

SnapDrive® 3.0.1 for UNIX® **Release Notes (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 215-03510_A0
October 2007

Copyright and trademark information

Copyright information

Copyright © 1994–2007 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, Data ONTAP, FAServer, FilerView, FlexVol, Manage ONTAP, MultiStore, NearStore, NetCache, SecureShare, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, Spinnaker Networks, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, SyncMirror, Topio, VFM, and WAFL are registered trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries. Cryptainer, Cryptoshred, Datafort, and Decru are registered trademarks, and Lifetime Key Management and OpenKey are trademarks, of Decru, a Network Appliance, Inc. company, in the U.S.A. and/or other countries. gFiler, Network Appliance, SnapCopy, Snapshot, and The evolution of storage are trademarks of Network Appliance, Inc. in the U.S.A. and/or other countries and registered trademarks in some other countries. ApplianceWatch, BareMetal, Camera-to-Viewer, ComplianceClock, ComplianceJournal, ContentDirector, ContentFabric, EdgeFiler, FlexClone, FlexShare, FPolicy, HyperSAN, InfoFabric, LockVault, NOW, NOW NetApp on the Web, ONTAPI, RAID-DP, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, SharedStorage, Simplicore, Simulate ONTAP, Smart SAN, SnapCache, SnapDirector, SnapFilter, SnapMigrator, SnapSuite, SohoFiler, SpinMirror, SpinRestore, SpinShot, SpinStor,

StoreVault, vFiler, 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 U.S.A.

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.

Contents of Release Notes

Topics covered

The following topics are covered in this Release Notes:

- ◆ [“Overview”](#) on page 2
- ◆ [“Limitations”](#) on page 5
- ◆ [“System requirements”](#) on page 6
- ◆ [“Finding information on known issues”](#) on page 7
- ◆ [“Known issues and troubleshooting tips affecting all hosts”](#) on page 10
- ◆ [“Solaris-specific issues and troubleshooting”](#) on page 24
- ◆ [“Obtaining the latest version of SnapDrive for UNIX software”](#) on page 28
- ◆ [“Obtaining other NetApp documentation”](#) on page 29

Overview

Overview

SnapDrive® for UNIX® is a Network Appliance™ application that you can use to back up and restore host data stored on NetApp® storage systems. You can also use SnapDrive for UNIX to manage the Snapshot™ copies that you create with it and to create and manage storage. SnapDrive for UNIX interacts with the host operating system (OS) and volume manager when it performs these tasks.

About this release

The SnapDrive 3.0.1 for UNIX release is specific to Solaris™ platform.

This release does not include the following:

- ◆ Guides

- ❖ *Installation and Administration Guide*

- ❖ *Quick Start Guide*

For the Installation and Administration Guide and the Quick Start Guide, see the SnapDrive 3.0 for UNIX documentation at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

- ◆ Media Kit

You will need to download the SnapDrive 3.0.1 for UNIX software from the NOW™ NetApp on the Web™ at <http://now.netapp.com>.

Bug fixes

The following bug is fixed in the SnapDrive 3.0.1 for UNIX release:

Bug ID	Description
260011	SnapDrive 3.0 for Solaris does not work with FCP without Veritas stack.

With SnapDrive for UNIX 3.0 on Solaris over FCP, the `snapdrive storage create` command gives the following error:

```
ADMINERROR:Could not find appropriate multipathing type, When UFS is not present.
```

When UFS is present the `snapdrive storage create` command gives the following error:

ADMINERROR: Transport Type FCP and File System UFS cannot be used together

SnapDrive for UNIX 3.0.1 solves the above problem. It allows SnapDrive for Unix 3.0.1 to work on non Veritas environments and use RAW disks for Solaris.

For more information on supported configuration, see the SnapDrive for UNIX *Compatibility Matrix* at

<http://now.netapp.com/NOW/knowledge/docs/docs.shtml>

Environments not affected: Following are the environments which are not affected by the bug fix:

- ◆ SnapDrive for UNIX 3.0 used in environments where Veritas software is installed
- ◆ SnapDrive for UNIX 3.0 used in iSCSI environments

Features

SnapDrive 3.0.1 for UNIX also includes all the features on the Solaris platform of SnapDrive 3.0 for UNIX:

- ◆ vFiler support on all host platforms

Note

The preceding feature is available on storage systems with Data ONTAP® 7.2.2.

