

SnapDrive® 5.2.1 for UNIX®

Release Notes

(AIX®, Linux®, and Solaris®)

NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089 U.S.
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-08631_A0_ur001
February 2014

Contents

Release notes	5
New and enhanced features in SnapDrive 5.2.1 for UNIX	6
Fixed issues in SnapDrive 5.2 for UNIX	7
Restore operations in SnapDrive for UNIX fail when the IP address is directly used in the Snapshot copy name	7
Restore operation fails when multiple security styles are configured	7
Some SnapDrive operations fail in clustered Data ONTAP	8
SnapDrive for UNIX daemon stops responding and fills up file system space in AIX	8
SnapManager for Oracle or SnapManager for SAP restore operation fails when clustered Data ONTAP 8.1.x is used with SnapDrive 5.2 for UNIX	9
SnapDrive for UNIX commands involving volume delete operations might fail	9
Cloning operation fails in SnapDrive for UNIX due to a volume group name conflict	10
SnapDrive for UNIX restore and cloning operations fail while mounting	10
SnapDrive for UNIX operations involving concurrent volume groups fail in certain AIX versions	10
SnapDrive for UNIX floods a trace log message	11
Limitations	12
Known issues	14
Known issues: Related to all hosts	14
SnapDrive dataset operations fail when performed from tertiary storage	14
Automatic detection does not work for more than one LUN in a secondary volume	14
If the Operations Manager server fails to discover the LUNs, SnapDrive for UNIX commands fail	15
Access control permissions file does not work as intended	15
Unable to clear a newly created export policy when the policy is deleted from a remote host	16
Storage provisioning and snap management operations might fail intermittently on the guest operating system	16

Unable to use either Telnet or Secure Shell in SFRAC environment in certain scenarios	17
Known issues in SnapDrive for UNIX on Linux	17
SnapDrive for UNIX does not work as expected when 32-bit versions of the required libraries are not installed	17
The restore operation for NFS directories that contain symbolic links fails in clustered Data ONTAP	18
SnapDrive operations might fail due to a known Red Hat Enterprise Linux issue	18
SnapDrive operations might hang intermittently in Red Hat Enterprise Linux 5.8	18
The SnapDrive daemon fails to start when you use Emulex OneCommand Manager 6.1 or later	19
The snap disconnect operation on the NFS file specification creates a stale export policy in clustered Data ONTAP	19
Known issues in SnapDrive for UNIX on AIX	20
The storage show commands for LUNs fail if volume group is unavailable	20
Known issues in SnapDrive for UNIX on Solaris	20
SnapDrive for UNIX delete and disconnect of disk groups fails without the -full option	20
Where to find product documentation and other information	22
Copyright information	23
Trademark information	24
How to send your comments	25

Release notes

The release notes describe new features, limitations, fixed issues, and known issues in SnapDrive 5.2.1 for UNIX.

New and enhanced features in SnapDrive 5.2.1 for UNIX

SnapDrive 5.2.1 for UNIX contains new features and existing features that have been enhanced.

The new and enhanced features are as follows:

- Support for Storage Area Network (SAN) and Network File System (NFS) in clustered Data ONTAP on the Solaris operating system
- Allows override of the SnapMirror or SnapVault existence check, which is one of the mandatory checks performed during volume-based SnapRestore (VBSR)

SnapDrive does not allow a VBSR if a SnapMirror or SnapVault relationship exists; this is because a VBSR that uses a Snapshot copy made before the base Snapshot copy breaks the SnapMirror or SnapVault relationship. You can override this and perform a VBSR by enabling the *override-vbsr-snapmirror-check* or *override-vbsr-snapvault-check* configuration variables. Enabling the override variables allows VBSR to be performed, but breaks the relationship. This variable is applicable only in Data ONTAP operating in 7-Mode and when OnCommand Data Fabric Manager is not configured.

- Supports the Mount Guard feature of AIX, which prevents simultaneous or concurrent mounts. If a file system is mounted on one node and the Mount Guard feature is enabled, AIX prevents the same file system from being mounted on another node.
- Supports paravirtual SCSI controlled devices (PVSCSI) on Linux guest operating systems

Note: For updated information about SnapDrive and its requirements, see the Interoperability Matrix.

Related information

Interoperability Matrix: support.netapp.com/matrix

Fixed issues in SnapDrive 5.2.1 for UNIX

Some known issues are fixed in SnapDrive 5.2.1 for UNIX.

