

SnapDrive® 4.1 for UNIX®

Release Notes (IBM® AIX®, HP-UX®, Linux®, Solaris®)

NetApp, 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: 215-04373_A0
December 2008

Copyright and trademark information

Copyright information

Copyright © 1994–2008 NetApp, Inc. All rights reserved. Printed in the U.S.A.

No part of this file 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 material of NetApp, Inc. 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 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, Cryptainer, Cryptoshred, DataFabric, DataFort, Data ONTAP, Decru, FAServer, FilerView, FlexClone, FlexVol, Manage ONTAP, MultiStore, NearStore, NetCache, NOW NetApp on the Web, SANscreen, SecureShare, SnapDrive, SnapLock, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, Spinnaker Networks, SpinCluster, SpinFS, SpinHA, SpinMove, SpinServer, StoreVault, SyncMirror, Topio, VFM, and WAFL are registered trademarks of NetApp, Inc. in the U.S.A. and/or other countries. gFiler, Network Appliance, SnapCopy, 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 NetApp arch logo; the StoreVault logo; ApplianceWatch; BareMetal; Camera-to-Viewer; ComplianceClock; ComplianceJournal; ContentDirector; ContentFabric; EdgeFiler; FlexShare; FPolicy; Go Further, Faster; HyperSAN; InfoFabric; Lifetime Key Management, LockVault; NOW; ONTAPI; OpenKey, RAID-DP; ReplicatorX; RoboCache; RoboFiler; SecureAdmin; Serving Data by Design; SharedStorage; Simplicore; Simulate ONTAP; Smart SAN; SnapCache; SnapDirector; SnapFilter; SnapMigrator; SnapSuite; SohoFiler; SpinMirror; SpinRestore; SpinShot; SpinStor; vFiler; VFM Virtual File

Manager; VPolicy; and Web Filer are trademarks of NetApp, Inc. in the U.S.A. and other countries. NetApp Availability Assurance and NetApp ProTech Expert are service marks of NetApp, Inc. in the U.S.A.

IBM, the IBM logo, AIX, and System Storage are trademarks and/or registered trademarks of International Business Machines Corporation.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. 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 U.S.A. and/or other countries.

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.

Contents of Release Notes

Topics covered

The following topics are covered in this Release Notes:

- ◆ [“Overview”](#) on page 2
- ◆ [“Limitations”](#) on page 4
- ◆ [“System requirements”](#) on page 6
- ◆ [“Finding information on known issues and troubleshooting”](#) on page 7
- ◆ [“Known issues affecting all hosts, and troubleshooting tips”](#) on page 8
- ◆ [“AIX specific issues and troubleshooting”](#) on page 22
- ◆ [“HP-UX specific issues and troubleshooting”](#) on page 24
- ◆ [“Linux specific issues and troubleshooting”](#) on page 26
- ◆ [“Solaris specific issues and troubleshooting”](#) on page 35
- ◆ [“Obtaining the latest version of SnapDrive for UNIX software”](#) on page 40
- ◆ [“Obtaining other NetApp documentation”](#) on page 41

Overview

Overview

SnapDrive® for UNIX® is a NetApp® 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 following feature is new in this release:

- ◆ Wizard CLI

For more information on wizard-supported commands, see “[Commands supported by the SnapDrive for UNIX wizard](#)” on page 3

Support for new platforms: Following is the list of new platforms supported with SnapDrive 4.1 for UNIX:

- ◆ Support for Sun® Native stack
- ◆ Support for Red Hat® Enterprise Linux® 4 Update 7, Red Hat Enterprise Linux 5 Update 1 and 2, and Oracle® Enterprise Linux 5 Update 2
- ◆ Support for Solaris® 10 Update 5
- ◆ Support for VMware® guest operating system in Linux with NFS and iSCSI protocols
- ◆ Support for SUSE® Linux Enterprise Server 10 with NFS protocol
- ◆ Support for Asymmetric Logical Unit Access (ALUA) in Solaris MPIO and AIX® Native MPIO
- ◆ Support for HP-UX vPAR
- ◆ Support for HP-UX 11iv3

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 *Interoperability Matrix*. This matrix is on the NOW™ NetApp on the Web™ site at

<http://now.netapp.com/NOW/products/interoperability/>

SnapDrive for UNIX is supported on the following host platforms:

- ◆ IBM® AIX
- ◆ HP-UX

- ◆ Linux

Note

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

- ◆ Solaris® on SPARC® and x86 architecture

With minor exceptions in functionality, the same feature set is available on each host. Also, many of the known problems and limitations of SnapDrive for UNIX are the same. As a result, this document contains information relevant to all platforms of SnapDrive for UNIX.

