



## **Restoring a Snapshot copy**

Snapdrive for Unix

NetApp

June 20, 2025

# Table of Contents

Restoring a Snapshot copy .....	1
Command to use to restore Snapshot copies .....	1
Restoring Snapshot copies on a destination storage system .....	1
Restoring multiple storage entities .....	2
Considerations for restoring a Snapshot copy .....	2
Information required for using the snapdrive snap restore command .....	2
Restoring a Snapshot copy .....	5
Restoring a Snapshot copy from a different host .....	6
Volume-based SnapRestore .....	6
What volume-based SnapRestore is .....	6
Considerations for using volume-based SnapRestore .....	7
Mandatory checks for volume-based SnapRestore .....	7
Checks that can be overridden by the user .....	8
Volume-based SnapRestore command .....	8
Information about LUNs mapped to local or remote hosts .....	10
Host filespec information for a particular volume .....	10
Volume-based SnapRestore for space reservation .....	12

# Restoring a Snapshot copy

You can restore a Snapshot copy of a single storage entity or multiple storage entities.

## Command to use to restore Snapshot copies

Use the command `snapdrive snap restore` to restore a Snapshot copy.

The `snapdrive snap restore` command restores data from the Snapshot copy you specify at the command prompt to the storage system. This operation replaces the contents of the `file_spec` arguments (for example disk groups, NFS files, NFS directory trees, file systems created directly on LUNs) that you specified on the `snapdrive snap restore` command with the contents of the `file_spec` arguments located in the specified Snapshot copy.

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

Normally, you restore Snapshot copies from the host where you took the Snapshot copies (in other words, the originating host).

- In a NFS entity, when the same volume is mounted on two different IPs, `snap restore` operations can be performed only for one file specification at a time for these volumes.
- If you are a non-root user then, from SnapDrive 4.1 for UNIX, you should have storage write capability on global in addition to GlobalSDSnapshot capability for Snapshot restore to work.

## Restoring Snapshot copies on a destination storage system

You can restore Snapshot copies on the storage system from which it is created.

When you create a Snapshot copy on a destination storage system, the Snapshot copy is automatically replicated, from the source system, where it is created to the destination storage system. SnapDrive for UNIX allows you to restore the Snapshot copy on the source storage system. You can also restore the Snapshot copy on the destination storage system.

If you are performing a single file snap restore in a SnapMirror relationship, then the source and destination volume names should not be the same. SnapDrive for UNIX displays the following error message if the source and destination volume name are the same:

```
0001-636 Command error: Snapdrive cannot restore LUNs on SnapMirror
destination filer volumes: <filer-vol-name>
```

In SnapDrive 4.0 for UNIX and later, if Role Based Access Control is enabled, then you can perform snap restore on vFiler unit only when you have the Snapshot Restore capability on the vFiler unit.

### Related information

[Role-based access control in SnapDrive for UNIX](#)

## Restoring multiple storage entities

You can restore a Snapshot copy that contains multiple storage entities.

To restore a Snapshot copy that contains storage entities that reside on multiple destination storage systems, you must meet the following requirements:

- The storage entities you specify at the command prompt must reside on a single storage system, or on a HA pair.
- The name of the volume of the source storage system must match the name of the volume of the destination storage system.
- You must set the `snapmirror-dest-multiple-filervolumes-enabled` argument in the `snapdrive.conf` file to on.

You can use one command to restore storage entities that reside on a single storage system or on a HA pair.

## Considerations for restoring a Snapshot copy

Before restoring Snapshot copies, ensure that you are not using the file system, and that you do not interrupt the restore process.

Before restoring a Snapshot copy, consider the following important information:

- Ensure you are not in any directory on a file system that you want to restore. You can perform the `snapdrive snap restore` command from any directory except the one on a file system to which you want to restore the information.
- Do not interrupt the restore operation by entering Ctrl-C. Doing so could leave your system in an unusable configuration. If that happens, you might have to contact the NetApp technical support to recover.
- When exporting the NFS entities to a volume, set the Anonymous User ID option to “0” for the `snapdrive snap restore` command to work successfully.

