attention notice contains instructions that you must follow to avoid a system crash, loss of data, or damage to the equipment.

DANGER

<p>A danger notice warns you of conditions or procedures that can result in death or severe personal injury.</p>

CAUTION:

A caution notice warns you of conditions or procedures that can cause personal injury that is neither lethal nor extremely hazardous.

How to send your comments

Your feedback is important in helping us provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, send us your comments by e-mail to starpubs@us.ibm.com. Be sure to include the following:

- Exact publication title
- Publication form number (for example, GC26-1234-02)
- Page, table, or illustration numbers
- A detailed description of any information that should be changed

Relevant documentation

The SnapDrive for UNIX *Release Notes*, Interoperability Matrix, and the manual pages are some of the documents that provide you with more information for SnapDrive for UNIX.

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

If you want to...	Go to...
Information about how to install and configure the attach kits and support kits	Host attach kit/support kit documentation for your specific host
General product information	The IBM NAS product Web page at http://www.ibm.com/storage/nas/
Product support information	The IBM NAS support Web page at http://www.ibm.com/storage/support/nas/

A quick look at SnapDrive for UNIX

This provides a brief overview of SnapDrive for UNIX and how you install and configure it.

What SnapDrive for UNIX is

SnapDrive for UNIX simplifies the storage management tasks for you. SnapDrive for UNIX uses the Snapshot technology to create an image of data stored on a shared or unshared storage system. SnapDrive also helps you with storage provisioning.

Following are the few important tasks you can accomplish using SnapDrive for UNIX:

- **Back up and restore data:** SnapDrive for UNIX lets you use Snapshot technology to make an image (Snapshot copy) of host data that is stored on a storage system. This Snapshot copy provides you with a copy of that data, which you can restore later. The data in the Snapshot copy can exist on one storage system or span multiple storage systems and their volumes. These storage systems can be in cluster-wide shared or node-local file systems or disk groups or LUNs in a cluster environment.
- **Manage storage:** You can create and delete storage components, including disk groups, host volumes, file systems, and LUNs on cluster and noncluster environment. SnapDrive for UNIX enables you to manage this storage by expanding it, connecting it to a host, and disconnecting it.
- **Role-Based Access Control:** SnapDrive 4.0 for UNIX and later provides Role-Based Access Control (RBAC). RBAC allows a SnapDrive administrator to restrict access to a storage system for various SnapDrive operations. This limited or full access for storage operations depends on the role that is assigned to the user. This allows the storage administrators to limit the operations that SnapDrive users can perform depending on their assigned roles.

Supported host platforms

The most current list of host platforms that SnapDrive for UNIX supports is in the online SnapDrive for UNIX Interoperability Matrix.

SnapDrive for UNIX is supported on the following host platforms:

- IBM AIX
- HP-UX

- Linux

Note: Linux includes Red Hat Enterprise Linux, SUSE Linux Enterprise Server, and Oracle Enterprise Linux.

- Solaris SPARC and Solaris for x86 architecture

For information about the requirements for running SnapDrive for UNIX on each of its platforms, see the Interoperability Matrix and the Compatibility and Configuration Guide for FC and iSCSI Products.

Terms used in this guide

Before you use SnapDrive for UNIX, you must know about disk groups, volume groups, logical volumes, host volumes, Snapshot copies, `file_spec`, and LUNs.

Term	Description
Disk group and volume group	Some of the platforms that SnapDrive for UNIX supports use the term <i>disk group</i> while others use the term <i>volume group</i> . As a result, some commands include both <code>-dg</code> and <code>-vg</code> . In general, you can use either argument and get the same result. To make this guide easier to follow, it uses the term disk group and <code>-dg</code> to refer to both disk and volume groups.
Logical volume and host volume	Some of the platforms that SnapDrive for UNIX supports use the term volume or host volume while others use the term logical volume. To simplify the terminology and make a distinction between volumes on the host and volumes on the storage system, this guide uses the term host volume and the command line argument <code>-hostvol</code> to refer to all volumes on the host.
Snapshot copy	Snapshot copy is an image of data on one or more storage systems that was made at a specific point in time using Snapshot technology. You can use this Snapshot copy to recover the data that has been accidentally deleted or modified since you took the Snapshot copy.
<code>file_spec</code>	<code>file_spec</code> is an object, like a host volume, disk group, file system or Network File System (NFS) file, or directory tree, which SnapDrive for UNIX uses to create a Snapshot copy.
LUN	Logical Unit Number (LUN) refers to a logical unit of storage identified by a number.

Installing SnapDrive for UNIX

You can quickly install and set up SnapDrive for UNIX. The following steps takes you through the basic installation and configuration steps. It also provides pointers to where you can find more information about a step.

1. Ensure that your host and storage systems are set up correctly.

For the host:

- If you have an FC or iSCSI configuration, make sure the FC Host Utilities or iSCSI Host Utilities is installed and you have performed the following tasks. These tasks are explained in the FC Host Utilities and iSCSI Host Utilities documentation.

Note: In SnapDrive 4.2 for UNIX and later versions, the products FC Host Attach Kit and iSCSI Support Kit are named FC Host Utilities and iSCSI Host Utilities.

- Set up the multipathing software, if it is supported on your host platform.
- Verify that the host and storage system can communicate.
- Check that the host meets the SnapDrive for UNIX system requirements.

Note: SnapDrive for UNIX requirements are a subset of the FC Host Utilities or iSCSI Host Utilities system requirements.