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* specific to your platform.

SnapDrive for UNIX components

SnapDrive for UNIX is available either on CD-ROM or as a software package that you can download from the NetApp NOW site (<http://now.netapp.com/>). The media kit includes the following items:

- ◆ CD-ROM

The CD-ROM for your host platform contains the NetApp software package in a standard installable package format that can be either of the following:

- ❖ An installation package or script for your host operating system
- ❖ The SnapDrive for UNIX software package

- ◆ License agreement

Commands supported by the SnapDrive for UNIX wizard

SnapDrive 4.1 for UNIX introduces a wizard to guide you through some of the commands. The wizard takes you through a set of questions, helping you complete the task in an interactive manner. Following is a list of commands that have an equivalent wizard:

- ◆ `snap connect`
- ◆ `storage create`
- ◆ `storage delete`
- ◆ `snap disconnect`
- ◆ `snap restore`

For more information on the SnapDrive for UNIX wizard, see the SnapDrive for UNIX *Installation and Administration Guide* specific to your platform.

Limitations

The SnapDrive 4.1 for UNIX software does not support the following functionality.

FC and iSCSI configuration

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

Volume-based SnapRestore

Volume-based SnapRestore is not supported in a multipathing setup using SANPath with an AIX host. Also, volume-based SnapRestore does not work with an SFRAC setup.

Multipathing

Multipathing is not supported on the following hosts:

- ◆ AIX hosts using iSCSI protocol
- ◆ Linux hosts using iSCSI protocol

AIX JFS

On AIX hosts, the JFS file system type is supported only for Snapshot operations and not for storage operations.

Thin provisioning

For the following Data ONTAP features SnapDrive for UNIX adheres to the volume options set by the storage administrator:

- ◆ Fractional reserve
- ◆ Snapshot reserve

The following Data ONTAP features are not supported:

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

vFiler®

Following are the vFiler unit limitations:

- ◆ SnapDrive for UNIX does not support access to vFiler units through FC, 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 and storage operations

Following are the limitations that are applicable to Snapshot and storage operations:

AIX: On an AIX originating host, the Snapshot connect operation is supported with NativeMPIO multipathing solution if the following patch level of AIX is used:

- ◆ 5.2 TL10 SP5
- ◆ 5.3 TL6 SP5
- ◆ 5.3 TL7 SP2

HP-UX: While creating a file system on RAW file system DMP multipathing, SnapDrive for UNIX mounts the file system on `/dev/dsk/<device-name>` device file and not on `/dev/vx/dmp/<device-name>` device file.

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 4.1 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 Interoperability Matrix* is at <http://now.netapp.com/NOW/products/interoperability/>. 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 Interoperability Matrix*.

Finding information on known issues and troubleshooting

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 affecting all hosts, and troubleshooting tips”](#) on page 8

Then check the section for your host platform:

- ◆ [“AIX specific issues and troubleshooting”](#) on page 22
- ◆ [“HP-UX specific issues and troubleshooting”](#) on page 24
- ◆ [“Linux specific issues and troubleshooting”](#) on page 26
- ◆ [“Solaris specific issues and troubleshooting”](#) on page 35
- ◆ The “Troubleshooting” chapter of the SnapDrive for UNIX *Installation and Administration Guide* for your specific OS platform.

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 clicking <http://now.netapp.com/NOW/cgi-bin/bol>.

Known issues affecting all hosts, and troubleshooting tips

If a SnapDrive for UNIX daemon stop is followed by a daemon start, the start operation might fail (299847)

Symptom: While running SnapDrive for UNIX commands in scripts, it is possible that a child process might be in time wait state on a SnapDrive for UNIX port for a long period. This period usually depends on the TCP timeouts and the life of that child process. In such cases, the daemon cannot use that port and daemon start might fail.

