



Storage VM administration

Cloud Volumes ONTAP

NetApp
January 13, 2026

Table of Contents

Storage VM administration	1
Manage storage VMs for Cloud Volumes ONTAP	1
Supported number of storage VMs	1
Work with multiple storage VMs	1
Modify the name of the default storage VM	2
Manage data-serving storage VMs for Cloud Volumes ONTAP in AWS	2
Supported number of storage VMs	2
Verify limits for your configuration	2
Allocate IP addresses in AWS	3
Create a storage VM on a single node system	4
Create a storage VM on an HA pair in a single AZ	5
Create a storage VM on an HA pair in multiple AZs	7
Manage data-serving storage VMs for Cloud Volumes ONTAP in Azure	10
Supported number of storage VMs	11
Create a storage VM	11
Manage storage VMs on single node systems and HA pairs	12
Manage data-serving storage VMs for Cloud Volumes ONTAP in Google Cloud	13
Supported number of storage VMs	13
Create a storage VM	13
Manage storage VMs	14
Set up storage VM disaster recovery for Cloud Volumes ONTAP	16

Storage VM administration

Manage storage VMs for Cloud Volumes ONTAP

A storage VM is a virtual machine running within ONTAP that provides storage and data services to your clients. You might know this as an *SVM* or a *vserver*. Cloud Volumes ONTAP is configured with one storage VM by default, but some configurations support additional storage VMs.


Supported number of storage VMs

Multiple storage VMs are supported with certain configurations. Go to the [Cloud Volumes ONTAP Release Notes](#) to verify the supported number of storage VMs for your version of Cloud Volumes ONTAP.

Work with multiple storage VMs

The NetApp Console supports any additional storage VMs that you create from ONTAP System Manager or the ONTAP CLI.

For example, the following image shows how you can choose a storage VM when you create a volume.



The screenshot shows the 'Details & Protection' section of a volume creation dialog. It includes fields for 'Storage VM Name' (set to 'svm_name1'), 'Volume Name' (empty), 'Size (GiB)' (empty), 'Snapshot Policy' (set to 'default'), and a note about 'Default Policy'.

Setting	Value
Storage VM Name	svm_name1
Volume Name	
Size (GiB)	
Snapshot Policy	default
Default Policy	(Note)

And the following image shows how you can choose a storage VM when replicating a volume to another system.



The screenshot shows a configuration interface with three main fields. The first field, 'Destination Volume Name', contains the value 'volume_copy'. The second field, 'Destination Storage VM Name', contains the value 'svm_name1'. The third field, 'Destination Aggregate', contains the value 'Automatically select the best aggregate'.

Destination Volume Name	volume_copy
Destination Storage VM Name	svm_name1
Destination Aggregate	Automatically select the best aggregate

Modify the name of the default storage VM

The Console automatically names the single storage VM that it creates for Cloud Volumes ONTAP. From ONTAP System Manager, the ONTAP CLI, or API, you can modify the name of the storage VM if you have strict naming standards. For example, you might want the name to match how you name the storage VMs for your ONTAP clusters.

Manage data-serving storage VMs for Cloud Volumes ONTAP in AWS

A storage VM is a virtual machine running within ONTAP that provides storage and data services to your clients. You might know this as an *SVM* or a *vserver*. Cloud Volumes ONTAP is configured with one storage VM by default, but some configurations support additional storage VMs.

To create additional data-serving storage VMs, you need to allocate IP addresses in AWS and then run ONTAP commands based on your Cloud Volumes ONTAP configuration.

Supported number of storage VMs

Multiple storage VMs are supported with specific Cloud Volumes ONTAP configurations starting with the 9.7 release. Go to the [Cloud Volumes ONTAP Release Notes](#) to verify the supported number of storage VMs for your version of Cloud Volumes ONTAP.

All other Cloud Volumes ONTAP configurations support one data-serving storage VM and one destination storage VM used for disaster recovery. You can activate the destination storage VM for data access if there's an outage on the source storage VM.

Verify limits for your configuration

Each EC2 instance supports a maximum number of private IPv4 addresses per network interface. You need to verify the limit before you allocate IP addresses in AWS for the new storage VM.

Steps

1. Go to the [Storage limits section in the Cloud Volumes ONTAP Release Notes](#).
2. Identify the maximum number of IP addresses per interface for your instance type.
3. Make note of this number because you'll need it in the next section when you allocate IP addresses in AWS.


Allocate IP addresses in AWS

Private IPv4 addresses must be assigned to port e0a in AWS before you create LIFs for the new storage VM.