Restore operations in SnapDrive for UNIX fail when the IP address is directly used in the Snapshot copy name

Restore operations in SnapDrive for UNIX fail if Internet Protocol (IP) address is directly mentioned in the Snapshot copy name.

The issue occurs because SnapDrive for UNIX treats the IP address as the host name attempts to resolve it, which fails.

For more information, see bug ID 679920 on Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

Restore operation fails when multiple security styles are configured

The restore operation in SnapDrive for UNIX fails when the storage system has multiple security styles configured in the exportfs rules for Data ONTAP operating in 7-mode.

Note: The restore operation works when multiple security styles are configured for the NFS export rule in clustered Data ONTAP.

For more information, see bug ID 715661 on Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

Some SnapDrive operations fail in clustered Data ONTAP

When you are using clustered Data ONTAP, some of the SnapDrive operations fail if performed on a volume that has a junctioned volume in its nested junction path. This is a known issue in Data ONTAP 8.2 and has been fixed in Data ONTAP 8.2.1.

The following error message is displayed: 0009-071 Admin error: Failed to end file list iteration on storage system sduvs2 : Stale NFS file handle.

For more information, see bug ID 685429 at Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

SnapDrive for UNIX daemon stops responding and fills up file system space in AIX

In some situations, the SnapDrive for UNIX daemon can fill up the file system space. This scenario usually occurs when SnapDrive for UNIX waits for free threads to be available.

Similarly, SnapDrive for UNIX might also stop responding during such operations as a job cleanup.

The following error message is displayed:

```
0002-890 Admin error: Error mounting temp directory
<filename>: Unexpected error from mount: Usage:
mount -V : print version
mount -h : print this help
mount : list mounted filesystems
mount -l : idem, including volume labels
```

For more information, see bug ID 694340 at Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

SnapManager for Oracle or SnapManager for SAP restore operation fails when clustered Data ONTAP 8.1.x is used with SnapDrive 5.2 for UNIX

SnapManager for Oracle or SnapManager for SAP restore operation fails if you are using clustered Data ONTAP 8.1.x with SnapDrive 5.2 for UNIX.

SnapDrive for UNIX always attempts to create a rollback Snapshot copy before proceeding with the restore operation. Due to the non-availability of a Consistency Group Snapshot (CGS) copy in clustered Data ONTAP 8.1.x, when SnapDrive 5.2 for UNIX attempts to create a CGS copy, SnapDrive for UNIX falls back to the best-effort mode and returns the following warning message:

```
Attempting to take a crash consistent snapshot. Switching to the best-effort mode
```

This issue occurs while using Network File System (NFS) file specifications.

For more information, see bug ID 707639 on Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

SnapDrive for UNIX commands involving volume delete operations might fail

Certain SnapDrive for UNIX commands that involve volume delete operations might fail when the storage system is in a jumpahead state. This can occur as the volume might not be online.

For more information, see bug ID 740182 at Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

Cloning operation fails in SnapDrive for UNIX due to a volume group name conflict

A cloning operation in SnapDrive for UNIX on Linux operating systems fails when there is a volume group name conflict on the host. When cloning is initiated, the volume group name is not renamed automatically and can lead to conflicts with existing entities on the host.

For more information, see bug ID 710158 at Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

SnapDrive for UNIX restore and cloning operations fail while mounting

SnapDrive for UNIX restore and cloning operations can fail while mounting due to a failure to detect the new mount guard feature that is available in IBM AIX.

For more information, see bug ID 650466 at Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

SnapDrive for UNIX operations involving concurrent volume groups fail in certain AIX versions

SnapDrive for UNIX operations involving concurrent volume groups fail in IBM AIX versions 6.1 Technology Level (TL) 8 and 7.1, due to a known AIX issue. This issue causes the Recreatevg command to fail in AIX versions 6.1 TL8 and 7.1.

For more information, see bug ID 770067 at Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

SnapDrive for UNIX floods a trace log message

SnapDrive for UNIX floods a trace log message during some Network File System (NFS) restore scenarios.

When a trace log message floods, it can take up space in the `/var/log` directory .

For more information, see bug ID 451449 at Bugs Online.

Related information

Bugs Online: support.netapp.com/NOW/cgi-bin/bol

Limitations

When working with SnapDrive, you must be aware of certain limitations that might affect your environment.

