



Product overview

Snapdrive for Unix

NetApp

February 12, 2024

This PDF was generated from https://docs.netapp.com/us-en/snapdrive-unix/linux-administration/concept_clustered_data_ontap_features_available_in_snapdrive_for_unix.html on February 12, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Product overview 1
 - SnapDrive highlights 1
 - Clustered Data ONTAP features available in SnapDrive for UNIX 1
 - Support for Vserver 3
 - Support for NetApp DataMotion for vFiler 3
 - Support for volume migration 4
 - vFiler unit support issues 4
 - Considerations for using SnapDrive for UNIX 5
 - Management of an LVM and raw entities 5
 - How to manage Snapshot copies of LVM, raw devices, and NFS entities 6
 - Security considerations 6
 - Access permissions on a storage system 7
 - Stack requirements 7
 - Supported FC, iSCSI, or NFS configurations 9
 - Limitations 9
 - Limitations of LUNs managed by SnapDrive 10
 - Limitations of RDM LUNs managed by SnapDrive 11
 - SnapDrive limitations for clustered Data ONTAP 12
 - Limitations for NFS files or directory trees 12
 - Thin provisioning in SnapDrive for UNIX 12
 - Volume managers on SnapDrive for UNIX 13

Product overview

SnapDrive for UNIX is a host-based storage and data management solution for UNIX environments. SnapDrive for UNIX simplifies data management, improves data availability, and increases the reliability of application data through simplified storage provisioning and file system consistent Snapshot copies.

You can use SnapDrive for UNIX as a stand-alone product or along with other SnapManager products that run on UNIX. SnapDrive for UNIX deployed along with SnapManager enables you to create application-consistent data backup.

SnapDrive highlights

SnapDrive for UNIX enables you to simplify storage management tasks. Some of the tasks that you can perform using SnapDrive for UNIX are:

- Creating Snapshot copies using clustered Data ONTAPNetApp Snapshot technology.

SnapDrive for UNIX enables you to create and delete Snapshot copies of a file system, volume group, host volume, or LUN and to clone storage entities from Snapshot copies.

- Provisioning and managing storage with little or no application or host downtime.

SnapDrive for UNIX enables you to create and delete storage components, including disk groups, host volumes, file systems, and LUNs on host cluster environments. SnapDrive for UNIX also enables you to expand storage, connect storage to a host, and disconnect storage from a host.

- Providing role-based access control, which enables storage administrators to restrict the operations that SnapDrive users can perform, depending on their assigned roles.

In addition to these major features, SnapDrive offers the following advantages:

- Enhanced clone-split operation
- Volume-based SnapRestore (VBSR)
- Command-line interface (CLI) options
- Data collection utility
- Support for storage system renaming
- Support for Fibre Channel RDM LUNs in Linux guest operating systems

Clustered Data ONTAP features available in SnapDrive for UNIX

SnapDrive for UNIX 5.0 and later versions support some features of clustered Data ONTAP.

The following features are supported:

- Port sets

- Vserver
- Handling errors caused by volume migration using the appropriate configuration variables
- Crash-consistent (consistency groups) Snapshot copies
- AutoSupport
- Raw Device Mapping (RDM) logical unit number (LUN) for Vservers
- Role-based access control (RBAC) with Operations Manager console
- Load-sharing mirrors (LSM)

A load-sharing mirror reduces the network traffic to a FlexVol volume by providing additional read-only access to the clients. You can create and manage the load-sharing mirrors to distribute read-only traffic away from a FlexVol volume. Load-sharing mirrors do not support Infinite Volumes. A set of load-sharing mirrors consists of a source volume that can connect with one or more destination volumes. Each load-sharing mirror in the set must belong to the same Storage Virtual Machine (SVM) as the source volume of the set. The load-sharing mirrors should also be created on different aggregates and accessed by different nodes in the cluster to achieve proper load balancing of client requests. For more information, see the [Clustered Data ONTAP Logical Storage Management Guide](#).



- LSM is supported with clustered Data ONTAP 8.2 and higher version of ONTAP. But you will have to configure cluster administrator if you are using clustered Data ONTAP 8.2.2 and below version using `snapdrive config set-cserver` command.
- LSM is not supported on any versions earlier than ONTAP 8.2 even if the cluster administrator is configured.