Note that an optional management LIF for a storage VM requires a private IP address on a single node system and on an HA pair in a single AZ. This management LIF provides a connection to management tools like SnapCenter.

Steps

1. Log in to AWS and open the EC2 service.
2. Select the Cloud Volumes ONTAP instance and click **Networking**.
3. Scroll down to **Network interfaces** and click the **Interface ID** for port e0a.



4. Select the network interface and click **Actions > Manage IP addresses**.
5. Expand the list of IP addresses for e0a.
6. Verify the IP addresses:
 - a. Count the number of allocated IP addresses to confirm that the port has room for additional IPs.

You should have identified the maximum number of supported IP addresses per interface in the previous section of this page.
 - b. Optional: Go to the ONTAP CLI for Cloud Volumes ONTAP and run **network interface show** to confirm that each of these IP addresses are in use.

If an IP address isn't in use, then you can use it with the new storage VM.

7. Back in the AWS Console, click **Assign new IP address** to assign additional IP addresses based on the amount that you need for the new storage VM.
 - Single node system: One unused secondary private IP is required.
An optional secondary private IP is required if you want to create a management LIF on the storage VM.
 - HA pair in a single AZ: One unused secondary private IP is required on node 1.
An optional secondary private IP is required if you want to create a management LIF on the storage VM.
 - HA pair in multiple AZs: One unused secondary private IP is required on each node.
8. If you're allocating the IP address on an HA pair in a single AZ, enable **Allow secondary private IPv4 addresses to be reassigned**.
9. Click **Save**.
10. If you have an HA pair in multiple AZs, then you'll need to repeat these steps for node 2.

Create a storage VM on a single node system

These steps create a new storage VM on a single node system. One private IP address is required to create a NAS LIF and another optional private IP address is needed if you want to create a management LIF.

Steps

1. Create the storage VM and a route to the storage VM.

```
vserver create -rootvolume-security-style unix -rootvolume root_svm_2  
-snapshot-policy default -vserver svm_2 -aggregate aggr1
```

```
network route create -destination 0.0.0.0/0 -vserver svm_2 -gateway  
subnet_gateway
```

2. Create a NAS LIF.

```
network interface create -auto-revert true -vserver svm_2 -service  
-policy default-data-files -home-port e0a -address private_ip_x -netmask  
node1Mask -lif ip_nas_2 -home-node cvo-node
```

Where *private_ip_x* is an unused secondary private IP on e0a.

3. Optional: Create a storage VM management LIF.

```
network interface create -auto-revert true -vserver svm_2 -service -policy default-management -home-port e0a -address private_ip_y -netmask node1Mask -lif ip_svm_mgmt_2 -home-node cvo-node
```

Where *private_ip_y* is another unused secondary private IP on e0a.

4. Assign one or more aggregates to the storage VM.

```
vserver add-aggregates -vserver svm_2 -aggregates aggr1,aggr2
```

This step is required because the new storage VM needs access to at least one aggregate before you can create volumes on the storage VM.

Create a storage VM on an HA pair in a single AZ

These steps create a new storage VM on an HA pair in a single AZ. One private IP address is required to create a NAS LIF and another optional private IP address is needed if you want to create a management LIF.

Both of these LIFs get allocated on node 1. The private IP addresses can move between nodes if failures occur.

Steps

1. Create the storage VM and a route to the storage VM.

```
vserver create -rootvolume-security-style unix -rootvolume root_svm_2 -snapshot-policy default -vserver svm_2 -aggregate aggr1
```

```
network route create -destination 0.0.0.0/0 -vserver svm_2 -gateway subnet_gateway
```

2. Create a NAS LIF on node 1.

```
network interface create -auto-revert true -vserver svm_2 -service -policy default-data-files -home-port e0a -address private_ip_x -netmask node1Mask -lif ip_nas_2 -home-node cvo-node1
```

Where *private_ip_x* is an unused secondary private IP on e0a of cvo-node1. This IP address can be relocated to the e0a of cvo-node2 in case of takeover because the service policy default-data-files indicates that IPs can migrate to the partner node.

3. Optional: Create a storage VM management LIF on node 1.

```
network interface create -auto-revert true -vserver svm_2 -service
-policy default-management -home-port e0a -address private_ip_y -netmask
node1Mask -lif ip_svm_mgmt_2 -home-node cvo-node1
```

Where *private_ip_y* is another unused secondary private IP on e0a.

4. Assign one or more aggregates to the storage VM.

```
vserver add-aggregates -vserver svm_2 -aggregates aggr1,aggr2
```

This step is required because the new storage VM needs access to at least one aggregate before you can create volumes on the storage VM.