For the storage system:

- Check that it meets the system requirements, including having the following setup:
 - Data ONTAP
 - The IP address on the storage system cluster (used if a takeover occurs)
 - The correct protocol license for your platform (FCP or iSCSI)
 - The SnapRestore license
 - Secure HTTP access to the storage system
 - MultiStore software on your storage system for vFiler unit setup
 - For NFS configurations:
 - The NFS service is on and operational
 - You have configured the exports file
 - You have a FlexClone license, for configurations with Data ONTAP 7.1 and FlexVol volumes
2. You can install the product software from software updates available for download.

Downloads are available only to entitled IBM N series customers who have completed the registration process on the IBM NAS support web site at www.ibm.com/storage/support/nas.

- a. Check the publication Matrix page at www.ibm.com/storage/support/nas for important alerts, news, interoperability details, and other information about the product before beginning the installation.
 - b. Sign in with your IBM ID and password.
If you do not have an IBM ID or password, or if you are adding new N series machines and serial numbers to an existing registration, click the Register link, follow the online instructions, and then sign in.
 - c. Select the N series software you want to download, and then select the Download view.
 - d. Use the “Software Packages” link on the Web page presented and follow the online instructions to download the software.
 - e. Launch the software installation program and then follow instruction on the prompts.
3. Log in as root and go to the directory where you placed the SnapDrive for UNIX software package.

Execute the install script or command to install SnapDrive for UNIX on your host:

- AIX: Uncompress the `ontap.snapdrive_aix_4_2.Z` file; then use SMIT to install `ontap.snapdrive_aix_4_2`
- HP-UX: `#swinstall -s pathname /ontapsnapdrive_hpux_4_2.depot snapdrive`
- Linux: `#rpm -U -v pathname/ontap.snapdrive.linux_4_2.rpm`

Note: Linux includes Red Hat Enterprise Linux, SUSE Linux, and Oracle Enterprise Linux.

- Solaris: Complete the following steps:
 - a. Uncompress the `ONTAPsnapdrive.tar.Z` file.
 - b. Extract the `ONTAPsnapdrive.tar` file.
 - c. Enter the `./install` command to install the software.
4. Specify the current login information for each storage system by using the `snapdrive config set user` command. SnapDrive for UNIX needs this information to access the storage system.

To get the user names for storage systems attached to the host, execute the `snapdrive config list` command. This command does not display the storage system passwords.

5. Note that the AutoSupport option is enabled by default, the `autosupport-enabled` on variable is set in the `snapdrive.conf` file. In SnapDrive 4.2 for UNIX and later versions, the `autosupport filer` variable

is not available in the `snapdrive.conf` file. The `snapdrive.conf` file is located in your installation directory. Use a text editor to edit it.

- For AIX, HP-UX, and Linux, `/opt/ontap/snapdrive/snapdrive.conf`
- For Solaris, `/opt/ontap/snapdrive/snapdrive.conf`

Note: The `snapdrive.conf` file contains several settings you can change, including the locations for the log files. In most cases, it is recommended that you accept the default values.

6. Specify the protocol that SnapDrive for UNIX uses as the transport type, when creating storage, if a decision is required. Enter `iscsi` or `fcp` value for the `default-transport` variable.
7. Specify the multipathing software to use by using the `multipathingtype` variable. The default value depends on the host operating system. This option applies only if one of the following is true:
 - There are more than one multipathing solutions available.
 - The configurations that include LUNs.
8. Specify the type of file system that you want to use in SnapDrive for UNIX operations by using the `fstype` variable.
9. Specify the type of volume manager that you want to use in SnapDrive for UNIX operations by using the `vmtype` variable.

Some configuration commands

There are few commands that help you in configuring SnapDrive for UNIX.

The following table summarizes additional commands that are helpful in configuring SnapDrive for UNIX.

Description	Command or action
Configuring and verifying your version of SnapDrive for UNIX software	
Check the values in the <code>snapdrive.conf</code> file.	<code>snapdrive config show</code>
Check the version of SnapDrive for UNIX.	<code>snapdrive version</code>
Start the SnapDrive for UNIX daemon.	<code>snapdrived start</code>
Prepare the host for creating a specific number of LUNs as well as to determine how many LUNs you can create. Note: You only need these commands on Linux hosts using FCP and iSCSI protocols, and Solaris hosts using FCP protocol. AIX and HP-UX hosts do not require any preparation prior to creating LUNs.	<code>snapdrive config prepare luns -count count snapdrive config check luns</code>

Description	Command or action
Prepare the guest for creating RDM LUN on Linux and Solaris x86 guest OS. Note: You only use these commands on Linux and Solaris x86 hosts using iSCSI protocols and FCP protocol with no host utilities.	<code>snapdrive config set -viadmin <user> <viadmin_name></code>
Create, estimate clone split for volume and LUN clone.	<code>snapdrive clone split</code>
Check for the following in the SFRAC cluster environment on a Solaris host: <ul style="list-style-type: none"> • SnapDrive for UNIX version • Group Membership Services/Atomic Broadcast (GAB) configuration • Cluster status • Cluster Volume Manager (CVM) status • Usage of rsh or ssh for a secure communication within the cluster nodes • Differences in setting the following configuration variable values in the <code>snapdrive.conf</code> file: <ul style="list-style-type: none"> – <code>default-transport= FCP</code> – <code>multipathing-type=DMP</code> 	<code>snapdrive config check cluster</code>
Change the path name and options for the log files. SnapDrive for UNIX logs information to three files: <ul style="list-style-type: none"> • An audit file • A recovery file • A trace file 	Edit the variables in the <code>snapdrive.conf</code> file.
Set and display access control permissions between a host and a storage system.	