## Information required for using the `snapdrive snap restore` command

To restore a Snapshot copy, determine which storage entity you want to restore, specify the name of the Snapshot copy, and so on.

The following table gives the information you need to supply when you use the `snapdrive snap restore` command.

Requirement/Option	Argument
Decide the type of storage entity that you wish to restore and enter that entity's name with the appropriate argument.	
<ul style="list-style-type: none"> <li>• If you specify a host volume or file system to be restored, the argument you give is translated to the disk group containing it. SnapDrive for UNIX then restores the entire disk group. SnapDrive for UNIX freezes any file systems in host volumes in those disk groups and takes a Snapshot copy of all storage system volumes containing LUNs in those disk groups.</li> <li>• If you specify a file specification that is an NFS mountpoint, the argument translates to a directory tree. SnapDrive for UNIX restores only the NFS directory tree or file. Within the directory tree, SnapDrive for UNIX deletes any new NFS files or directories that you created after you created the Snapshot copy. This ensures that the state of the restored directory tree will be the same as when the Snapshot copy of the tree was made.</li> <li>• If you restore a LUN, SnapDrive for UNIX restores the LUN you specify.</li> <li>• If you restore a file system that is created directly on a LUN, SnapDrive for UNIX restores the LUN and the file system.</li> <li>• If the Snapshot copy contains storage entities that span multiple storage system volumes, you can restore any of the entities in that Snapshot copy.</li> </ul>	
LUN (-lun <i>file_spec</i> )	<i>name of the LUN. You must include the name of the storage system, volume, and LUN.</i>
Disk group (-dg <i>file_spec</i> c) or volume group (-vg <i>file_spec</i> )	<i>name of the disk or volume group</i>
File system (-fs <i>file_spec</i> )	<i>name of the file system</i>
File (-file <i>file_spec</i> )	<i>name of the NFS file</i>
Host volume (-hostvol <i>file_spec</i> ) or logical volume (-lvol <i>file_spec</i> )	<i>name of the host or logical volume. You must supply both the requested volume and the disk group containing it; for example, -hostvol dg3/acct_volume.</i>

Requirement/Option	Argument
Specify the name for the Snapshot copy. If any of the <i>file_spec</i> arguments you supply at the command prompt currently exist on the local host, you can use a short form of the Snapshot copy name.	
If none of the <i>file_spec</i> arguments exist on the host, you must use a long form of the Snapshot copy name where you enter the storage system name, volume, and Snapshot copy name. If you use a long name for the Snapshot copy and the path name does not match the storage system and/or storage volume information at the command prompt, SnapDrive for UNIX fails. The following is an example of a long Snapshot copy name: big_filer:/vol/account_vol:snap_20031115	
Sometimes, the value supplied with the <i>file_spec</i> argument might not exist on the host. For example, if you had unmounted a file system or removed a disk group by exporting, deporting, or destroying, that file system or disk group could still be a value for the <i>file_spec</i> argument. It would, however, be considered a non-existent value. SnapDrive for UNIX can restore Snapshot copies for such a non-existent <i>file_spec</i> , but you must use the long Snapshot copy name.	
When you restore Snapshot copies that span multiple storage systems and volumes, and contain a nonexistent <i>file_spec</i> argument, SnapDrive for UNIX permits an inconsistency in the command line. It does not allow for existing <i>file_spec</i> arguments. If you want to restore only one storage entity from a multiple storage system Snapshot copy, the Snapshot copy you specify does not need to be on the same storage system as the storage system containing the storage entity.	
The short form of the same Snapshot copy name would omit the storage system and storage system volume name, so it would appear as: <i>snap_20031115</i>	
Snapshot copy name (-snapname)	<i>snap_name</i>
It can be either a short name, such as <i>mysnap1</i> , or a long name that includes the storage system name, volume, and Snapshot copy name.	
Generally, NetApp recommends that you use the short name. If the <i>file_spec</i> argument is non-existent: that is, it no longer exists on the host; see the explanation of the <i>file_spec</i> argument. Then you must use the long name for the Snapshot copy.	
-reserve   -noreserve	
Optional: If you want SnapDrive for UNIX to create a space reservation when you restore the Snapshot copy.	
-force	~
-noprompt	~

