



## **Proxmox**

### **ONTAP SAN Host Utilities**

NetApp  
January 30, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap-sanhost/hu-proxmox-ve-9x.html> on January 30, 2026. Always check docs.netapp.com for the latest.

# Table of Contents

- Proxmox ..... 1
  - Configure Proxmox VE 9.x for FCP and iSCSI with ONTAP storage ..... 1
    - Step 1: Confirm the multipath configuration for your host ..... 1
    - Step 2: Confirm the iSCSI configuration for your host ..... 4
    - Step 3: Optionally, exclude a device from multipathing ..... 7
    - Step 4: Customize multipath parameters for ONTAP LUNs. .... 8
    - Step 5: Review the known issues ..... 9
  - Configure Proxmox VE 8.x for FCP and iSCSI with ONTAP storage. .... 9
    - Step 1: Confirm the multipath configuration for your host ..... 9
    - Step 2: Confirm the iSCSI configuration for your host ..... 12
    - Step 3: Optionally, exclude a device from multipathing ..... 15
    - Step 4: Customize multipath parameters for ONTAP LUNs. .... 16
    - Step 5: Review the known issues ..... 17

# Proxmox

## Configure Proxmox VE 9.x for FCP and iSCSI with ONTAP storage

Configure Proxmox VE 9.x for multipathing and with specific parameters and settings for FCP and iSCSI protocol operations with ONTAP storage.

FCP and iSCSI with Proxmox VE 9.x have the following known limitations:

- The Linux Host Utilities do not support Proxmox VE 9.x operating systems.
- The SAN boot configuration is not supported.

### Step 1: Confirm the multipath configuration for your host

You can use multipathing with Proxmox VE 9.x to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

#### Steps

1. Verify that the `/etc/multipath.conf` file exists. If the file doesn't exist, create an empty, zero-byte file:

```
touch /etc/multipath.conf
```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

```
systemctl enable multipathd
```

```
systemctl start multipathd
```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

The following table shows the Linux OS native compiled multipath parameter settings for ONTAP LUNs.

### Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"always"
hardware_handler	"1"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Optionally, override the default value for the `find_multipaths` parameter to ensure that ONTAP LUNs are correctly discovered and managed by `multipathd`:

- a. Set `find_multipaths` to "no" in the defaults section of `/etc/multipath.conf`:

```
defaults {
    find_multipaths "no"
}
```

- b. Reload the multipath service:

```
systemctl reload multipathd
```



By default, the Proxmox OS-native multipath configuration sets `find_multipaths` to **"strict"** with the empty zero byte `/etc/multipath.conf` configuration file each time you boot the host. This can prevent the host discovering newly presented ONTAP LUNs as multipath devices, which means they do not appear under multipath control automatically. Existing ONTAP LUNs remain discovered and under multipath control after each reboot.

5. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn't require more than four paths. If there are more than four paths, it might cause issues with the paths during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

## ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

### Show example

```
multipath -ll
3600a098038315071592b59713261566d dm-38 NETAPP,LUN C-Mode
size=100G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 8:0:0:7 sdbv 68:144 active ready running
  |- 9:0:0:7 sdbx 68:176 active ready running
  |- 6:0:0:7 sdbz 68:80 active ready running
  `-- 7:0:0:7 sdbt 68:112 active ready running
```

## AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren't available.

The following example shows the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

### Show example

```
multipath -ll
3600a0980383149764b5d567257516273 dm-0 NETAPP,LUN C-Mode
size=150G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 16:0:3:0 sdcg 69:64 active ready running
| `-- 10:0:0:0 sdb 8:16 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 10:0:1:0 sdc 8:32 active ready running
  `-- 16:0:2:0 sdcf 69:48 active ready running
```

## Step 2: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

### About this task

You perform the following steps on the iSCSI host.

### Steps

1. Verify that the iSCSI initiator package (open-iscsi) is installed:

```
$apt list |grep open-iscsi
```

You should see an output similar to the following example:

```
open-iscsi/noble-updates,noble-updates,now 2.1.9-3ubuntu5.4 amd64
```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

```
InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

```
node.session.timeo.replacement_timeout = 5
```

The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

```
$systemctl enable iscsid
```

5. Start the iSCSI service:

```
$systemctl start iscsid
```

6. Verify that the iSCSI service is running:

```
$systemctl status iscsid
```

### Show example

```
●iscsid.service - iSCSI initiator daemon (iscsid)
   Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
   enabled; preset: disabled)
   Active: active (running) since Mon 2026-01-12 12:53:18 IST; 2
   days ago
   TriggeredBy: ● iscsid.socket
     Docs: man:iscsid(8)
    Main PID: 1127419 (iscsid)
      Tasks: 2 (limit: 76557)
     Memory: 4.3M (peak: 8.8M)
        CPU: 1.657s
     CGroup: /system.slice/iscsid.service
            └─1127418 /usr/sbin/iscsid
            └─1127419 /usr/sbin/iscsid
```

### 7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

### show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.100.197
192.168.100.197:3260,1046 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.200.199:3260,1049 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.100.199:3260,1048 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.200.197:3260,1047 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
```

### 8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

### 9. Set iSCSI to log in automatically when the host boots:



```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n  
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-  
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 -p  
192.168.100.197:3260 -o update -n node.startup -v automatic
```

#### 10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

#### Show example

```
iscsiadm --mode session  
tcp: [1] 192.168.200.197:3260,1047 iqn.1992-  
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)  
tcp: [2] 192.168.100.197:3260,1046 iqn.1992-  
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)  
tcp: [3] 192.168.100.199:3260,1048 iqn.1992-  
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)  
tcp: [4] 192.168.200.199:3260,1049 iqn.1992-  
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
```

### Step 3: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

#### Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is 360030057024d0730239134810c0cb833.