Description	Command or action
<p>Specify the access control permissions a host has on a storage system.</p> <p>You can set the following access levels for a host on a storage system:</p> <ul style="list-style-type: none"> • NONE—No access. • SNAP CREATE—Create Snapshot copies. • SNAP USE—Delete and rename Snapshot copies. • SNAP ALL—Create, restore, delete, and rename Snapshot copies. • STORAGE CREATE DELETE—Create, resize, and delete storage. • STORAGE USE—Connect and disconnect storage. • STORAGE ALL—Create, delete, connect, and disconnect storage. • ALL ACCESS—All operations. 	<p>Edit the access control permissions file (<i>sd hostname.prbac</i>) on the storage system associated with that host.</p>
<p>Display information about the access control permissions a host has to a storage system.</p>	<p><code>snapdrive config access show filer_name</code></p>
<p>If you want to remove the specified user name-password pair from SnapDrive for UNIX.</p>	<p><code>snapdrive config delete filename [filename . . .]</code></p>

Snapshot copies and storage in SnapDrive for UNIX

You can view high-level steps that describe creating and using Snapshot copies and provisioning storage. It also lets you to use a sample script to work with Snapshot copies and storage.

SnapDrive for UNIX to create Snapshot copies

You can quickly start using SnapDrive for UNIX to take Snapshot copies. The steps provide the basic information needed to work with Snapshot copies as well as to set up a sample Snapshot copy and work with it.

Note: The following steps focus on the Snapdrive for UNIX Snapshot copy commands; they do not use the storage commands.

These steps and their examples show how you can easily perform the following tasks:

- Create a Snapshot copy
- Display information about the Snapshot copy
- Rename a Snapshot copy
- Restore disk groups, file systems, or LUNs from a Snapshot copy
- Connect a Snapshot copy to a different location on the host
- Disconnect a Snapshot copy from a location on the host
- Delete a Snapshot copy

It is assumed that you have experience setting up and running applications on your host operating system. It is recommended that you must read and understand the guidelines of Snapshot operations before you perform them. For more information about the guidelines, see the guidelines section of each Snapshot operation in the *SnapDrive for UNIX Installation and Administration Guide* of your specific platform.

Performing Snapshot operations

SnapDrive for UNIX enables you to execute Snapshot commands. It is assumed that you have a volume group named `vg1` and a storage system named `toaster`. When performing this exercise, it is recommended that you do it on a test system to ensure that you do not accidentally create a problem on a production system.

Creating a Snapshot copy

You can create a Snapshot copy using SnapDrive for UNIX.

Create a Snapshot copy of a volume group called vg1 using the snapdrive snap create command, and then provide a Snapshot copy names as snap1. # snapdrive snap create -vg vg1 -snapname snap1

Example 1:

The following command lines use the snapdrive snap create command with disk groups dg1 and dg2, which have host volumes myvol1 and myvol2. The host volume dg1/myvol2 has a file system mounted on /myfs/fs2. Each command line creates a Snapshot copy called toaster:/vol/vol1:snap1.

```
# snapdrive snap create -fs /myfs/fs2 -snapname snap1
# snapdrive snap create -dg dg1 dg2 -snapname snap1
# snapdrive snap create -hostvol dg1/myvol1
-snapname snap1 -force
# snapdrive snap create -hostvol dg1/myvol2 -fs /fs2
-snapname snap1
```

The following example creates a Snapshot copy of storage system entities that do not have dependent writes during Snapshot copy creation. SnapDrive for UNIX creates a Snapshot copy in which the file system /mnt/fs1 and the disk group dg1 are crash consistent as individual storage entities, but are not treated as a group:

```
# snapdrive snap create -fs /mnt/fs1 -dg dg1 -unrelated -snapname fs1_dg1
```

The following example creates a Snapshot copy of a shared file system on dg testdg:

```
# snapdrive snap create -dg testdg -snapname sfractestdgsnap
```

The following example creates a Snapshot copy of NFS volumes and directories:

```
# snapdrive snap create -fs /mnt/nfs -snapname nfs_snap
Successfully created snapshot nfs_snap on abc:/vol/vol1
```

You can also create Snapshot copies of volume groups that span multiple storage system volumes or multiple storage systems. SnapDrive for UNIX checks the reads or writes against the storage entities in the Snapshot copy to ensure that all Snapshot copy data is crash consistent. SnapDrive for UNIX will not create a Snapshot copy unless the data is crash consistent.

Displaying information about a Snapshot copy

You can see the information about the Snapshot copies.

Display information about the Snapshot copy you just created using the `snapdrive snap show` command.

```
# snapdrive snap show
-snapname toaster:/vol/vol1:snap1
```

Note: You can use either `snapdrive snap show` or `snapdrive snap list` in the command line. These commands are synonyms.

Examples of additional command lines:

These are valid `snapdrive snap show` command lines. If you are not sure of a Snapshot copy name, you can include a wildcard character (*), but only at the end of a Snapshot copy name.

```
# snapdrive snap show
-snapname toaster:/vol/vol2:snapA snapX snapY
# snapdrive snap show -verbose -snapname toaster:/vol/vol2:snapA /vol/vol3:snapB snapC
# snapdrive snap show -snapname toaster:/vol/vol2:snapA
# snapdrive snap show toaster:/vol/vol2:snapA
# snapdrive snap show toaster:/vol/vol0:*
# snapdrive snap show -verbose -hostvol hp1vol23
# snapdrive snap show -verbose -snapname toaster:/vol/vol2:snapA
# snapdrive snap show -verbose -dg dg21
# snapdrive snap show -verbose toaster:/vol/vol2:snapA /vol/vol3:snapB snapC
```

Restoring a Snapshot copy

This section describes how to restore a Snapshot copy.

