



Upgrade AFF A250 to AFF A400 by converting to a drive shelf

Upgrade controllers

NetApp
August 29, 2024

Table of Contents

- Upgrade AFF A250 to AFF A400 by converting to a drive shelf 1
 - Workflow 1
 - Migrate LIFs and data aggregates on node2 to node1 2
 - Convert node2 to a drive shelf and connect to node4 4
 - Reassign drives from node2 to node4 5
 - Migrate data aggregates, epsilon, and LIFs on node1 to node4 7
 - Convert node1 to a drive shelf and connect to node3 9
 - Reassign drives from node1 to node3 10
 - Migrate LIFs and data aggregates on node4 to node3 11

Upgrade AFF A250 to AFF A400 by converting to a drive shelf

Workflow

You can perform a nondisruptive upgrade from a NetApp AFF A250 system to an NetApp AFF A400 system by converting each AFF A250 node to an NS224 drive shelf and then connecting to the AFF A400 replacement nodes.

About this task

In this procedure, the AFF A250 high-availability (HA) pair controllers are called node1 and node2, and the replacement AFF A400 HA pair controllers are called node3 and node4.

1

Migrate LIFs and data aggregates on node2 to node1

Before converting AFF A250 node2 to a drive shelf, you migrate the logical interfaces (LIFs) and data aggregates on node2 to node1.

2

Convert node2 to a drive shelf and connect to node4

You convert AFF A250 node2 to an NS224 drive shelf and then connect to AFF A400 node4 before reassigning drives from node2 to node4.

3

Reassign drives from node2 to node4

After converting AFF A250 node2 to an NS224 drive shelf and connecting to AFF A400 node4, you reassign the drives that previously belonged to node2 to node4

4

Migrate data aggregates, epsilon, and LIFs on node1 to node4

Before converting AFF A250 node1 to a drive shelf, you migrate the data aggregates, epsilon, and LIFs on node1 to AFF A400 node4.

5

Convert node1 to a drive shelf and connect to node3

Convert AFF A250 node1 to an NS224 drive shelf and then connect to AFF A400 node3 before reassigning drives from node1 to node3.

6

Reassign drives from node1 to node3

After converting AFF A250 node1 to an NS224 drive shelf and connecting to AFF A400 node3, you reassign the drives that previously belonged to node1 to node3.

7

Migrate LIFs and data aggregates on node4 to node3

To complete the upgrade, you connect node3 to node4 and then migrate the data LIFs and data aggregates on node4 to node3.

Migrate LIFs and data aggregates on node2 to node1

Before converting AFF A250 node2 to a drive shelf, you migrate the logical interfaces (LIFs) and data aggregates on node2 to node1.

Before you begin

Verify that you meet the following requirements:

- The AFF A250 and AFF A400 controllers are running the same ONTAP release and patch version.
- You must netboot and install the ONTAP version on each AFF A400 that is identical to the version running on the AFF A250.
- Both the primary and backup boot image of each AFF A400 must have the same ONTAP version.
- If the AFF A400 clusters were previously configured, you must clear any residual cluster configuration by performing a `wipeconfig` from the boot menu.
- Both AFF A400 controllers are on standby at the LOADER prompt.
- You have all the appropriate cabling on-hand.



About this task

The following steps are performed on AFF A250 node1.

Steps

1. Access the advanced privilege level:

```
set -privilege advanced
```

2. Disable storage failover automatic giveback:

```
storage failover modify -node node1 -auto-giveback false
```

3. Disable auto-revert of the LIFs across both nodes of the HA pair:

```
network interface modify -lif * -auto-revert false
```

4. Display the status of all data network LIFs:

```
network interface show -role data
```

5. Display the status of the cluster management LIFs:

```
network interface show -role cluster_mgmt
```

6. Migrate all data LIFs from the storage virtual machines hosted on node2:

```
network interface migrate -vserver vserver_name -lif lif_name -destination  
-node node1 -destination-port port_name
```



This command only migrates non-SAN LIFs. You cannot use it to migrate iSCSI and FCP LIFs.