What to do: Wait until processes in time wait state are timed out. This wait period can be as long as 5 minutes.

Note

This problem might occur more frequently on Solaris platforms.

If Operations Manager server fails to discover LUNs, SnapDrive for UNIX commands fail with error "Could not check"

Symptom: SnapDrive for UNIX command displays the following type of error message:

```
0002-332 Admin error: Could not check SD.Storage.Write access on
lun storage_array1:/vol/vol1/lun1 for user unix-host\root on
Operations Manager server(s) ops-mngr-server1
```

What to do :

1. Verify that sd-admin is granted {Global, DFM.Core.AccessCheck} capability.
2. Verify that the storage resource is refreshed on Operations Manager. If storage resource is not refreshed on Operations Manager; then do one of the following:
 - Manually execute `dfm host discover <storage-system>`.
 - Assign {Global, DFM.Database.Write} capability to sd-admin. In this case SnapDrive for UNIX will automatically refresh Operations Manager and reissue access check.

The default refresh interval used by SnapDrive for UNIX is 60 seconds (`mgmt-retry-sleep-secs * 30 seconds`). If you still receive the above error message, then increase the value of `mgmt-retry-sleep-secs` configuration variable to 4, 6, 8, and so on to match your environment.

The refresh interval required by Operations Manager depends on the number of LUNs and volumes you have in your storage system.

**AIX, HP-UX only:
storage show for
LUNs fails if volume
group unavailable
(161395)**

Symptom: On AIX and HP-UX hosts, the `snapdrive storage show -lun lun_name` and the `snapdrive storage show -all` commands fail if any volume group on the system is inactive (HP-UX) or varied-off (AIX).

Other forms of the `snapdrive storage show` command work without problems as long as you do not specify an inactive or varied off volume group.

Explanation: SnapDrive for UNIX cannot determine the volume groups to which the LUNs belong. As a result, if the command specifies one or more LUNs, as both the `-lun` and `-all` options do, and the volume group is unavailable, the command fails. If you use one of the other command options (`-dg`, `-hostvol`, or `-fs`), the command captures the information for that entity and returns it.

What to do: When it is possible, activate or vary on all your host volume groups. If you have a host cluster, activating all the volumes might not be practical. In this case, use `snapdrive storage show` with the `-dg`, `-hostvol`, or `-fs` option rather than the `-lun` or `-all` option.

**Access control
permissions file not
working as intended**

Symptom: The local access control mechanism is not working as it should. 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`.

**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 is not able to access the storage system. 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 storage system option `httpd.admin.access` is set to “legacy”, 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 storage system.

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 storage system.

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

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

Note

You can set `trusted.hosts` or `httpd.admin.access` to the 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_snapshot
```

```

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 fail to remove Snapshot copies on the storage system 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 storage system 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 that are not connected to a host	Do one of the following to delete the Snapshot copy: ◆ Split the cloned LUNs from the storage system by using the <code>filer split clone</code> command. ◆ Delete the cloned LUNs.
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: ◆ Cloned LUNs in a Snapshot copy that are also present in other Snapshot copies ◆ Backup 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 contain the Snapshot copies of the cloned LUNs, as shown in the following example: <pre>tonic*> lun snap usage vol1 james_lun (filer view) mpiotest2:</pre> LUN: <pre>/vol/vol1/.snapshot/mpiotest2/d hilip_0</pre> Backed By: <pre>/vol/vol1/.snapshot/dhilip_lun/ james</pre> 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, as shown in the following 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 is being accessed, you will not 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*.
- ◆ 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
AIX	etc/filesystems
HP-UX or Linux	/etc/fstab
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.

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/um1 (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:

- ◆ NTAPasl
- ◆ 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
AIX	1. cfmgr 2. vxdisk scandisk
HP-UX	1. ioscan -fnC disk 2. ioinit -i 3. vxdisk scandisks

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.

AIX specific issues and troubleshooting

Message about clustering appears during snap restore

Symptom: The following message appears when you execute the `snapdrive snap restore` command:

```
# snapdrive snap restore -vg clustervg -snapname mysnap
It appears you are using high-availability host clustering
software. To avoid a possible cluster failover, please
disable cluster monitoring for the resources you want to
restore before restoring them.
Do you want to continue? (y/n)
```

Explanation: This is a generic message that applies to any SnapDrive for UNIX host operating system with host cluster monitoring. It appears even if your system does not provide host cluster monitoring or you have it turned off.