General limitations of SnapDrive for UNIX

- SnapDrive for UNIX does not support snap operations on a Network File System (NFS) mount point when the volume is exported with Kerberos authentication security types such as krb5, krb5i, or krb5p.
- In clustered Data ONTAP, SnapDrive for UNIX does not support symbolic links for any mount points.
- Snap operations might be inconsistent if you perform SnapRestore on a mount point where a different entity is mounted other than what was created by the Snapshot copy.
- SnapDrive modifies the mount point permissions from non-root user to root user for a qtree after volume-based SnapRestore (VBSR) operations.
- From SnapDrive for UNIX 5.1, only asynchronous mode is supported for clone-split operations. The *split-clone-async* configuration variable and *-splitmode* command-line options are removed.
- The clone-split operation is not supported on a filespec, which is a combination of a volume clone and LUN clone, LUN clone and LUN, or volume clone and LUN.
- Restoring the storage stack on the host fails if you remove a LUN by using an external command.
- SnapDrive for UNIX does not allow you to restore a Snapshot copy of a qtree if the Snapshot copy is mounted on an NFS in the SnapMirror destination.
- SnapDrive for UNIX does not support Broadcom Converged Network Adapters (CNA).
- In clustered Data ONTAP, to configure data Vserver with SnapDrive for UNIX, you must ensure that the IP address of the data Vserver's management logical interface (LIF) is mapped to the data Vserver name either in the Domain Name System (DNS) or in the */etc/hosts* file. You must also ensure that the data Vserver name is configured in SnapDrive by using the *snapdrive config set <vsadmin> <Vserver name>* command.
- In clustered Data ONTAP, aggregates that contain data Vserver volumes must be assigned to the data Vserver's aggregate list to configure the data Vserver and perform SnapDrive operations.
- In clustered Data ONTAP, SnapDrive for UNIX does not support the snap connect operation with the *-readonly* option from the secondary host, which does not have export permissions on the parent volume).
In clustered Data ONTAP, the export policy is defined only at the volume level. Therefore, the *.snapshot* directory cannot be exported to the secondary host, which does not have export permission on the parent volume.

Limitations of SnapDrive for UNIX on Solaris only

- LUNs greater than 1 TB with the Storage Management Initiative (SMI) label are not supported.

- SnapDrive for UNIX does not support Extensible Firmware Interface (EFI) and SMI LUNs in the same disk group.

Limitations of SnapDrive for UNIX on AIX only

- The SnapDrive daemon takes a long time to restart if one or more Fibre Channel (FC) port links are down or not available.
However, the SnapDrive daemon starts irrespective of the state of the FC ports.
- By default, SnapDrive for UNIX creates disk groups with non-concurrent settings, and you need to change or provide storage provisioning operations manually.
- While upgrading SnapDrive for UNIX on AIX, if you get a warning message `Already installed`, perform a force upgrade.

Known issues

You must be aware of the known issues and the workaround for each issue so that you can use SnapDrive 5.2.1 for UNIX more effectively.

Known issues related to all hosts

You must be aware of some potentially undesirable behaviors that impact all supported host platforms and the workarounds.

SnapDrive dataset operations fail when performed from tertiary storage

The dataset operations, such as restore and cloning, fail from tertiary storage, if the protection policy is set to “Mirror, then Back up”. This issue occurs if you are using 5.1 and 5.2.

Workaround

None.

For more information, see bug ID 682608 on Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

Automatic detection does not work for more than one LUN in a secondary volume

If you delete more than one LUN in a secondary volume, and run the `snapdrive storage delete` command, SnapDrive for UNIX displays the following error message: `Could not detect type of the entity....`

The `snapdrive storage delete` command deletes only the first LUN in the secondary volume.

```
# snapdrive storage delete storagesystem:/vol/user/test1 test2
/vol/user/test3 test4
Auto detection of file_spec(s) in progress ...
- 'storagesystem:/vol/user/test1' detected as LUN
- 'test2' detected as LUN
- '/vol/user/test3' detected as LUN
9000-305 Command error: Could not detect type of the entity test4.
Provide a specific option (-lun, -dg, -fs or -lvol) if you
know the type of the entity.
```

Workaround

None.

For more information, see bug ID 540314 on Bugs Online.