5. If you're running Cloud Volumes ONTAP 9.11.1 or later, modify the network service policies for the storage VM.

Modifying the services is required because it ensures that Cloud Volumes ONTAP can use the iSCSI LIF for outbound management connections.

```
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service data-fpolicy-client
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service management-ad-client
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service management-dns-client
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service management-ldap-client
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service management-nis-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service data-fpolicy-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service management-ad-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service management-dns-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service management-ldap-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service management-nis-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service data-fpolicy-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service management-ad-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service management-dns-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service management-ldap-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service management-nis-client
```

Create a storage VM on an HA pair in multiple AZs

These steps create a new storage VM on an HA pair in multiple AZs.

A *floating* IP address is required for a NAS LIF and is optional for a management LIF. These floating IP addresses don't require you to allocate private IPs in AWS. Instead, the floating IPs are automatically configured in the AWS route table to point to a specific node's ENI in the same VPC.

In order for floating IPs to work with ONTAP, a private IP address must be configured on every storage VM on each node. This is reflected in the steps below where an iSCSI LIF is created on node 1 and on node 2.

Steps

1. Create the storage VM and a route to the storage VM.

```
vserver create -rootvolume-security-style unix -rootvolume root_svm_2  
-snapshot-policy default -vserver svm_2 -aggregate aggr1
```

```
network route create -destination 0.0.0.0/0 -vserver svm_2 -gateway  
subnet_gateway
```

2. Create a NAS LIF on node 1.

```
network interface create -auto-revert true -vserver svm_2 -service  
-policy default-data-files -home-port e0a -address floating_ip -netmask  
node1Mask -lif ip_nas_floating_2 -home-node cvo-node1
```

- The floating IP address must be outside of the CIDR blocks for all VPCs in the AWS region in which you deploy the HA configuration. 192.168.209.27 is an example floating IP address. [Learn more about choosing a floating IP address](#).
- `-service-policy default-data-files` indicates that IPs can migrate to the partner node.

3. Optional: Create a storage VM management LIF on node 1.

```
network interface create -auto-revert true -vserver svm_2 -service  
-policy default-management -home-port e0a -address floating_ip -netmask  
node1Mask -lif ip_svm_mgmt_2 -home-node cvo-node1
```

4. Create an iSCSI LIF on node 1.

```
network interface create -vserver svm_2 -service-policy default-data-  
blocks -home-port e0a -address private_ip -netmask node1Mask -lif  
ip_node1_iscsi_2 -home-node cvo-node1
```

- This iSCSI LIF is required to support LIF migration of the floating IPs in the storage VM. It doesn't have to be an iSCSI LIF, but it can't be configured to migrate between nodes.
- `-service-policy default-data-block` indicates that an IP address does not migrate between nodes.
- `private_ip` is an unused secondary private IP address on eth0 (e0a) of cvo_node1.

5. Create an iSCSI LIF on node 2.

```
network interface create -vserver svm_2 -service-policy default-data-  
blocks -home-port e0a -address private_ip -netmaskNode2Mask -lif  
ip_node2_iscsi_2 -home-node cvo-node2
```

- This iSCSI LIF is required to support LIF migration of the floating IPs in the storage VM. It doesn't have to be an iSCSI LIF, but it can't be configured to migrate between nodes.
- `-service-policy default-data-block` indicates that an IP address does not migrate between nodes.
- *private_ip* is an unused secondary private IP address on eth0 (e0a) of cvo_node2.

6. Assign one or more aggregates to the storage VM.

```
vserver add-aggregates -vserver svm_2 -aggregates aggr1,aggr2
```

This step is required because the new storage VM needs access to at least one aggregate before you can create volumes on the storage VM.

7. If you're running Cloud Volumes ONTAP 9.11.1 or later, modify the network service policies for the storage VM.

Modifying the services is required because it ensures that Cloud Volumes ONTAP can use the iSCSI LIF for outbound management connections.

```
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service data-fpolicy-client
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service management-ad-client
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service management-dns-client
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service management-ldap-client
network interface service-policy remove-service -vserver <svm-name>
-policy default-data-files -service management-nis-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service data-fpolicy-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service management-ad-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service management-dns-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service management-ldap-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-blocks -service management-nis-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service data-fpolicy-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service management-ad-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service management-dns-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service management-ldap-client
network interface service-policy add-service -vserver <svm-name> -policy
default-data-iscsi -service management-nis-client
```

Manage data-serving storage VMs for Cloud Volumes ONTAP in Azure

A storage VM is a virtual machine running within ONTAP that provides storage and data services to your clients. You might know this as an *SVM* or a *vserver*. Cloud Volumes ONTAP is configured with one storage VM by default, but you can create additional storage VMs when running Cloud Volumes ONTAP in Azure.