Restore the Snapshot copy using the `snapdrive snap restore` command.

```
# snapdrive snap restore -dg vg1 -snapname toaster:/vol/vol1:bkup40105
```

Make sure that you are not in the directory in which you want to restore this command.

Note: The `snapdrive snapshot restore` command can only restore Snapshot copies that is created by SnapDrive for UNIX

Attention: This command replaces the current contents of the LUNs that make up the host disk groups with the contents of the LUNs in the specified Snapshot copy. It can take several minutes. Do not press Ctrl-C during this procedure. If you try to halt the procedure, you risk leaving the disk groups in an incomplete state.

Example 1:

These are examples of `snapdrive snap restore` command lines:

```
# snapdrive snap restore -fs /mnt/dir -snapname toaster:/vol/vol1:NewSnap33
# snapdrive snap restore -dg dg1 dg2 -snapname toaster:/vol/vol1:Tuesday
```

The following set of command lines uses a host that has a disk group `dg1` with the 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`. All of the following commands restore the same data:

```
# snapdrive snap restore -dg dg1 -snapname toaster:/vol/vol1:snap1
# snapdrive snap restore -vg dg1 -snapname toaster:/vol/vol1:snap1
# snapdrive snap restore -hostvol dg1/myvol2 dg1/myvol1 -snapname toaster:/vol/vol1:snap1
# snapdrive snap restore -hostvol dg1/myvol2 -fs /fs2 -snapname toaster:/vol/vol1:snap1
```

The following example shows to restore a storage entity in a host cluster environment:

```
# snapdrive snap restore -fs /mnt/sfortesting -snapname f270-197-109:/vol/vol2:testsfarcsnap
-devicetype shared
```

The following command line restores an NFS entity:

```
# snapdrive snap restore -fs /mnt/nfs -snapname toaster:/vol/vol1:tru1
```

Note: Use this command to restore failed and takeover filespecs in a cluster environment.

Connecting a Snapshot copy to a different location

You can connect a Snapshot copy to a different location.

Connect a Snapshot copy to a new location on a host in which you took the Snapshot copy.

```
# snapdrive snap connect -fs /db2/datafiles2 /db2_bkup40105 -snapname toaster:/vol/vol1:bkup40105
```

Example 1:

The following command can be used to connect to a Snapshot copy to a new location on the same host or a different host. The following are sample `snapdrive snap connect` command lines:

```
# snapdrive snap connect -hostvol srcdg/myvol destdg/myvol -snapname toaster:/vol/vol1:snap1
# snapdrive snap connect -vg vg2 -snapname toaster:/vol/vol1:snap1
# snapdrive snap connect -fs /db2/datafiles2 /db2_snap1 -snapname toaster:/vol/vol1:snap1
# snapdrive snap connect -fs /db2/datafiles2 -snapname toaster:/vol/vol1:snap1 -autoexpand -autorename
```

The following command can be used to connect to a Snapshot copy that contains shared storage entities on an originating cluster:

```
# snapdrive snap connect -fs /mnt/sfortesting /mnt/sfortesting2 -snapname f270-197109:/vol/vol2:testsfarcsnap
-autorename
```

The following command can be used to connect to a Snapshot copy that contains shared storage entities on a nonoriginating cluster:

```
# snapdrive snap connect -fs /mnt/sfortesting -snapname f270-197-109:/vol/vol2:testsfarcsnap
-devicetype shared
```

The following command can be used to connect to a Snapshot copy that contains the NFS volume and directory:

```
# snapdrive snap connect -fs /mnt/nfs -snapname toaster:/vol/vol1:nfs_snap -autorename
```

Note: When you identify a disk group to be connected by specifying a logical volume or file system, you must specify all the logical volumes contained in that disk group in order to connect the entire disk group. The `-autoexpand` option simplifies this process by letting you name only a subset of the logical volumes or file systems in the disk group. The option then expands the connection to the rest of the logical volumes or file systems in the disk group. The `-autorename` option tells SnapDrive for UNIX to rename the destination entities as needed to avoid naming conflicts.

Disconnecting a Snapshot copy from a location on a host

You can disconnect a Snapshot copy from a location on a host.

Remove the mappings between the host and the LUNs in the Snapshot copy `bkup40105` by using the `snapdrive snap disconnect` command.

```
snapdrive snap disconnect -fs /bkup40105
```

Note: SnapDrive for UNIX does not modify the Snapshot copy during a disconnect except to delete the temporary LUNs that were created during the `snapdrive snap connect` operation.

Example 1:

The following are sample snapdrive snap disconnect command lines:

```
# snapdrive snap disconnect -hostvol dg5/myvolume
# snapdrive snap disconnect -vg vg1 vg10
# snapdrive snap disconnect -dg dg1 dg2 dg3
```

The following command disconnects a Snapshot copy that contains shared storage entities:

```
# snapdrive snap disconnect -lun long_lun_name [lun_name...] [-devicetype {shared | dedicated}]
```

Deleting a Snapshot copy

You can delete a Snapshot copy by using SnapDrive for UNIX.

Delete the Snapshot copy bkup40105 using the snapdrive snap delete command.

```
# snapdrive snap delete toaster:/vol/vol11:bkup40105
```

Example of additional command lines:

The snapdrive snap delete command displays a list of Snapshot copies that it is deleted. The following is an example of the type of output it provides:

```
# snapdrive snap delete -v toaster:/vol/vol10:snap1 snap2 snap3
snapdrive: deleting
toaster:/vol/vol10:snap1
toaster:/vol/vol10:snap2
toaster:/vol/vol10:snap3
```

SnapDrive for UNIX for storage provisioning

SnapDrive for UNIX provides end-to-end storage management. It provides a number of storage options that enable you to automate storage provisioning tasks on the storage system to manage the entire storage hierarchy.