Related information

Bugs Online: support.netapp.com/NOW/cgi-bin/bol

If the Operations Manager server fails to discover the LUNs, SnapDrive for UNIX commands fail

SnapDrive for UNIX commands fail, if the Operations Manager server fails to discover the LUNs.

The following error message is displayed: 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

Workaround

None.

For more information, see bug ID 304724 on Bugs Online.

Related information

Bugs Online: support.netapp.com/NOW/cgi-bin/bol

Access control permissions file does not work as intended

The local access control mechanism does not work as intended when the storage system does not have a vol0 volume.

By default, SnapDrive for UNIX can find the permissions file for local access control at `/vol/vol0/sdprbac/sdhostname.prbac`, where *hostname* is the name of the local host. This issue does not occur if the vol0 volume is on the storage system.

Workaround

Use the `rbac-vol-name` parameter in the `snapdrive.conf` file to specify a different volume name where SnapDrive for UNIX looks for the permissions file. Do not specify the name of the volume with the `/vol` suffix. For example, you can specify `rbac-vol-name="vol1"`.

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

For more information, see bug ID 161689 on Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](https://support.netapp.com/NOW/cgi-bin/bol)

Unable to clear a newly created export policy when the policy is deleted from a remote host

A new *exportfs* policy is created when a clone is created on the host, where it does not have *exportfs* permission. The newly created *exportfs* policy cannot be cleared even if you delete the clone.

Workaround

You must ensure that you have the appropriate permissions on the remote host, and then manually delete the *exportfs* policy from the filer.

For more information, see bug ID 773433 at Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](https://support.netapp.com/NOW/cgi-bin/bol)

Storage provisioning and snap management operations might fail intermittently on the guest operating system

SnapDrive for UNIX (SDU) leverages Virtual Storage Control (VSC) for storage provisioning on the guest operating system (OS). Storage provisioning and snap management operations might fail intermittently on the guest OS because of VSC scalability limitations.

Workaround

None.

For more information, see bug ID 778241 at Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](https://support.netapp.com/NOW/cgi-bin/bol)

Unable to use either Telnet or Secure Shell in SFRAC environment in certain scenarios

While using SnapDrive for UNIX in the Symantec Storage Foundation for Oracle RAC (SFRAC) environment, if you restart the host after performing a clone operation, you cannot use either Telnet or Secure Shell (SSH) on the host.

Workaround

Before restarting the host, you must comment the mount point corresponding to the cloned file specification in the `/etc/vfstab` file. After restarting the host, you can remount the cloned file specification.

For more information, see bug ID 725028 on Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

Known issues in SnapDrive for UNIX on Linux

You must be aware of some potentially undesirable behaviors that impact the Linux host platform and the workarounds.

SnapDrive for UNIX does not work as expected when 32-bit versions of the required libraries are not installed

SnapDrive for UNIX does not work as expected if the 32-bit versions of the required libraries are not installed. This limitation persists because SnapDrive for UNIX is a 32-bit binary.

The required binaries include 32-bit versions of `libgcc` and `libn1` packages.

Additionally, in Qlogic environments, the 32-bit `QConvergeConsoleCLI` package is required.

Workaround

None.

For more information, see bug ID 720020 on Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

The restore operation for NFS directories that contain symbolic links fails in clustered Data ONTAP

In clustered Data ONTAP, if you use SnapDrive for UNIX to restore an NFS directory that contains symbolic links, the operation fails.

Workaround

None.

For more information, see bug ID 573044 on Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

SnapDrive operations might fail due to a known Red Hat Enterprise Linux issue

SnapDrive for UNIX operations in Linux might fail, when `sg_*` commands that run against Device Mapper (DM) multipath devices fail, due to a path group switch with insufficient input or output operations to activate the path group.

The failure might occur in SnapDrive 5.0 for UNIX and later versions, for the hosts on Red Hat Enterprise Linux 5.6 and later, using Qlogic Fibre Channel adapters.

Workaround

Run a few input and output operations to read the DM multipath device. This activates the path group. Then retry the SnapDrive for UNIX operation that failed by running the following command :

```
dd if= /dev/mapper/<device> of=/dev/null bs=4096 count=100 iflag=direct
```

For more information, see bug IDs 689501, 521233, and 554518 on Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

SnapDrive operations might hang intermittently in Red Hat Enterprise Linux 5.8

SnapDrive operations might hang intermittently for the hosts on Red Hat Enterprise Linux 5.8, because the `vgdisplay` command fails.

