



SnapDrive for UNIX overview

Snapdrive for Unix

NetApp
March 24, 2021

Table of Contents

- SnapDrive for UNIX overview 1
 - What SnapDrive for UNIX does 1
 - Features supported in SnapDrive for UNIX 1
 - Operations performed on stand-alone UNIX systems 2
 - Support for Vserver 2
 - vFiler unit support issues 3
 - Considerations for using SnapDrive for UNIX 3
 - Management of an LVM and raw entities 4
 - How to manage Snapshot copies of LVM, raw devices, and NFS entities 5
 - Security considerations 5
 - Access permissions on a storage system 5
 - Requirements for storage systems 6
 - Stack requirements 7
 - Supported FC, iSCSI, or NFS configurations 9
 - Limitations 9
 - Limitations of LUNs managed by SnapDrive 10
 - Support for NetAppDataMotion for vFiler 10
 - Support for volume migration 11
 - What configuration checker is 12
 - Limitations for NFS files or directory trees 12
 - Multipathing in SnapDrive for UNIX 12
 - Thin provisioning in SnapDrive for UNIX 12
 - Volume managers on SnapDrive for UNIX 12

SnapDrive for UNIX overview

SnapDrive for UNIX simplifies data management and increases the availability and reliability of application data through simplified storage provisioning and file system consistent Snapshot copies.

SnapDrive for UNIX simplifies data backup so that you can recover data even if it is deleted or modified. SnapDrive for UNIX uses Snapshot technology to create an image of the data on a shared or unshared storage system attached to a UNIX host. When you restore a Snapshot copy, it replaces the current data on the storage system with the data in the Snapshot copy.

SnapDrive for UNIX provides storage features that enable you to manage the entire storage hierarchy, such as the host-side application-visible file, the volume manager, and the storage-system-side logical unit number (LUN).

SnapDrive for UNIX provides support for role-based access control. A storage administrator uses RBAC to restrict a user's access to the storage system based on the role and task that the user performs.



You must use Operations Manager console 3.7 or later to use RBAC.

What SnapDrive for UNIX does

SnapDrive for UNIX simplifies the storage management tasks for you. SnapDrive for UNIX uses Snapshot technology to create an image of data stored on a shared or unshared storage system. SnapDrive also helps you with storage provisioning.

You can use SnapDrive for UNIX for performing the following tasks:

- **Back up and restore data:** SnapDrive for UNIX enables you to use Snapshot technology to create an image (Snapshot copy) of the host data that is stored on a storage system. This Snapshot copy provides you with a copy of that data, which you can restore later. The data in the Snapshot copy can exist on one storage system or span multiple storage systems and their volumes. These storage systems can be in HA pair or node-local file systems or disk groups, or LUNs in a host cluster environment.
- **Manage storage:** You can create and delete storage components, including disk groups, host volumes, file systems, and LUNs on host cluster and non-cluster environments. SnapDrive for UNIX enables you to manage this storage by expanding it, connecting it to a host, and disconnecting it.
- **Role-based access control:** SnapDrive for UNIX provides role-based access control (RBAC). RBAC allows a SnapDrive administrator to restrict access to a storage system for various SnapDrive operations. This access for storage operations depends on the role that is assigned to the user. RBAC allows the storage administrators to limit the operations that SnapDrive users can perform depending on their assigned roles.

Features supported in SnapDrive for UNIX

You can find information about the features and supported configurations in SnapDrive for UNIX.

SnapDrive for Unix has the following capabilities and supports the following features:

- SnapDrive for UNIX configuration checker
- Role-based access control (RBAC) permissions
- SnapDrive for UNIX wizard
- Enhanced clone-split operation
- Creation and usage of Snapshot copies
- Volume-based SnapRestore (VBSR)
- Command-line interface (CLI) options
- Data collection utility
- Modified system name support for a storage system in SnapDrive operations
- Transparent handling of errors caused by volume migration and NetAppDataMotion for vFiler; controls the same using the appropriate configuration variables



The latest information about SnapDrive for UNIX and its requirements is available in the Interoperability Matrix.

Related information

[NetApp Interoperability](#)

Operations performed on stand-alone UNIX systems

SnapDrive for UNIX on stand-alone UNIX systems enable you to create storage and manage Snapshot copy of LUNs, file systems, logical volumes, and disk groups.

- Create storage that include LUNs, file systems, logical volumes, and disk groups.

After creating the storage, you can increase or reduce the storage capacity, connect the storage to a host or disconnect it, and display configuration information about the storage.