The problem occurs because the `snapdrive snap restore` command brings down the volume group temporarily. If cluster monitoring is going on at this time, performing a `snapdrive snap restore` operation could cause a cluster failover.

What to do: This should not be a problem on AIX hosts running HACMP because HACMP does not provide default resource monitoring for volume groups and applications.

If you customized an application to provide storage monitoring under HACMP, you must disable the application and monitoring before you execute a `snap restore` command.

Intermittent failure of storage show operation

Symptom: On an AIX host which uses the DotHill SANpath® for mutipathing, the `snapdrive storage show -all` operation fails with the following error message:

```
LUN information not available. the LUN may by busy or offline
```

Explanation: This happens due to a bug in the DotHill SANpath product, which incorrectly populates the physical volume name of the AIX (native) logical volume manager, after a `snapdrive snap connect` operation on an originating host.

This condition often manifests itself on a cluster setup which has more than one path for the FCP SAN. This does not affect any I/O on the cloned or the original file system.

What to do: Do one of the following:

- ◆ Run the `setsp -W0` command.
- ◆ Reboot the host.

Intermittent failure of Snapshot connect operation

Symptom: Intermittent failure of Snapshot connect operation on an AIX host.

Explanation: On the AIX 5.2 host the Snapshot connect operation fails, if the following are true:

- ◆ SANpath 4.0.5l software installed
- ◆ The `failback` attribute for `spdctrl` is set to “3”

What to do: The workaround is to disable SANpath’s PVID feature. To do this, enter the following command:

```
setsp -W0
```

Note

SANpath 4.0.5m fixes this issue.

MPIO-specific issue

The MPIO version of FC IBM AIX Host Utilities has limited support for cfmode types. If you are running Data ONTAP 7.1.x, only single-image or SSI cfmode is supported. With Data ONTAP 7.2 or later, both single-image cfmode and partner cfmode are supported. MPIO is not supported in Data ONTAP 7.0.x or earlier.

HP-UX specific issues and troubleshooting

Command fails with /dev/ in path name (116375)

Symptom: The SnapDrive for UNIX commands fail when the name of the volume group includes the full HP-UX specification. The following error message appears:

```
# snapdrive snap create -vg /dev/tester_vg_2 -snapname qa2
Error (Command error 0001-058): Diskgroup /dev/tester_vg_2 is not
found on Volume Manager
```

Explanation: The SnapDrive for UNIX commands do not accept the full HP-UX specification when you supply the name of a volume group as an argument.

What to do: Enter the volume group name without preceding it with /dev/. For example, the following SnapDrive for UNIX command line executes without producing an error message:

```
# snapdrive snap create -vg tester_vg_2 -snapname qa2
```

Failure of Snapshot restore operation (129741)

Symptom: The Snapshot restore operation fails if you manually activate the volume group during the restore operation.

Explanation: The Snapshot restore operation automatically performs the following tasks:

- ◆ Deactivates a volume group.
- ◆ Exports it.
- ◆ Restores the LUNs.
- ◆ Re-imports and re-activates the volume group.

If you try to activate the volume group manually during the Snapshot restore operation, you must then perform a manual recovery.

What to do: Do not manually activate or import a volume group during a snap restore operation.

If you do and this problem occurs, provide NetApp technical support with the recovery log and use their help to complete the Snapshot restore operation.

**Failure of Snapshot
and host
disconnect
operations (248871)**

Symptom: On an HP-UX host with a DMP multipathing solution, the `snapdrive snap disconnect` and `snapdrive host disconnect` commands fail with the following error:

0001-023 Admin error: Unable to discover all LUNs in file system

Explanation: On an HP-UX host with a DMP multipathing solution, SnapDrive for UNIX does not support a file system on a raw LUN.

Linux specific issues and troubleshooting

Data disappears when you add LUNs to a volume group

Symptom: You add an existing LUN to a volume group, but the data on the LUN disappears.

Explanation: If you add an existing LUN to a volume group that is managed by the LVM and then execute a `pvcreeate` command, the data is overwritten.