7. Display the status of all data LIFs in the cluster:

```
network interface show -role data
```

8. If any LIFs are down, set the administrative status of the LIFs to `up` by entering the following command, once for each LIF:

```
network interface modify -vserver vserver_name -lif lif_name -status-admin up
```

9. Display the status of all data aggregates in the cluster:

```
storage aggregate show
```

10. Display failover eligibility:

```
storage failover show
```

11. Migrate the data aggregates on `node2` to `node1`:

```
storage aggregate relocation start -aggregate aggregate_name -node node2  
-destination node1
```

12. Display the status of all data aggregates in the cluster:

```
storage aggregate show
```

13. Display the status of all data volumes in the cluster:

```
volume show
```

14. Display the `ha` status and ownership of `epsilon`:

```
cluster show
```

15. Disable `cluster ha`:

```
cluster ha modify -configured false
```

16. Display the `ha` status and ownership of `epsilon`:

```
cluster show
```

17. Halt `node2`:

```
halt -node node2 -inhibit-takeover true -ignore-quorum-warnings true
```

What's next?

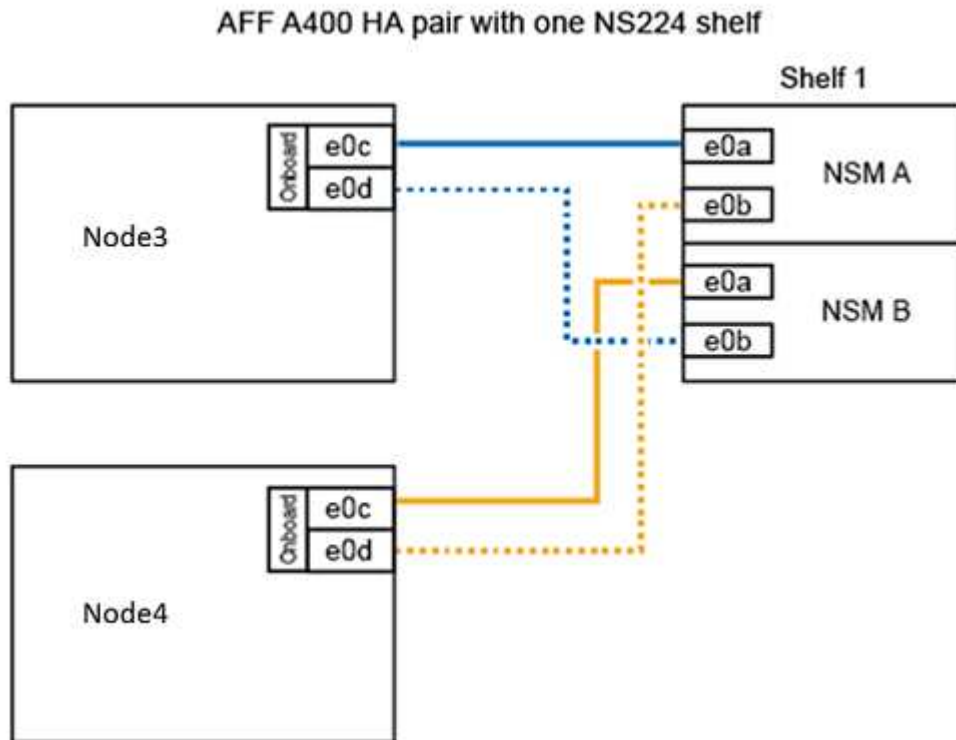
[Convert `node2` to a drive shelf and connect to `node4`](#)

Convert node2 to a drive shelf and connect to node4

Convert AFF A250 node2 to an NS224 drive shelf and then connect to AFF A400 node4 before reassigning drives from node2 to node4.