2. Add the WWID to the "blacklist" stanza:

```

blacklist {
    wwid      360030057024d0730239134810c0cb833
    devnode   "^(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9] *"
    devnode   "^hd[a-z] *"
    devnode   "^cciss.*"
}

```

## Step 4: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the [default settings for multipath parameters](#).



You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

### Show example

```

defaults {
    path_checker      readsector0
    no_path_retry     fail
}

devices {
    device {
        vendor        "NETAPP"
        product        "LUN"
        no_path_retry  queue
        path_checker   tur
    }
}

```

## Step 5: Review the known issues

There are no known issues.

# Configure Proxmox VE 8.x for FCP and iSCSI with ONTAP storage

Configure Proxmox VE 8.x for multipathing and with specific parameters and settings for FCP and iSCSI protocol operations with ONTAP storage.

FCP and iSCSI with Proxmox VE 8.x have the following known limitations:

- The Linux Host Utilities do not support Proxmox VE 8.x operating systems.
- The SAN boot configuration is not supported.

## Step 1: Confirm the multipath configuration for your host

You can use multipathing with Proxmox VE 8.x to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

### Steps

1. Verify that the `/etc/multipath.conf` file exists. If the file doesn't exist, create an empty, zero-byte file:

```
touch /etc/multipath.conf
```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

```
systemctl enable multipathd
```

```
systemctl start multipathd
```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

The following table shows the Linux OS native compiled multipath parameter settings for ONTAP LUNs.

## Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"always"
hardware_handler	"1"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Optionally, override the default value for the `find_multipaths` parameter to ensure that ONTAP LUNs are correctly discovered and managed by `multipathd`:

- a. Set `find_multipaths` to "no" in the defaults section of `/etc/multipath.conf`:

```
defaults {
    find_multipaths "no"
}
```

- b. Reload the multipath service:

```
systemctl reload multipathd
```



By default, the Proxmox OS-native multipath configuration sets `find_multipaths` to **"strict"** with the empty zero byte `/etc/multipath.conf` configuration file each time you boot the host. This can prevent the host discovering newly presented ONTAP LUNs as multipath devices, which means they do not appear under multipath control automatically. Existing ONTAP LUNs remain discovered and under multipath control after each reboot.

5. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn't require more than four paths. If there are more than four paths, it might cause issues with the paths during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

### ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

#### Show example

```
multipath -ll
3600a098038315071592b59713261566d dm-38 NETAPP,LUN C-Mode
size=100G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 8:0:0:7 sdbv 68:144 active ready running
  |- 9:0:0:7 sdbx 68:176 active ready running
  |- 6:0:0:7 sdbz 68:80 active ready running
  `-- 7:0:0:7 sdbt 68:112 active ready running
```

### AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren't available.

The following example shows the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

#### Show example

```
multipath -ll
3600a0980383149764b5d567257516273 dm-0 NETAPP,LUN C-Mode
size=150G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 16:0:3:0 sdcg 69:64 active ready running
| `-- 10:0:0:0 sdb 8:16 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 10:0:1:0 sdc 8:32 active ready running
  `-- 16:0:2:0 sdcf 69:48 active ready running
```

## Step 2: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

### About this task

You perform the following steps on the iSCSI host.

### Steps

1. Verify that the iSCSI initiator package (open-iscsi) is installed:

```
$apt list |grep open-iscsi
```

You should see an output similar to the following example:

```
open-iscsi/noble-updates,noble-updates,now 2.1.9-3ubuntu5.4 amd64
```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

```
InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

```
node.session.timeo.replacement_timeout = 5
```

The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

```
$systemctl enable iscsid
```

5. Start the iSCSI service:

```
$systemctl start iscsid
```

6. Verify that the iSCSI service is running:

```
$systemctl status iscsid
```

### Show example

```
●iscsid.service - iSCSI initiator daemon (iscsid)
   Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
   enabled; preset: disabled)
   Active: active (running) since Mon 2026-01-12 12:53:18 IST; 2
   days ago
   TriggeredBy: ● iscsid.socket
     Docs: man:iscsid(8)
    Main PID: 1127419 (iscsid)
      Tasks: 2 (limit: 76557)
     Memory: 4.3M (peak: 8.8M)
        CPU: 1.657s
     CGroup: /system.slice/iscsid.service
            └─1127418 /usr/sbin/iscsid
            └─1127419 /usr/sbin/iscsid
```

### 7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

### show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.100.197
192.168.100.197:3260,1046 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.200.199:3260,1049 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.100.199:3260,1048 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.200.197:3260,1047 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
```

### 8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

### 9. Set iSCSI to log in automatically when the host boots:



```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 -p
192.168.100.197:3260 -o update -n node.startup -v automatic
```

#### 10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

#### Show example

```
iscsiadm --mode session
tcp: [1] 192.168.200.197:3260,1047 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [2] 192.168.100.197:3260,1046 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [3] 192.168.100.199:3260,1048 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [4] 192.168.200.199:3260,1049 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
```

### Step 3: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

#### Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is 360030057024d0730239134810c0cb833.

2. Add the WWID to the "blacklist" stanza:

```

blacklist {
    wwid      360030057024d0730239134810c0cb833
    devnode   "^(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9] *"
    devnode   "^hd[a-z] *"
    devnode   "^cciss.*"
}

```

## Step 4: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the [default settings for multipath parameters](#).



You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

### Show example

```

defaults {
    path_checker      readsector0
    no_path_retry     fail
}

devices {
    device {
        vendor        "NETAPP"
        product        "LUN"
        no_path_retry  queue
        path_checker   tur
    }
}

```

## **Step 5: Review the known issues**

There are no known issues.

## 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.