To create and manage additional data-serving storage VMs in Azure, you should use the APIs. This is because the APIs automate the process of creating the storage VMs and configuring the required network interfaces. When creating the storage VMs, the NetApp Console configures the required LIF services, as well as an iSCSI LIF that's required for outbound SMB/CIFS communications from the storage VM.

For information about running Cloud Volumes ONTAP API calls, refer to [Your first API call](#).

Supported number of storage VMs

Beginning with Cloud Volumes ONTAP 9.9.0, based on your license, multiple storage VMs are supported with specific configurations. Refer to the [Cloud Volumes ONTAP Release Notes](#) to verify the supported number of storage VMs for your version of Cloud Volumes ONTAP.

All versions of Cloud Volumes ONTAP prior to 9.9.0 support one data-serving storage VM and one destination storage VM used for disaster recovery. You can activate the destination storage VM for data access if there's an outage on the source storage VM.

Create a storage VM

Based on your configuration and license type, you can create multiple storage VMs on a single node system or in a high-availability (HA) configuration by using the APIs for the NetApp Console.

About this task

When you create storage VMs using the APIs, along with configuring the required network interfaces, the Console also modifies the `default-data-files` policies on the data storage VMs by removing the following services from the NAS data LIF and adding them to the iSCSI data LIF that's used for outbound management connections:

- `data-fpolicy-client`
- `management-ad-client`
- `management-dns-client`
- `management-ldap-client`
- `management-nis-client`

Before you begin

The Console agent requires specific permissions to create storage VMs for Cloud Volumes ONTAP. The required permissions are included in [the policies provided by NetApp](#).

Single node system

Use the following API call to create a storage VM on a single node system.

```
POST /azure/vsa/working-environments/{workingEnvironmentId}/svm
```

Include the following parameters in the request body:

```
{ "svmName": "myNewSvm1"  
  "svmPassword": "optional, the API takes the cluster password if not  
  provided"  
  "mgmtLif": "optional, to create an additional management LIF, if you  
  want to use the storage VM for management purposes"}
```

HA pair

Use the following API call to create a storage VM on an HA pair:

```
POST /azure/ha/working-environments/{workingEnvironmentId}/svm
```

Include the following parameters in the request body:

```
{ "svmName": "NewSvmName"  
  "svmPassword": "optional value, the API takes the cluster password if  
  not provided"  
  "mgmtLif": "optional value, to create an additional management LIF, if  
  you want to use the storage VM for management purposes"}
```

Manage storage VMs on single node systems and HA pairs

Using the APIs, you can rename and delete storage VMs in both single node and HA configurations.

Before you begin

The Console agent requires specific permissions to manage storage VMs for Cloud Volumes ONTAP. The required permissions are included in [the policies provided by NetApp](#).

Rename a storage VM

To rename a storage VM, you should provide the names of the existing storage VM and new storage VM as parameters.

Steps

- Use the following API call to rename a storage VM on a single node system:

```
PUT /azure/vsa/working-environments/{workingEnvironmentId}/svm
```

Include the following parameters in the request body:

```
{  
  "svmNewName": "NewSvmName",  
  "svmName": "OldSvmName"  
}
```

- Use the following API call to rename a storage VM on an HA pair:

```
PUT /azure/ha/working-environments/{workingEnvironmentId}/svm
```

Include the following parameters in the request body:

```
{  
  "svmNewName": "NewSvmName",  
  "svmName": "OldSvmName"  
}
```

Delete a storage VM

In a single node or HA configuration, you can remove a storage VM if it doesn't have any active volumes.

Steps

- Use the following API call to delete a storage VM on a single node system:

```
DELETE /azure/vsa/working-environments/{workingEnvironmentId}/svm/{svmName}
```

- Use the following API call to delete a storage VM on an HA pair:

```
DELETE /azure/ha/working-environments/{workingEnvironmentId}/svm/{svmName}
```

Related information

- [Prepare to use the API](#)
- [Cloud Volumes ONTAP workflows](#)
- [Get required identifiers](#)
- [Use the REST APIs for NetApp Console](#)

Manage data-serving storage VMs for Cloud Volumes ONTAP in Google Cloud

A storage VM is a virtual machine running within ONTAP that provides storage and data services to your clients. You might know this as an *SVM* or a *vserver*. Cloud Volumes ONTAP is configured with one storage VM by default, but some configurations support additional storage VMs.