SnapDrive for UNIX lets you perform the following storage tasks:

- Create storage by creating LUNs, file systems, logical volumes, and disk groups
- Display information about storage entities
- Connect LUNs and storage entities to the host
- Resize storage by increasing the size of the storage
- Disconnect LUN and storage entities' mappings from the host

- Connect to host-side storage entities
- Disconnect from host-side storage entities
- Delete storage

In addition to helping you to set up a sample storage system and work with it, each step also includes the examples of additional SnapDrive for UNIX command lines.

These steps assume you have experience setting up and running applications on your host operating system. You should read and understand the guidelines of storage operations before you perform them. For more information on guidelines, see the guidelines section of each storage operation in the *SnapDrive for UNIX Installation and Administration Guide* for your platform.

Note: To make the exercises in these steps easier to follow, they do not use all the options available for each SnapDrive for UNIX command. For a complete list of these commands, their options, and detailed explanations of how they work, see the *SnapDrive for UNIX Installation and Administration Guide* for your platform.

Creating a storage entity

This section describes how to create a storage entity using SnapDrive for UNIX.

Create a host volume and file system on a 1-GB LUN using the `snapdrive storage create` command.

```
snapdrive storage create -fs /mnt/qa/sdufs1 -fstype jfs2 -lun f270c-198-147:/vol/hacmpvol/sdulun1 -lunsize 1g
LUN f270c-198-147:/vol/hacmpvol/sdulun1 ... created
mapping new lun(s) ... done
discovering new lun(s) ... done
LUN to device file mappings:
- f270c-198-147:/vol/hacmpvol/sdulun1 => /dev/hdisk4
disk group sdufs1_SdDg created
host volume sdufs1_SdHv created
file system /mnt/qa/sdufs1 created
```

In this example, the `-fstype` option is used to specify the file system type. This option applies only if there is more than one file system type available on the host.

The `snapdrive storage create` command automatically performs all the tasks needed to set up LUNs, including preparing the host, performing discovery mapping, and connecting to each LUN you create.

Displaying information about a storage entity

This section describes how to display information about a storage entity.

Display the information about the storage entity you just created using the `snapdrive storage show` command.

```
snapdrive storage show -all

dg: sdufs1_SdDg dgtype LVM_AIX
hostvol: /dev/sdufs1_SdHv state: AVAIL

fs: /dev/sdufs1_SdHv mount point: /mnt/qa/sdufs1
(persistent) fstype jfs2

device filename adapter path size proto state clone lun path backing snapshot
-----
/dev/hdisk4 - P 1g FC online No f270c-198-147:/vol/hacmpvol/sdulun1 -
```

Disconnecting a storage entity that is mapped to a host

This section describes how to disconnect a storage entity that is mapped to a host.

Disconnect a disk group called `sdufs1_SdDg` from the host using the `snapdrive storage disconnect` command.

```
snapdrive storage disconnect -dg sdufs1_SdDg -full disconnecting disk group sdufs1_SdDg
- hostvol sdufs1_SdDg/sdufs1_SdHv ...
disconnected
- dg sdufs1_SdDg ... disconnected
- LUN f270c-198-147:/vol/hacmpvol/sdulun1 ...

disconnected
0001-669 Warning:

Please save information provided by this command.
You will need it to re-connect disconnected filespecs.
```

The storage disconnect operation removes the LUNs, or the LUNs and storage entities that were mapped to the host. When SnapDrive for UNIX removes the LUN mappings, it exports the disk groups or file systems that the LUNs contain.

Connecting a storage entity to a host

This section describes how to connect a storage entity using SnapDrive for UNIX.

Connect a file system and host volume to the LUN on the host using the `snapdrive storage connect` command.

```

snapdrive storage connect -fs /mnt/qa/sdufs1 -hostvol sdufs1_SdDg/sdufs1_SdHv -lun f270c-198-147:/vol/hacmpvol/sdulun1

mapping lun(s) ... done
discovering lun(s) ... done
LUN f270c-198-147:/vol/hacmpvol/sdulun1 connected
- device filename(s): /dev/hdisk4
Importing sdufs1_SdDg
Connected fs /mnt/qa/sdufs1

```

You can use the `snapdrive storage connect` command to map storage to a new location, and to access existing storage from a different host than the one used to create it. It also lets you make existing LUNs, file systems, disk groups, and logical volumes accessible on a new host.

Disconnecting a storage entity from the host side

This section describes how to disconnect a storage entity from the host side.

Disconnect `sdufs1_SdDg` from the host without unmapping the LUN using the `snapdrive host disconnect` command.

```

snapdrive host disconnect -dg sdufs1_SdDg -full

disconnecting disk group sdufs1_SdDg
- fs /mnt/qa/sdufs1 ... disconnected
- hostvol sdufs1_SdDg/sdufs1_SdHv ...

disconnected
- dg sdufs1_SdDg ... disconnected

```

Connecting a storage entity from the host side

This section describes how to connect to a storage entity from the host side.

Connect the file system and host volume to the host using the `snapdrive host connect` command.

```

snapdrive host connect -fs /mnt/qa/sdufs1 -hostvol sdufs1_SdDg/sdufs1_SdHv -lun f270c-198-147:/vol/hacmpvol/sdulun1

discovering lun(s) ... done
LUN f270c-198-147:/vol/hacmpvol/sdulun1 connected
- device filename(s): /dev/hdisk4
Importing sdufs1_SdDg
Connected fs /mnt/qa/sdufs1

```

Increasing the storage entity size

This section describes how to increase the size of a storage entity.