- ◆ Thin provisioning of LUNs

Configuration variables:

- ◆ Changed configuration variables: The default value of the following configuration variables is changed in the snapdrive.conf file:
 - ❖ use-https-to-filer=on
 - ❖ enable-implicit-host-preparation=on
- ◆ New configuration variables: The following new configuration variables are added to the snapdrive.conf file:
 - ❖ enable-split-clone="off"
To enable to split the cloned volumes or LUNs during Snapshot connect and Snapshot disconnect operations
 - ❖ fstype="vxfs"
To specify the default file system type you want to use for SnapDrive for UNIX operations
 - ❖ space-reservations-enabled=on
To enable space reservation when creating LUNs

- ❖ `vmtype="vxvm"`
To specify the default volume manager type you want to use for SnapDrive for UNIX operations
- ◆ **Command-line options:** The following new command-line options are added:
 - ❖ `-fstype`
The type of file system you want to use for the SnapDrive for UNIX operations
 - ❖ `-vmtype`
The type of volume manager you want to use for the SnapDrive for UNIX operations
 - ❖ `-split`
Enables to split the cloned volumes or LUNs during Snapshot connect and Snapshot disconnect operations
 - ❖ `-status`
Used with the `snapdrive storage show` command to know if the volume or logical unit number (LUN) is cloned

Current host platforms for SnapDrive for UNIX

The most current list of host platforms that SnapDrive for UNIX supports is in the online SnapDrive for UNIX *Compatibility Matrix*. This matrix is on the NOW™ NetApp on the Web™ site at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml>.

SnapDrive for UNIX is supported on the following host platforms:

- ◆ Solaris

For more information about the differences between the SnapDrive for UNIX hosts, see the overview chapter in the SnapDrive for UNIX *Installation and Administration Guide*. All platforms use the same manual.

Limitations

The SnapDrive 3.0.1for UNIX software does not support the following:

FCP/iSCSI configuration

FCP and iSCSI configurations are not supported on the same host.

Thin provisioning

The following Data ONTAP features are not supported:

- ◆ Fractional reserve
- ◆ Space reclamation (Hole punching)
- ◆ Snapshot reserve
- ◆ Space monitoring command-line interface (CLI)
- ◆ LUN fill on demand
- ◆ Volume fill on demand
- ◆ Volume autogrow/Snapshot autodelete

vFiler

Following are the vFiler limitations:

- ◆ SnapDrive for UNIX does not support access to vFiler units through FCP, because this process is not supported by Data ONTAP.
- ◆ SnapDrive for UNIX provides storage provisioning operations, Snapshot operations, host-side operations, and configuration operations on a vFiler unit only if the vFiler is created on a FlexVol volume. These operations are not supported on a vFiler unit that is created on a qtree because Snapshot operations are disallowed unless the vFiler unit owns the entire storage volume.
- ◆ Snapshot operations are not supported on a vFiler unit if the root of the vFiler unit is a qtree.

System requirements

System requirements

For FCP and iSCSI configurations, SnapDrive for UNIX requires that you install the FCP or iSCSI Host Utility for your host system. Your host system must meet the requirements specified by that utility, regardless of whether you are using an FCP or iSCSI Host Utility.

Note

In SnapDrive 3.0.1 for UNIX, the products FCP Host Attach Kit and iSCSI Support Kit are named as FCP Host Utilities and iSCSI Host Utilities.

Because NetApp adds new kits and components on an ongoing basis, you need to check the following on-line matrices to determine the system requirements for SnapDrive for UNIX on your host platform. These matrices contain the most up-to-date information for NetApp products in a SAN environment.

- ◆ The *SnapDrive for UNIX Compatibility Matrix* is at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml/>. It provides a high-level list of the SnapDrive for UNIX requirements for each host platform. If a feature (for example, a file system) is not listed in the matrix, then SnapDrive for UNIX does not support it on that host platform, even if the utility does. To avoid duplicating detailed information, this matrix refers you to the *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* for specific information on requirements, such as any patches that are required for the operating system.
- ◆ The *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* guide is at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/. It provides a detailed list of the requirements for setting up a NetApp SAN environment. It provides specific information on components, such as the patch levels required for the operating systems. You need to check the information in this matrix to determine the details on the requirements listed in the *SnapDrive for UNIX Compatibility Matrix*.