To create and manage additional data-serving storage VMs in Google Cloud, you should use the APIs. This is because the APIs automate the process of creating the storage VMs and configuring the required network interfaces. When creating the storage VMs, the NetApp Console configures the required LIF services, as well as an iSCSI LIF that's required for outbound SMB/CIFS communications from the storage VM.

For information about running Cloud Volumes ONTAP API calls, refer to [Your first API call](#).

Supported number of storage VMs

Beginning with Cloud Volumes ONTAP 9.11.1, based on your license, multiple storage VMs are supported with specific configurations. Refer to the [Cloud Volumes ONTAP Release Notes](#) to verify the supported number of storage VMs for your version of Cloud Volumes ONTAP.

All versions of Cloud Volumes ONTAP prior to 9.11.1 support one data-serving storage VM and one destination storage VM used for disaster recovery. You can activate the destination storage VM for data access if there's an outage on the source storage VM.

Create a storage VM

Based on your configuration and license type, you can create multiple storage VMs on a single node system or in a high-availability (HA) configuration by using the APIs.

About this task

When you create storage VMs using the APIs, along with configuring the required network interfaces, the Console also modifies the `default-data-files` policies on the data storage VMs by removing the following services from the NAS data LIF and adding them to the iSCSI data LIF that's used for outbound management connections:

- `data-fpolicy-client`
- `management-ad-client`
- `management-dns-client`
- `management-ldap-client`
- `management-nis-client`

Before you begin

The Console agent requires specific permissions to create storage VMs for Cloud Volumes ONTAP HA pairs. The required permissions are included in [the policies provided by NetApp](#).

Single node system

Use the following API call to create a storage VM on a single node system.

```
POST /gcp/vsa/working-environments/{workingEnvironmentId}/svm
```

Include the following parameters in the request body:

```
{ "svmName": "NewSvmName"  
  "svmPassword": "optional value, the API takes the cluster password if  
not provided"  
  "mgmtLif": "optional value, to create an additional management LIF, if  
you want to use the storage VM for management purposes"}
```

HA pair

Use the following API call to create a storage VM on an HA pair:

```
POST /gcp/ha/working-environments/{workingEnvironmentId}/svm/
```

Include the following parameters in the request body:

```
{ "svmName": "NewSvmName"  
  "svmPassword": "optional value, the API takes the cluster password if  
not provided"  
 }
```

Manage storage VMs

Using the APIs, you can rename and delete storage VMs in both single node and HA configurations.

Before you begin

The Console agent requires specific permissions to manage storage VMs for Cloud Volumes ONTAP HA pairs. The required permissions are included in [the policies provided by NetApp](#).

Rename a storage VM

To rename a storage VM, you should provide the names of the existing storage VM and new storage VM as parameters.

Steps

- Use the following API call to rename a storage VM on a single node system:

```
PUT /gcp/vsa/working-environments/{workingEnvironmentId}/svm
```

Include the following parameters in the request body:

```
{  
  "svmNewName": "NewSvmName",  
  "svmName": "OldSvmName"  
}
```

- Use the following API call to rename a storage VM on an HA pair:

```
PUT /gcp/ha/working-environments/{workingEnvironmentId}/svm
```

Include the following parameters in the request body:

```
{  
  "svmNewName": "NewSvmName",  
  "svmName": "OldSvmName"  
}
```

Delete a storage VM

In a single node or HA configuration, you can remove a storage VM if it doesn't have any active volumes.

Steps

- Use the following API call to delete a storage VM on a single node system:

```
DELETE /gcp/vsa/working-environments/{workingEnvironmentId}/svm/{svmName}
```

- Use the following API call to delete a storage VM on an HA pair:

```
DELETE /gcp/ha/working-environments/{workingEnvironmentId}/svm/{svmName}
```

Related information

- [Prepare to use the API](#)
- [Cloud Volumes ONTAP workflows](#)

- [Get required identifiers](#)
- [Use the REST APIs for the NetApp Console](#)

Set up storage VM disaster recovery for Cloud Volumes ONTAP

The NetApp Console does not offer setup or orchestration support for storage VM (SVM) disaster recovery. To perform these tasks, use ONTAP System Manager or the ONTAP CLI.

If you set up SnapMirror SVM replication between two Cloud Volumes ONTAP systems, the replication must be between two HA pair systems or two single node systems. You can't set up SnapMirror SVM replication between an HA pair and a single node system.

Refer to the following documents for the ONTAP CLI instructions.

- [SVM Disaster Recovery Preparation Express Guide](#)
- [SVM Disaster Recovery Express Guide](#)

Copyright information

Copyright © 2026 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.