Steps

1. Disconnect all the network cables from node2.
2. Remove node2 from the AFF A250 chassis.
3. Insert the NVMe shelf module (NSM) into the bay of node2.
4. Connect the NSM to node4 by cabling the node4 100GbE port e0c to NSM B port e0a.



5. Connect the 25GbE cabling from node2 ports e0c and e0d to any two 25GbE onboard ports (e0e, e0f, e0g, or e0h) on node4 to create temporary cluster connections.



If the AFF A400 system uses FC ports as onboard ports, install a 25Gb Ethernet adapter into each node for cluster connectivity during the migration.

6. Connect the 25GbE HA interconnect cables between the AFF A400 nodes using ports e0a and e0b. Do not cross-connect the ports.
7. Connect the 100GbE cluster interconnect cables between the AFF A400 nodes using ports e3a and e3b. Do not cross-connect the ports.

What's next?

[Reassign drives from node2 to node4](#)

Reassign drives from node2 to node4

After converting AFF A250 node2 to an NS224 drive shelf and connecting to AFF A400 node4, you must reassign the drives that previously belonged to node2 to node4.

Before you begin

Verify that both node3 and node4 are on standby at the LOADER prompt.

About this task

You perform the following steps on node4.

Steps

1. At the LOADER prompt, boot node4 into Maintenance Mode:

```
boot_ontap maint
```

2. Show the state of the 100GbE interfaces:

```
storage port show
```

3. Set the 100GbE interfaces to storage ports:

```
storage port modify -p e0c -m storage
```

```
storage port modify -p e0d -m storage
```

4. Verify the mode changes to the 100GbE interfaces:

```
storage port show
```

Output like the following example should display:

```

*> storage port modify -p e0c -m storage
Nov 10 16:27:23 [localhost:nvmeof.port.modify:notice]: Changing NVMe-oF
port e0c to storage mode.

Nov 10 16:27:29 [localhost:nvmeof.subsystem.add:notice]: NVMe-oF
subsystem added at address fe80::2a0:98ff:fefa:8885.

*> storage port modify -p e0d -m storage
Nov 10 16:27:34 [localhost:nvmeof.port.modify:notice]: Changing NVMe-oF
port e0d to storage mode.

Nov 10 16:27:38 [localhost:nvmeof.subsystem.add:notice]: NVMe-oF
subsystem added at address fe80::2a0:98ff:fefa:8886.

*> storage port show
Port Type Mode      Speed(Gb/s) State      Status  VLAN ID
---- ---- -
e0c  ENET storage 100 Gb/s    enabled  online  30
e0d  ENET storage 100 Gb/s    enabled  offline 30

```

5. Display all attached drives:

```
disk show -v
```

6. Record the local system ID value; this is the system ID of node4. Also record the system IDs of node1 and node2 from the "OWNER" column.

7. Reassign all drives from node2 to node4:

```
disk reassign -s node2_system_ID -d node4_system_ID -p node1_system_ID
```

8. Verify that all reassigned drives are viewable to the new system ID:

```
disk show -s node4_System_ID
```



If drives are not viewable, **STOP** and contact technical support for assistance.

9. Verify that the root aggregate of node2 is reported in the output and the aggregate is online:

```
aggr status
```

10. Exit maintenance mode:

```
halt
```

What's next?

[Migrate data aggregates, epsilon, and LIFs on node1 to node4](#)

Migrate data aggregates, epsilon, and LIFs on node1 to node4

Before converting AFF A250 node1 to a drive shelf, you migrate the data aggregates, epsilon, and logical interfaces (LIFs) on node1 to AFF A400 node4.

Steps

1. At the LOADER prompt for node4, boot the node into the boot menu:

```
boot_ontap menu
```

2. Select option 6 Update flash from backup config to restore the /var file system to node4.

This replaces all flash-based configuration with the last backup to disks.

3. Enter *y* to continue.



The node automatically reboots to load the new copy of the /var file system.

The node reports a system ID mismatch warning. Enter *y* to override the system ID.