- IPspaces

An IPspace defines a distinct IP address space in which Storage Virtual Machines (SVMs) can be created. A routing table is maintained for each SVM within an IPspace; no cross-SVM or cross-IPspace traffic routing occurs. For more information, see the [Clustered Data ONTAP Network Management Guide](#).



This feature is supported only if you are using clustered Data ONTAP 8.3.

- MetroCluster setup

This feature is supported only if you are using clustered Data ONTAP 8.3. For information about MetroCluster setup, see the [Clustered Data ONTAP MetroCluster Installation and Configuration Guide](#).

- Tokens to reserve space on node to create FlexClone files and FlexClone LUNs

You must have the FlexClone license if you want to use this feature. For more information, see the [Clustered Data ONTAP Logical Storage Management Guide](#).

- Custom qtree exports

The qtree can have different exportfs policy than the parent volume.



This feature is supported only if you are using clustered Data ONTAP 8.2.1 and later.

- Federal Information Processing Standard (FIPS) compliance

Related information

[Support for storage system rename](#)

[Using port set in SnapDrive for UNIX](#)

[Volume migration in SnapDrive for UNIX](#)

[Support for Vserver](#)

Support for Vserver

SnapDrive for UNIX supports Vserver. Vserver is a secure virtual storage server that supports multiple protocols and unified storage. A Vserver contains data volumes and one or more LIFs, which it uses to serve data to clients.

The Vserver securely isolates the shared virtualized data storage and network, and appears as a single dedicated server to the clients. Each Vserver has a separate administrator authentication domain and can be managed independently by a Vserver administrator.

The volumes of each Vserver are related through junctions and are mounted on junction paths. The file system of each volume appears to be mounted at the junctions. The root volume of the Vserver is found at the top level of the namespace hierarchy; additional volumes are mounted to the Vserver's root volume to extend the global namespace. The Vserver's data volumes contain files and LUNs.

- You can use SnapDrive for UNIX to perform storage provisioning, Snapshot copy operations, and configuration operations on a Vserver.
- Application data is not stored in the root volume of the Vserver.
- If the root of the Vserver is a qtree, Snapshot copy operations are not supported.
- Every volume that is created on the Vserver must be mounted on a junction path.

Related information

[Configuration information for Vserver](#)

[Verifying login information for Vserver](#)

[Specifying login information for the Vserver](#)

[Deleting a user from a Vserver](#)

[ONTAP 9 Software Setup Guide](#)

Support for NetApp DataMotion for vFiler

SnapDrive for UNIX supports DataMotion for vFiler. When you perform DataMotion for vFiler, few SnapDrive for UNIX operations might fail during the cutover phase.



If the SnapDrive for UNIX operations fail during the cutover phase of vFiler migration, you must perform the SnapDrive for UNIX operations after the DataMotion for vFiler operations are completed.

You must set the `datamotion-cutover-wait` variable in the `snapdrive.conf` file to execute SnapDrive

for UNIX operations.



If you run a volume-based `snap restore` command during the cutover phase of a vFiler migration, the `snap restore` operation might fail. After the migration is complete and the vFiler is available at the destination, running a volume-based `snap restore` operation places the LUNs offline. You must manually put the LUNs back online.

Support for volume migration

SnapDrive for UNIX supports volume migration, which enables you to nondisruptively move a volume from one aggregate to another within the same controller for capacity utilization, improved performance, and to satisfy service-level agreements. In a SAN environment, FlexVol volumes and the LUNs in the volumes are moved nondisruptively from one aggregate to another.

You must set the `volmove-cutover-retry` and the `volmove-cutover-retry-sleep` variables in the `snapdrive.conf` file to execute SnapDrive operations.

Related information

[ONTAP 9 Logical Storage Management Guide](#)

Volume migration in SnapDrive for UNIX

You can execute SnapDrive operations during volume migration.

The volume migration consists of the following three phases:

- Setup
- Move
- Cutover

SnapDrive for UNIX operations function smoothly in setup and move phases.

When you execute any SnapDrive commands during the cutover phase, SnapDrive for UNIX can retry the operation as defined in the variables `volmove-cutover-retry` and `volmove-cutover-retry-sleep` in the `snapdrive.conf` file.