Requirement/Option	Argument
Optional: Decide if you want to overwrite an existing Snapshot copy. Without this option, this operation halts if you supply the name of an existing Snapshot copy. When you supply this option and specify the name of an existing Snapshot copy, it prompts you to confirm that you want to overwrite the Snapshot copy. To prevent SnapDrive for UNIX from displaying the prompt, include the <code>-noprompt</code> option also. (You must always include the <code>-force</code> option if you want to use the <code>-noprompt</code> option.)	
You must include the <code>-force</code> option at the command prompt if you attempt to restore a disk group where the configuration has changed since the last Snapshot copy. For example, if you changed the way data is striped on the disks since you took a Snapshot copy, you would need to include the <code>-force</code> option. Without the <code>-force</code> option, this operation fails. This option asks you to confirm that you want to continue the operation unless you include the <code>-noprompt</code> option with it.	
 If you added or deleted a LUN, the restore operation fails, even if you include the <code>-force</code> option.	
<code>mntopts</code>	<code>~</code>
<b>Optional:</b> If you are creating a file system, you can specify the following options:	
<ul style="list-style-type: none"> <li>Use <code>-mntopts</code> to specify options that you want to pass to the host mount command (for example, to specify host system logging behavior). The options you specify are stored in the host file system table file. Allowed options depend on the host file system type.</li> <li>The <code>-mntopts</code> argument is a file system <code>-type</code> option that is specified using the mount command <code>-o</code> flag. Do not include the <code>-o</code> flag in the <code>-mntopts</code> argument. For example, the sequence <code>-mntopts tmpplog</code> passes the string <code>-o tmpplog</code> to the <code>mount</code> command, and inserts the text <code>tmpplog</code> on a new command line.</li> </ul>	
 If you pass any invalid <code>-mntopts</code> options for storage and snap operations, SnapDrive for UNIX does not validate those invalid mount options.	

## Restoring a Snapshot copy

You can restore a Snapshot copy by using the `snapdrive snap restore` command. The restore operation can take several minutes, depending on the type, and amount of data being restored.

### Steps

1. Enter the following command to restore a Snapshot copy:

```
snapdrive snap restore -snapname snap_name {-lun | -dg | -fs | -hostvol | -file
} file_spec [file_spec...][{ -lun | -dg | -fs | -hostvol | -file } file_spec
file_spec ...] [-force [-noprompt]] [-mntopts options][{-reserve | -noreserve}]
```

SnapDrive for UNIX replaces the contents of the LUNs you specify in the `snapdrive snap restore` command line with the contents of the LUNs in the Snapshot copy you specify. This operation can take several

minutes. When the operation is complete, SnapDrive for UNIX displays a message similar to the following:

```
snap restore <filespec list> succeeded
```

In the following example, file system 1 (fs1) resides on storage system1, and file system 2 (fs2) resides on storage system1 and also on storage system 2, which is the partner storage system. File system 3 (fs3) resides on storage system1, partner storage system 2, and storage system3, which is not part of the HA pair. An additional file system, fs4, resides entirely on storage system 4.

The following command creates a Snapshot copy of fs1, fs2, fs3, and, fs4:

```
# snapdrive snap create -fs /mnt/fs1 /mnt/fs2 /mnt/fs3 /mnt/fs4  
-snapname fs_all_snap
```

The next command restores fs1 and fs2 on the destination storage system. Both fs1 and fs2 reside on a HA pair, so you can restore them with one command:

```
# snapdrive snap restore -fs /mnt/fs1 /mt/fs2 -snapname fs_all_snap
```

The following command restores fs4:

```
snapdrive snap restore -fs /mnt/fs4 -snapname fs_all_snap
```

SnapDrive for UNIX cannot restore fs3 on the destination storage system, because this file system resides on storage system1, storage system 2, and storage system 3.