4. Migrate the cluster LIFs:

```
set -privilege advanced
```

```
network port show
```



If the system cluster ports are not similar when upgrading an AFF A250 to an AFF A400, you might have to temporarily change the interfaces on node4 into cluster ports:

```
network port modify -node node4 -port port_name -mtu 9000 -ipSpace  
Cluster
```

```
network interface migrate -vserver Cluster -lif cluster_LIF -destination-node  
node4 -destination-port port_name
```

5. Wait for the cluster to come into quorum and then verify that the cluster nodes are healthy:

```
- cluster show
```



The HA pair and storage failover remain disabled in the current state.

6. Move the cluster LIFs to the temporary 25G cluster ports on node4:

```
network interface modify
```

7. If interface groups and data VLANs are in use on the AFF A250 cluster you are upgrading, perform this step. If not, go to [Step 8](#).

The physical network port names differ between AFF A250 and AFF A400 systems. As a result, there might be incorrectly configured interface groups and displaced VLANs on node4. Check, and if required, fix

any incorrectly configured interface groups and displaced VLANs.

8. Migrate the data aggregates on node1 to node4:

```
storage aggregate relocation start -aggregate-list aggregate_list_name -node
node1 -destination node4 -ndo-controller-upgrade true -override-destination
-checks true
```

9. Display the status of all data aggregates in the cluster:

```
storage aggregate show
```

10. Migrate the epsilon by removing it from node1 and instead moving it to node4.

a. Remove epsilon from node1:

```
cluster modify -epsilon false -node node1
```

b. Move epsilon to node4:

```
cluster modify -epsilon true -node node4
```

11. Display the cluster status:

```
cluster show
```

12. Display all data network LIFs:

```
network interface show -role data
```

13. Migrate all data LIFs to node4:

```
network interface migrate -vserver vserver_name -lif lif_name -destination
-node node4 -destination-port port_name
```

14. Display the status of all data LIFs in the cluster:

```
network interface show -role data
```

15. If any LIFs are down, set the administrative status of the LIFs to up by entering the following command, once for each LIF:

```
network interface modify -vserver vserver_name -lif lif_name -status-admin up
```

16. Migrate the cluster management LIF:

```
network interface migrate -vserver vserver_name -lif cluster_mgmt -destination
-node node4 -destination-port port_name
```

17. Display the status of the cluster management LIF:

```
network interface show cluster_mgmt
```

18. Halt node1:

```
halt -node node1 -inhibit-takeover true -ignore-quorum-warnings true
```

What's next?

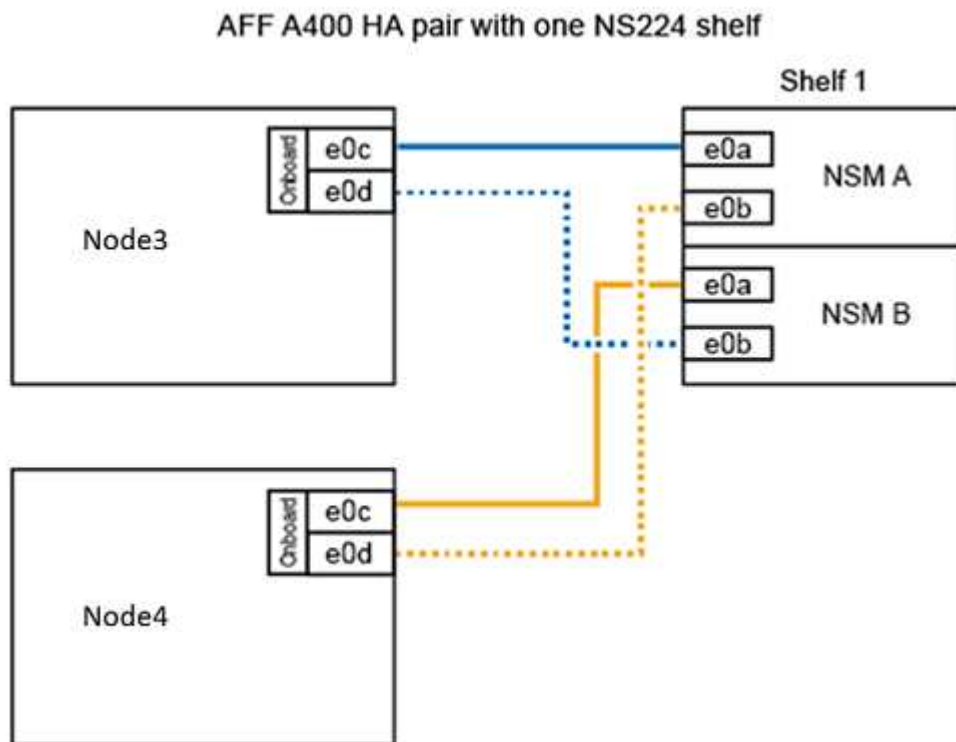
[Convert node1 to a drive shelf and connect to node3](#)