Increase the size of the disk group by adding a LUN using the `snapdrive storage resize` command.

```

snapdrive storage resize -dg sdufs1_SdDg -growto 2g -
addlun
discovering filer LUNs in disk group sdufs1_SdDg...done
LUN f270c-198-147:/vol/hacmpvol/sdufs1_SdLun ...
created
mapping new lun(s) ... done
discovering new lun(s) ... done.
initializing LUN(s) and adding to disk group
sdufs1_SdDg...done
Disk group sdufs1_SdDg has been resized
Desired resize of host volumes or file systems
contained in disk group must be done manually

```

The snapdrive storage resize command lets you increase the size of the 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

Note: The storage resize operation can only increase the size of a disk group and not the size of a file system. You have to manually increase the file system size.

Deleting a storage entity

This section describes how to delete a storage entity.

Delete the disk group sdufs1_SdDg using the snapdrive storage delete command.

```

snapdrive storage delete -dg sdufs1_SdDg -full

deleting disk group sdufs1_SdDg
- fs /mnt/qa/sdufs1 ... deleted
- hostvol sdufs1_SdDg/sdufs1_SdHv ... deleted
- dg sdufs1_SdDg ... deleted
- LUN f270c-198-147:/vol/hacmpvol/sdulun1 ...
deleted
- LUN f270c-198-147:/vol/hacmpvol/sdufs1_SdLun
... deleted

```

The snapdrive storage delete command removes all the specified host side entities, their underlying entities, and the LUNs associated with them.

Snapshot copies and storage

The example in this section allows you to perform simplified versions of many tasks you might want to perform in a production environment. By going through these examples, you can quickly start using SnapDrive for UNIX to manage Snapshot copies and storage on your host system.

Note: You can execute the commands in the following steps. When performing an exercise such as this, it is recommended that you do it on a test system to ensure that you do not accidentally create a problem for a production system.

The steps that follow show you how to perform the following actions:

- Create a volume group and file system to be used as storage.
- Display information about the storage.
- Increase the size of the storage.
- Take a Snapshot copy of the volume group.
- Destroy the data in the volume group and then restore it using a Snapshot copy.
- Disconnect the LUNs from the host.
- Reconnect the Snapshot copy to the host.

Each step also includes the examples of additional SnapDrive for UNIX commands so that you can see other ways to use the command.

These steps assume you have experience setting up and running applications on your host operating system. If the steps or the sample command lines are not clear to you, look in the SnapDrive for UNIX documentation.

Note: To make the exercises in these steps easier to follow, they do not use all the options available for each SnapDrive for UNIX command. For a complete list of the SnapDrive for UNIX commands, their options, and detailed explanations of how they work, see the *SnapDrive for UNIX Installation and Administration Guide*.

Setting up a file system on a host entity

This section describes how to set up a file system on a host entity.

Create a volume group vg1 with three LUNs (lunA, lunB, and lunC) on a storage system toaster and a file system /db1/datafiles1. Assign each LUN a size of 100 MB.

Enter the following snapdrive storage create command:

```
# snapdrive storage create -vg vg1 -lun toaster:/vol/vol1/lunA lunB lunC -lunsize 100m -fs /db1/datafiles1
```

Now create a stand-alone LUN called lunD:

```
# snapdrive storage create -lun toaster:/vol/vol1/lunD -lunsize 100m
```

The following are snapdrive storage create command. These commands illustrate how you can specify different information when you are creating storage.

Example 1:

This command line provides the minimal information required for setting up a file system. It creates a file system in a 100 MB disk group by specifying only the file system mountpoint, the disk group size, and the storage system volume. Snapdrive for UNIX supplies the default values for the other parameters, including the volume group name, number of LUNs, LUN size, and file system type.

```
# snapdrive storage create -fs /mnt/myfs -dsize 100m -filervol toaster:/vol/vol1
```

This command line specifies a VxFS file system type:

```
# snapdrive storage create -fs /mnt/acctfs -fstype vxfs -filervol toaster:/vol/acct -dsize 1g
```

This command line specifies a shared storage system on a Solaris host:

```
# snapdrive storage create -fs /mnt/sfortesting -dsize 300m -filervol f270-197-109:/vol/vol2  
-devicetype shared
```

This command line specifies an AIX file system type:

```
# snapdrive storage create -fs /mnt/acctfs -fstype jfs2 -lun toaster:/vol/vol1/lunA lunB lunC  
-lunsize 1g
```

Displaying information about the storage

This section describes how to display information about the storage system.

Confirm that SnapDrive for UNIX created all the necessary parameters by executing the snapdrive storage show command.

```
# snapdrive storage show -vg vgl
```

Note: You can use either snapdrive storage show or snapdrive storage list in the command line. These commands are synonyms.

Examples of additional command lines:

The following is a list of sample command lines that you can use with the storage show command:

```
# snapdrive storage show -filer toaster1 toaster2 -verbose
# snapdrive storage list -dg dg1 dg2
# snapdrive storage show -fs /mnt/myfs /mnt/myfs2
# snapdrive storage show -hostvol dg2/vol1 dg4/myvol3
# snapdrive storage list -dg dg2 -fs /mnt/myfs
# snapdrive storage show -all -devicetype shared
```

Increasing the size of storage

This section describes how to increase the size of storage.