Finding information on known issues

Finding information on known issues for SnapDrive for UNIX

At the time of this release, there were some known problems, limitations for SnapDrive for UNIX. While some issues affect all SnapDrive for UNIX host platforms, others affect only a specific host platform. To locate information on known issues and troubleshooting tips, first check the following section:

- ◆ “[Known issues and troubleshooting tips affecting all hosts](#)” on page 10

Then check the section for your host platform:

- ◆ “[Solaris-specific issues and troubleshooting](#)” on page 24
- ◆ The “Troubleshooting” chapter of the SnapDrive for UNIX *Installation and Administration Guide*

Finding information on public bugs

In some cases, problems and limitations have a public bug ID assigned to them. Public burts are not included in these release notes. For information on checking the status of these issues, see “[Accessing information about bug fixes](#)” on page 7. You can also use Bugs Online to determine if any new issues have been found since this document was published.

Note

The term *disk group* in this documentation refers to either a disk or volume group.

Accessing information about bug fixes

Some issues with SnapDrive for UNIX are listed in an online system called Bugs Online. For a list of fixed bugs, access Bugs Online by completing the following steps.

Step	Action
1	Use a Web browser to go to http://now.netapp.com .
2	Log into NOW and click Service & Support.
3	Click the link for Bugs Online & Release Tools.

Step	Action	
4	If you want to check by using the...	Then...
	Quick Search	a. Enter either the bug ID in the ID box or the phrase “SnapDrive for UNIX” in the Key Word Search box. Note — You can also include an operating system keyword with “SnapDrive for UNIX.” If you do, the search finds all information related to the operating system, not just the SnapDrive for UNIX information. b. Click the Go button next to the text box for your search.
	Advanced Search: Bug Type	a. For results specific to SnapDrive for UNIX, select the following two types in the Bug Types box: ❖ SnapDrive - UNIX This choice displays the bugs that pertain to all SnapDrive for UNIX hosts. ❖ SnapDrive - <i>host-operating-system</i>; for example SnapDrive-AIX This choice displays all the SnapDrive for UNIX bugs that are specific to that host operating system.

Step	Action
	You should select both bug types to ensure that you see all the SnapDrive for UNIX bugs that apply to your host. Bugs that are common to all platforms do not specify an operating system type. To select multiple types, highlight the first type. Then press the CTRL key as you highlight subsequent types. To select a range of types, highlight the first one and then press the Shift key when you highlight the next one. b. Click the Go button in the Advanced Search section.
	Advanced Search: Product Type a. For results about all platforms of SnapDrive, select SnapDrive in the Product Types box. b. Click the Go button in the Advanced Search section. Bugs Online displays the bugs for all platforms of SnapDrive, not just the SnapDrive for UNIX platforms. c. Check the Product Type category of each bug displayed to see whether the bug applies to SnapDrive for UNIX.

Known issues and troubleshooting tips affecting all hosts

Access control permissions file not working as intended

Symptom: The local access control mechanism does not work as expected. This happens when the storage system does not have a vol0.

Explanation: By default, SnapDrive for UNIX expects to find the permissions file for local access control at the pathname `/vol/vol0/sdprbac/sd<hostname>.prbac`, where `<hostname>` is the name of the local host. This does not work if your storage system does not contain a vol0.

What to do: You can use the parameter `rbac-vol-name` in the `snapdrive.conf` file to specify a different volume name where SnapDrive for UNIX looks for the permissions file. Specify the name of the volume with no `/vol` suffix, as in:

```
rbac-vol-name="vol1"
```

SnapDrive for UNIX then uses that volume name for all the storage systems accessed from the local host.

Restoring the storage stack on the host after destroying a LUN

You should always use SnapDrive storage commands to remove LUNs. Using commands outside of SnapDrive to remove LUNs can have an unpredictable effect on the host system.

If you remove a LUN without using a SnapDrive storage command, you should reboot the host to restore the storage system stack to a known state. You can also restore the host system state by entering the `snapdrive snap restore` command and specifying a Snapshot copy that contains an instance of the destroyed LUN.

Restoring qtrees in NFS-mounted file systems on SnapMirror

SnapDrive for UNIX does not allow you to restore a Snapshot copy of qtree that is Network File System (NFS)-mounted on a SnapMirror destination. SnapDrive for UNIX does allow you to restore LUNs on Qtrees on a SnapMirror destination.

Restoring NFS directories with symbolic links

SnapDrive for UNIX fails if you use it to restore an NFS directory that contains symbolic links.

SnapDrive for Unix produces an error 18 in trace log (225765)

Symptom: The host from which SnapDrive for UNIX is executed will not be able to access the filer. The following error message appears:

SnapDrive for UNIX cannot contact the storage system needed for this command. Check the connectivity to the storage system indicated in the error message.