What to do: If you need data that is on an exiting LUN, you must migrate that data to a LUN that you create using the LVM. Complete the following steps.

Step	Action
1	Create new LUNs.
2	Place the LUNs under LVM control by using the <code>pvcreeate</code> , the <code>vgcreate</code> , and the <code>lvcreate</code> commands.
3	Copy data from the existing LUNs to the LUNs you just created and placed under LVM control.

SnapDrive for UNIX commands fails if the transport protocol connectivity fails (241596)

Symptom: SnapDrive for UNIX commands fail.

Explanation: This happens if the iSCSI or FCP connection between the host and storage system fails, or if iSCSI or FCP on the storage system fails.

What to do: To successfully use SnapDrive for UNIX commands, you must first restore the iSCSI or FCP service.

Autodetection does not work with MPIO setup for an FS created over a raw LUN

Symptom: Autodetection does not work with file system on raw LUN on a Linux MPIO setup with SnapDrive for UNIX.

What to do: Use the option `-fs` with all the snapdrive operations.

Attempt to connect a Snapshot copy fails

Symptom: SnapDrive for UNIX fails to connect to a Snapshot copy when a Linux host reboots during the SnapDrive for UNIX operation. The following error message appears:

Volume manager configuration file has tags that are added by SnapDrive for UNIX. Clean the file manually

Explanation: This happens on Linux LVM2 during a `snap connect` operation on an originating host that reboots during the SnapDrive for UNIX operation.

What to do: Complete the following steps.

Step	Action
1	Ensure that the LUN clones of devices belonging to the disk groups are not mapped to the host. If they are mapped to the host, then unmap and delete them. They might have been present from a previously failed SnapDrive operation.
2	Edit the <code>lvm.config</code> file (located at <code>/etc/lvm/</code>) in the following manner: a. Remove all lines with suffix <code>#SDUTEMP</code>. b. Restore all lines that are preceded by the comment tag <code>#SDUEXISTING</code>.

If none of the above works, contact NetApp technical support.

Attempt to delete LUNs with exported disk groups on Linux LVM 2 fails

Symptom: SnapDrive for UNIX does not delete LUNs with exported disk groups on Linux LVM2.

Explanation: If you export a disk group (dg) on Linux LVM2 and try to delete the LUN, the delete operation fails.

What to do: You have to use `vgimport <dgname>` to import the disk group and then use the following command to delete the disk group and LUN:

```
# snapdrive storage delete - dg <dgname> [-full]
```

SnapDrive for UNIX occasionally fails to discover LUNs

Symptom: SnapDrive for UNIX occasionally fails to discover LUNs on a host running SUSE 9, Service Pack 3.

Explanation: The iSCSI stack sometimes gets corrupted and causes SnapDrive for UNIX to fail to discover LUNs or devices during SnapDrive for UNIX operations.

What to do: Reboot the host.

After each reboot snap connect yields "0001-680 Admin error: Host OS requires an update to internal data" (226889)

Symptom: On Red Hat Enterprise Linux 4 and Oracle Enterprise Linux 4 hosts, after every reboot the Snapshot connect and storage create operations yield the following error:

0001-680 Admin error: Host OS requires an update to internal data

Explanation: This happens because sg driver deletes all the inactive device /dev/sgXX files after a reboot.

What to do: Complete the following steps.

Step	Action
1	Check the value of the configuration variable <code>enable-implicit-host-preparation</code> in the <code>/opt/NetApp/snapdrive/snapdrive.conf</code> file. It should be set to "on". After rebooting the host and before using SnapDrive for UNIX, complete the following steps.
2	Verify that the sg and iSCSI services are loaded by entering the following command: <pre>lsmod</pre>
3	If the sg and iSCSI (if iSCSI setup is used) services are not loaded, enter the following commands: ◆ <code>modprobe sg</code> ◆ <code>service iscsi start</code>
4	(Optional) If you do not want to perform steps 2 and 3 manually, add the following commands in the host startup script: ◆ <code>modprobe sg</code> ◆ <code>service iscsi start</code>

SnapDrive errors out if /lib/libgcc_s.so.1 is not properly installed.