## Restoring a Snapshot copy from a different host

Use the `snapdrive snap restore` command to restore a Snapshot copy from a different host.

Usually, you can restore a Snapshot copy from the host where you took the Snapshot copy. Occasionally, you might need to restore a Snapshot copy using a different or non-originating host. To restore a Snapshot copy using a non-originating host, use the same `snapdrive snap restore` command that you would normally use. If the Snapshot copy you restore contains NFS entities, the non-originating host must have permission to access the NFS directory.

## Volume-based SnapRestore

SnapDrive 4.0 for UNIX and later provides Snapshot restore capability at a volume level. This explains the various storage operations you can carry out using volume-based Snapshot restore.

### What volume-based SnapRestore is

Volume-based SnapRestore (VBSR) restores the volume with all its storage objects. The volume-based restore is faster than each storage object restored individually. VBSR also

works with FlexClone volumes and vFiler configurations. VBSR for vFiler is available for Data ONTAP 7.3 and later.

SnapDrive 3.0 for UNIX and earlier can restore LUNs for a host-side entity like file system, disk groups, and host volumes, or normal files created over NFS from an application consistent snapshot. SnapDrive for UNIX uses Single File Snap Restore (SFSR) implemented in Data ONTAP. SFSR works as following:

- For normal files while the SFSR is proceeding, any operation which tries to change the file is suspended until SFSR is complete.
- For LUNs, when SFSR is proceeding the LUN is available and I/Os (both reads and writes) are allowed. SFSR for normal files as well as LUNs may take a long time depending on the size of the LUN or the file being restored.

Therefore, for some environments SFSR is an obstruction.

SnapDrive 4.0 for UNIX and later enables you to take volume-based Snapshot copies. VBSR is much faster and requires less CPU and storage resources. It restores all the data on the active file system. This functionality can be used if a user wants to restore all LUNs or normal files on a volume from the same Snapshot copy.

Volume-based Snapshot copy should be used with caution because all Snapshot copies taken after volume Snapshot copy that is used for the restore operation are deleted. All the new files and new LUNs created on this volume must be deleted.

## **Considerations for using volume-based SnapRestore**

You can use volume-based SnapRestore (VBSR) with keeping few points in mind. Keeping these points in mind helps you to use the VBSR feature safely.

You must keep the following points in mind:

- VBSR reverts the entire volume to the state of the time when the Snapshot copy is created, which is used for VBSR. VBSR includes the following:
  - All the files and LUNs for the specified host filespec during `snap create` operation.
  - All the files and LUNs that are part of the application consistent Snapshot copy during `snap create` operation.
- VBSR removes all the newer files and LUNs created on the current volume after the Snapshot copy used for restore.
- VBSR removes all the newer Snapshot copies that are taken after the Snapshot copy which is being used for restore.
- It is recommended that you run `-vbsr preview` command before using `-vbsr execute` command.

## **Mandatory checks for volume-based SnapRestore**

Before volume based SnapRestore is carried out, SnapDrive for UNIX does some mandatory checks with the system. These checks are required so that volume based SnapRestore is used safely. The mandatory checks cannot be over-ridden by the user.

The following are the mandatory checks which SnapDrive for UNIX makes before implementing volume-based SnapRestore:

- Volume-based SnapRestore works only with Snapshots which are created using SnapDrive for UNIX.
- The volume in volume-based SnapRestore should not be a storage system's root volume.
- SnapDrive for UNIX checks for volume clones. It does not allow the volume restore operation if there are any volume clones from new Snapshot copies. This is a limitation imposed by Data ONTAP.
- The volume in volume-based SnapRestore should not have any mapped LUNs apart from the LUNs specified (raw LUN or LUNs present in file system, disk group or host volume) for restore.
- SnapDrive for UNIX checks whether the volume exists in a SnapMirror relationship.
- SnapDrive for UNIX checks whether the volume exists in a SnapVault relationship.