Explanation: If the filer option `httpd.admin.access` is set to "legacy" value, then the host from which SnapDrive for UNIX is executed has to be entered in the `trusted.hosts` list, for the host to access the filer.

What to do: Add all the hosts from which SnapDrive for UNIX is executed to the `trusted.hosts` list. This enables SnapDrive for UNIX to access the filer.

Example: In the following example, SnapDrive for UNIX is executed from the ruby and 10.88.88.191 hosts, and each of these hosts require access to filer X:

```
options trusted.hosts 10.88.88.191, ruby
```

Note

You can set `trusted.hosts` or `httpd.admin.access` to * value to give HTTP access to all the hosts in the network.

Restore fails with message that file system is mounted

Symptom: You unmount the file system, but when you attempt to restore the Snapshot copy the operation fails and says the file system is mounted.

Explanation: With a Linux host, you can mount the same file system more than once. You can then take a Snapshot copy using SnapDrive for UNIX. The Snapshot copy, though, contains only the first entry in the mount table for the file system. When you issue an unmount command, it unmounts the most recent mountpoint, not the one shown in the Snapshot copy. Then, when you attempt to restore the Snapshot copy, it fails with a message that the volume is already mounted.

The following example illustrates what can happen if a file system is mounted more than once:

```
# mount /dev/dg3/vol1 /var/tmp/mnts/dg3/vol1
# mount /dev/dg3/vol4 /mnt/uml
# mount /dev/dg3/vol4 /var/tmp/mnts/dg3/vol4
# mount |grep dg3
/dev/dg3/vol1 on /var/tmp/mnts/dg3/vol1 type ext3 (rw)
/dev/dg3/vol4 on /mnt/uml type ext3 (rw)
/dev/dg3/vol4 on /var/tmp/mnts/dg3/vol4 type ext3 (rw)
# snapdrive snap create -dg dg3 -snapname dg3_vol2_snapshot
snap create: snapshot dg3_vol2_snapshot contains:
    disk group dg3 containing host volumes
        vol1 (filesystem: /var/tmp/mnts/dg3/vol1)
        vol2
        vol3
        vol4 (filesystem: /mnt/uml)
snap create: created snapshot eccentric:/vol/vol1:dg3_vol2_snapshot
# snapdrive snap show -v -snapname eccentric:/vol/vol1:dg3_vol2_sn
```

snap name	host	date	snapped

eccentric:/vol/vol1:dg3_vol2_snapshot	childs	Sep 1 14:14	dg3
host OS: Linux 2.4.21-9.ELsmp #1 SMP Thu Jan 8 17:08:56 EST 2004			
snapshot name: dg3_vol2_snapshot			
Volume Manager: linuxlvm 1.0.3			
disk group: dg3			
host volume: vol1			
host volume: vol2			
host volume: vol3			
host volume: vol4			
file system: vol1	file system type: ext3	mountpoint:	
/var/tmp/mnts/dg3/vol1			
file system: vol4	file system type: ext3	mountpoint: /mnt/uml	
lun path dev paths			

```
eccentric:/vol/vol1/swl14 /dev/sds
# snapdrive snap restore -dg dg3 -snapname dg3_vol2_snapshot
Starting to restore LUNs in disk group dg3
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 dg3
Error (Command error 0001-034): mount failed: mount: /dev/dg3/vol4 already mounted or
/mnt/uml busy
mount: according to mtab, /dev/dg3/vol4 is already mounted on /mnt/uml
```

What to do: Make sure that the file system is not mounted more than once.

Failure of Snapshot create and delete operations (210958)

Symptom: The Snapshot create and delete operations fails to remove Snapshot copies on the filer because the LUNs are busy. The following error message appears:

```
0001-124 Admin error: Failed to remove snapshot <snap-name> on  
filer <filer-name>: LUN clone
```

Explanation: You might encounter this error and not be able to directly delete a Snapshot copy if the following are true:

- ◆ A disk group spans multiple storage systems.
- ◆ The following SnapDrive commands are executed in a loop for more than 50 iterations (the number of iterations depends on the host operating system):
 - ❖ `snapdrive storage create`
 - ❖ `snapdrive snap create`
 - ❖ `snapdrive storage delete`
 - ❖ `snapdrive snap restore`
 - ❖ `snapdrive snap connect`
 - ❖ `snapdrive storage delete`
 - ❖ `snapdrive snap delete`

At this point, the Snapshot copy status on the filer shows `LUNs-Busy`, and this status will be prolonged for a few hours.