Symptom: SnapDrive for UNIX errors out if the /lib/libgcc_s.so.1 path is not properly installed on 64-bit Red Hat Enterprise Linux 4 and Oracle Enterprise Linux 4 hosts.

Explanation: If the path /lib/libgcc_s.so.1 is not properly installed on a 64-bit Red Hat Enterprise Linux 4 and Oracle Enterprise Linux 4 hosts, SnapDrive for UNIX displays the following error message:

libgcc_s.so.1 must be installed for pthread_cancel to work

What to do: Follow the procedure that is explained in the kb21402 article at <http://now.netapp.com/Knowledgebase/solutionarea.asp?id=kb21402>.

HBA_LoadLibrary: Encountered and error loading: /usr/lib/libqlsdrm.so (240502)

Symptom: On all updates of Red Hat Enterprise Linux 4 and Oracle Enterprise Linux 4 hosts with iSCSI setup, all SnapDrive for UNIX commands error out and display the following message:

HBA_LoadLibrary: Encountered and error loading:
/usr/lib/libqlsdrm.so

Explanation: During unloading or uninstalling the QLogic driver, the libHBaAPI.so library remains in the /usr/lib file and is not removed by the QLogic ./install scripts.

What to do: Rename or move (mv) the libqlsdrm.so library. For example,

mv /usr/lib/libqlsdrm.so /usr/lib/libqlsdrm.so.old.

Restoring a file system from a Snapshot copy fails (243636)

Symptom: After migration from Red Hat Enterprise Linux 3 updates to Red Hat Enterprise Linux 4 updates, that is, from LVM 1 to LVM 2, restoring a file system from a Snapshot copy fails.

Explanation: During restoring a file system (that is taken on LVM1) from a Snapshot copy using LVM2, the snapdrive snap restore -fs command does not mount the file system. However, it successfully restores the disk group and host volume associated with it.

This happens because Red Hat Enterprise Linux 4 update 3 hosts use LVM2, which uses a device mapper to interface with the kernel, whereas Red Hat Enterprise Linux 3 update 5 hosts use LVM1, which does not use a device mapper.

For example, the host volume created by LVM1 will have a device file at /dev/Mydg_SDDg/Mydg_Sdhv; and the host volume created by LVM2 will have a device file at /dev/mapper/Mydg_SdDg-Mydg_Sdhv.

What to do: Manually mount the file system.

**SnapDrive for UNIX
commands fail to
discover LUNs
because of iSCSI
daemon limitation
(243917)**

Symptom: SnapDrive for UNIX commands fail to discover LUNs after continuous multiple iterations. The following error message is displayed:

```
0000-001 Admin error: iscsi transport protocol has dropped sessions
in the host. Please ensure the connection and the service in the
filer
```

Explanation: This happens because of the iSCSI daemon limitation in the following hosts:

- ◆ Red Hat Enterprise Linux 4.0
- ◆ SLES 9, Service Pack 3
- ◆ Oracle Enterprise Linux 4.0

What to do: Complete the following steps.

Note

Before modifying the following system parameters, contact the OS vendor.

Step	Action
1	Set the following parameter to “65536” value: <pre># sysctl -w fs.file-max=65536</pre> Note_____ NetApp recommends to set it to 65536, but you have to set the value depending on the applications on your host.
2	Add the fs.file-max = 65536 parameter to /etc/sysctl.conf file.

Step	Action
3	Update the default shell start-up file for the root user's UNIX account. ◆ For the Bourne, Bash, or Korn shell, add the following lines to the /etc/profile file by running the following command: <pre>cat >> /etc/profile <<EOF if [\\${USER} = "root"]; then if [\\${SHELL} = "/bin/ksh"]; then ulimit -p 16384 ulimit -n 65536 else ulimit -u 16384 -n 65536 fi umask 022 fi EOF</pre> ◆ For the C shell (csh or tcsh), add the following lines to the /etc/csh.login file by running the following command: <pre>cat >> /etc/csh.login <<EOF if (\\${USER} == "root") then limit maxproc 16384 limit descriptors 65536 endif EOF</pre>
4	Add the following parameters to the /etc/security/limits.conf file: <pre>root soft nofile 65536 root hard nofile 65536</pre>
5	Activate all the kernel parameters for the system by executing the following command: <pre>sysctl -p</pre>

Snapshot restore operation fails to restore file system on raw LUN (251215)

Symptom: The Snapshot restore operation of file system on raw LUN fails after performing a Snapshot connect operation on a non-originating host.