If SnapDrive for UNIX operations fail during the volume migration, you must perform the SnapDrive for UNIX operations after the volume migration operations are completed.

vFiler unit support issues

SnapDrive for UNIX supports SnapDrive operations on a vFiler unit created on a FlexVol volume. However, SnapDrive for UNIX does not support vFiler units if you are using Fibre Channel (FC).

You must be aware of some considerations related to SnapDrive for UNIX supporting vFiler units:

- SnapDrive operations are not supported on a vFiler unit that was created on a qtree.

These operations are allowed if the vFiler unit owns the entire storage volume.

- When configuring SnapDrive to support vFiler units, you must ensure that the management and data paths are not configured for an interface on vFiler0.
- In Data ONTAP operating in 7-Mode, you must ensure that the Data ONTAP configuration variable `vfiler.vol_clone_zapi_allow` is set to `on` to connect to a Snapshot copy for a volume or LUN in a vFiler unit.

Considerations for using SnapDrive for UNIX

You must be aware of various considerations for using SnapDrive for UNIX.

- You must use the default value for the space reservation setting for any LUN managed by SnapDrive for UNIX.
- In FC and iSCSI configurations, set snap reserve on the storage system to zero percent for each volume.
- Place all LUNs connected to the same host on a dedicated storage system volume accessible only by that host.
- If you use Snapshot copies, you cannot use the entire space on a storage system volume to store your LUNs.

The storage system volume that hosts the LUNs should be at least twice the combined size of all the LUNs on the storage system volume.

- Data ONTAP uses `/vol/vol0` (root volume) to administer the storage system.

Do not use this volume to store data. If you have configured any other volume (other than `/vol/vol0`) as root volume to administer the storage system, do not use it to store data.

Management of an LVM and raw entities

SnapDrive for UNIX enables you to manage LVM (Logical Volume Manager) and raw entities. SnapDrive for UNIX also provides commands that help in provisioning and managing storage when you create storage entities.

SnapDrive for UNIX provisioning in an LVM environment

SnapDrive for UNIX storage commands provision LVM entities by creating LVM objects.

If you request a SnapDrive for UNIX storage operation that provisions an LVM entity—for example, a disk group that includes host volumes or file systems—the `snapdrive storage` command works with the LVM to create the LVM objects and file systems that use the storage.

During storage provisioning operation, the following actions occur:

- The host LVM combines LUNs from a storage system into disks, or volume groups.

The storage is then divided into logical volumes, which are used as if they were raw disk devices to hold file systems or raw data.

- SnapDrive for UNIX integrates with the host LVM to determine which NetApp LUNs make up each disk group, host volume, and file system requested for a Snapshot copy.

Because the data from any specified host volume can be distributed across all disks in the disk group, Snapshot copies can be made and restored only for whole disk groups.

SnapDrive for UNIX and raw entities

SnapDrive for UNIX enables storage operation for a raw entity such as a LUN, or a file system that can be created directly on a LUN, and performs the storage operation without using the host system, LVM.

SnapDrive for UNIX storage commands manage raw entities such as LUNs without activating the LVM. SnapDrive for UNIX enables you to create, delete, connect, and disconnect LUNs, and the file systems that they contain, without activating the LVM.

How to manage Snapshot copies of LVM, raw devices, and NFS entities

You can use SnapDrive commands to create, restore, and manage Snapshot copies of LVM, raw devices and NFS entities.

You must run the commands on the host to create, restore, and manage Snapshot copies of storage entities.

- Volume manager entities

The volume manager entities are disk groups with host volumes and file systems that you created using the host volume manager.

- Raw entities

The raw entities are either LUNs or LUNs that contain file systems without creating any volumes or disk groups and are mapped directly to the host.

- NFS entities

The NFS entities are NFS files and directory trees.

The Snapshot copy that you create can exist on multiple storage systems and storage system volumes. SnapDrive checks the read or write privilege against the storage entities in the Snapshot copy to ensure that all Snapshot copy data is crash-consistent. SnapDrive does not create a Snapshot copy unless the data is crash-consistent.

Security considerations

You can enable SnapDrive for UNIX to access the storage systems connected to the host, and must configure the host to use the login names and passwords assigned to the storage systems. If you do not provide this information, SnapDrive for UNIX cannot communicate with the storage system.