Use the `snapdrive storage resize` command with the `-growby` option to expand the size of the storage. The `-growby` option combined with the `-addlun` option tells SnapDrive for UNIX to increase the size of the volume group by adding a LUN that is say, 155 MB, to the volume group.

```
# snapdrive storage resize -vg vg1 -addlun -growby 155m
```

Example 1:

This command line includes the `-growto` option with the `-addlun` option to tell SnapDrive for UNIX to increase the size of the storage until it reaches a total of 155 MB. SnapDrive for UNIX adds a LUN to the volume group and bases its size on the difference between the current size of the storage and 155 MB.

```
# snapdrive storage resize -vg myvg -addlun -growby 155m
```

The next command line is for a Solaris host and uses the `-growto` option to tell SnapDrive for UNIX to increase the size of the storage until it reaches 200 MB. SnapDrive for UNIX performs the necessary calculations to determine how much the storage needs to grow by before it reaches 200 MB. It adds the new storage to the last LUN in the volume group.

```
# snapdrive storage resize -dg my_dg -growto 200m
```

Note: When you use this command, make sure that the total size of the volume or disk group does not exceed the operating system limits.

This command allows to increase the size of the disk group either by increasing the size of a LUN or by adding a new LUN in a cluster environment, on a Solaris host.

```
# snapdrive storage resize -dg shared -growby 100m -addlun
discovering filer LUNs in disk group shared...done
LUN f270-197-109:/vol/vol1/lunShared_SdLun ...
created
mapping new lun(s) ... done
discovering new lun(s) ... done.
Connecting cluster node: sfrac-58
mapping lun(s) ... done
discovering lun(s) ... done
LUN f270-197-109:/vol/vol1/lunShared_SdLun
connected
- device filename(s): /dev/vx/dmp/c3t0d4s2
initializing LUN(s) and adding to disk group
shared...done
Disk group shared has been resized
Desired resize of host volumes or file systems
contained in disk group must be done manually
```

Making a Snapshot copy

This section describes how to make a Snapshot copy.

1. Add three files. Make each file of size 1 MB and name the files file1, file2, and file3.

```
# cd /db1/datafiles1
# mkfile 1m file1 file2 file3
# cd ..
```

Note: After you have created the files, leave that directory (that is, it should not be your current directory). Certain SnapDrive for UNIX commands, such as the `snapdrive snap restore` command, do not work if you are in the directory you want to restore.

2. Create a Snapshot copy called snap1 using the `snapdrive snap create` command.

```
# snapdrive snap create -vg vg1 -snapname snap1
```

Examples of additional command lines:

The following `snapdrive snap create` command lines assume you have disk groups dg1 and dg2 with host volumes myvol1 and myvol2. The host volume dg1/myvol2 has a file system mounted on /myfs/fs2. The disk groups have three LUNs: toaster:/vol/vol1/lun0, toaster:/vol/vol1/lun1, and toaster:/vol/vol1/lun2. These command lines all create a Snapshot copy called toaster:/vol/vol1:snap1.

```
# snapdrive snap create -fs /myfs/fs2 -snapname snap1
# snapdrive snap create -dg dg1 dg2 -snapname snap1
# snapdrive snap create -vg dg1 -snapname snap1
# snapdrive snap create -hostvol dg1/myvol1 -snapname snap1 -force
# snapdrive snap create -hostvol dg1/myvol2 -fs /fs2 -snapname snap1
```

Restoring a Snapshot copy

This section describes how to restore a Snapshot copy.

1. Remove the files you earlier created by using the command `# rm -rf /db1/datafiles1/*`
2. Restore a Snapshot copy using the `snapdrive snap restore` command. Make sure you are not in the directory where you want to restore the files, or you will receive an error message.

```
# snapdrive snap restore -vg vg1 -snapname snap1
```

Attention: This procedure might take several minutes. Do not press Ctrl-C during this procedure. If you try to halt the procedure, you risk leaving the disk groups in an incomplete state.

Examples of additional command lines:

The following command lines assume you have a Snapshot copy that includes data from `vg1` and `vg2`. The first command line performs a partial restore. It restores the Snapshot copy data from `vg2` only. The second one restores data for both `vg2` and `vg1`.

```
# snapdrive snap restore -vg vg2 -snapname snap1
# snapdrive snap restore -vg vg1 vg2 -snapname snap1
```

Disconnecting the LUN mappings from the host

This section describes how to disconnect a LUN from a host.

Disconnect `lunD` from the host where you created it using the `snapdrive storage disconnect` command.

```
# snapdrive storage disconnect -lun toaster:/vol/vol1/lunD
```

Note: If you want to disconnect a host-side entity that contains other entities, you must include the `-full` option on the command line.

Examples of additional command lines:

These are sample storage disconnect command lines.

```
# snapdrive storage disconnect -lun toaster:/vol/vol1/lunA muffin:/vol/vol1/lunB
# snapdrive storage disconnect -dg dg1 dg2 dg3 -full
# snapdrive storage disconnect -hostvol dg1/mylvoll
# snapdrive storage disconnect -fs /mnt/shared -devicetype shared
```

Connecting an existing LUN to the host

This section describes how to connect an existing LUN to a host.

Reconnect a lunD to the host by using the snapdrive storage connect command.

```
# snapdrive storage connect -lun toaster:/vol/vol1/lunD
```

Note: You can use this command to connect LUNs not created with SnapDrive for UNIX as well as LUNs created with it.

Examples of additional command lines:

The following are examples of snapdrive storage connect command lines:

```
# snapdrive storage connect -lun filer1:/vol/vol1/ntap_lun2
# snapdrive storage connect -lun filer1:/vol/vol1/ntap_lun2 ntap_lun3 ntap_lun4
# snapdrive storage connect -fs /mnt/shared_fs - hostvol shared_SdDg/shared_SdHv -lun f270-197-109:/vol/vol1/shared_SdLun -devicetype shared
```

Deleting storage from the host

This section describes how to delete storage from a host.