What to do: Do one of the following.

If...	Then...
There are cloned LUNs in a Snapshot copy which are not connected to a host	Do one of the following to delete the Snapshot copy: ◆ Split the cloned LUNs from the filer by using the <code>filer split clone</code> command. ◆ Delete the cloned LUNs.

If...	Then...
There are cloned LUNs in a Snapshot copy which are connected to a host	Complete the following steps: a. Disconnect the cloned LUNs from the host by using the <code>snapdrive snap disconnect</code> command. b. Delete the cloned LUNs. c. Delete the Snapshot copy.

If...	Then...
There are any of the following: ◆ There are cloned LUNs in a Snapshot copy which are also present in other Snapshot copies ◆ There are back up cloned LUNs in a Snapshot copy	Complete the following steps: a. Run the following command at the storage system command prompt: <pre>lun snap usage -s <volume-name><snap-name></pre> This will list the name of the Snapshot copies that contains the Snapshot copies of the cloned LUNs. For example, <pre>tonic*> lun snap usage vol1 dhilip_lun (filer view) mpiotest2:</pre> LUN: /vol/vol1/.snapshot/mpiotest2/dhilip_0 Backed By: /vol/vol1/.snapshot/dhilip_lun/dhilip b. Delete the Snapshot copies from the storage system in the same order as they are listed in the <code>lun snap usage</code> command output. For example, <pre>snap delete vol1 mpiotest2</pre> Note _____ Ensure that the Snapshot copy you are deleting is not being used by any other host. If it does, you will be able to delete it. _____
There are no cloned LUNs	Wait until the <code>LUNs-Busy</code> status of the Snapshot copy on the storage system disappears.

Attempts to create a Snapshot copy fail

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

- ◆ 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 through SCSI commands.
- ◆ The LUNs in the disk group must be online and mapped to the disk group before you take the Snapshot copy. 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 copy.
- ◆ The access permissions do not allow the host to create a Snapshot copy of information on that storage system. For more information on access permissions, see section on setting up access control in the *SnapDrive for UNIX Installation and Administration Guide*.
- ◆ 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 copy fail

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

- ◆ 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 must allow the host to create and restore a Snapshot copy of information on that storage system. For more information on access permissions, see section on setting up access control in the *SnapDrive for UNIX Installation and Administration Guide*.
- ◆ The NFS mounted directory must be exported correctly to the host, so that it can be mounted. For additional information, see section on NFS considerations in the *SnapDrive for UNIX Installation and Administration Guide*.

Caution

You should always take another Snapshot copy 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 copy and then try to restore the Snapshot copy, 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 copy are no longer accessible when you create and restore the Snapshot copy.

- ◆ If you rename a host volume or a file system or change the mountpoint after taking a Snapshot copy 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 do not accidentally overwrite something that you did not intend to overwrite.

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

Caution

After you start a Snapshot 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.

How SnapDrive commands mount file systems and modify system files

Certain SnapDrive for UNIX commands can cause file systems to be mounted or unmounted. When a command performs an operation that mounts a file system, SnapDrive for UNIX adds the file system's name to the standard system file for your host platform. If an operation unmounts a file system, SnapDrive for UNIX removes the name from the system file. This should not pose a problem; however, it is useful to know when system files are modified.

The name of the system file varies depending on your host platform. The following table lists the host platforms and their system files.

Host	System file
Solaris	/etc/vfstab

Any of the following commands can mount file systems:

- ◆ `snapdrive snap restore`
- ◆ `snapdrive snap connect`
- ◆ `snapdrive storage create`
- ◆ `snapdrive host connect`
- ◆ `snapdrive storage connect`

The following commands perform operations that can unmount a file system:

- ◆ `snapdrive snap restore`
- ◆ `snapdrive snap disconnect`
- ◆ `snapdrive storage disconnect`
- ◆ `snapdrive storage delete`
- ◆ `snapdrive host disconnect`

Some commands are delayed when the storage cluster is in failover mode

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

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

This problem affects `snapdrive 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 copy.