A root user can allow other users to run specific commands, depending on the roles assigned to them. You need not be a root user to perform the storage and snap management operations.

SnapDrive for UNIX stores user authentication information about the host in an encrypted file. On Linux hosts, by default, SnapDrive encrypts the password information and sends it across the network. SnapDrive for UNIX communicates using HTTPS over the standard IP connection.

Access permissions on a storage system

Access permissions indicate whether a host can perform certain Snapshot copy and storage operations. Access permissions do not affect any of the SnapDrive storage show or storage list operations. SnapDrive enables you to specify the access permissions for each host in a file that resides on the storage system.

You can also specify the action that SnapDrive must take when it does not find a permission file for a specified host. You can specify the action by setting the value in the `snapdrive.conf` configuration file for `all-access-if-rbac-unspecified`. You can also enable or disable the access to the storage system by modifying the access permissions.



In SnapDrive 4.0 for UNIX and later, you can perform storage operations depending on the role-based access control capabilities.

Stack requirements

SnapDrive for UNIX requires a host operating system, host file systems, NFS, volume managers, FC or iSCSI Host Utilities, storage system licenses, ONTAP software, MultiStore software, and Internet Protocol (IP) access. SnapDrive for UNIX also has certain stack requirements which it must fulfill.

Host-side entities

The following is a list of host-side entities:

- The host operating system
- A volume manager
- File system
- Linux Host Utilities

Guest-side entities for RDM LUN support

The following is a list of guest-side entities:

- The guest operating system
- A volume manager
- File system
- Linux iSCSI Host Utility alone is adequate, if the protocol is iSCSI

SnapDrive for UNIX stack

You must enter acceptable values for the *multipathing-type*, *fstype*, *default-transport*, and *vmtype* variables in the *snapdrive.conf* file as provided in the matrix stack. You must verify that the values entered are installed and running in your host system.

Host platform	Default transport type	Multipathing type	fstype	vmtype
Linux	FCP	none	ext4	lvm
	iSCSI	none	ext4	lvm
	FCP	nativempio	ext4	lvm
	iSCSI	nativempio	ext4	lvm
	FCP	none	ext3	lvm
	iSCSI	none	ext3	lvm
	FCP	nativempio	ext3	lvm
	iSCSI	nativempio	ext3	lvm



SnapDrive for UNIX does not support the Ext2 file system.

- If you have FCP and iSCSI storage stacks, SnapDrive will only support FCP storage stacks.

SnapDrive does not support iSCSI storage stacks for AIX.

- The host operating system and appropriate patches for Linux are installed .
- The volume manager for Linux is LVM2 .
- Host Utilities are installed in Linux
- ONTAP software is installed on your storage system.
- MultiStore software is installed on your storage system for a vFiler unit setup.
- Internet Protocol (IP) access is available between the host and storage system.

NetApp modifies host utilities and components on an ongoing basis. You can track this changes by using the Interoperability Matrix, which contains up-to-date information for using NetApp products in a SAN environment.

The storage system license and the MultiStore license constitute the storage system entities.

Storage system licensing requirements

- An FC, iSCSI, or NFS license, depending on your configuration
- A FlexClone license

- A SnapRestore license on the storage system

Related information

[Unable to select a storage stack](#)

[SnapDrive configuration wizard](#)

[NetApp Interoperability](#)

Supported FC, iSCSI, or NFS configurations

SnapDrive for UNIX supports host cluster and HA pair topologies. FC or iSCSI configurations support the same host cluster and HA pair configurations that the FC Host Utilities or iSCSI Host Utilities supports.

SnapDrive for UNIX supports the following host cluster and HA pair topologies:

- A stand-alone configuration in which a single host is connected to a single storage system
- Any topology that involves HA pair failover of a storage system
- Any topology that has host clusters supported by NetApp

For more information about the recommended configurations for your host and the storage systems you are using, see the Linux Host Utilities documentation.



If you need a SnapDrive for UNIX configuration that is not mentioned in the utilities documentation, contact technical support.