Delete the volume group, LUNs, and file system using the snapdrive storage delete command. Because you are deleting the file system, you do not need to specify the volume group and LUN names. (Do not unmount the file system before you delete the storage, or SnapDrive for UNIX will not be able to find it.)

```
# snapdrive storage delete -fs /dbl/datafiles1
```

Note: If you want to delete a host-side entity that contains other entities, you must include the `-full` option on the command line.

Examples of additional command lines:

The following are examples of the snapdrive storage delete command lines.

```
# snapdrive storage delete -lun toaster:/vol/vol1/lunC
# snapdrive storage delete -vg vg1
# snapdrive storage delete -dg dg1 dg2 dg3
# snapdrive storage delete -hostvol mydg/vol3 mydg/vol5
# snapdrive storage delete -dg dg1 dg2 dg3 -full
# snapdrive storage delete -dg testdg -devicetype shared -full
```

Splitting the volume clone or LUN clone

SnapDrive for UNIX enables you to split a volume clone or LUN clone. After the clone split operation is complete, the parent volume and the parent LUN have independent clones that get detached from the parent volume and parent LUN, and have their own individual storage space.

SnapDrive for UNIX has the following clone split operations:

- Estimate the disk space (in MB) for a volume clone or LUN clone.
- Split a volume clone or LUN clone in a synchronous or in an asynchronous mode.

By default, the clone split is in asynchronous mode.

- Stop the clone split operation
- View the status of the clone split that is in progress, completed, or failed
- View the result of clone split that succeeds or failed using the job ID or file specification

Note: The LUN clone split operation provides support for Snapshot copies taken in SnapDrive 4.2 for UNIX and later.

Some SnapDrive for UNIX features in brief

Web services and daemon

SnapDrive for UNIX Web service provides a uniform interface for all the SnapManager and third party products to help them integrate seamlessly with SnapDrive for UNIX.

Various SnapManager products communicate with SnapDrive for UNIX using CLI. It puts a constraint on SnapManager and SnapDrive for UNIX in terms of performance and manageability.

Using SnapDrive for UNIX daemon, all the commands will work as a unique process in background. Daemon service does not affect the way SnapDrive for UNIX commands are used. SnapDrive for UNIX Web service allows third party applications to integrate with SnapDrive for UNIX seamlessly. They interact with SnapDrive for UNIX using APIs. For any SnapDrive command to work, you need to make sure that daemon is running.

To start the daemon, enter the `snapdrived start` command in the CLI.

Note: Only root user can start the daemon service. If a user other than a root user tries to start the daemon, SnapDrive for UNIX gives an error message.

SnapDrive for UNIX configuration checker

SnapDrive 4.2 for UNIX configuration checker tool is bundled with SnapDrive for UNIX. The configuration checker tool helps you to verify the various configurations that is required for proper working of SnapDrive for UNIX. This tool helps in transitioning the task of checking the configuration from the end user to the person who support and maintain SnapDrive for UNIX.

Automatic detection of host entities

SnapDrive 4.1 for UNIX, for some commands, automatically detects the type of host entity. Purpose of automatic detection is to minimize the number of inputs to a command by the user. SnapDrive for UNIX automatically detects the type of the host entity if you provide the name without the type specification.

The following commands are enabled for automatic detection:

- storage delete

- storage resize
- snap create
- snap restore
- snap connect
- snap list

SnapDrive for UNIX can automatically detect LUNs, disk groups, file systems, and host volumes. You can perform automatic detection on a file_spec that already exists on the hosts. SnapDrive 4.0 for UNIX and later does not support automatic detection of Volume Manager.

Following is one example of automatic detection syntax and usage:

```
snapdrive storage delete /mnt/test_fs_snapdrive
Auto detection of file_spec(s) in progress ...
- '/mnt/test_fs_snapdrive' detected as file system.
delete file system /mnt/test_fs_snapdrive
- fs /mnt/test_fs_snapdrive ... deleted
- hostvol test_fs_snapdrive_SdDg/test_fs_snapdrive_SdHv ... deleted
- dg test_fs_snapdrive_SdDg ... deleted
- LUN bart:/vol/voldm/test_fs_snapdrive_SdLun ... deleted
- LUN bart:/vol/voldm/test_fs_snapdrive-1_SdLun ... deleted
```

FlexClone volumes in SnapDrive for UNIX

A flexible volume clone, FlexClone volume, is a point-in-time copy of a parent flexible volume. Changes made to the parent volume after the clone is created are not inherited by the clone.

FlexClone volumes provide substantial space savings with minimal overhead. Using this you can manage many more data set variations in less time and with less risk. FlexClone volumes do not have write reservation by default. SnapDrive for UNIX uses FlexClone volumes in NFS and SAN environments.

Following are few operations which you perform on FlexClone volumes:

- Snap disconnect operation
- Reuse of FlexClone volumes
- Splitting the FlexClone volumes
- Space reservation on FlexClone volumes
- Splitting a volume clone or LUN clone

Volume-based SnapRestore

SnapDrive for UNIX 4.0 and later provides Snapshot capability at a volume level. Volume-based SnapRestore restores the volume with all its storage objects. The volume-based restore is faster than each storage object restored individually.

Volume-based SnapRestore is also supported for FlexClones and vFiler unit configurations. Volume-based SnapRestore for vFiler units is available for Data ONTAP 7.3 and later.



NA 210-04890_A0, Printed in USA

GC26-7879-07