Explanation: The restore operation replaces the data on a disk based on the way it was when you took the Snapshot copy. If you take a Snapshot copy and then change the way data is distributed among disks in the disk group, the restored Snapshot copy 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: Do one of the following:

- ◆ If you change the data distribution on the disks in the disk group, take a new Snapshot copy.
- ◆ You can use the parameter `rbac-vol-name` in the `snapdrive.conf` file to specify a different volume name where SnapDrive for UNIX looks for the permissions file. Specify the name of the volume with no `/vol` suffix, as in:
`rbac-vol-name="vol1"`

SnapDrive for UNIX then uses that volume name for all the storage systems accessed from the local host.

Commands appear blocked (116277)

Symptom: The SnapDrive for UNIX commands are blocked for as much as 10 minutes.

Explanation: This behavior occurs in rare cases when VxVM/DMP needs time to manage the volume configuration daemon. Triggers for this behavior include the following:

- ◆ A storage system that has just recovered from a panic
- ◆ A storage system target adapter that has transition from offline to online

What to do: Wait for the command to complete. The Veritas configuration daemon needs several minutes to run and discover the current state of the devices.

LUN is online, but not seen (116590)

Symptom: SnapDrive for UNIX reports that a LUN is offline even though it is online.

Explanation: The LUN was taken offline and then put back online. Veritas is not aware of the LUN's transition back to an online state until `vxconfig` is run.

What to do: Run `vxconfig`. After `vxconfig` has been run, SnapDrive for UNIX works as it did before.

Snap connect operation fails if #SDUTEMP tag is present in the lvm.conf file

Symptom: The Snapshot connect operation fails and the original volume group will not be loaded after rebooting the host.

Explanation: On a Linux host, the Snapshot connect operation fails on an originating host if SnapDrive for UNIX is interrupted, either by rebooting the host or by using Ctrl-C. The interruption results in the `#SDUTEMP` tag being added to the `/etc/lvm/lvm.conf` file; any subsequent SnapDrive for UNIX operations would fail with the following error:

0001-945 Admin error: Volume Manager's configuration file contains tags from previous SDU operation. Please cleanup the cloned lun(s), if any, existing from a previous unsuccessful operation and cleanup the configuration file before proceeding.

What to do: Complete the following steps.

Step	Action
1	Remove any cloned LUNs from the storage system that might exist from a previously failed or interrupted Snapshot connect operation. Note_____ You can use the <code>snapdrive storage show</code> command with the <code>-all</code> option to verify. _____
2	Check if the LUN has been removed from the host by entering the following command: ◆ For iSCSI, enter the following command: <code>/etc/init.d/iscsi reload</code> ◆ For FCP, enter the following command: <code>snapdrive test os linux rescanLuns</code>
3	Remove the volume group from the <code>/etc/lvm/lvm.conf</code> file. After removing the cloned LUNs from the <code>/etc/lvm.conf</code> file, the file might look as follows: <pre>#SDUEXISTING volume_list = ["vgX" , "vgY"] volume_list = ["vgX"] #SDUTEMP</pre>
4	Delete the entire line containing the <code>#SDUTEMP</code> tag at the end. If a line exists that begins with the <code>#SDUEXISTING</code> tag, remove the tag only, and retain the line.
5	Run the <code>vgscan</code> command to retain the changes made in the configuration file.
6	Enter the <code>vgchange -a y</code> command to bring back the original volume group.

Step	Action
7	Execute the appropriate mount command to mount the logical volume group after reviewing the /etc/fstab file in the host to remount the file system.

Directory access and ownership rights change during the Snapshot restore operation

Symptom: NFS directory access and ownership rights change during the Snapshot restore operation.

Explanation: For non-existing NFS directories on Data ONTAP 7.0 and 7.1, and for existing and non-existing NFS directories on Data ONTAP 7.2, the directory access and permission rights change from user to root user.

What to do: To change the ownership right, enter the following command at the command prompt:

```
chown <username> <directory name>
```

To change the access right, enter the following command at the command prompt:

```
chmod <permission> <directory name>
```

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 “HBA assistant not found”. This happens after one of the following situations occurs:

- ◆ An iSCSI stop was issued.
- ◆ A takeover has happened and the iSCSI service has not started yet.
- ◆ A giveback is happening and the iSCSI service has not 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 execute the `snapdrive storage show` command when the iSCSI service is not running, you will see output similar to the following:

```
# snapdrive storage show -dg toaster1
0001-877 Admin error: HBA assistant not found. Commands involving
LUNs should fail.
```