- Create a Snapshot copy of one or more volume groups on a storage system.

The Snapshot copy can contain file systems, logical volumes, disk groups, LUNs, and NFS directory trees. After you create a Snapshot copy, you can rename, restore, or delete the Snapshot copy. You can also connect a Snapshot copy to a different location on the same or different host. After you connect the Snapshot copy, you can view and modify the content, or disconnect. You can also display information about Snapshot copies that you create.

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](#)

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.

Related information

[Data ONTAP 8.2 MultiStore Management Guide For 7-Mode](#)

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 AIX 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.

Requirements for storage systems

You must consider the storage system requirements before setting it up.

Component	Requirement
Operating system	<p>Data ONTAP 7.3.5 or later.</p> <ul style="list-style-type: none">• SnapDrive for UNIX supports FlexVol volumes but does not use all FlexVol volume features.• Configurations that use NFS must use Data ONTAP 7.3.5 or later and FlexVol volumes must use <code>snapdrive snap connect</code> to read and write to a connected NFS file or directory tree. <p>Configurations with traditional volumes are provided with read-only access to NFS files and directory trees.</p>
Storage system setup	<p>You must specify the partner IP address in the HA pair that can be used if a storage system failover occurs.</p> <div data-bbox="867 1192 932 1257"></div> <p>You specify the IP address when you run the setup program on the storage system.</p>

Component	Requirement
Licenses	<ul style="list-style-type: none"> • FC, iSCSI, or NFS, depending on the host platform • FlexClone license <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;">  <p>You must have appropriate protocols running on the storage system for SnapDrive for UNIX to execute.</p> </div> <ul style="list-style-type: none"> • SnapRestore software • MultiStore software <p>You should set the SnapRestore and MultiStore licenses when you set up the storage system. You need a MultiStore license if you want to set up a vFiler environment.</p> <ul style="list-style-type: none"> • Secure HTTP access to the storage system.

SnapDrive for UNIX operations are not case-sensitive with respect to storage system host name, you must ensure that the IP addresses are unique when you configure the storage system.



For the latest SnapDrive requirements, see the Interoperability Matrix.

Related information

[NetApp Interoperability](#)

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
- AIX Host Utilities

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
AIX	FCP	none	jfs2	lvm
	FCP	none	jfs	lvm
	FCP	nativempio	jfs2	lvm
	FCP	nativempio	jfs	lvm
	FCP	DMP	vxfs	vxvm

- 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 AIX are installed .
- The volume manager for AIX is LVM or VxVM .
- Host Utilities are installed in AIX

For example, if multipathing is in use with SnapDrive for UNIX with an AIX host, you must set up the features required by the AIX Host Utilities for that host.

- 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](#)

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 AIX 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 set vsadmin Vserver 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 AIX

- The SnapDrive for UNIX daemon takes longer to restart if one or more FC port links are down or unavailable.

However, the SnapDrive for UNIX daemon can start regardless of the state of the FC ports.

- By default, SnapDrive for UNIX creates disk groups with non-concurrent settings; you must manually change or provide storage provisioning operations.
- SnapDrive for UNIX does not support iSCSI transport protocol on AIX.

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.
- 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.

Support for NetAppDataMotion 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](#)

[Data ONTAP 8.2 SAN Administration Guide for 7-Mode](#)

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.

What configuration checker is

The SnapDrive configuration checker tool helps you to identify and verify the configurations supported by SnapDrive for UNIX.

The configuration checker tool addresses the problem of multiple support configurations by verifying the user configuration in both NAS and SAN environments.

The SnapDrive configuration checker tool supports the following configuration checks in SnapDrive for UNIX:

- Checks the configuration specified in the `snapdrive.conf` file, which is available in the default location.
- Checks the configuration specified using the CLI.

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.

Multipathing in SnapDrive for UNIX

SnapDrive for UNIX supports multipathing. SnapDrive for UNIX does not support the use of aliases for devices, created using the `mknod` and/or `rendev` commands.

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 AIX, the volume manager is termed as Native LVM and Veritas Volume Manager (VxVM).

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
AIX	Native LVM	Volume groups <code>vg</code>	<code>dev/lvol</code> All logical volumes share the same namespace	<code>/dev/hdisk</code> (FC only) multipathing is not supported with ISCSI
	Veritas Volume Manager (VxVM)	Volume groups <code>vg</code>	<code>/dev/vx/dsk/dg/lvol</code>	<code>/dev/vx/dmp/Disk_1</code>

Copyright Information

Copyright © 2021 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, 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.