The SnapMirror and SnapVault checks can be overridden if SnapDrive for UNIX is using Operations Manager for RBAC and user has SD.SnapShot.DisruptBaseline capability on volume. For more information about the specific RBAC capability for user to override these checks, you can see Role Based Access Control in SnapDrive for UNIX.

## Related information

[Role-based access control in SnapDrive for UNIX](#)

## Checks that can be overridden by the user

Before volume-based SnapRestore is carried out, SnapDrive for UNIX performs some checks that a user can override using `-force` option. These checks are required so that volume-based SnapRestore is used safely.

It is recommended that you follow the various checks that the system performs, but you can override these checks by using the `-force` option.

You can override the following SnapDrive for UNIX for volume-based SnapRestore checks:

- SnapDrive for UNIX checks for LUNs in the Snapshot copy which are not crash consistent. If it finds an application inconsistent LUN in the Snapshot copy, it warns you about it.
- SnapDrive for UNIX checks whether there are additional LUNs in the active volume which were created after the Snapshot copy was taken. If SnapDrive for UNIX finds additional LUNs, it warns you that those additional LUNs in the active volume is lost.
- SnapDrive for UNIX checks for new Snapshot copies. These new Snapshot copies do not get restored and are lost.
- SnapDrive for UNIX checks for normal file (files that are visible from the local host) in a volume.
- SnapDrive for UNIX checks for NFS exports.
- SnapDrive for UNIX checks for CIFS shares.

If Snapshot copies were created using SnapDrive 3.0 for UNIX or earlier, volume-based SnapRestore preview is not able to perform the first two checks in the earlier mentioned list. If you have provided `-force` option, then during volume-based SnapRestore execute, a prompt message is displayed to over-ride these checks and proceed.

## Volume-based SnapRestore command

This section describes the commands and the options to use volume based

## SnapRestore.

-vbsr option is added in snap restore CLI to select volume based SnapRestore. Use the following command syntax to perform restore using volume based SnapRestore:

```
snapdrive snap restore {-lun | -dg | -vg | -hostvol |  
-lvol | -fs |-file} file_spec [file_spec ...] [{-lun | -dg |  
-vg | -hostvol | -lvol | -fs -file} file_spec [file_spec ...]  
... ] -snapname snap_name [-force [-noprompt]] [{-reserve |  
-noreserve}]  
[-vbsr [execute | preview]]
```

If no argument is supplied with the -vbsr, the default output is that of the preview option. A -verbose option is used which enables detailed output of all the mandatory checks that can be over-ridden by the user. The default output when -verbose option is not provided displays the results of the checks that is failed.

If you do not want any confirmation message to prompt, while executing volume based SnapRestore, you can use -noprompt and -force option with snap restore -vbsr execute command. The following table describes SnapDrive for UNIX behavior depending on the options provided by you.

S.N.	<b>-vbsr execute</b>	<b>-force</b>	<b>-noprompt</b>	<b>Result</b>
1.	No	NA	NA	Preview mode is the default mode. All the checks are done and report for each check is generated.
2.	Yes	No	No	All checks are done. If any mandatory checks that a user can override fails, SnapDrive for UNIX displays an error message.
3.	Yes	Yes	No	All the checks are done. If any mandatory checks fail, SnapDrive for UNIX displays an error message. If any check that a user can override fails, SnapDrive for UNIX prompts you.

S.N.	<b>-vbsr execute</b>	<b>-force</b>	<b>-noprompt</b>	<b>Result</b>
4.	Yes	Yes	Yes	All the checks are done. If any mandatory checks fail, SnapDrive for UNIX displays an error message. If any check that a user can override fails, SnapDrive for UNIX does not prompt you.