Limitations

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

Generic limitations

- SnapDrive for UNIX requires VMs to use BIOS during boot up for SnapManager for Virtual Infrastructure (SMVI) support. Using the Unified Extensible Firmware Interface (UEFI) is not supported.
- SnapDrive for UNIX does not support MetroCluster configuration in an RDM (Raw Device Mapping) environment in a guest operating system because the MetroCluster configuration is not supported by Virtual Storage Console (VSC).
- SnapDrive for UNIX does not support Snapshot operations on an NFS mount point when the volume is exported with Kerberos security authentication types krb5, krb5i, or krb5p.
- Snapshot operations might be inconsistent if you perform a snap restore operation on a mount point where a different entity other than the one created in the Snapshot copy is mounted.
- SnapDrive for UNIX does not support operations on file specifications or LUNs if they are located across Data ONTAP operating in 7-Mode and clustered Data ONTAP storage systems.
- If you are using clustered DATA ONTAP to configure a storage virtual machine (SVM) with SnapDrive for UNIX, check that the IP address of the SVM's LIF is mapped to the SVM name either in the DNS or in the

/etc/hosts file.

You must also check that the SVM name is configured in SnapDrive for UNIX by using the `snapdrive config setvsadminVserver name` command.

- SnapDrive for UNIX modifies the mount point permissions from a nonroot user to a root user for a qtree after VBSR operations.
- SnapDrive for UNIX does not support non-English language environments.
- The snap restore operation fails if it is restored from the Snapshot copy that was created before the LUNs were moved to another volume.
- If you are using ONTAP 8.2 or later, the Snapshot operations on a Snapshot copy might fail if a cloning operation that is in progress uses the same Snapshot copy.

You must retry the operation later.

- OnCommand Unified Manager 6.0 or later does not support Protection Manager on clustered Data ONTAP. As a result, the integration between OnCommand Unified Manager 6.0 or later and Snapdrive for UNIX is not supported, and the following SnapDrive for UNIX features are not supported:
 - Role-based access control (RBAC) integration with OnCommand Unified Manager 6.0 or later on ONTAP
 - Protection Manager integration with OnCommand Unified Manager 6.0 or later on ONTAP
- You must assign the aggregates that contain SVM volumes to the SVM's aggregate list to configure the SVM and execute SnapDrive for UNIX operations.
- SnapDrive for UNIX does not support automount, using any type of automount might cause SnapDrive for UNIX operations to fail.

Limitations on Linux

- In a multipath environment, the `snapdrive snap list-v` or `snapdrive snap show-v` operation does not show the development path in a raw LUN, and the `snapdrive storage show-all` operation does not show a raw device and mount point while in case of host LVM not been involved.
- SnapDrive for UNIX does not show the complete alias name if the alias name has the special character " - `". SnapDrive for UNIX supports only " _ " special character in alias name.
- The `snapdrive config prepare luns` command is not supported in the RDM LUN environment.
- The `snapdrive lun fixpaths` command is not supported in a guest operating system.

Limitations of LUNs managed by SnapDrive

When working with SnapDrive, you must be aware of the limitations related to LUNs.

- A LUN managed by SnapDrive cannot serve either as a boot disk or a system disk.
- The Linux hosts have operating system limits on how many LUNs you can create.

You can run the `snapdrive config check luns` command when you create the LUNs on these hosts. This command enables you to determine how many LUNs you can create.

- SnapDrive does not support the colon (:) in the long forms of the names for LUNs and Snapshot copies.

The colon is allowed between the components of a long Snapshot copy name or between the storage system name and storage system volume name of a LUN. For example, `toaster:/vol/vol1:snap1` is a typical long Snapshot name, while `toaster:/vol/vol1/lunA` is a typical long LUN name.

Limitations of RDM LUNs managed by SnapDrive

SnapDrive has a few limitations for provisioning RDM LUNs. You must be aware of the limitations that might affect your environment.

- An RDM LUN cannot serve either as a boot disk or system disk.
- SnapDrive does not support RDM LUNs in Internet Small Computer System Interface (iSCSI) environments.
- SnapDrive does not support MPIO in the guest operating system, although VMware ESX server supports MPIO.
- When the transport protocol is *FC*, the igroup that is specified in the CLI command is ignored by SnapDrive, and the igroup is automatically created by the virtual interface.
- You can rename, move, or delete the `/usr/bin/rescan-scsi-bus.sh` script that is bundled as part of `sg3_utils` to avoid limiting the number of RDM LUNs to eight.