What to do: Make sure the iSCSI service is running. If the storage system is in the process of a takeover or giveback, wait until it starts the iSCSI service. If the storage system is in normal running mode, restart the iSCSI service. After 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          dgtype lvm
hostvol: /dev/mapper/toaster1-lvol1    state: AVAIL
hostvol: /dev/mapper/toaster1-lvol2    state: AVAIL
fs: /dev/mapper/toaster1-lvol1  mount point: /mnt/uml (non-
persistent) fstype ext3
fs: /dev/mapper/toaster1-lvol2  mount point: NOT MOUNTED

device filename adapter path    size    proto    state    clone    lun
path backing snapshot
-----
---
/dev/sde                -      P      64m     iscsi    online   No
gin:/vol/vol1/Lun1
/dev/sdf                -      P      64m     iscsi    online   No
gin:/vol/vol1/Lun2
```

**DMP configuration:
snapdrive storage
delete -lun
command hangs
(247942)**

Symptom: On hosts with FCP and DMP configurations, the `snapdrive storage delete -lun` command hangs for more than 10 minutes.

Explanation: This happens due to wrong, improper installation or configuration of the following components:

- ◆ NTAPasI
- ◆ Veritas stack (storage foundation)
- ◆ Multipathing licenses

What to do: Refresh the DMP paths information properly after any FCP path is enabled, disabled, or added. To refresh the DMP paths, execute the following commands in the same sequence as they are mentioned below.

Host	Commands
Solaris	1. <code>devfsadm -Cv</code> 2. <code>vxdisk scandisk</code>

Discrepancy in the error code message (247976)

If you execute the `snapdrive storage delete` command when both the FCP and iSCSI services are up or when they are down, SnapDrive for UNIX displays a wrong error message. The following table gives more details.

Transport protocol	Message displayed (The incorrect message)	Message that should be displayed instead (The correct message)
FCP and iSCSI both are configured	0002-181 Command error: Could not select appropriate File System. Please check the value of <code>fstype</code> in config file, and ensure proper file system is configured in the system.	0002-143 Admin error: Coexistence of <code>linuxiscsi</code> <code>linuxfc</code> drivers is not supported. Please ensure only one of the drivers is loaded in the host and retry.
FCP and iSCSI both are not configured	0001-465 Command error: The following filespecs do not exist and cannot be deleted:	0001-877 Admin error: HBA assistant not found. Commands involving LUNs should fail.

Solaris-specific issues and troubleshooting

Solaris platform limits concurrent LUN paths to 128

Symptom: An unexpected problem occurs on Solaris hosts when you attempt to use more than 128 concurrent paths to LUNs.

Explanation: SnapDrive for UNIX does not place limits on the number of concurrent paths to the LUNs; however, in some cases, a component of the software stack below it does.

What to do: Current testing indicates that you can avoid such unexpected problems by reducing the number of concurrent paths you use to be fewer than 128.

Storage connect operation fails to connect to a LUN created using filer command

Symptom: The `snapdrive storage connect` command with `-lun` option gives the following error:

```
0001-638 Admin error: Unable to identify device nodes corresponding to the following LUN(s). The LUN(s) on the filer may be offline or unmapped or may not be completely initialized on the host: for lun create manually on the filer.
```

Explanation: The storage connect fails to connect to a LUN if it is created on a filer by using the `lun create` command or by using FilerView. This error appears because the LUN is not formatted and Solaris does not recognize it, but the LUN gets discovered and mapped to the host.

What to do: Do not create LUNs manually. Use the `snapdrive storage create` command with the `-lun` option to create LUNs. If you have to use an existing LUN after executing the `snapdrive storage connect` command, complete the following steps.

Step	Action
1	Format the LUN manually.
2	Create a file system on it using the native volume manager and operating system commands. After creating a file system, you can use SnapDrive for UNIX operations for this LUN.

SnapDrive for UNIX does not show the supported file system types for storage create operation

Symptom: On a Solaris 10 update 1 host, the snapdrive storage create operation fails to create a UFS file system over raw LUNs, using the -fstype option and the iSCSI transport. The following error message appears:

0001-539 Command error: Snapdrive cannot create a file system of type ufs Supported types for storage create are:

Explanation: The snapdrive storage create operation fails to create a UFS file system over raw LUNs (LUNs without any volume manager components), using the -fstype option and the iSCSI transport, if the following are true:

- ◆ UFS file system is present
- ◆ VxVM components (VxVM, VxFS and DMP) are not installed
- ◆ HBA driver for Emulex is loaded

What to do: Unload the Emulex driver (lpfc) on the host using the modunload command.

Creating a file system directly on a LUN fails (244689)

Symptom: On Solaris hosts, creating a UFS file system directly on a LUN fails.

Explanation: If both UFS and Veritas stack are installed on the same Solaris host, and after setting the default-transport configuration variable as "iscsi", SnapDrive for UNIX does not pick UFS as the default file system for creating a file system directly on a LUN. The operation fails with the following error message:

0002-138 Admin error: Transport Type: iscsi Volume Manager: VxVM is not supported combination.

What to do: Use the "ufs" value for the -fstype command to create a UFS file system directly on a LUN.

After rebooting the host, OS and cluster services fail to restart (249086)

Symptom: On a Solaris 10 host, if the file system fails to mount after reboot, OS and cluster services will not start.

Explanation: After rebooting the Solaris 10 host, if the service svc:/system/filesystem/local fails to mount any of the following file system, the dependent OS or cluster services will not start.

- ◆ Local file system
- ◆ Shared file system

Note

The preceding file system may have been created by SnapDrive for UNIX.

What to do: Complete the following steps.

Step	Action
1	Find the service which failed to start after the reboot by using the following command: <pre>svcs -vx /network/inetd</pre> Example: <pre># svcs -vx /network/inetd svc:/network/inetd:default (inetd) State: offline since Fri 25 May 2007 04:06:59 PM IST Reason: Service svc:/system/filesystem/local:default is not running because a method failed.</pre>
2	If the service /system/filesystem/local:default is failed, find the persistent file system which failed to mount by using the following command: <pre>tail /var/svc/log/system-filesystem-local:default.log</pre> Example: <pre># tail /var/svc/log/system-filesystem-local:default.log [May 25 15:58:35 Method "start" exited with status 95] [May 25 16:07:17 Executing start method ("/lib/svc/method/fs-local")] UX:vxfs fsck: ERROR: V-3-20693: cannot access device /dev/vx/rdisk/dg2/fs1_SdHv: UX:vxfs mount: ERROR: V-3-20003: Cannot open /dev/vx/dsk/dg2/fs1_SdHv: No such device or address UX:vxfs mount: ERROR: V-3-24996: Unable to get disk layout version WARNING: /sbin/mountall -l failed: exit status 1 [May 25 16:07:18 Method "start" exited with status 95]</pre>
3	Comment the entries in /etc/vfstab for the file system which failed to mount after rebooting the host.

Step	Action
4	Restart the OS service: <code>svc:/system/filesystem/local:default</code> by executing the following commands: a. <code>svcadm disable</code> <code>svc:/system/filesystem/local:default</code> b. <code>svcadm enable</code> <code>svc:/system/filesystem/local:default</code> Result: The OS and cluster services will start.
5	For local file system that is failed to mount, complete the following steps: a. Manually mount the file system. b. Uncomment the <code>/etc/vfstab</code> entry.
6	For cluster services, complete the following steps: a. Use the <code>cfsmount</code> command to check if the shared file systems are properly mounted. b. Uncomment the <code>/etc/vfstab</code> entry on all cluster nodes for this file system.
7	If local or cluster file system fails to mount even after completing the preceding steps, contact NetApp technical support.

Obtaining the latest version of SnapDrive for UNIX software

What you need

The SnapDrive for UNIX software package that you download from the NOW site is the most recent version of the NetApp software. You can check the NOW site to see whether the software available there is more recent than the file you have on CD-ROM. If so, download the software instead of using the CD-ROM version of the software.

Obtaining the software from NOW

To download the software from the NOW site, complete the following steps.

Step	Action
1	Log in to the NOW site at 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 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	For instructions about extracting and installing this file, see the SnapDrive for UNIX <i>Installation and Administration Guide</i> .

Obtaining other NetApp documentation

Obtaining additional NetApp documentation

NetApp's NOW Web site contains the most current documentation for SnapDrive for UNIX, FCP and iSCSI Host Utilities, SAN storage system components, and various interconnect methods. For information, see the following:

- ◆ The online SnapDrive for UNIX *Compatibility Matrix* at <http://now.netapp.com/NOW/knowledge/docs/docs.shtml/>
- ◆ The online *Compatibility and Configuration Guide for NetApp's FCP and iSCSI Products* at http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/
- ◆ The online NOW Description Page for SnapDrive for UNIX on your host operating system
- ◆ The SnapDrive for UNIX *Installation and Administration Guide*
- ◆ SnapDrive for UNIX *Quick Start Guide*