Explanation: The following error occurs when the cloned LUN is already connected to the same filespec present in Snapshot copy and then you try to execute the snapdrive snap restore command.

"0001-034 Command error: mount failed: mount: <device name> is not a valid block device"

The command fails because the iSCSI daemon remaps the device entry for the restored LUN when you delete the cloned LUN.

What to do: Do either of the following:

- ◆ Execute the `snapdrive snap restore` command again.
- ◆ Delete the connected LUN (if it is mounted on the same filespec as in Snapshot copy) before trying to restore a Snapshot copy of a original LUN.

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. _____

Step	Action
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> Note_____ <code>/etc/init.d/iscsi reload</code> option is applicable for versions RHEL 4 update 6 and prior. ◆ For RHEL 5 and later, if you are using iSCSI you have to remove LUNs manually. Also, you have to perform host rescan manually. If you are using multipathing solution, then perform rescan for multipathing also manually. You can do so by running the following command: <pre>multipath</pre> Note_____ If you are using FCP, then the above process applies to all the versions prior to RHEL 4 update 6 and RHEL 5 and later.
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 vgscan command to retain the changes made in the configuration file.
6	Enter the vgchange -a y 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.

Using QLogic rescan script in multipathing environment

If you are working on Linux host with Qlogic FC setup, then download the Qlogic rescan script from the Qlogic website and copy it to /root path and rename it as `dynamic-lun-rescan.sh` for SnapDrive for UNIX to use it. The QLogic rescan script asks for user confirmation; SnapDrive for UNIX may hang at this confirmation stage indefinitely. You can avoid this confirmation by setting the variable `QL_DISABLE_WARNING` to 1 in the script.

Solaris specific issues and troubleshooting

The snap restore cmd fails to mount the filesystem (298282)

Known issue: With SFRAC 5.0 on Solaris, the `snap restore` and `snap connect` commands fail to mount the filesystem.

What to do: Before taking a Snapshot user has to execute the following command on the cluster node where the user is going to issue `snapdrive snap create` command:

```
/opt/VRTS/bin/fsclustadm setprimary mount_point
```

LUN discovery fails when using SnapDrive for UNIX in single path Environment on Solaris 10 update 4 (298551)

Symptom: On Solaris 10 update 4 on x86 architecture, `sd` driver leaves behind stale entries on the host after LUN deletion in a single path environment with FCP. Later this device gets associated with a new LUN and leads to I/O error.

What to do: Creating a dummy LUN on storage system and mapping it to the host solves the problem mentioned above. Hence on a fresh install, a dummy LUN is required to be mapped to the host before any SnapDrive for UNIX operations are started. The above dummy LUN can be created using SnapDrive for UNIX as follows:

```
bash-3.00# snapdrive storage create -lun f3050-221-178:/vol/vol5_test/sun221-50_rrt_lun0 -lunsize 40m
```

```
LUN f3050-221-178:/vol/vol5_test/sun221-50_rrt_lun0 ... created
```

```
mapping new lun(s) ... done
discovering new lun(s) ... done
```

```
LUN to device file mappings: - f3050-221-178:/vol/vol5_test/sun221-50_rrt_lun0 => /dev/dsk/c4t500A0982964800C5d0s2
```

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 storage system 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 storage system 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 might 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 2008 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-filessystem-local:default.log</pre> Example: <pre># tail /var/svc/log/system-filessystem-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: a. Enter your user name and password. b. From the Select Start Page list box, click Download Software. c. 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 <i>SnapDrive for UNIX Installation and Administration Guide</i> for your specific platform.

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 *Interoperability Matrix* at <http://now.netapp.com/NOW/products/interoperability/>
- ◆ 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* for your specific platform (available as a downloadable PDF file from the NOW site and in printed form with the SnapDrive for UNIX media kit)
- ◆ SnapDrive for UNIX *Quick Start Guide* (available as a downloadable PDF file from the NOW site and in printed form with the SnapDrive for UNIX media kit)