## Information about LUNs mapped to local or remote hosts

The only mapping information available from the volume to SnapDrive for UNIX during volume based SnapRestore is the initiator group information for a LUN. If the initiator groups used by you are always created by SnapDrive for UNIX, then the fully qualified domain name of the host is part of the initiator group name.

If SnapDrive for UNIX administrator specifies the **-igroup** CLI option or if you use manually created initiator groups, then the **igroup** name need not necessarily have the host name. For all the earlier reasons, SnapDrive for UNIX cannot reliably detect local or remote bindings for a LUN. Therefore, SnapDrive for UNIX displays the full LUN initiator group and initiator information as part of the volume based SnapRestore.

## Host filespec information for a particular volume

SnapDrive for UNIX as part of volume restore preview report displays the LUN mapping information. This displayed information is relevant for the checks and the normal files are reverted. Finding out all the host filespecs based on LUNs on a particular volume is a time consuming process and slows down the volume restore process.

If you want to know that the host filespec information for the local host mapped to a particular storage system volume, you can use `snapdrive storage show -filervol <full-volume-name>`. An example of this is shown in the following.

```
#snapdrive storage show -filervol bart:/vol/volusecase2
```

Connected LUNs and devices:

device lun path	filename	adapter path	size	proto	state	clone
	backing	snapshot	-----	-----	-----	-----
/dev/sdg		- P	100m	iscsi	online	No
bart:/vol/volusecase2/lun5		-				

Host devices and file systems:

dg: vbsrfs_1_SdDg	dgtype lvm
hostvol: /dev/mapper/vbsrfs_1_SdDg-vbsrfs_1_SdHv	state: AVAIL
fs: /dev/mapper/vbsrfs_1_SdDg-vbsrfs_1_SdHv	mount point: /mnt/vbsrfs_1
(persistent)	fstype jfs2

device lun path	filename	adapter path	size	proto	state	clone	lun
	backing	snapshot	-----	-----	-----	-----	-----
/dev/sdbe		- P	100m	iscsi	online	No	
bart:/vol/volusecase1/lun9_0		-					
/dev/sdbf		- P	100m	iscsi	online	No	
bart:/vol/volusecase2/lun4_0		-					

raw device: /dev/sdbr1 mount point: /mnt/fs11 (persistent) fstype jfs2

device lun path	filename	adapter path	size	proto	state	clone	lun
	backing	snapshot	-----	-----	-----	-----	-----
/dev/sdbr		- P	200m	iscsi	online	No	
bart:/vol/volusecase2/fs11_SdLun		-					

NFS device: bart:/vol/volusecase1 mount point: /mnt/volusecase1 (non-persistent)

LUNs not connected to this host:

lun path	size	state
-----	-----	-----
bart:/vol/volusecase2/lunotherhost	20m	online

## Volume-based SnapRestore for space reservation

For space reservation with volume based snap reserve you need to set *space-reservations-volume-enabled* option in the `snapdrive.conf` file.

The *space-reservations-volume-enabled* option is used to set space guarantee policy on volume and can take the following three values.

- **Snapshot:** This is the default value. Space guarantee on the volume does not change.
- **volume:** Space guarantee on the volume is at the volume level.
- **none:** Space guarantee is set as none.

Following table describes the behavior of volume-based snap reserve for space reservation.

No space reserve CLI option used; <code>-vbsr</code> execute is specified	<code>space-reservations-volume-enabled=</code>	Result
none	snapshot	Space guarantee on the volume does not change.
none	none	Attempt to set space guarantee as "none" for the volumes.
<code>-reserve</code>	configuration value is over-ridden	Attempt to set space guarantee for volumes as "volume"
<code>-noreserve</code>	configuration value is over-ridden	Attempt to set space guarantee for volumes as "none"
none	volume	Attempt to set space guarantee for volumes as "volume"



`-vbsr` preview does not check for any of the space reservation options.

## Copyright information

Copyright © 2025 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.

**LIMITED RIGHTS LEGEND:** Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

## Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.