If you want to retain `/usr/bin/rescan-scsi-bus.sh` in `sg3_utils` and avoid limiting the number of RDM LUNs to eight, then you must create a wrapper script `/root/dynamic-lun-rescan.sh` and from that script run `/usr/bin/rescan-scsi-bus.sh`, with the options `-w`, `-c`, and `-r` and assign full permissions.

The following is an example of the modified content of `/root/dynamic-lun-rescan.sh`:

```
#cat /root/dynamic-lun-rescan.sh
#Wrapper script used to call the actual rescan script.
/usr/bin/rescan-scsi-bus.sh -w -c -r
```

Limitations related to VMware ESX server

- Each guest operating system can be configured with four SCSI controllers, and each SCSI controller can be mapped to 16 devices.

However, one device is reserved per controller, and therefore a total of 60 ($16 * 4 - 4$) RDM LUNs can be mapped to the guest operating system.

- Each ESX server can be mapped to a maximum of 256 RDM LUNs.

Related information

[VMware VMotion support in SnapDrive for UNIX](#)

[Configuring Virtual Storage Console for SnapDrive for UNIX](#)

[Considerations for provisioning RDM LUNs](#)

SnapDrive limitations for clustered Data ONTAP

Some SnapDrive features and operations are not supported for clustered Data ONTAP.

- Storage provisioning and Snapshot management operations are not supported on the Vserver's root volume.
- Operations on file specification or LUNs that span Data ONTAP operating in 7-Mode and storage systems running on clustered Data ONTAP are not supported.
- The symbolic links that reside inside a mount point are not supported.
- Aggregates that contain Vserver volumes must be assigned to the Vserver's aggregate list to configure the Vserver and execute SnapDrive operations.
- SnapDrive for UNIX does not support migration of the Vserver because it is not supported by clustered Data ONTAP.
- Snap connect operation with the `-readonly` option mounts only the `.snapshot` directory on the host and does not create the clone volume.

In clustered Data ONTAP, the export policy is defined only at the volume level and not at the directory level. Therefore, the `.snapshot` directory cannot be exported to the secondary host (the host that does not have the export permission on the parent volume).

- SnapDrive for UNIX does not support volume-based SnapRestore (VBSR) operations from a Snapshot copy created prior to the base Snapshot copy, when the volume is in a SnapMirror or SnapVault operation.
- SnapDrive for UNIX does not support the `snapdelete` operation of the Snapshot copy created before SnapVault update, when the volume is in SnapVault (XDP) relationship.

Limitations for NFS files or directory trees

SnapDrive does not provide storage provisioning commands for NFS files or directory trees. SnapDrive supports `snapdrive snap create` and `snapdrive snap restore` commands only if you are using Data ONTAP 7.3.5 and later.

The `snapdrive snap connect` and `snapdrive snap disconnect` commands involve NFS and use the Data ONTAP FlexVol volumes feature for read and write access. Therefore you can run these commands only if you are using Data ONTAP 7.3.5 or later. The configurations with Data ONTAP 7.3.5 or later and traditional volumes enable you to create and restore Snapshot copies, but the Snapshot connect operation is restricted to only read-only access.

Thin provisioning in SnapDrive for UNIX

You cannot set the fractional reserve value, and there is no integration with Data ONTAP capabilities such as `autodelete` and `autosize` in SnapDrive for UNIX. Although you can safely use the Data ONTAP capabilities with SnapDrive for UNIX, SnapDrive for UNIX does not register `autodelete` or `autosize` events.

Volume managers on SnapDrive for UNIX

On Linux, the volume manager is Native LVM2.

The following table describes the volume managers on your host platform:

Host	Volume manager	Volume or disk groups	Location of logical volumes	Location of multipathing devices
RHEL 4 and RHEL 5	Native LVM2	Volume groups <i>vg</i>	/dev/mapper/dgname-lvolname	/dev/mpath /dev/mapper /dev/dm
RHEL 6	Native LVM2	Volume groups <i>vg</i>	/dev/mapper/dgname-lvolname	/dev/mapper /dev/dm
SLES 10 and SLES 11	Native LVM2	Volume groups <i>vg</i>	/dev/mapper/dgname-volname	/dev/mapper /dev/dm

Copyright information

Copyright © 2024 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.