Workaround

Restart the multipath daemon.

Note: If the SnapDrive operations hang intermittently in the Fibre Channel SAN boot environment, then restart the host.

For more information, see bug ID 678909 on Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

The SnapDrive daemon fails to start when you use Emulex OneCommand Manager 6.1 or later

The SnapDrive daemon fails to start when you use Emulex OneCommand Manager 6.1 or later with Emulex LightPulse host bus adapters (HBAs). The following error message is displayed: Unable to start daemon...check daemon log file for more details.

Workaround

Use Emulex OneCommand Manager 6.0 or earlier.

For more information, see bug ID 668199 on Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

The snap disconnect operation on the NFS file specification creates a stale export policy in clustered Data ONTAP

When you perform the snap connect operation on the Network File System (NFS) file specification from the secondary host, which does not have export permissions on the parent volume in clustered Data ONTAP, SnapDrive creates a new export policy. In the newly created export policy, SnapDrive inherits the parent export policy rules as well as appends the rules for the secondary host, if it does not have the export permission on the parent volume.

When you perform the snap disconnect operation on the NFS file specification from the secondary host, clustered Data ONTAP deletes the newly created export policy and creates `unexport_policy`, which is a stale export policy. `unexport_policy` does not contain any rules.

Workaround

None.

For more information, see bug ID 601895 on Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

Known issues in SnapDrive for UNIX on AIX

You must be aware of some potentially undesirable behaviors that impact the AIX host platform and the workarounds.

The storage show commands for LUNs fail if volume group is unavailable

On AIX hosts, the `snapdrive storage show -lun lun_name` and `snapdrive storage show -all` commands fail if any volume group on the system is varied off.

SnapDrive for UNIX cannot determine the volume groups to which the LUNs belong. This issue might occur because the storage show commands specify one or more LUNs using the options `-lun` and `-all`, and the volume to which the LUN belongs may be inactive. If you use one of the other command options such as `-dg`, `-hostvol`, or `-fs`, the command captures the information for that entity. Other forms of the `snapdrive storage show` command work without issues if you do not specify an inactive or varied-off volume group.

Workaround

Activate the inactive volume group by using the `varyon` command. If you have a host cluster, activating all the volumes might not be practical. In this case, use the `snapdrive storage show` command with the `-dg`, `-hostvol`, or `-fs` option rather than the `-lun` or `-all` option.

For more information, see bug ID 161395 on Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

Known issues in SnapDrive for UNIX on Solaris

You must be aware of some potentially undesirable behaviors that impact SnapDrive for UNIX operating on the Solaris host platform, and their workarounds.

SnapDrive for UNIX delete and disconnect of disk groups fails without the -full option

SnapDrive for UNIX delete and disconnect of disk groups operations fail because the host volumes are not completely disconnected or deleted.

Workaround

Run the `-full` option to delete or disconnect all file systems and the associated host volumes.

For more information, see bug ID 756309 on Bugs Online.

Related information

Bugs Online: support.netapp.com/NOW/cgi-bin/bol

Where to find product documentation and other information

The NetApp Support Site contains the most current documentation for SnapDrive for UNIX, FC and iSCSI Host Utilities, SAN storage system components, and various interconnects.

The following table lists some of the documents on the NetApp Support Site that contain current information about host requirements, supported configurations, and procedures for installing and managing SnapDrive for UNIX:

If you need more information about...	Go to...
The latest supported configurations	The Interoperability Matrix
Supported Fibre Channel SAN topologies	The <i>SAN Configuration Guide</i> (called <i>Fibre Channel and iSCSI Configuration Guide in Data ONTAP 8.1 and earlier</i>)
How to download and install SnapDrive for UNIX	The <i>SnapDrive for UNIX Installation and Administration Guide</i>
Configuring your host for the Host Utilities	The Product A-Z page

Related information

Product A-Z page: support.netapp.com/documentation/productsatoz/index.html
Interoperability Matrix: support.netapp.com/matrix

Copyright information

Copyright © 1994–2014 NetApp, Inc. All rights reserved. Printed in the U.S.

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

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

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

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

The product described in this manual may be protected by one or more U.S. 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 NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Customer Fitness, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, 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.

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

NetApp, Inc. NetCache is certified RealSystem compatible.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277