Convert node1 to a drive shelf and connect to node3

Convert AFF A250 node1 to an NS224 drive shelf and then connect to AFF A400 node3 before reassigning drives from node1 to node3.

Steps

1. Disconnect all network cables from node1.
2. Remove node1 from the AFF A250 chassis.
3. Insert the NVMe shelf module (NSM) into the bay of node1.
4. Connect the NSM to node3 by cabling node3 100GbE port e0c to NSM A port e0a.



5. Move the temporary cluster connections to node3 by moving the 25GbE cabling from node1 ports e0c and e0d to any two 25GbE onboard ports (e0e, e0f, e0g, or e0h) on node3.



If the AFF A400 system uses FC ports as onboard ports, install a 25Gb Ethernet adapter into each node for cluster connectivity during the migration.

What's next?

[Reassign drives from node1 to node3](#)

Reassign drives from node1 to node3

After converting AFF A250 node1 to an NS224 drive shelf and connecting to AFF A400 node3, you must reassign the drives that previously belonged to node1 to node3.

Steps

1. At the LOADER prompt, boot node3 into Maintenance Mode:

```
boot_ontap maint
```

2. Show the state of the 100GbE interfaces:

```
storage port show
```

3. Set the 100GbE interfaces to storage ports:

```
storage port modify -p e0c -m storage
```

```
storage port modify -p e0d -m storage
```

4. Verify the mode changes to the 100GbE interfaces:

```
storage port show
```

Output like the following example should display:

```
*> storage port modify -p e0c -m storage
Nov 10 16:27:23 [localhost:nvmeof.port.modify:notice]: Changing NVMe-oF
port e0c to storage mode.

Nov 10 16:27:29 [localhost:nvmeof.subsystem.add:notice]: NVMe-oF
subsystem added at address fe80::2a0:98ff:fefa:8885.

*> storage port modify -p e0d -m storage
Nov 10 16:27:34 [localhost:nvmeof.port.modify:notice]: Changing NVMe-oF
port e0d to storage mode.

Nov 10 16:27:38 [localhost:nvmeof.subsystem.add:notice]: NVMe-oF
subsystem added at address fe80::2a0:98ff:fefa:8886.

*> storage port show
Port Type Mode      Speed (Gb/s) State      Status  VLAN ID
---- ---- -
e0c  ENET storage 100 Gb/s   enabled  online  30
e0d  ENET storage 100 Gb/s   enabled  offline 30
```

5. Display all attached drives:

```
disk show -v
```

- Record the local system ID value; this is the system ID of node3. Also record the system IDs of node1 and node2 from the "OWNER" column.
- Reassign all drives from node1 to node3:

```
disk reassign -s node1_system_ID -d node3_system_ID -p node4_system_ID
```

- Verify that all reassigned drives are viewable to the new system ID:

```
disk show -s node3_system_ID
```



If drives are not viewable, **STOP** and contact technical support for assistance.

- Exit Maintenance Mode:

```
halt
```

What's next?

[Migrate LIFs and data aggregates on node4 to node3](#)

Migrate LIFs and data aggregates on node4 to node3

To complete the upgrade, you connect node3 to node4 and then migrate the data logical interfaces (LIFs) and data aggregates on node4 to node3.

Steps

- At the LOADER prompt for node3, boot the node into the boot menu:

```
boot_ontap menu
```

- Select option 6 Update flash from backup config to restore the /var file system to node3.

This replaces all flash-based configuration with the last backup to disks.

- Enter `y` to continue.
- Allow the node to boot as normal.



The node automatically reboots to load the new copy of the /var file system.

The node reports a warning that there is a system ID mismatch. Enter `y` to override the system ID.

- Connect node3 to node4:
 - Attach multipath high availability (MPHA) cables to the NS224 shelf to ensure redundancy. Connect node3 100GbE port e0d to NSM B port e0b and connect node4 100GbE port e0d to NSM A port e0a.
 - Verify that HA ports e0a and e0b are connected between the nodes.
 - Verify that cluster ports e3a and e3b are connected between the nodes.

6. Migrate the cluster LIFs:

```
set -privilege advanced  
  
network port show
```

7. Modify the cluster broadcast domain to include the desired cluster ports:

```
network port broadcast-domain remove-ports -broadcast-domain  
broadcast_domain_name -ports port_names  
  
network port broadcast-domain add-ports -broadcast-domain Cluster -ports  
port_names
```



Beginning with ONTAP 9.8, new IPspaces and one or more broadcast domains might be designated to existing physical ports that are intended for cluster connectivity.

8. Modify the cluster IPspace to include the desired cluster ports and set the maximum transmission unit to 9000 if not already set:

```
network port modify -node node_name -port port_name -mtu 9000 -ipspace Cluster
```

9. Display all cluster network LIFs:

```
network interface show -role cluster
```

10. Migrate all cluster network LIFs on both nodes to their home ports:

```
network interface migrate -vserver vservice_name -lif lif_name -destination  
-node node_name -destination-port port_name
```

11. Display all cluster network LIFs:

```
network interface show -role cluster
```

12. Verify the home ports for the cluster network LIFs:

```
network interface modify -vserver vservice_name -lif lif_name -home-port  
port_name
```

13. Migrate all data LIFs to node3:

```
network interface migrate -vserver vservice_name -lif lif_name -destination  
-node node_name -destination-port port_name
```

14. Display all data network LIFs:

```
network interface show -role data
```

15. Configure the home node and home port for all data LIFs. If any LIFs are down, set the administrative status of the LIFs to `up` by entering the following command, once for each LIF:

```
network interface modify -vserver vservice_name -lif lif_name -home-node
```

```
node_name -home-port port_name -status-admin up
```

16. Migrate the cluster management LIF:

```
network interface migrate -vserver vservice_name -lif cluster_mgmt -destination  
-node node3 -destination-port port_name
```

17. Display the status of the cluster management LIF:

```
network interface show cluster_mgmt
```

18. Display the status of all data aggregates in the cluster:

```
storage aggregate show
```

19. Enable the HA pair, storage failover, and auto-giveback:

```
cluster ha modify -configured true
```

20. Migrate data aggregates owned by node4 to node3:

```
storage aggregate relocation start -aggregate aggregate_name -node node4  
-destination node3
```

21. Display the status of all data aggregates in the cluster:

```
storage aggregate show
```

22. Enable auto-revert of the network LIFs across the nodes:

```
network interface modify -lif * -auto-revert true
```

23. Enable storage failover automatic giveback:

```
storage failover modify -node * -auto-giveback true
```

24. Display the cluster status:

```
cluster show
```

25. Display failover eligibility:

```
storage failover show
```



In the cluster report output, a node might incorrectly own aggregates that belong to another node. If this occurs, normalize by performing a takeover and giveback from both sides of the cluster.

26. Display the status of all data aggregates in the cluster:

```
storage aggregate show
```

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.